

Managing MC/ServiceGuard Extension for SAP R/3



Manufacturing Part Number: B7885-90011

June 2001

Legal Notices

The information contained in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Copyright © 2000-2001 Hewlett-Packard Company.

This document contains information which is protected by copyright. All rights are reserved. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws.

Corporate Offices:

*Hewlett-Packard Co.
3000 Hanover St.
Palo Alto, CA 94304*

Use, duplication or disclosure by the U.S. Government Department of Defense is subject to restrictions as set forth in paragraph (b)(3)(ii) of the Rights in Technical Data and Software clause in FAR 52.227-7013.

Rights for non-DOD U.S. Government Departments and Agencies are as set forth in FAR 52.227-19(c)(1,2).

Use of this manual and flexible disc(s), compact disc(s), or tape cartridge(s) supplied for this pack is restricted to this product only. Additional copies of the programs may be made for security and back-up purposes only. Resale of the programs in their present form or with alterations, is expressly prohibited.

A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

EMC², Symmetrix, and SRDF are trademarks of EMC Corporation. Windows NT is a trademark of Microsoft Corporation. SAP and R/3 are trademarks of SAP AG.

Contents

1. Understanding the MC/ServiceGuard SAP R/3 Integration

Overview of SGeSAP	12
One-Package Configuration Model	12
Two-Package Configuration Model	13
Application Servers	15
SAP Application Server packages	17
Highly Available NFS	19
Networking	20
Directory Structures	21
Common Directories that are Kept Locally	21
Directories that Reside on Shared Disks	23
Directories that are Handled by the Automounter	23
SAP R/3 Local Executables	25
MC/ServiceGuard Package Concepts	26
Extension Files for the One Package Model	28
Extension Files For the Two-Package Model	29
SGeSAP and HP Somersault	31
HP Somersault Machine Sets	32
Server Consolidation	33
Dialog Instance Handling	33
Shared Memory Resource Handling	34
SGeSAP Product Files	35
SGeSAP and SCM High Availability- APO/LiveCache Configuration	37

Contents

2. Step by Step Installation Guide

Planning the LVM Setup	43
Preparing the Integration	46
Steps if SAP R/3 is Already Installed	46
Steps if SAP R/3 is Not Installed Yet	48
HP-UX Configuration.	52
Cluster Node Synchronization	53
Cluster Node Configuration	59
External Application Server Host Configuration	62
SGeSAP Files Configuration	63
Standard Parameters of Release 3.x	64
Additional Parameters in sap.conf.	67
Using Database Reconnect and Transaction Reset.	70
Advanced Options of sap.conf	74
SGeSAP configuration for SCM High Availability - APO/LiveCache	77
Service configuration	79
HP Somersault Integration with SGeSAP.	83
MC/ServiceGuard Configuration.	84
Automounter Configuration	90
Database Configuration.	93
Additional Steps for Informix	93
Additional Steps for Oracle	96
SAP R/3 System Configuration	102

Contents

3. SGeSAP Administration

Switching SGeSAP Off and On	108
Switch Off SGeSAP	108
Switch On SGeSAP	113
SGeSAP Administration Issues	116
Hardware Aspects	116
HP-UX Software Aspects	117
SAP R/3 Administration Aspects	119
MC/ServiceGuard Administration Aspects	121
SGeSAP Administration Aspects	121
SGeSAP Administration With HP Somersault Enabled	123

Contents

Printing History

Table 1

Printing Date	Part Number	Edition	SAP R/3 Release	Operating System Releases
June 2000	B7885-90004	Edition 1	SAP R/3 3.0D and up	HP-UX 10.20 HP-UX 11.00
March 2001	B7885-90009	Edition 2	SAP R/3 3.03D and up	HP-UX 10.20 HP-UX 11.00 HP-UX 11i
June 2001	B7885-90011	Edition 3	SAP R/3 3.04 and up	HP-UX 10.20 HP-UX 11.00 HP-UX 11i

The printing date and part number indicate the current edition. The printing date changes when a new edition is printed. (Minor corrections and updates which are incorporated at reprint do not cause the date to change.) The part number changes when extensive technical changes are incorporated.

New editions of this manual will incorporate all material updated since the previous edition.

HP Printing Division:

*Business Critical Computing Business Unit
Hewlett-Packard Co.
19111 Pruneridge Ave.
Cupertino, CA 95014*

About this Manual...

This document describes how to configure and install highly available R/3 systems on HP-UX 10.x and 11.x using MC/ServiceGuard. It refers to the HP product B7885BA – MC/SG Extension for SAP. To get further information regarding SAP and High Availability in general, please refer

to the SAP documents:

- *SAP R/3 in Switchover Environments*
- *SAP BC High Availability*

To understand this document you have to have knowledge of the basic MC/ServiceGuard concepts and commands. Experience in the Basis Components of SAP R/3 will also be helpful.

This manual consists of three chapters:

- Chapter 1, “Understanding the MC/ServiceGuard SAP R/3 Integration,” describes MC/ServiceGuard Extension for SAP R/3 (SGeSAP) concepts. It also explains the current approach and why it was taken. It can be used like a FAQ list. If you have technical questions about the installation, especially during the scoping phase, refer to this section to get the information you need.
- Chapter 2, “Step by Step Installation Guide,” describes the installation of SGeSAP step-by-step down to the HP-UX command level.
- Chapter 3, “SGeSAP Administration,” contains a summary of how to switch off SGeSAP. This is useful for very complex R/3 reconfiguration tasks like *putting* the R/3-System.

Table 2 lists abbreviations used in this manual:

Table 2

Abbreviations

Abbreviation	Meaning
<SID>, <sid>	System ID of the SAP R/3 system in uppercase/lowercase
<INSTNR>	instance number of the SAP R/3 system
<pri_host>, <sec_host>, <local_host>	names mapped to the local hostnames (output of <code>uname -n</code>)
<primary>, <secondary>, <local>	names mapped to local IP addresses of the client LAN
<relocdb>, <reloci>, <relocdbci>	names mapped to relocatable IP addresses of MC/SG packages in the client LAN

Table 2**Abbreviations**

Abbreviation	Meaning
<primary_s>,<secondary_s>, <local_s>	names mapped to local IP addresses of the server LAN
<relocdb_s>,<reloci_s>, <relocdbci_s>	names mapped to relocatable IP addresses of MC/ServiceGuard packages in the server LAN

Related Documentation

The following documents contain additional related information:

- *Managing Systems and Workgroups* (Part Number B2355-90157)
- *Managing MC/ServiceGuard* (Part Number B3936-90026) for HP-UX 10.20
- *Managing MC/ServiceGuard* (Part Number B3936-90045) for HP-UX 11.0 and 11i
- *Managing Highly Available NFS* (Part Number B5125-90001)
- *ServiceGuard Extension for SAP Release Notes* (For your version of the product)

1

Understanding the MC/ServiceGuard SAP R/3 Integration

HP MC/ServiceGuard Extension for SAP R/3 (SGeSAP) extends MC/ServiceGuard's powerful failover capabilities to SAP R/3 environments. It continuously monitors the health of each SAP R/3 node and automatically responds to failures or threshold violations. It can also minimize planned downtime when performing SAP R/3 upgrades. MC/ServiceGuard protects the SAP R/3 central instance and database by defining them in MC/ServiceGuard packages.

This chapter provides an overview of how SGeSAP works. Topics are as follows:

- Overview of SGeSAP
- Application Servers
- Highly Available NFS
- Networking
- Directory Structures
- SAP R/3 Local Executables
- MC/ServiceGuard Package Concepts
- SGeSAP and HP Somersault
- Server Consolidation
- SGeSAP Product Files

Overview of SGeSAP

SAP R/3 allows a great amount of flexibility in setup and configuration. The SGeSAP Integration Scripts preserve much of this flexibility through the use of two integration models:

- One Package Configuration Model
- Two Package Configuration Model

One-Package Configuration Model

In a one-package configuration, both the database (DB) and central instance (CI) run on the same node at all times and are configured in one MC/ServiceGuard package. Other nodes in the MC/ServiceGuard cluster function as backups for the primary node (on which the system runs during normal operation).

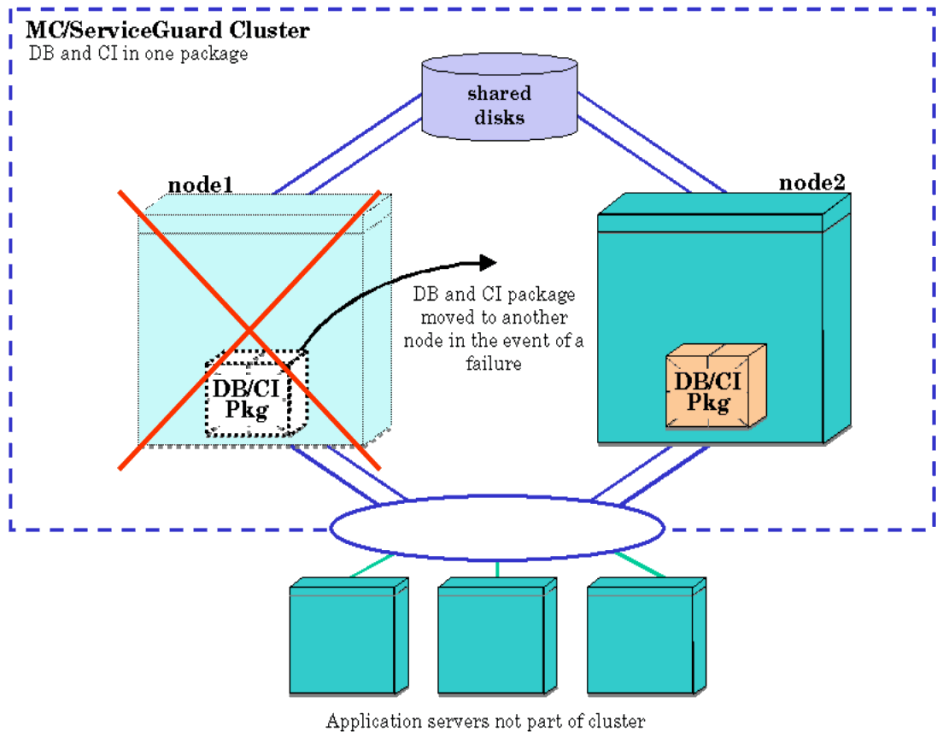
If the primary node fails, the database and the central instance fail over and continue functioning on an adoptive node. The process of failover results in downtime that can last several minutes, depending on the work in progress when the failover takes place. The main portion of this downtime is needed for the recovery of the database.

After failover, the system runs without any manual intervention needed. The application servers are not part of the cluster but can either stay up or be restarted during failover.

A sample configuration in Figure 1-1 shows *node1* with a failure, which causes the package containing the database and central instance to fail over to *node2*.

Figure 1-1

One-Package Failover Scenario



Use the one-package model for all configurations where you can put the database and central instance on one node and you have available an equal sized node as a backup. During normal operation, the backup node can be used as an idle standby, application server, or test system.

Two-Package Configuration Model

If you are planning to distribute the database and central instance between two nodes, use the two-package model. The SAP R/3 functionality is separated into two MC/ServiceGuard packages, one for the database (DB) and the other for the SAP R/3 central instance (CI). The database package contains the filesystems for the NFS mount points.

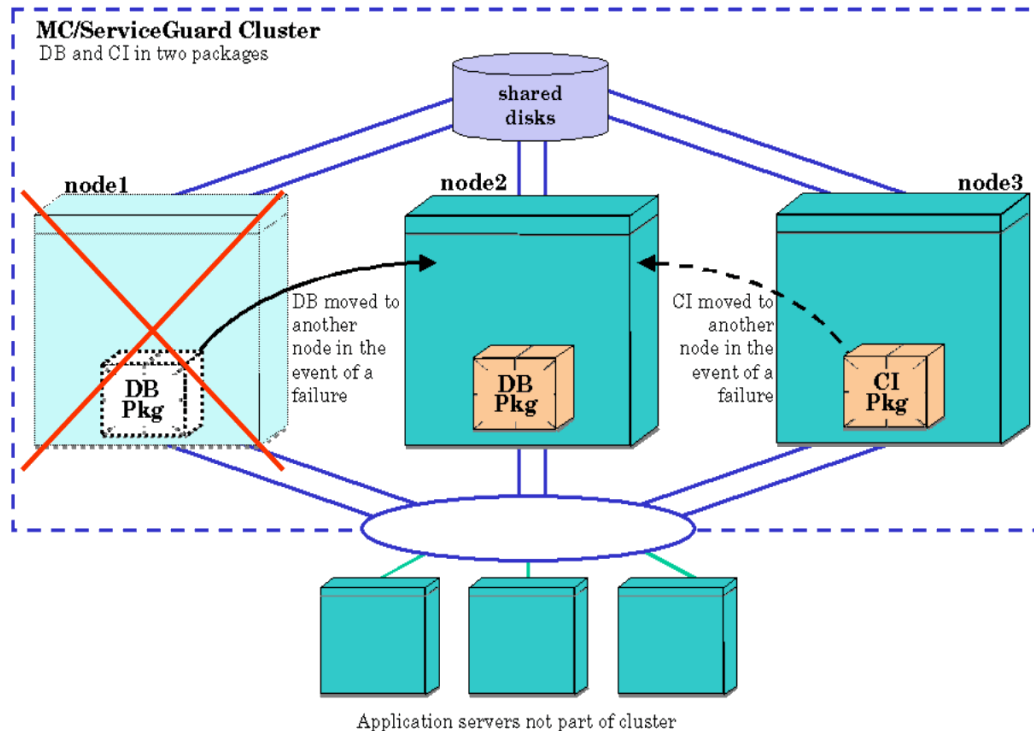
The cluster can be configured so that the two nodes back up each other, or so that one or more dedicated hosts back up the nodes running the SAP R/3 packages as illustrated in Figure 1-2.

Understanding the MC/ServiceGuard SAP R/3 Integration Overview of SGeSAP

If either *node1* or *node2* fails, the package can fail over to *node3*. If a failover to a node that is not idle takes place (for example if *node3* were used as an application server), SGeSAP, if requested, can bring down running instances to free the resources needed to get the system up.

Under normal conditions, all backup hosts can be used to run application servers or instances of different test or development systems, or they can be idle. If needed, additional application servers inside and outside the cluster can be restarted automatically. It is possible to define more than one highly available SAP R/3 system in one cluster.

Figure 1-2 Two-Package Failover Scenario with Third Node as Standby



Application Servers

The database and the SAP R/3 central instance are always running on nodes that are protected by MC/ServiceGuard. This is different for systems that have the enqueue service, SAP R/3's central lock handling server process, protected by HP Somersault (B8438). Refer to "SGeSAP and HP Somersault" on page 31 for additional information.

Other SAP R/3 services can be run on HP-UX application server hosts. These hosts do not need to be part of the MC/ServiceGuard cluster. Even if the additional SAP R/3 services are run on nodes in the MC/ServiceGuard cluster, they are not protected by MC/ServiceGuard. A combination of Windows NT/HP-UX application servers is technically possible but additional software is required to access HP-UX filesystems or HP-UX-like remote shells from the Windows NT system.

All application servers different from the central instance are called *additional application servers*. An additional application server that runs on a cluster node is called an *internal application server*. External application servers run on HP-UX- or Windows NT-hosts that are not part of the cluster. Even if application servers are external to the cluster, they are affected by package startup and shutdown.

Run the SAP R/3 services dialog, update, batch, spool, and gateway, on additional application servers. Do not run the message or enqueue services on additional application servers. Set up one message server and one enqueue server for each SAP R/3 system.

The message server, enqueue server, database and NFS server are all single points of failure (SPOF). To maintain high availability, all these SPOFs for the SAP R/3 system must be eliminated by configuring them in MC/ServiceGuard cluster nodes.

In standard SAP R/3 scenarios the SPOFs database, NFS, message and enqueue server are all protected. It is highly recommended that you also protect at least one instance of each additional SAP R/3 service in the MC/ServiceGuard cluster.

Understanding the MC/ServiceGuard SAP R/3 Integration

Application Servers

For all application server nodes outside the MC/ServiceGuard cluster, consider the following for each of the SAP R/3 services:

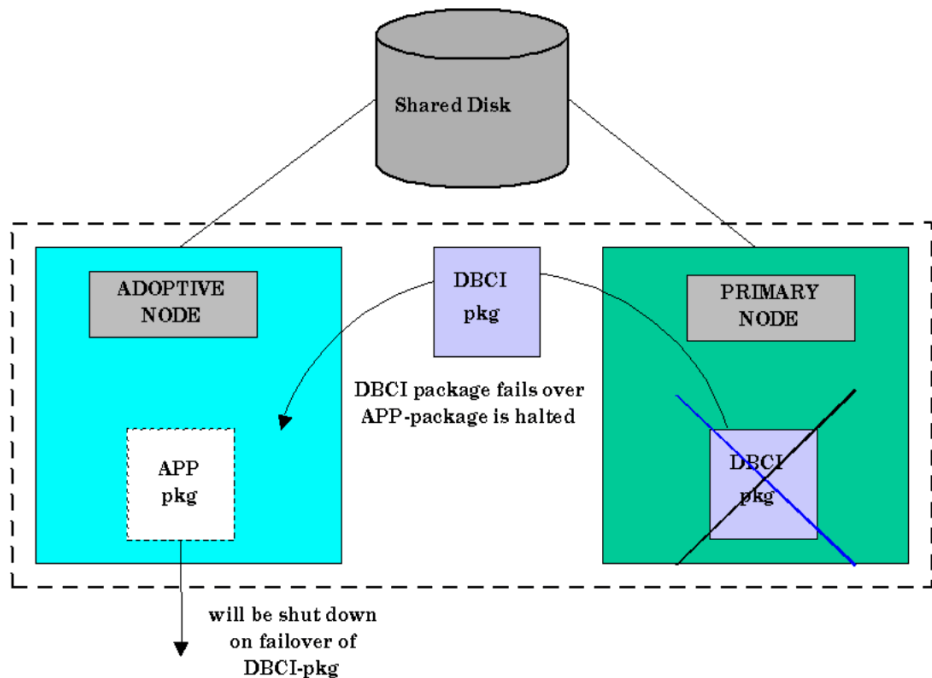
- **Dialog**—Logon using `saplogon` for an application server groups instead of `saptemu/sapgui` for each individual application server. When logging on to an application server group with two or more application servers, you not need a different login procedure even if one of the application servers of the group fails. Also, using login groups provides workload balancing between application servers.
- **Update**—Configure an update service on a node only for application services running on the same node. This ensures that the remaining SAP R/3 servers, running on different nodes, are affected if an outage occurs on the update server. However, if the update server is configured to be responsible for application servers running on different nodes, any failure of the update server will lead to subsequent outages at these nodes. Configuring the update service on the high available central instance is recommended. Consider using local update servers only if performance issues require it.
- **Batch**—Do not specify a *destination host* when defining a batch job. This allows the batch scheduler to choose a batch server that is available at the start time of the batch job. If you must specify a destination host, specify the batch server running on the highly available central instance.
- **Spool**—Print requests stay in the system until the node is available again and the spool server has been restarted. These requests could be moved manually to other spool servers if one spool server is unavailable for a long period of time. An alternative is to print all time critical documents through the highly available spool server of the central instance.
- **Gateway**—You don't have to worry about the gateway processes `gwr` and `gwr` with SAP R/3 3.0 or higher. Do not confuse this with the SNA gateway.

SGeSAP monitors the hardware within the cluster (CPU, Network cards, etc), but it does not monitor the health of the SAP R/3 processes or the database. This means that MC/ServiceGuard will not switch the packages to another node if you have problems with SAP R/3 or the database.

SAP Application Server packages

Beginning with SGeSAP Version 3.03 it is possible to configure SAP Application Server within a MC/ServiceGuard package. For example, Figure 1-3 illustrates that in a 2 node cluster it is common to configure the adoptive node to run as an additional application server.

Figure 1-3 Adoptive Node as Application Server



Whenever a failover of the DBCI-package happens this application server package will be shut down to free up resources for the DBCI-package. When the failed primary node is afterwards repaired the application server package can be easily brought up on this repaired node. The advantage is that the DBCI package can remain running on the former adoptive node without additional downtime for the application.

An additional example for the need of this SAP Application Server

Understanding the MC/ServiceGuard SAP R/3 Integration

SAP Application Server packages

package is to protect dedicated batch servers against hardware failures.

Refer to Chapter 2, page 59 to get detailed information how to configure Application Server packages.

Highly Available NFS

The SAP R/3 system requires the Network File System (NFS) to be highly available. Otherwise the NFS is one of the four single points of failure (SPOF) in the SAP R/3 environment. SGeSAP works with HP's Highly Available NFS product, the NFS Toolkit (B5140BA). This product is a separately purchasable product with a set of configuration and control files that can and must be customized for the SGeSAP environment. SGeSAP uses the NFS-Toolkit interface to make the NFS server in an SAP R/3 environment highly available.

Reminder when configuring HA NFS for HP-UX 11.i systems:

The default "NFS" transport mode in HP-UX 11.i is "TCP". The MC/ServiceGuard NFS Toolkit for HP-UX 11.i is currently supported only for the "UDP" transport mode. Therefore it is necessary to specify the "UDP" transport mode using the appropriate mount option "-o proto=udp" in the auto.direct configuration file for the automounter. See the <automount> manpage and the MC/ServiceGuard NFS Toolkit release notes for detailed information on this.

Networking

Communication with the highly available components takes place by using special IP addresses, called *relocatable IP addresses*. They are handled by MC/ServiceGuard. Think of them as addresses of a package rather than host machine addresses.

Some of the SAP R/3 processes require that hostnames are not longer than eight characters and do not contain special characters.

One additional IP address is the minimum address requirement for the one package concept, two IP addresses are needed for the two package solution (one for each package). If you are planning to use a frontend LAN and a server LAN as suggested by SAP R/3, double the number of IP addresses, that is, use two or four.

You can reduce the number of additional IP addresses with the two package concept if you do not want to access the database directly through the frontend LAN. You can save another IP address, if you do not have frontends that directly connect to the central instance, but only access your SAP R/3 system through different application servers.

The reason for relocatable IP addresses on both networks is, that a switchover must remain transparent to both the frontends and the additional application servers.

The complete system is installed if there is no server LAN at all, that is, all profiles only refer to addresses of the frontend LAN. Static routing of the relocatable addresses from the client LAN to the server LAN redirect the communication traffic of the application servers and the database to the backend LAN.

Directory Structures

This section provides a deeper understanding of the directory layout. For information about LVM setup, refer to the section, “Planning the LVM Setup” on page 43.

Depending upon your particular setup, the following important groups of directories need special treatment:

- Common Directories that are Kept Locally
- Directories that Reside on Shared Disks
- Directories that are Handled by the Automounter

Common Directories that are Kept Locally

Use a standard setup for the following common directories and their files that are kept locally on any host of the cluster:

- `/etc/cmcluster`—the directory in which MC/ServiceGuard keeps its configuration files and the runtime directory
- `/home/<SID>adm`—the home directory of the SAP R/3 system Administrator
- `/oracle/<SID>`—the directory in which the ORACLE client software resides

If you install SAP R/3 release 4.0A or higher with ORACLE database, you have this additional local directory on each internal application server.

The contents of this group of directories must be synchronized manually on all hosts of the cluster. Do not automatically make `/home/<SID>adm` the same on all of the hosts. For example, it is possible to install an additional application server on a host of the cluster that is not be part of any package. If it is local to its host, the SAP R/3 startup scripts are only needed on its dedicated host. You do not need to distribute them to other hosts.

Change the SAP R/3 startup scripts `startsap_<local_hostname>_<id>` individually in the home directory of `<SID>adm` on all cluster nodes, to configure different startup profiles. The central instance can be configured differently, depending on which node it is actually started. Be

Understanding the MC/ServiceGuard SAP R/3 Integration Directory Structures

careful with this option. You have to make sure by yourself that the central instance is capable of doing its work in any of the possible configurations.

The standard HP-UX configuration files are local. To prevent malfunction, never delete the mutual `.rhosts` entries of the root user and `<SID>adm` on any of the nodes.

The SAP R/3 option to introduce a local copy of the `ORACLE_HOME` directory `/oracle/<SID>` allows you to use older versions of the ORACLE client software in combination with a more recent ORACLE server release. This provides more flexibility and is useful if the operating system releases are supported by a different client than the server component.

In switchover environments only external application servers can take advantage of this. The server component has to be able to run on all nodes. You cannot use an operating system release with a server component that is not supported. This is true for any host within the cluster. Within a cluster, the client and server components must have the same release number.

This is important, because `/oracle/<SID>` also acts as a mountpoint for the directories of the server component. The client software residing on a local disk is no longer accessible after a mount takes place. This happens, if the database switches to a host on which an application server is running. The whole process remains transparent to the running application server because the database server directories deliver the same files at the same places. Every file is where the application server expects it. Already opened files can still be accessed, even though they are no longer visible in the directory tree. Please note, that the application server will only have read access. No writing will occur.

NOTE

Because `/oracle/<SID>` is needed as a free mountpoint by the package, it must not be a mountpoint for a logical volume containing the client ORACLE files. Alternatives include using `/oracle` as a mountpoint for the local filesystem or using a link. Using crossmounting with `/oracle/<SID>` impacts performance even during normal operation. Crossmounting is not needed since only read access occurs on the client side.

Directories that Reside on Shared Disks

Changing shared disk directory files on any host of the cluster affects the whole cluster. Files in the following directories and all subdirectories are normally shared:

- /usr/sap/<SID>/DVEBMGS<instance_id>
- /export/usr/sap/trans
- /export/sapmnt/<SID>
- /export/informix or /oracle/<SID>

The files in these directories are only available on a host if the package they belong to is running on it. MC/ServiceGuard switches them to another node with the package. Under a two package concept the three /export and /oracle directories belong to the database package. The /usr directories belong to the central instance package.

Never use /usr/sap/<SID> as mountpoint on a shared disk. You risk making the Instance Directory of an internal application server unreachable if a switchover takes place.

All filesystems mounted below /export are part of the crossmounting. Please note that /oracle/<SID> is treated differently for performance reasons.

Directories that are Handled by the Automounter

Directories handled by the automounter are mounted automatically as needed. This is true not only for the nodes of the cluster, if you use external application servers, they will also need them. Automounter directories are:

- /sapmnt/<SID>
- /usr/sap/trans
- /informix

These directories are NFS-mounted from their equivalents inside of the /export directory of the nodes which run the packages. The automounter uses the relocatable IP addresses. This ensures that the directories are quickly available again after a switchover.

There are two important issues concerning these directories:

Understanding the MC/ServiceGuard SAP R/3 Integration Directory Structures

- The directories below `/export` are exported without root permissions.

This happens according to the recommendations of SAP R/3 and enhances the security of the installation. The effect is, that the root user cannot modify these directories or their contents. With standard permissions set, the root user cannot even see the files. This is not an error and the system runs without problems.

If you want to modify anything as root, use the equivalent directory below `/export` on the host the package runs on.

- If the database package is halted, you cannot log in as `<SID>adm` unless you keep the binaries local.

The reason for this is, that `/usr/sap/<SID>/SYS/exe` is part of the path of `<SID>adm`. Without local binaries, this directory links to `/sapmnt/<SID>` which is handled by the automounter. The automounter cannot contact the host belonging to the relocatable address that is configured because the package is down. The system will hang. To prevent this, always keep local copies of the executables.

On a HP-UX system, all SAP R/3 executables should be kept locally on all nodes. This avoids problems during a failover. You have to setup `sapcpe` which allows SAP R/3 to distribute new executables after an SAP R/3 upgrade.

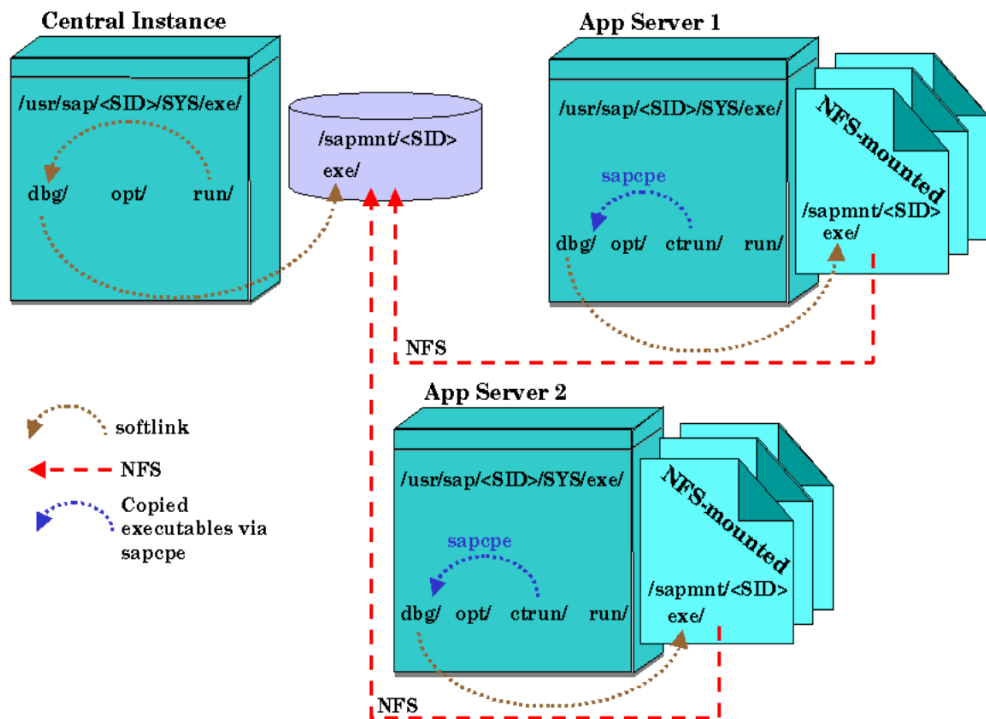
SAP R/3 Local Executables

In clustered SAP R/3 system environments it is required to install local executables. For example, starting SAP R/3 instances from local executables is quicker than starting them from centrally stored executables.

To automatically update the executables SAP R/3 includes the `sapepe` mechanism. With every startup of an SAP R/3 system, `sapepe` matches new executables stored centrally with those stored locally.

Figure 1-4 shows a filesystem layout with `sapepe`.

Figure 1-4 `sapepe` Mechanism for Executables



MC/ServiceGuard Package Concepts

Specific files are required to build MC/ServiceGuard packages. The files required apply to both the SGeSAP one package model and two package model. Figure 1-5 illustrates the required configuration files, and the questions they answer, for an SGeSAP application package in MC/ServiceGuard.

Figure 1-5 Configuration Files Needed to Build Cluster

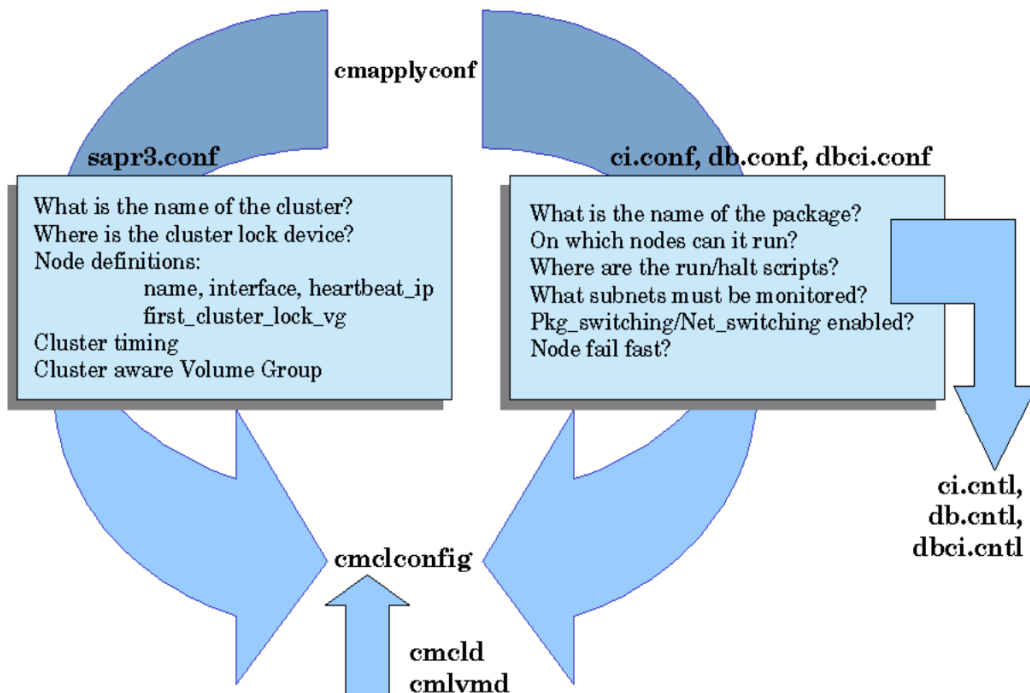
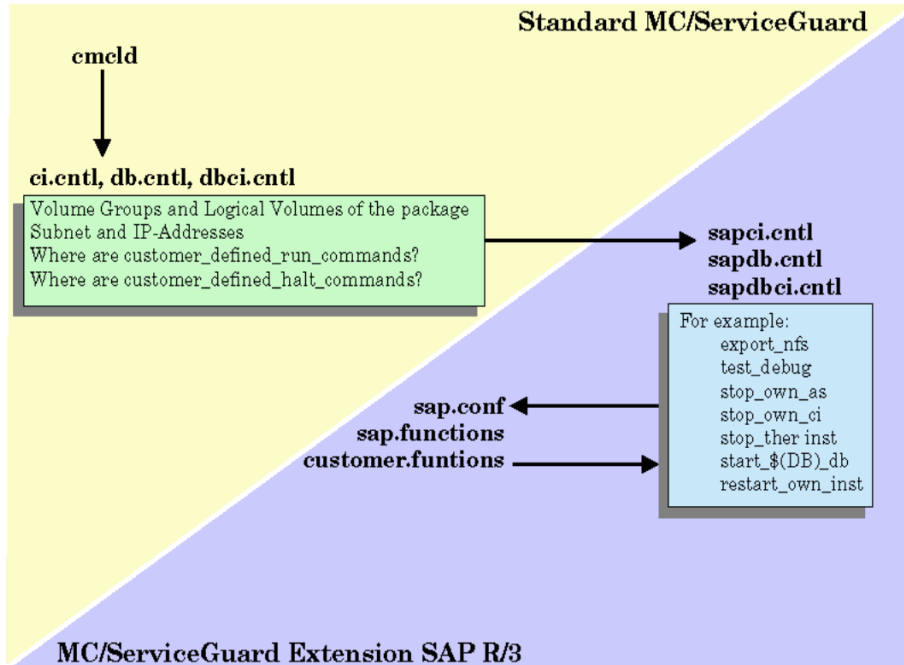


Figure 1-6 illustrates the required control files for an SGeSAP application package in MC/ServiceGuard.

Figure 1-6 Control Files Containing Runtime Logic



Extension Files for the One Package Model

In the one package model, the SAP R/3 functionality—database and central instance—along with the highly available NFS server are all been placed in one MC/ServiceGuard package.

The SGeSAP scripts of the one package concept are organized as follows:

- `/etc/cmcluster/sapr3.conf`—cluster configuration file of MC/ServiceGuard.
- `/etc/cmcluster/<SID>/dbci.conf`—package configuration file based on the NFS toolkit configuration file.
- `/etc/cmcluster/<SID>/dbci.cntl`—package control file based on the NFS toolkit configuration file. It must call `<SID>/sapdbci.cntl` in the `customer_defined_[run|halt]_cmds` section.
- `/etc/cmcluster/<SID>/sap.conf`—contains the actual configuration and options of the SGeSAP Integration Scripts. It needs to be customized.
- `/etc/cmcluster/sap.functions`—contains all standard runtime logic used by one package and two package configurations. Do not modify this file!
- `/etc/cmcluster/customer.functions`—contains all functions which can be modified for special setups. This file may also contain your own additional functions. If you have to change functions in `/etc/cmcluster/sap.functions`, copy them over to `/etc/cmcluster/customer.functions` and modify them there. Never modify `/etc/cmcluster/sap.functions` itself. Any function provided in `customer.functions` overrides its equivalent in `sap.functions`.
- `/etc/cmcluster/<SID>/sapdbci.cntl`—contains database and SAP R/3 specific control logic. This file is valid for one package configurations only. Sources include `<SID>/sap.conf`, `customer.functions`, and `sap.functions`.

Use the one pkg concept for all configurations where you can put the database and central instance on one node and have an equally sized node as a backup. During normal operation, backup nodes can be used as an idle standby, application server host, or test system.

If you are planning to distribute the database and central instance on two nodes in the near future apply the two package concept.

Extension Files For the Two-Package Model

In the two package model, the SAP R/3 functionality is separated into two MC/ServiceGuard packages. One for the database (DB) and the other for the SAP R/3 central instance (CI). The database package contains the filesystems for the NFS mount points. The SGeSAP scripts of the two package concept are organized as follows:

- `/etc/cmcluster/sapr3.conf`—cluster configuration file of MC/ServiceGuard.
- `/etc/cmcluster/db.conf`—database package configuration file based on the NFS toolkit configuration file. The provided file is for reference only, create your own.
- `/etc/cmcluster/ci.conf`—central instance package configuration file based on the NFS toolkit configuration file. The provided file is for reference only, create your own.
- `/etc/cmcluster/db.cntl`—database package control file based on the NFS toolkit configuration file. It must call `<SID>/sapdb.cntl` in the `customer_defined_[run|halt]_cmds` sections.
- `/etc/cmcluster/ci.cntl`—central instance package control file based on the NFS toolkit configuration file. It must call `<SID>/sapdbci.cntl` in the `customer_defined_[run|halt]_cmds` sections.
- `/etc/cmcluster/<SID>/sap.conf`—contains the actual configuration and options of the Integration Scripts. It needs to be customized.
- `/etc/cmcluster/sap.functions`—contains all standard runtime logic used by one package and two package configurations. Do not modify this file!
- `/etc/cmcluster/customer.functions`—contains all functions which can be modified for special setups. This file may also contain your own additional functions. If you have to change functions in `/etc/cmcluster/sap.functions` copy them over to `/etc/cmcluster/customer.functions` and modify them there. Never modify `/etc/cmcluster/sap.functions` itself. Any function provided in `customer.functions` overrides its equivalent

Understanding the MC/ServiceGuard SAP R/3 Integration

MC/ServiceGuard Package Concepts

in `sap.functions`.

- `/etc/cmcluster/<SID>/sapci.cntl`—contains SAP R/3 specific control logic. This file is valid for a two package configuration only. Sources include `<SID>/sap.conf`, `customer.functions`, and `sap.functions`.
- `/etc/cmcluster/<SID>/sapdb.cntl`—contains database control logic. This file is valid for a two package configuration only. Sources include `<SID>/sap.conf`, `customer.functions`, and `sap.functions`.

Extension Files for an application server package

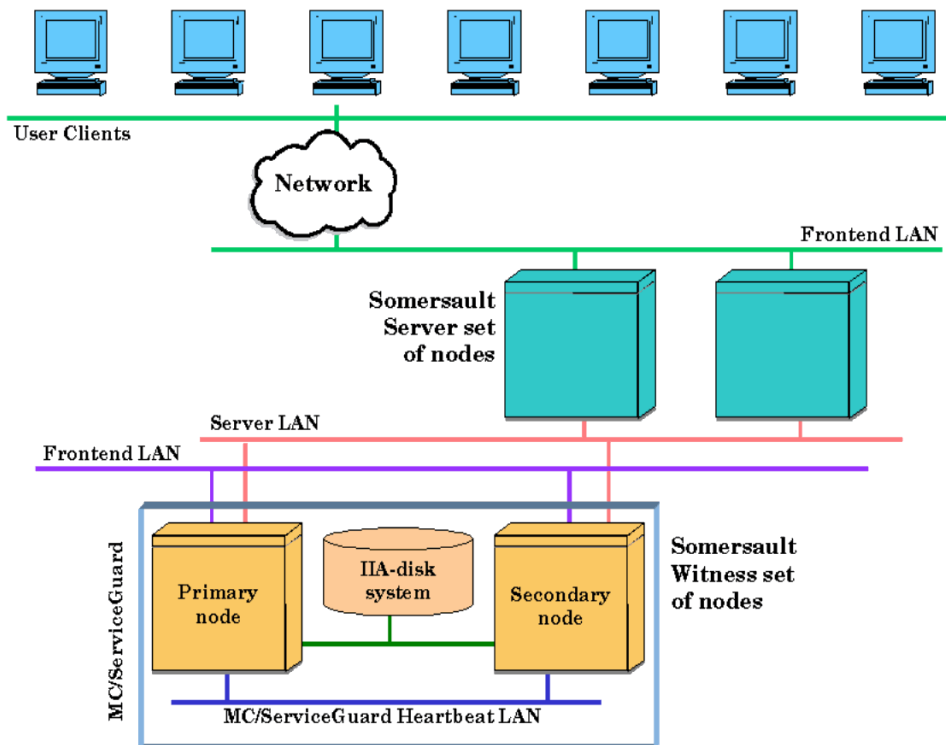
The creation of an Application server package only makes sense if a one-package or two-package model has already been implemented. The needed Application server files are organized as follows:

- `/etc/cmcluster/<SID><INSTNR>.conf`—application server package configuration file based on the NFS toolkit configuration file.
- `/etc/cmcluster/<SID>/ap<SID><INSTNR>.cntl`—application server package control file based on the NFS toolkit control file. It must call `/etc/cmcluster/<SID>/sapap.cntl` in the `customer_defined_[run | halt]_cmds` section.
- `/etc/cmcluster/<SID>/sapap.cntl`— contains SAP Application server specific control logic. This file applies to the an Application server package only and must not be edited.

SGeSAP and HP Somersault

HP Somersault (B8438) is the latest technology to guarantee a failsafe operation of the enqueue server, one of the four SPOF's in an SAP R/3 system. Figure 1-7 shows a sample HP Somersault system environment.

Figure 1-7 HP Somersault System Environment



SAP R/3 supports the usage of HP Somersault for protecting the enqueue server only in complete High Availability scenario with MC/ServiceGuard and SGeSAP eliminating the other SPOFs.

MC/ServiceGuard and SGeSAP protect against the loss of service due to database and message server failure, and HP Somersault protects against Enqueue failure.

The enqueue server is the central process that provides access to the Enqueue lock table in an SAP R/3 system. HP Somersault protects the Enqueue service by introducing the HP Somersault Enqueue process (`ensserv`) that replaces the standard SAP R/3 Enqueue service. An HP Somersault unit consists of three members:

- primary
- secondary
- witness

The primary and secondary members host the HP Somersault enqueue process (`ensserv`). The witness acts as an arbitrator to the unit in the event of a failure.

HP Somersault Machine Sets

An HP Somersault machine set is a group of machines that have been designated to fulfill a specific HP Somersault role. The two kinds of machine sets that are of interest when integrating HP Somersault with MC/ServiceGuard are:

- Server Set Machines

The server machine set is stretched over nodes outside the MC/ServiceGuard cluster only. These machines are able to host the primary and secondary members of the HP Somersault unit.

- Witness Set Machines

The witness machine set consists of the systems that host the witness member of the HP Somersault unit. All machines in the MC/ServiceGuard environment are members of the witness machine set. This means that every cluster node must have the HP Somersault installed and configured.

In an MC/ServiceGuard environment the handling of the HP Somersault unit is transparent to the administrator. HP Somersault starts and stops with the startup and shutdown of the Central Instance out of a MC/ServiceGuard package. Refer to Chapter 3, “SGeSAP Administration,” for additional information. For supported configurations and further details on HP Somersault refer to your HP Somersault documentation.

Server Consolidation

Server consolidation in an SAP R/3 environment is a strategy to put multiple SAP R/3 systems on a single physical server, a single host system. Since HP introduced HP-UX 11.00 with its 64-bit technology, addressing memory is no longer a limitation. The availability of 64-bit compiled SAP R/3 versions now enables a full blown 64-bit solution for SAP R/3. This results in having either multiple SAP R/3 clients running on a single SAP R/3 server or having multiple 2-tier SAP R/3 systems running on a single host. Additional SAP R/3 dialog instances can also run on the same server—in both scenarios.

Clustering a consolidated SAP R/3 server is different from clustering a single SAP R/3 instance. SGeSAP Version 3.1 supports dialog instance handling and shared memory resource handling on a consolidation host.

Dialog Instance Handling

With SGeSAP 3.1, SAP R/3 dialog instance treatment during CI / DB / DBCI package startup and stop has changed. SGeSAP now supports dialog instances on a host where CI / DB / DBCI packages are currently started as listed in Table 1-1.

Table 1-1 Dialog Instance Treatment During Startup and Shutdown

ASTREAT	Restart	Final_ Stop	Active	Runpkg DB	Haltpkg DB	Runpkg CI	Haltpkg CI
0	0	0	0	nop	nop	nop	nop
1	0	0	1	CI-start	nop	st_if_nl	nop
2	0	1	0	nop	nop	nop	stop
3	0	1	1	CI-start	nop	st_if_nl	stop
4	1	0	0	stop	nop	stop	nop
5	1	0	1	stop CI-start	nop	stop start	nop
6	1	1	0	stop	nop	stop	stop
7	1	1	1	stop CI-start	nop	stop st_if_nl	stop

Table 1-1 Dialog Instance Treatment During Startup and Shutdown

ASTREAT	Restart	Final_ Stop	Active	Runpkg DB	Haltpkg DB	Runpkg CI	Haltpkg CI
nop = no operation st_if_nl = start if <i>not</i> on local host							

Shared Memory Resource Handling

SGeSAP 3.1 provides shared memory resource cleanup policies as listed in Table 1-2.

Table 1-2 Shared Memory Cleanup Policies and their Effects

CLEANUP_ POLICY	CI Resources	DB Resources	AppServer Resources
lazy	no action	no action	no action
normal	use <code>cleanipc</code> to remove unused own resources	unused Oracle SGA is removed	use <code>cleanipc</code> to remove unused own resources
strict	remove any resource belonging to any SAP R/3 instance	remove any resource belonging to any database	user <code>cleanipc</code> to remove unused own resources

SGeSAP Product Files

All the files that are available after the installation of SGeSAP can be found in `/opt/cmcluster/sap`. The files that are installed are:

- The standard SGeSAP run-time functions and configurable functions repository container in `/opt/cmcluster/sap`:

`sap.functions`—contains standard SAP R/3 functions. Do not edit this file!

`customer.functions`—contains functions that can be modified for special setups.

- The run-time configuration options and control logic in `/opt/cmcluster/sap/SID`:

`sap.conf`—contains configuration and options for SGeSAP

`sapci.cntl`—contains SAP R/3 specific control logic for two package configurations

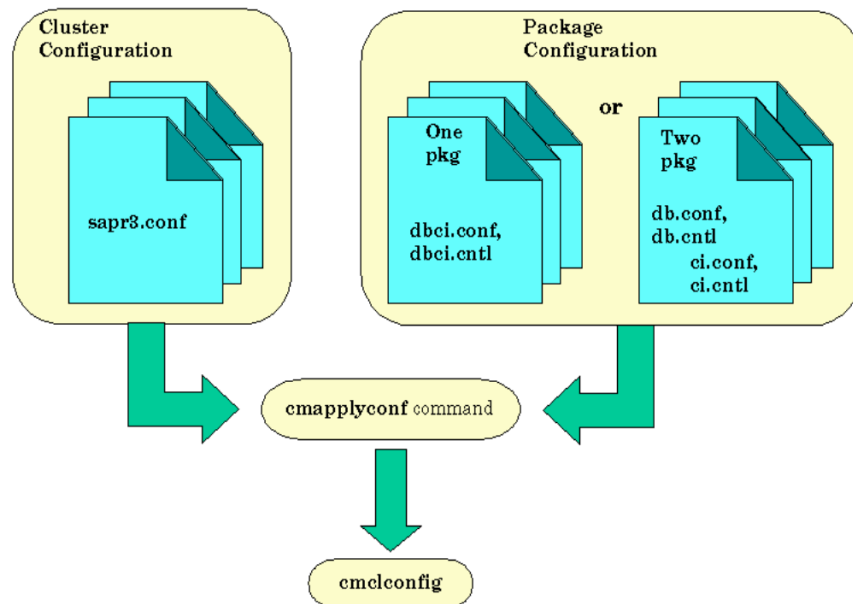
`sapdb.cntl`—contains database specific control logic for two package configurations

`sapdbci.cntl`—contains database and SAP R/3 specific control logic for one package configurations

`sapap.cntl`—contains the Application specific control logic for an application server package

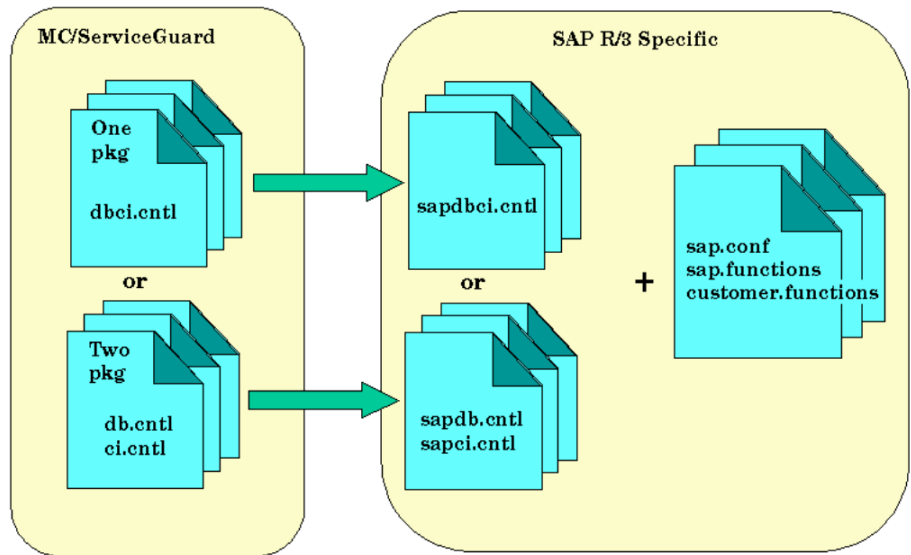
Figure 1-8 shows that the configuration files are used as input for the standard MC/ServiceGuard `cmapplyconf` command to generate the `cmclconfig` file that is used to manage the cluster.

Figure 1-8 Configuration Files



The SGeSAP files work in tandem with standard MC/ServiceGuard files to monitor the health of the cluster and packages, and to move packages to another node in case of failure. Figure 1-9 shows the relationship between the standard MC/ServiceGuard files and the SGeSAP scripts and control logic.

Figure 1-9 Control Files Containing Run-time Logic



SGeSAP and SCM High Availability- APO/LiveCache Configuration

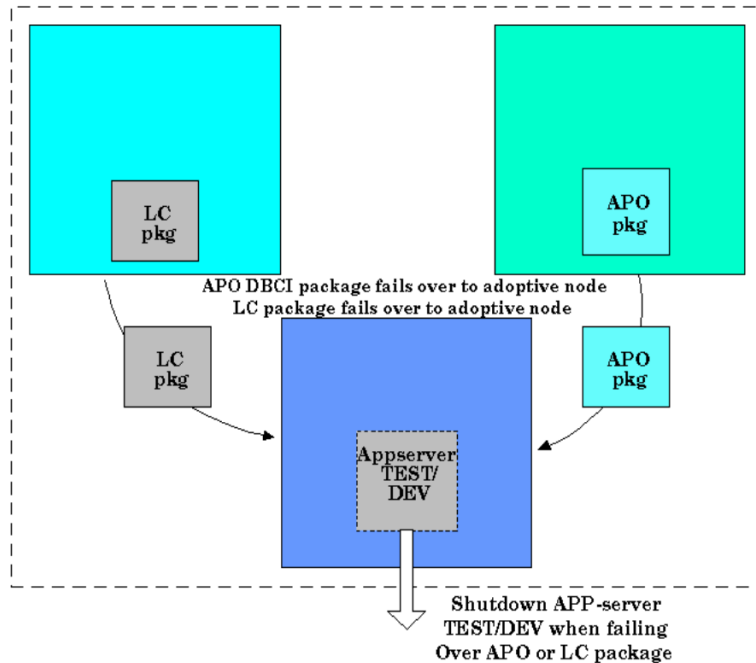
SGeSAP supports customers who choose to run their Supply Chain Management (SCM) systems in a High Available system environment. The SAP SCM environment consists of 2 systems: the APO system and the LiveCache. Considering only database and application server, an APO system looks like an standard R/3 system. It uses the same basic technology like R/3 (kernel, databases, executables) but with different business logic.

Additional components like the liveCache and optimization server further enhance the system. The liveCache server consists of an SAPDB database running in memory. Disk space is required to host the

operating system of the node liveCache is running, pagefiles, executables and dev-spaces for an filesystem image of a consistent DB image.

Figure 1-10 illustrates a sample scenario:

Figure 1-10 3 Node SGeSAP Cluster for APO/LiveCache



The LiveCache package is required to run on a separate server. Therefore the 3-node cluster is an appropriate setup to run a customer SCM system in a high available environment. The adoptive node that can run an additional Application server or a development system under normal condition is able to run either the LiveCache-package or the APO-package in case of a system failure. The additional APO Application-server or a development system will be shut down in case of a failover.

Additional scenarios are a 4 node cluster where both LiveCache primary and APO primary node have a dedicated backup node each. The requirement is that the LiveCache package always runs on a node where no other application like a database or an additional SAP system is running.

Setting up the APO package in SGeSAP is standard as the APO is based on a standard R/3 kernel. For details about the setup of a liveCache package refer to the step by step installation guide in chapter 2 “SGeSAP configuration for SCM High Availability - APO/LiveCache” of this book.

NOTE

For all High Availability configurations of APO and LiveCache it is required to follow SAP’s latest recommendations and requirements. Check all SAP OSS notes and all appropriate documentation about HA scenarios and Recovery for APO / LiveCache before setting up a HA scenario for APO / LiveCache! Configuring a HA LiveCache system requires some additional tasks that differ from a standard R/3 configuration.

The APO system itself can be treated as an additional SAP R/3 system and therefore can be configured in SGeSAP using the standard installation steps. The LiveCache concept described in this documentation is based on the SAP documentation “LiveCache and High Availability on Unix platforms” and “Backup and recovery for APO 3.0”. It is recommended to follow these documents before configuring LiveCache with the SGeSAP.

Understanding the MC/ServiceGuard SAP R/3 Integration
SGeSAP Product Files

2 Step by Step Installation Guide

This chapter describes how to configure and install highly available R/3 systems on HP-UX 10.x and 11.x using MC/ServiceGuard. The process is split into the following logical tasks:

- Planning the LVM Setup
- Preparing the Integration
- HP-UX Configuration
- SGeSAP Files Configuration
- MC/ServiceGuard Configuration
- Automounter Configuration
- Database Configuration
- SAP R/3 System Configuration

The tasks are presented as a sequence of steps. Each mandatory installation step is accompanied by a unique number of the format *XXnnn*, where *nnn* are incrementing values and *XX* indicates the step relationship, as follows:

- ISnnn—Installation Steps - mandatory
- OSnnn—Optional Steps
- ORnnn—ORacle database only steps
- IRnnn—InfoRmix database only steps

Whenever appropriate, HP-UX sample commands are given to guide you through the process in as detailed a manner as possible. It is assumed that hardware as well as the operating system and MC/ServiceGuard are already installed properly on all cluster hosts. Sometimes a condition is specified with the installation step. Follow the information presented *only* if the condition is true for your situation.

NOTE

For installation steps in this chapter that require the adjustment of SAP specific parameter in order to run the SAP R/3 system in a switchover environment usually example values are given. These values are for

reference ONLY and it is recommended to read and follow the appropriate SAP OSS notes for SAP's latest recommendation. Whenever possible the OSS note number is given.

Planning the LVM Setup

IS020 Installation Step:

Record all the minor numbers already in use on the cluster hosts. See Table 2-1.

The device minor numbers of the shared volume groups must be the same on all cluster hosts. They must differ from all other volume group device minor numbers used on the cluster hosts.

To find these minor device numbers on each cluster hosts, type:

```
ll /dev/vg*/group|cut -c 44-45
```

Table 2-1

Hosts and Device Minor Numbers

hostname	device minor numbers
...	
...	
...	
...	

IS030 Installation Step:

Create new special files for the volume groups as needed.

Refer to Table 2-1, check that all the minor numbers for one host are unique for that host. If any host uses the same number twice create new special files with a unique number.

IS040 Installation Step:

Specify the needed volume groups in another table.

For the standard setup there should be at least three cluster volume groups:

- one volume group for filesystems that are specific to the Central Instance
- one volume group for more general SAP R/3 filesystems

Step by Step Installation Guide
Planning the LVM Setup

- one volume group for the database specific filesystems

Two shared volume groups would also work but this is not recommended because of lack of flexibility. If you put all logical volumes in one volume group, you will be restricted to the one package concept. You should end up with a table similar to Table 2-2.

NOTE

If there is more than one SAP R/3 system in the cluster then /usr/sap/trans may be outside of this configuration.

Table 2-2 Volume Groups Needed for MC/ServiceGuard Packages

MC/SG package	VG Name	lvol Names	Later Mount Point	Device minor number
ci<SID>	vgsap<SID>	lvsap<SID>	/usr/sap/<SID>/DVEBMGS<INSTNR>	08
db<SID>	vgmnt<SID>	lvmnt<SID> lvtrans	/export/sapmnt/<SID> /export/usr/sap/trans	09
	vgdb<SID>	lvoracle lvstage_805 lvsaparch lvsapreorg lvsapdata1 lvsapdata2 lvsapdata3 lvsapdata4 lvsapdata5 lvsapdata6 lvoriglogA lvoriglogB lvmirrlogA lvmirrlogB or /lvinformix ...	/oracle/<SID> /oracle/stage/stage_805 /oracle/<SID>/saparch /oracle/<SID>/sapreorg /oracle/<SID>/sapdata1 /oracle/<SID>/sapdata2 /oracle/<SID>/sapdata3 /oracle/<SID>/sapdata4 /oracle/<SID>/sapdata5 /oracle/<SID>/sapdata6 /oracle/<SID>/origlogA /oracle/<SID>/origlogB /oracle/<SID>/mirrlogA /oracle/<SID>/mirrlogB or /export/informix ...	10
ap<SID><INSTNR>	vg<SID><INSTNR>	ivsap<SID>	/usr/sap/<SID>/D<INSTNR>	11

Using the one package concept, all the volume groups will belong to one package. For reference we call this package `dbci<SID>` or the DBCI package.

Using the two package concept, `vg<SID>` belongs to the Central Instance package (CI). The CI package contains all logical volumes that are exclusively needed by the Central Instance. `vg<SID>` and `vgdb<SID>` are configured to belong to the database package (DB).

If you use only two cluster volume groups, put `/export/sapmnt/<SID>` and `/export/usr/sap/trans` in the volume group which belongs to the node that runs the HA NFS server (this should be the database server). If you plan to use more than one SAP R/3 system using the same transport directory `/usr/sap/trans`, do not configure it with an SAP R/3 system package. If you stop the SAP R/3 system package, other systems will be affected, too. They cannot be restarted and transports are impossible. CCMS functionality should be used with care. For additional `/export/usr/sap/trans` security use the HP provided NFS toolkit. Using the Application server package you should setup a separate volume group for this package. This volume group `vg<SID><INSTNR>` consists of one logical volume called `lvsap<SAP>`. It holds the work, log and data directories for this specific instance.

If you have more than one system, place `/oracle/stage/stage_nnn` and `/usr/sap/put` on separate volume groups created on shared drives. These directories should not be added to any package. This ensures that they are independent from any SAP R/3 system and you can mount them on any host by hand if needed.

If SAP R/3 is not installed yet, specify values in the device minor column that are different from all values you identified in Step IS020. Use Table 2-1 as a template for recording the device minor values.

Table 2-3 Device Minor Values

MC/SG package	VG name	lvol names	later mount point	Device Minor Number
...				
...				
...				
...				

Preparing the Integration

If the SAP R/3 system you are going to make highly available is already installed, proceed section “Steps if SAP R/3 is Already Installed.” If the SAP R/3 system you are going to make highly available is not already installed, proceed to section “Steps if SAP R/3 is Not Installed Yet.”

Steps if SAP R/3 is Already Installed

IS050 Installation Step:

- Logon as root to the system where the SAP R/3 Central Instance is installed. If the database is installed on a different host, also open a shell as root on the database machine. Stop SAP R/3 and the database if they are not already.
- Verify that the existing volume group layout is compliant with the needs of the MC/ServiceGuard package(s) as specified in Table 2-3 of Step IS040. Test the following:
 - If you install using the two package concept, make sure that database specific filesystems and Central Instance specific filesystems are separated onto different volume groups.
 - Specify one mountpoint as `/usr/sap/<SID>`. For MC/ServiceGuard, change this to `/usr/sap/<SID>/DVEBMGS<INSTNR>`. If `/usr/sap/<SID>` is the mountpoint you have to move some files to a local logical volume and change the mountpoint. For example:

```
mkdir /usr/sap/<SID>.new
cd /usr/sap/<SID>
bdf . # Remember the filesystem column.
      # It will be referred to as <dev_path> later.
find . -depth -print|cpio -pd /usr/sap/<SID>.new
cd /
umount /usr/sap/<SID>
rmdir /usr/sap/<SID>
mv /usr/sap/<SID>.new /usr/sap/<SID>
chmod 751 /usr/sap/<SID>
chown <SID>adm:sapsys /usr/sap/<SID>
cd /usr/sap/<SID>/DVEBMGS<INSTNR>
rm -r * # be careful with this
```

```
cd ..
mount <dev_path> /usr/sap/<SID>/DVEBMGS<INSTNR>
cd DVEBMGS<INSTNR>
ls
# remove everything that is different from
DVEBMGS<INSTNR>,
# Example: rm -r SYS
#           rm -r D00
cd DVEBMGS<INSTNR>
find . -depth -print|cpio -pd
/usr/sap/<SID>/DVEBMGS<INSTNR>
rm -r * # be careful with this
cd ..
rmdir DVEBMGS<INSTNR>
```

- Mark all volume groups as members of the cluster. This only works if the cluster services are already available. For example:

```
cd /
# umount all logical volumes of the volume group
vgchange -a n <vg_name>
vgchange -c y <vg_name>
vgchange -a e <vg_name>
# remount the logical volumes.
```

- The device minor numbers must be different from all device minor numbers gathered on the other hosts. Verify this by comparing numbers listed in Table 2-2 in Step IS040 to the numbers listed in Table 2-1 in Step IS020.
- Create and distribute mapfiles for all shared volume groups.

NOTE

The logical volumes that are called `lvmnt<SID>` and `lvtrans` in the sample Table 2-2 of Step IS040 are not mounted below `/export`, but below the root directory. Keep it like this. It's OK.

IS060 Installation Step:

Comment out the references to any filesystem that is mentioned in Step IS040 from the `/etc/fstab`.

OS070 Optional Step:

If there are additional internal application servers on the cluster hosts:

Logon to all cluster hosts that have application servers installed. Verify the following:

- Application servers in the cluster must have an instance number which is different from the instance number of the Central Instance. Execute the command:

```
ls -d /usr/sap/<SID>/*<INSTNR>
```

It should reply with No Match. If it does not, reinstall the instance with another instance ID.

- Comment out `/etc/fstab` entries to filesystems mentioned in Step IS040.
- If you use SAP R/3 4.0B or higher, `/oracle/SID` must not be a mountpoint of a local filesystem. You can use `/oracle` as mountpoint.

IS080 Installation Step:

The usage of local executables with the SAPCPE mechanism is required for SGeSAP environments. You can do this now, because it prevents hangs when performing a `su to <sid>adm` later in this installation.

Check if the Central Instance host and all application servers have a directory called, type:

```
/usr/sap/<SID>/SYS/exe/ctrun.
```

If the directories exists, you are done. The system is already using local executables through `sapcpe`.

If the directories do not exist, setup `sapcpe`. Refer to your OSS notes and the SAP R/3 online documentation for information about how to do this.

Steps if SAP R/3 is Not Installed Yet

SAP R/3 is installed on one host. This remains true even if you want to distribute the database and Central Instance using the three-tier concept. MC/ServiceGuard allows you to remove one or the other later. The host you choose now will become the primary host.

IS090 Installation Step:

Logon to the primary host as root.

Create the shared volume groups using the minor numbers specified in Step IS040. Use the command line interface to create the whole volume group, and later, the distribution. Do not use SAM. SAM will scramble the minor numbers. For example:

```
mkdir /dev/vgsap<SID>
mknod /dev/vgsap<SID>/group c 64 0x080000
pvcreate <pv_path>
vgcreate vgsap<SID> <pv_path>
vgchange -c y vgsap<SID>
lvcreate -n lvsap<SID> vgsap<SID>
vgchange -a e vgsap<SID>
```

Refer to your SAP R/3 documentation for a description on how to distribute the filesystems on physical disks. Refer to the SAP R/3 installation documents to find out about the filesystem sizes you need. Use `lvextend` and `mkfs -F vxfs` commands to create the filesystems on the shared volume groups. Create and distribute the mapfiles with `vgexport (1m)`.

NOTE

R3INST/R3SETUP creates a file called `SAPMADB.PAR`, `SAPMAIN.PAR` or `SAPAPPL.PAR` in the installation directory. It shows needed filesystems and recommends sizes for them.

For the recommended file system refer to the latest OSS notes.

OS100 Optional Step:

If there is a different SAP R/3 system already installed on another cluster node, propagate the additional SAP R/3 HP-UX users and groups to all the nodes of the cluster. This prevents conflicts with the UIDs. The SAP R/3 installation process creates users, groups and services and it performs local consistency checks. It is possible, that UIDs, GIDs or service ports already exist on other cluster hosts.

Synchronize the `/etc/passwd`, `/etc/group` and `/etc/services` files of the cluster nodes so that they have the same contents. This prevents conflicts later in the installation process.

IS110 Installation Step:

Logon to the primary host as root.

Create the mountpoints as recommended in the SAP R/3 installation documents, except, do not use `/usr/sap/<SID>` as a mountpoint. Mount one level below the instance directory. For example:

```
mkdir -p /usr/sap/<SID>/DVEBMGS<INSTNR>  
...
```

The filesystems should be distributed on disks and volume groups as recommended in Step IS040. But instead of using the mountpoints below `/export` that are suggested in Step IS040, create mountpoints below `/` as recommended by the SAP R/3 documentation. For example:

```
mkdir -p /sapmnt/<SID>  
mkdir -p /usr/sap/trans  
mkdir -p /informix  
...
```

Mount all created filesystems manually. Do not create entries in `/etc/fstab`.

IS120 Installation Step:

Begin SAP R/3 installation now. Install the complete system on the primary machine. Refer to your standard SAP R/3 documentation and OSS notes on how to do this.

If you use `R3SETUP` to install, the space check might fail. Modify the `SPACECHECK` entries in the `.R3S` file in the installation directory to reflect the mountpoint changes. For example:

Search for `[Z_CENTRDBSPACECHECK_IND_IND]` in `CENTRDB.R3S`.

The `@SAPDIR@/@SAPSYSTEMNAME@` entries check the disk space of `/usr/sap/<SID>`.

Additional free space caused by the logical volume that is mounted below at the mountpoint `/usr/sap/<SID>/DVEBMGS<INSTNR>` is not counted.

Do not continue until SAP R/3 is properly installed on the primary node.

OS130 Optional Step:

If internal application servers are needed:

Setup application servers with R3INST/R3SETUP as usual. If you install an application server on a node within the cluster, consider the following:

- Application servers in the cluster must have an instance number which is different from the instance number of the Central Instance.
- Do not create `/etc/fstab` entries. Mount manually to avoid problems with the automounter later on.

IS140 Installation Step:

The usage of local executables with the SAPCPE mechanism is required for SGeSAP environments. Do this now to prevent hangs when performing a `su` to `<sid>adm` later in this installation process. Refer to your OSS notes and the SAP R/3 online documentation for information.

HP-UX Configuration

Correct HP-UX configuration ensures that all cluster nodes provide the environment and system configuration required to run SAP R/3. This section describes how to distribute the SAP R/3 installation configuration changes made in the previous section, “Preparing the Integration,” among the cluster nodes, that is, from all primary nodes to all other nodes.

Several of the following steps must be repeated on each node. Record the steps completed for each node, as you complete them. This helps identify errors in the event of a malfunction later in the integration process. The HP-UX configuration task is split into the following sections:

- Cluster Node Synchronization

This section consists of steps performed on the backup nodes. These ensure that the primary node and the backup nodes have a similar environment.

Repeat the steps in this section for each node of the cluster that is different from the primary.

The primary host is the host where the Central Instance was installed. If your database is currently running on a machine different from this, repeat all steps once again for each node with the database machine as primary host. This distributes the database configuration, too.

- Cluster Node Configuration

This section consists of steps performed on all the cluster nodes, regardless if the node is a primary node or a backup node.

Repeat the steps in this section for each node of the cluster.

- External Application Server Host Configuration

This section consists of steps performed on any host outside of the cluster that runs another instance of the SAP R/3 system.

Repeat this section for each host that has an external application server installed.

Performing some of the iterations in parallel is fine, just use caution in any complex setup situation.

Rather than using the cut and paste mechanism you can also fill out the tables provided by first analyzing the primary host(s). Afterwards you can use the tables to synchronize all nodes.

Cluster Node Synchronization

Repeat the steps in this section for each node of the cluster that is different than the primary.

- Logon as root to the primary host.
- Prepare a logon for each of its backup hosts.

IS150 Installation Step:

Look at the groupfile file, `/etc/group`, on the primary side.

If any of the groups listed in Table 2-4 exist on the primary node and they do not exist on the backup node, copy them from the primary node to the backup node. If any group exists, verify that it has the same GID on both the primary and backup nodes. Merge the group members lists.

Table 2-4

Groupfile File Groups

groups	GID	group members
sapsys		
dba		
oper		
informix		
super_archive		

IS160 Installation Step:

Look at the password file, `/etc/passwd`, on the primary side.

If any of the users listed in Table 2-5 exist on the primary node, recreate them on the backup node. Assign the users on the backup nodes the same user and group ID as the primary nodes.

INFORMIX users must have the same passwords, as well, on both the backup and primary nodes. Beware of copying over into `/etc/passwd` if

your HP-UX is running in Trusted System mode.

Table 2-5 Password File Users

username	UID	GID	home directory	shell
<sid>adm				
ora<sid>				
informix				
sapr3				

IS170 Installation Step:

Look at the service file, /etc/services, on the primary side.

Replicate all services listed in Table 2-6 that exist on the primary node onto the backup node.

Table 2-6 Service File Services

service name	service port
sapdp00	
sapdp01	
...	
sapdp<nn>	
...	
sapdp99	
sapdp<nn>s	
sapgw00	
sapgw01	
...	
sapgw<nn>	
...	

Table 2-6 Service File Services

service name	service port
sapgw99	
sapgw<nn>s	
sapms<SID>	
orasrv	
tlisrv	
sapinf<SID>	

IS180 Installation Step:

Change the HP-UX kernel on the backup node to meet the SAP R/3 requirements.

Compare the Tunable Parameters section of `/stand/system` on both hosts. All values on the backup host must reach or exceed the values of the primary host. A tool, `kinst`, configures the kernel automatically. It is bundled with SAP R/3 releases prior to 3.0F.

Install all HP-UX patches that are recommended for MC/ServiceGuard and patches recommended for SAP R/3.

If you are planning to configure multiple Application Servers that are configured to be started parallel (see IS510) make sure the kernel parameter “`tcp_conn_request_max`” of your HP-UX systems are configured appropriately to 1024. This will prevent closing tcp ports (for example the LISTENER port configured for Oracle DB).

IS200 Installation Step:

Build a new kernel with `mk_kernel (1m)` on the backup host if `/stand/system` was changed in either Step IS180 or .

IS210 Installation Step:

If the primary node has the Central Instance installed:

- Copy the `<sid>adm` home directory to the backup node. This is a local directory on each node.
- Rename the start, stop and environment scripts in `/home/<sid>adm` on the secondary node. Some of the environment scripts may not

Step by Step Installation Guide

HP-UX Configuration

exist. For example:

```
su - <sid>adm
mv startsap_<primary>_<INSTNR>
startsap_<secondary>_<INSTNR>
mv stopsap_<primary>_<INSTNR> stopsap_<secondary>_<INSTNR>
mv .sapenv_<primary>.csh .sapenv_<secondary>.csh
mv .sapenv_<primary>.sh .sapenv_<secondary>.sh
mv .dbenv_<primary>.csh .dbenv_<secondary>.csh
mv .dbenv_<primary>.sh .dbenv_<secondary>.sh
# Remove logfiles from <primary> if any
exit
```

Never use the relocatable address in these filenames. If an application server was already installed, do not overwrite any files which will start the application server. If the rc-files have been modified, correct any hardcoded references to the primary hostname.

IR220 Informix Database Step:

If the primary node has the INFORMIX database installed:

- Copy the home directories of `sapr3` and `informix` to the backup node.
- Rename the environment scripts in `/home/informix` on the backup node. For example:

```
su - informix
mv .dbenv_<primary>.csh .dbenv_<secondary>.csh
mv .dbenv_<primary>.sh .dbenv_<secondary>.sh
exit
```


OR230 Oracle Database Step:

If the primary node has the ORACLE database installed:

Create additional links in `/oracle/<SID>` on the primary node. For example:

```
su - ora<sid>  
ln .dbenv_<primary>.csh .dbenv_<secondary>.csh  
ln .dbenv_<primary>.sh .dbenv_<secondary>.sh  
exit
```

OR240 Oracle Database Step:

If you are using ORACLE:

Create a mountpoint for the Oracle files on the backup node if it is not already there. For example:

```
su - ora<sid>  
mkdir -p /oracle/<SID>  
exit
```

NOTE

Complete this step also, if you are using 4.0A, the two package concept, and the backup node will be configured to run the Central Instance package only.

OS250 Optional Step:

If you need different configurations of the Central Instance depending on the node it runs on:

Be careful if you create different instance profiles for the different machines. You have to make sure by yourself that the Central Instance is capable of doing the intended work with any of the possible configurations.

Perform the following steps to setup different configurations:

- Execute the command:

```
su - <sid>adm
```

- On each host the files
`/home/<sid>adm/startsap_<local>_<INSTNR>` and

Step by Step Installation Guide

HP-UX Configuration

`/home/<sid>adm/stopsap_<local>_<INSTNR>` contain a line that specifies the start profile. After a standard installation this line is similar to:

```
START_PROFILE="START_DVEBMGS<INSTNR>_<primary>"
```

Change the line individually on each host, MC/ServiceGuard in a two-node cluster:

— On the primary host keep:

```
START_PROFILE="START_DVEBMGS<INSTNR>_<primary>"
```

— On the secondary host change the value in both files to:

```
START_PROFILE="START_DVEBMGS<INSTNR>_<secondary>"
```

- The start profile and the instance profile can be found in the shared profile directory. Be careful if you rename the instance profile. You must change all references to this file afterwards. The following is an example of the steps to create two individual configurations in a two-node cluster:

```
cp /sapmnt/<SID>/profile/START_DVEBMGS<INSTNR>_<primary>\
/sapmnt/<SID>/profile/START_DVEBMGS<INSTNR>_<secondary>
cp /sapmnt/<SID>/profile/<SID>_DVEBMGS<INSTNR>_<primary>\
/sapmnt/<SID>/profile/<SID>_DVEBMGS<INSTNR>_<secondary>
```

- Now replace any reference to `<SID>_DVEBMGS<INSTNR>_<primary>`, type:

```
<SID>_DVEBMGS<INSTNR>_<secondary> in the file
/sapmnt/<SID>/profile/START_DVEBMGS<INSTNR>_<secondary>.
```

IS260 Installation Step:

If the primary node has the Central Instance installed and the other node has no internal application server installed:

Distribute the local directory tree `/usr/sap/<SID>/SYS`. Do not use `rcp(1)`, it will follow all links and copy a lot of files from the shared disks that are not needed.

For example:

On the primary node:

```
cd /usr/sap/<SID>/SYS
find . -depth -print | cpio -o >/tmp/SYS.cpio
```

use ftp(1) to copy the file over to the secondary node

On the secondary node:

```
su - <sid>adm
mkdir -p /usr/sap/<SID>/SYS
cd /usr/sap/<SID>/SYS
cpio -id </tmp/SYS.cpio
exit
```

IS270 Installation Step:

Import the shared volume groups using the minor numbers specified in Step IS040.

The whole volume group distribution should be done using the command line interface. Do not use SAM. SAM will scramble the minor numbers. Specify the device minor numbers explicitly by creating the groupfile manually. For example:

```
mkdir /dev/vgsap<SID>
mknod /dev/vgsap<SID>/group c 64 0x080000
```

Now you can use vgimport(1m) with the mapfile created on the primary host during Step IS050 or Step IS090.

IS275 Installation Step:

Create a mountpoint for the Central Instance directory so the node can run the Central Instance. For example:

```
mkdir -p /usr/sap/<SID>/DVEBMGS<INSTNR>
```

Cluster Node Configuration

Repeat the steps in this section for each node of the cluster.

Logon as root.

IS280 Installation Step:

Check that /etc/loggingroup is a link to /etc/group.

IS290 Installation Step:

Create an .rhosts file in the home directories of the HP-UX users root, <sid>adm and (if applicable) informix. Allow login for root as root from all nodes including the node you are logged into. Allow login for root and

Step by Step Installation Guide

HP-UX Configuration

<sid>adm as <sid>adm from all nodes including the node you are logged into. Be careful with this step, many problems result from an incorrect setup of remote access.

Check the setup with remsh commands. If you have to provide a password, the .rhosts does not work.

IS300 Installation Step:

Create all directories below /export as specified in Step IS040. For example:

```
su - <sid>adm
mkdir -p /export/sapmnt/<SID>
mkdir -p /export/usr/sap/trans
exit
su - informix
mkdir -p /export/informix (Informix only)
exit
```

IS310 Installation Step:

Add all relocatable IP address information to /etc/hosts. Do not forget heartbeat IP addresses. Use Table 2-7 to record the addresses.

Table 2-7

Relocatable IP Address Information

name/aliases	IP address
	. . .
	. . .
	. . .
	. . .
	. . .
	. . .
	. . .
	. . .
	. . .
	. . .
	. . .

IS320 Installation Step:

If you use DNS:

Configure `/etc/nsswitch.conf` to avoid problems. For example:

```
hosts: files[NOTFOUND=continue UNAVAIL=continue \  
TRYAGAIN=continue] dns
```

IS330 Installation Step:

If you establish frontend and server LANs to separate network traffic:

Add routing entries to the internet routing configurations of `/etc/rc.config.d/netconf`. This is the only phase of the whole installation in which you will need to specify addresses of the server LAN. Route all relocatable client LAN addresses to the local server LAN addresses. For example, a two package concept:

```
ROUTE_DESTINATION [n]=<relocdb>  
ROUTE_MASK [n]=""  
ROUTE_GATEWAY [n]=<relocdb_s>  
ROUTE_COUNT [n]=1  
ROUTE_ARGS [n]=""  
ROUTE_DESTINATION [n+1]=<relocci>  
ROUTE_MASK [n+1]=""  
ROUTE_GATEWAY [n+1]=<relocci_s>  
ROUTE_COUNT [n+1]=1  
ROUTE_ARGS [n+1]=""
```

IS335 Installation Step:

If you have SAP R/3 4.0A or higher:

Beginning with SAP R/3 4.0A, during installation SAP R/3 appends some entries to the standard `.profile` files in the user home directories instead of using a new file defined by SAP R/3. On HP-UX, by default, there is the following in the given profiles:

```
set -u
```

This confuses the `.dbenv*.sh` and `.sapenv*.sh` files of SAP R/3. They fail during execution if the environment is not setup properly. Using SGeSAP the package startup fails trying to bring up the database.

Search the `.profile` of `<sid>adm` and remove the `set -u`, if found.

External Application Server Host Configuration

Repeat the steps in this section for each host that has an external application server installed.

Logon as root.

IS340 Installation Step:

Create an `.rhosts` file in the home directories of the HP-UX users `root`, `<sid>adm` and (if applicable) `informix`. Allow login for users `root` and `<sid>adm` from all cluster nodes.

For reasons described in Step IS335, search `.profile` in the home directory of `<sid>adm` and remove the `set -u`, if found.

IS345 Installation Step:

If you have SAP R/3 4.0A or higher:

For reasons described in Step IS335, search `.profile` in the home directory of `<sid>adm` and remove the `set -u`, if found.

IS350 Installation Step:

Add all relocatable IP address information to `/etc/hosts`.

IS360 Installation Step:

If you establish frontend and server LANs to separate network traffic:

Add routing entries to the internet routing configurations of `/etc/rc.config.d/netconf`. This is the only phase of the whole installation in which you will need to specify addresses of the server LAN. Route all relocatable client LAN addresses to the local server LAN addresses. For example, a one package concept:

```
ROUTE_DESTINATION [n] = <relocdbci>  
ROUTE_MASK [n] = "  
ROUTE_GATEWAY [n] = <relocdbci_s>  
ROUTE_COUNT [n] = 1  
ROUTE_ARGS [n] = "
```

SGeSAP Files Configuration

The MC/ServiceGuard Extension for SAP R/3 (SGeSAP) integration needs information about the specific setup at the customer site. It gathers this information from a file that is called `sap.conf`. You have to modify this file manually. Refer to the example provided with the integration files. It can be used as a template.

Logon to the primary host as root.

IS370 Installation Step:

Install the product depot file SGeSAP B7885BA using the `swinstall` tool.

B7885BA depends on B5140BA, the NFS toolkit which will also be installed if it is not already available.

IS375 Installation Step:

The `swinstall` process copied relevant files to `/opt/cmcluster` for reference. For your installation to work, copy the following files:

```
cp /opt/cmcluster/sap/SID/sap.conf /etc/cmcluster/<SID>/sap.conf
cp /opt/cmcluster/sap/sap.functions /etc/cmcluster/sap.functions
cp /opt/cmcluster/sap/customer.functions /etc/cmcluster/customer.functions
```

IS377 Installation Step:

In addition, for a one package installation you need the following package control and configuration files out of the NFS toolkit repository:

For the 1 package installation you need:

```
cp /opt/cmcluster/sap/SID/sapdbci.cnt1/etc/cmcluster/<SID>/sapdbci.cnt1
cp /opt/cmcluster/nfs/nfs.cnt1/etc/cmcluster/<SID>/dbci.cnt1
cp /opt/cmcluster/nfs/nfs.conf/etc/cmcluster/<SID>/dbci.conf
```

For the two package installation you need:

```
cp /opt/cmcluster/sap/SID/sapdb.cnt1 /etc/cmcluster/<SID>/sapdb.cnt1
cp /opt/cmcluster/sap/SID/sapci.cnt1 /etc/cmcluster/<SID>/sapci.cnt1
cp /opt/cmcluster/nfs/nfs.cnt1/etc/cmcluster/<SID>/db.cnt1
cp /opt/cmcluster/nfs/nfs.cnt1/etc/cmcluster/<SID>/ci.cnt1
cp /opt/cmcluster/nfs/nfs.conf/etc/cmcluster/<SID>/db.conf
cp /opt/cmcluster/nfs/nfs.conf/etc/cmcluster/<SID>/ci.conf
```

For the a application server package installation you need:

```
cp /opt/cmcluster/sapap.cnt1 /etc/cmcluster/<SID>/sapap.cnt1
```

Step by Step Installation Guide

SGeSAP Files Configuration

```
cp /opt/cmcluster/nfs/nfs.cnt1 /etc/cmcluster/<SID>/ap<SID><ASNR>.conf
cp /opt/cmcluster/nfs/nfs.cnt1 /etc/cmcluster/<SID>/ap<SID><ASNR>.cnt1
```

IS380 Installation Step:

Customize the package control and configuration files as described in the standard MC/ServiceGuard manual *Managing MC/ServiceGuard* (Part Number B3936-90026) and *Managing Highly Available NFS* (Part Number B5125-90001).

IS385 Installation Step:

Customization of `sap.conf`.

The customization of `sap.conf` is divided into subsections as follows:

- Standard Parameters of Release 3.x
- Additional Parameters in `sap.conf`
- Using Database Reconnect and Transaction Reset
- Advanced Options of `sap.conf`
- HP Somersault Integration with SGeSAP

Standard Parameters of Release 3.x

Open the file `/etc/cmcluster/<SID>/sap.conf` with a text editor.

IS390 Installation Step:

SGeSAP performs activities specific to the database you use. Specify the underlying database vendor using the `DB` parameter. The compatible options are: `ORACLE` or `INFORMIX`.

IS400 Installation Step:

Provide information about the Central Instance that will be protected by a MC/ServiceGuard package. Set the parameter `CINR` to the *Instance ID* of your Central Instance.

IS410 Installation Step:

If you are using the two package concept you can use the `DB_RECONNECT` functionality that is provided by SAP R/3.

If you want to use `DB_RECONNECT` with any instance of your SAP R/3 system, you *must* disable the automatic restart of the Central Instance, type:


```
CIRESTART=0
```

Setting this prevents the Central Instance from restarting in the event of a database failover. This means that the Central Instance profiles must be configured to use `DB_RECONNECT`. Remember that if you want to use `DB_RECONNECT` with any application server you must use it with the Central Instance, too. A reconnection of the application server to the switched database is only possible if the Central Instance has not been restarted.

The `CIRESTART` parameter has no meaning with a Central System that is secured by the one package concept and should not be modified. A failure of the database with a one package concept always means that the Central Instance has failed. The Central Instance is always restarted and reconnection of application servers is not allowed. The same conditions apply if you plan to use a two package approach with both packages on the same node all the time. When configuring the `DBRECONNECT` mechanism follow the appropriate OSS notes 109036, 98051 and 24806

IS420 Installation Step:

The file `sap.conf` contains four arrays that describe the additional application servers of your system.

Additional application servers are all application servers that belong to the SAP R/3 system and are different from the Central Instance. Each array has one entry for each application server. This is true whether the application server is running on a node in the cluster or outside of the cluster.

- In the `ASHOST[*]` array, specify the hostnames on which the application servers reside. Never use a relocatable name, even if the host is part of the cluster!
- Specify the Instance ID for each application server using the `ASNR[*]` array.

You already specified the Instance ID of the Central Instance with the parameter `CINR`.

Each index value n , `ASNR[n]` refers to the same application server as `ASHOST[n]`. If the corresponding `ASHOST` entry specifies a host that is part of the cluster, provide an ID that is different from the ID used by the Central Instance.

Make sure that there is no other SAP R/3 Instance on the same host

using the same ID.

- In the third array called `ASTREAT [*]`, you can define the way the application server acts if the status of the package changes. `ASTREAT [*] = 0` means that the application server is not affected by any changes that happen to the package status.
 - Add 1 to `ASTREAT [*]` if the application server should be started automatically during startup of a package.
 - Add 2 to `ASTREAT [*]` if the application server should be stopped automatically after initiating `cmhaltpkg` for `ci-pkg`.
 - Add 4 to `ASTREAT [*]` if the application server should be restarted automatically if a DB-switchover caused by a failure takes place. If you do not use the restart option you have to configure the instance to use `DB-RECONNECT`.
 - In the fourth array called `ASPLATFORM [*]` you specify the platform on which the Application Server runs. Supported values are:
 - “HP-UX”: standard SAP Application server running on an HP-UX server
 - “LINUX”: standard SAP Application server running on a linux server
 - “NT”: standard SAP Application server running on a NT server. The NT Application server handling is not standardized as there is no way to open a remote DOS shell that starts R/3 Application servers on a windows platform. SGeSAP right now contains examples of functions using the ATAMAN™ TCP Remote Logon.
 - “SG-PACKAGE”: MC/ServiceGuard packaged SAP Application server. Specify this value is you want to run the Application Server within a ServiceGuard Cluster package. Refer to IS 377 for prerequisites.

IS430 Installation Step:

Specify the relocatable hostnames of the database and the Central Instance in the parameters `DBRELOC` and `CIRELOC`.

They will be the same if you use the one package concept.

Specify the relocatable hostname in `TRANSRELOC`, the host from which

the mount to the transport directory `/usr/sap/trans` is initiated. In a multiple SAP R/3 system environment, specify a separate relocatable hostname for `TRANSRELOC`. In single system environments `TRANSRELOC` can be set equal `DBRELOC`.

OS435 Installation Step:

Specify `AS_PSTART=0` if you want the Application Server startup to run sequentially. The default value here is `AS_PSTART=1` for parallel startup. Setting `AS_PSTART=0` will slow down your total failover time.

OS450 Optional Step:

Specify `SAPOSCOL_STOP=1` if `saposcol` should be stopped together with each instance that is stopped.

SGeSAP makes sure that the collector only stops if there is no instance of an SAP R/3 system running on the host.

OS460 Optional Step:

It is possible to failover both packages of the two package concept to the same node.

If both packages try to come up on a single node after a failover at the same time, it is likely that the Central Instance package wants to start up the Central Instance before the database is fully recovered. This would lead to a failure because SAP R/3 cannot connect to the database.

To deal with this situation, there is a loop implemented that polls the database in increasing intervals of time. The startup of the CI package is delayed until the database is reached. After the first poll, the script waits 10 seconds before initiating another one. After the second poll, the waiting time is increased to 20 seconds, and so forth. This continues until the database responds or up to a maximum of `DELAY_INTERVALS` polling attempts. You can modify `DELAY_INTERVALS` if you expect long recovery times.

Additional Parameters in `sap.conf`

Starting with release 3.0.03 additional parameters were introduced. You can add them on an *as needed* basis to already existing `sap.conf` files. The `sap.functions` file and `customer.functions` of release 3.0.03 or higher works with `sap.conf` files of previous releases without changing them.

OS470 Optional Step:

Step by Step Installation Guide

SGeSAP Files Configuration

If your setup consists of application servers that are significantly slower than the Central Instance host, it is possible that the Central Instance shuts down before application server shutdown is completed. This can lead to unsafe shutdowns and Instance crash.

To be safe, specify one of the following:

- `WAIT_OWN_AS=1`
the shutdown of all application servers takes place in parallel, but the scripts do not continue before all of these shutdown processes have come to an end.
- `WAIT_OWN_AS=2`
if the package should also wait for all application servers to come up successfully. You have to use this value if you want to prevent the integration from temporarily opening a new process group for each application server during startup.
- `WAIT_OWN_AS=0`
can significantly speed up the package start and stop, especially if Windows NT application servers are used. Use this value only if you have carefully tested and verified that timing issues will not occur.

OS480 Optional Step:

Control the handling of resources.

On 32bit HP-UX, a shared memory shortage can occur if you install more than one instance on a host. Prior to any instance startup the SGeSAP tries to free up unused or unimportant resources to make the startup more likely to succeed. A database package only frees up database related resources, a Central Instance package only removes IPCs belonging to SAP R/3 administrators. Table 2-8 summarizes how the behavior of SGeSAP is affected by changing the `CLEANUP_POLICY` parameter:

Table 2-8

Resource Handling and the `CLEANUP_POLICY` Parameter

<code>CLEANUP_POLICY =</code>	CI resources	DB resources	AppServer resources
lazy	no action	no action	no action

Table 2-8 Resource Handling and the CLEANUP_POLICY Parameter

CLEANUP_POLICY =	CI resources	DB resources	AppServer resources
normal	use cleanipc to remove unused own resources	unused ORACLE SGA is removed	use cleanipc to remove unused own resources
strict	remove any resource belonging to any SAP R/3-Instance	remove any resource belonging to any database	use cleanipc to remove unused own resources

- Integration versions prior to release 3.0.03 use the lazy policy: No additional cleanup takes place apart from SAP R/3 standard shutdown activities.
- `cleanipc` is an SAP R/3 tool used to free up the IPC resources of specific SAP R/3 Instances. It is used to free up resources of an Instance that is to be started soon on HP-UX. This prevents `shmem`-problems due to a prior crash of the Instance. Accordingly, an obsolete ORACLE SGA is also removed if a database crash occurred.
- The `strict` policy uses HP-UX commands to free up all semaphores and shared memory segments that belong to *any* SAP R/3 Instance of *any* SAP R/3 system on the host if the Central Instance is to be started soon. It uses HP-UX to free up all semaphores and shared memory segments that belong to *any* database if the SAP R/3 database is to be started soon.

Do not use the strict policy unless it is critical that you do. Be aware that the strict option can crash running instances of different SAP R/3 systems on the backup host!

- Use this value only if you have a productive system that is much more important than any other SAP R/3 system you have. In this case a switchover of the productive system is more robust, but additional SAP R/3 systems will crash.
- You can also use `strict` policy, if your SAP R/3 system is the only one running at the site and you are low on memory. `Strict` policy frees up more of its own shared memory segments than the normal policy does.

OS485 Optional Step:

SERVER_CONSOLIDATION is used to identify the usage of a consolidated system environment that has multiple SAP R/3 systems on the same node.

To enable the handling of shared memory resources and application servers in a consolidated environment, specify:

```
SERVER_CONSOLIDATION=1
```

To specify the Dialog instances handling refer to Step OS510.

Using Database Reconnect and Transaction Reset

Starting with V3.0 of the SGeSAP Integration Scripts you can configure the DB-Reconnect functionality. DB-Reconnect is a feature provided by SAP R/3. It allows SAP R/3 to remain running if the database switches and the Central Instance stays up. DB-Reconnect can be used with the two package concept. Transaction Reset can be used with either the one package or two package concept. Both DB-Reconnect and Transaction Reset can be combined in the two package concept.

SAP R/3 V4.0B kernel patch level 85 introduced the Transaction Reset functionality. It allows additional SAP R/3 application servers to remain running if the Central Instance switches. The first implementation of Transaction Reset was unreliable and caused problems with some SAP R/3 transactions. Use at least kernel patch level 269.

The DB-Reconnect can be different for each application server. If you want to use DB-Reconnect for any application server, the Central Instance must also be configured to reconnect. Make sure that the data remains consistent under all conditions. If the Central Instance fails, all application servers are always restarted.

OS490 Optional Step:

If you want to use the DB-Reconnect functionality:

Add entries to:

- /sapmnt/<SID>/profile/<SID>_DVEBMGS<INSTANCENR>
- Instance Profiles of all application servers that use DB-Reconnect

When configuring the DBRECONNECT feature follow the appropriate OSS notes 109036, 98051 and 24806.

For example:

```
rsdb/reco_trials = 15
rsdb/reco_sleep_time = 60
rsdb/reco_sosw_for_db = off (based on OSS #109036)
rsdb/reco_sync_all_server = on
```

OS500 Optional Step:

If you want to use the DB-Reconnect functionality:

Make sure that you configured CIRESTART=0 as described in Step IS410 the default setting of CIRESTART is 0 starting with SGeSAP Release 3.01.

OS510 Optional Step:

If you want to use the DB-Reconnect functionality:

Use ASTREAT[] values that are <4 for any application server that uses DB-Reconnect.

The ASTREAT[] array defines how the application servers are treated during switchover.

ASHOST[] and ASTREAT[] values with the same index belong to the same application server:

```
ASTREAT [X] = Y
```

where —

X specifies the corresponding application server.

Y is an integer value that defines the way the application server is treated.

Y = 0 means that the application server is not affected by any changes that happen to the package status. At the moment there are three triggers you can use to customize the way the scripts affect the application server. They are:

- ACTIVE—Add 1 to Y if the application server should be started automatically during the startup of a package.

NOTE

Not during failover – only during package startup.

Under normal circumstances all application servers should be configured to be active.

Under special conditions, for example, MC/ServiceGuard maintenance, you can stop an application server manually. In situations such as these, also configure the application server to be inactive. For now you need to change the scripts manually.

To deactivate an application server, you have to decrease his `ASTREAT []` value by one in the `sap.conf` files on all nodes. This ensures the number is even. Otherwise a switchover would accidentally cause an unwanted startup of the previously stopped application server.

To reactivate the application server Instance, increase the `ASTREAT []` value on all nodes.

If you deactivate an application server by changing the `sap.conf` files, the application server Instance does *not* stop automatically or immediately. When you deactivate this way, the Instance is not (re-)started automatically by MC/ServiceGuard. You can still have manual startups and MC/ServiceGuard-triggered shutdowns of the Instance.

- **FINAL STOP**—Add 2 to `Y` if the application server should automatically be stopped after initiating `cmhaltpkg`.
- **RESTART**—Add 4 to `Y` if the Application Instance processes should be stopped and restarted automatically if a DB-switchover caused by a failure takes place.

Use this for extremely critical environments, in which safety is the most important concern. The restart does not take place if the application server is configured to be inactive.

Table 2-9 provides a summary of values for Y:

Table 2-9 Y Values for ASTREAT[]

Value	Action
0	This application server configured is never touched by the Extension Scripts. <i>Note:</i> This means the application server has to be configured for DB-RECONNECT.
1	This application server is an active server. This means the scripts try to start the application server if the package is coming up. <i>Note:</i> Configure the application server to use DB-RECONNECT.
2	This application server stops if the ci-package halts. All startups have to be done manually. <i>Note:</i> This means the application server has to be configured for DB-RECONNECT.
3	This application server always starts if the package is coming up. It stops if the package is halted manually. In case of a failover the application server reconnects by itself. Use this value if you want to use RECONNECT and additionally want to control the SAP R/3 system as a whole by using the package commands. <i>Note:</i> This means the application server has to be configured to use DB-RECONNECT.
4,5 or 6	Do not configure.
7	This application server stops and starts in case of a switchover (no DB-RECONNECT). It stops if the ci-package halts. It always starts if the package is coming up.

The Central Instance is treated the same as any of the additional application servers. Use CIHOST and CIRESTART instead of ASHOST [] and ASTREAT [] .

Table 2-10 provides an overview of the possible configurations and the corresponding actions.

Table 2-10 **ASTREAT Configuration Options**

ASTREAT	Restart	Final_ Stop	Active	Runpkg DB	Haltpkg DB	Runpkg CI	Haltpkg CI
0	0	0	0	nop	nop	nop	nop
1	0	0	1	CI-start	nopstart	start	nop
2	0	1	0	nop	nop	nop	stop
3	0	1	1	CI-start	nop	start	stop
4	1	0	0	stop	nop	stop	nop
5	1	0	1	stop CI-start	nop	stop start	nop
6	1	1	0	stop	nop	stop	stop
7	1	1	1	stop CI-start	nop	stop start	stop

nop = no operation
st_if_nl = start if *not* on local host

The Central Instance should always be ACTIVE and configured for FINAL_STOP.

The RECONNECT behavior can be specified.

Advanced Options of sap.conf

In `/etc/cmcluster` there is a file called `customer.functions`. Do not change the `sap.functions`. The `customer.functions` templates that are delivered with SGeSAP work with additional parameters in the second part of `sap.conf`. Use the `customer.functions` as needed.

OS520 Optional Step:

SAPROUTER is a sample additional program that always starts on the Central Instance host.

To start an `saprouter` on the CI host automatically, specify:

```
SAPROUTER_START=1
```

You can also provide an option string that is passed to the `saprouter`, for example to set the path to the `saprouttab` file to reside on a shared volume.

Typically, set a maximum of one `saprouter` inside of the cluster. `saprouters` cannot share a service port. So make sure that you use different ones, if a failover to the same node is possible for the `saprouters`.

OS530 Optional Step:

Sometimes, when a failover occurs, you want to stop other SAP R/3 Instances running on backup nodes.

This is useful if you have to free up limited resources before you are able to start the failed Instance again. You have to configure two groups of arrays to use this functionality:

- The first group consists of the arrays `RMNR[*]`, `RMADM[*]` and `RMDEP[*]`.
 - Specify the Instance-IDs and the name of the System Administrators of each instance that shall be stopped in the arrays `RMNR[*]` and `RMADM[*]`.
 - For Dialog-Instances and other additional application servers also specify:

```
RMDEP[*] = -1
```

This means, that there are no instances which depend on the services offered by them.

- For Central Instances, specify the index into the second group of arrays in `RMDEP[*]`.

If you allow stopping of Central Instances, you have to specify the list of related additional application servers, too. The application servers are stopped first. It does not matter on which host they are running.

- The second group of arrays consists of `RMDEPNR[*]` and `RMDEPHOST[*]`.

Specify the application servers depending on the Central Instances which are allowed to be stopped. Fill in the Instance-IDs and the local

Step by Step Installation Guide

SGeSAP Files Configuration

hostnames on which the application servers run. The list of application servers which belong to the same SAP R/3 system start at the index which is specified in the `RMDEP[*]` entry of their Central Instance and the list continues until an `RMDEPNR[*] = -1` is found. Never forget this entry.

The `RMNR[*]` specified instances are halted only if they are running on the host you currently fail over to. This functionality never halts any MC/ServiceGuard package. It is possible to stop a Central Instance that is using MC/ServiceGuard packages, but the associated package remains in the running state. Application servers belonging to the same SAP R/3 system as the failing package do not need to be specified here. They stop automatically.

OS540 Optional Step:

If there is a special demand to use values different from the default, it is possible to redefine some of the following values:

- name of SAP R/3 System Administration User
- home directory of SIDADM
- name of the database package
- name of the Central Instance package
- name of Oracle Administrator
- SID of Oracle Database
- database home directory
- path of the startdb logfile
- path of the stopdb logfile
- path of the SAP R/3 default profile
- path of the SAP R/3 global transport directory
- name of the database listener

OS550 Optional Step:

If you want to use a database different from ORACLE or INFORMIX:

Currently only ORACLE and INFORMIX databases are supported, but there is a mechanism to easily integrate another database, provided its handling is not significantly different.

Specify a database vendor different from ORACLE and INFORMIX in

the DB parameter of `sap.conf`: `DB=<VENDOR>`

SGeSAP Integration calls the script functions:

```
start_<VENDOR>_db, stop_<VENDOR>_db, and stop_<VENDOR>_as
```

These functions start and stop the database as well as stop an application server gracefully. The third function is needed, in case there are specific tasks required to stop an application server after the underlying database has failed. These functions must be specified in the database dependent part of the `customer.functions` file.

SGeSAP configuration for SCM High Availability - APO/LiveCache

Configuring a HA LiveCache system requires some additional task that differ from a standard R/3 configuration. The APO system itself can be treated as an additional SAP R/3 system and therefore can be configured in SGeSAP using the standard installation steps.

For the setup of LiveCache on SGeSAP it is required to follow the attached special installation steps for LiveCache. These steps are based on the SAP documentation “LiveCache and High Availability on Unix platforms” and “Backup and recovery for APO 3.0”. It is recommended to read these documents before proceeding with the SGeSAP configuration of LiveCache.

It is required to follow the listed SAP OSS notes appropriately: 383312, 383468, 154997, 111865

OS552 LC Installation Step:

Customization control/config files.

The LiveCache is configured similar to a standard db-package. Execute the following steps for the configuration of a LiveCache package in appropriate order:

- Follow installation step IS370 and IS375.cp
/opt/cmcluster/sap/SID/saplc.cntl
/etc/cmcluster/<SID>/saplc.cntlcp
/opt/cmcluster/sap/SID/saplc.mon
/etc/cmcluster/<SID>/saplc.mon
- Customize the package control and configuration files according to IS380

OS553 LC Installation Step:

Customize sap.conf. Customize sap.conf appropriately. Only the following parameters in sap.conf are relevant to set for a LiveCache package:

- DB=LC
- DBRELOC=<reloc1c>
- TRANRELOC=<reloctrans>
- CLEANUP_POLICY=strict-normal

The strict option should be set in order to guarantee the LiveCache package gets all resources required on the adoptive node. Remember that this option can crash a system running on the adoptive node.

- RMDEP[*] array Systems – an additional APO Application server or a complete development system – that are running on the adoptive node need to be shut down before the LiveCache package can start. Specify systems running on the adoptive node using the RMDEP[*] array. Make sure the remsh access to this system is working fine:

```
RMNR [0] =<SYSTEMNUMBER>; RMDEP [0] =0; RMADM [0] =<SID>adm
```

```
RMDEPNR [0] =<SYSTEMNUMBER>; RMDEPHOST [0] =<adoptive_node>; RMDEP  
PLATFORM=HP-UX
```

```
RMDEPNR [1] =-1
```

See the examples given in sap.conf for more details on how to set this array.

OS554 LC Installation Step:

Set up for LiveCache Monitor

SAP recommends the use of a monitor in order to test the availability of LiveCache periodically. This monitor periodically checks the availability of the LiveCache system. If the monitor recognizes the LiveCache to be unavailable, it will try to restart LiveCache several times on the node it is currently running. It will switch the package when the LiveCache is still unavailable.

The monitor program is shipped with SGeSAP in the saplc.mon file. Follow the attached steps on how to integrate the monitor into the LiveCache package configuration. The monitor actually runs as a Service attached to a MC/ServiceGuard package. Find additional information about services in MC/ServiceGuard environments in the

“Managing MC/ServiceGuard” documentation.

Service configuration

Configure Service]<SID> for the monitor in file
/etc/cmcluster/<SID>/lc.conf:

- SERVICE_NAME <SID> SERVICE_FAIL_FAST_ENABLED
NOSERVICE_HALT_TIMEOUT 300
- Configure Service Name <SID> for the monitor in file
/etc/cmcluster/<SID>/lc.conf:SERVICE_NAME [0] = "<SID>
"SERVICE_CMD [0] = "/etc/cmcluster/LHP/saplc.mon
monitor"SERVICE_RESTART [0] = "-r 2"

Refer to the comments in the conf and cntl files to get more information about setting the parameters appropriately. The examples show reasonable values.

The SERVICE_RESTART [0] parameter sets the number of restarts MC/ServiceGuard will try to restart the service on the node it was running. When this number is met the MC/ServiceGuard will switch the package to the adoptive node. The number of the actual service restarts can be found when executing “cmviewcl -v”.

The monitor can be paused/resumed by placing/removing a lc.debug file in /etc/cmcluster/<SID>. This is helpful when configuring LiveCache and running the LiveCache package in debug-mode.

The file /sapdb/LHP/db/sap/lccluster is used by the LiveCache program lcinit to indicate that the LiveCache is running in a cluster environment.

OS555 LC Installation Step:

Configure LiveCache

LiveCache need to be aware it is running in a cluster environment. Therefore several configurations are required.

Inform LiveCache about the presence of MC/ServiceGuard:

- Create /sapdb/LHP/db/sap/lccluster file :
ln -s /etc/cmcluster/LHP/saplc.mon
/sapdb/LHP/db/sap/lccluster

LiveCache Configuration

Volume groups and filesystems

The LiveCache dev-spaces (data, log, system) need to be placed on the MC/ServiceGuard-shared disks. Usually these dev-spaces are placed on raw devices. Additionally /sapdb/programs and /sapdb/data need to be on shared devices. See attached table for more information about the volume group needed for the LiveCache:

Table 2-11

LiveCache package	VG Name	lv Names	Mount Point Usage	Device Number
lc<SID>	vg<SID>	rlvLHP1	raw – data 1	06
		rlvLHP2	raw – data 2	
		rlvLHP3	raw – data 3	
		rlvLOG	raw - log	
		lvSYS	/sapdb/<SID>/ dbsys	
		lvSAPDB	/sapdb	

Cluster node synchronization

- Synchronize /usr/spool/sql recursively to the adoptive node:
- Synchronize the following ports to /etc/services:
 - sql307200/tcp
 - sql67210/tcp
- Synchronize /etc/passwd with the <SID>adm
- Synchronize /etc/group with group “sapsys”

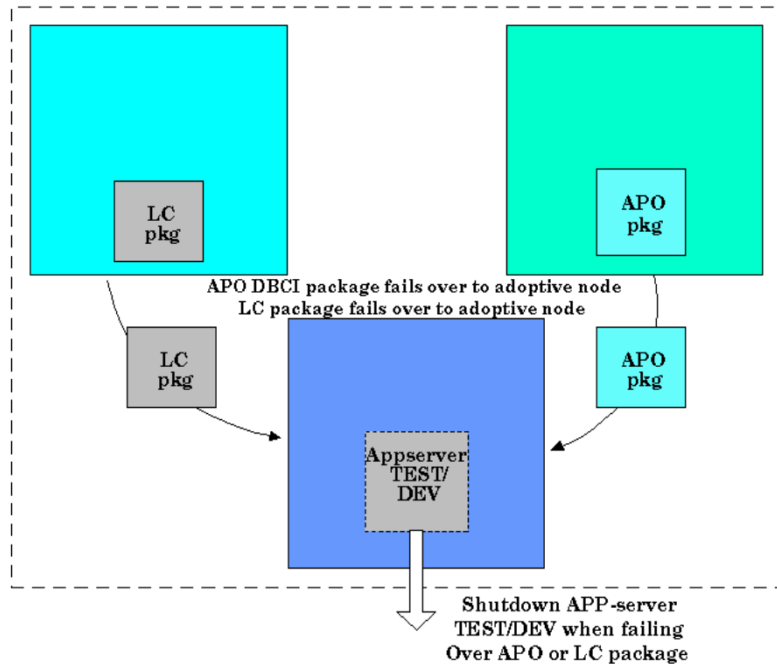
IS556 LC Installation Step:

Configure LiveCache connections within APO:

Running LiveCache within a MC/ServiceGuard cluster package means that the LiveCache instance is now configured for the relocatable IP of the package. This configuration needs to be adopted in the APO system that connects to this LiveCache.

Figure 2-1 shows an example for configuring “LCA” to “reloc8”.

Figure 2-1 **LiveCache Connections Within APO**



Run SAP transaction “lc10” and configure the logical LiveCache names “LCA” and “LCD” to listen to the relocatable IP, which the LiveCache-package is running under.

IS557 LC Installation Step:

Configure XUSER file in the APO-user homedirectory

The XUSER file in the home directory of the APO-user is the file that holds the connection information and grant information for a client connecting to LiveCache. The XUSER content needs to be adopted to the relocatable IP the LiveCache in running on.

Run the following commands on every APO-Application server:

```
# su - <SID>adm(<SID> is there the SID of the APO-system)
# dbmcli -d <LIVECACHE_SID> -n <LC_relocatable_IP> \
  -us control,control
```

To verify that the relocatable IP is in the XUSER file, run the following command as <SID>adm of the APO user:

```
# dbmcli -ux SAPR3,SAP -ul
```

You should get output like the following for a relocatable hostname “reloc8” and a physical hostname “hpcc072”:

```
OK
DEFAULT SAPR3
1LHPhpcc072 CONTROL
1LHPreloc8 CONTROL
```

IS558 LC Installation Step:

Verify LiveCache package setup

Verify your LiveCache installation following the attached steps:

- Touch a /etc/cmcluster/<SID>/debug file and /etc/cmcluster/<SID>/lc.debug file and startup the LiveCache package
- Manually start LiveCache as APO <SID>adm using transaction LC10
- Shutdown LiveCache package
- Remove both debug files
- Startup/shutdown the LiveCache package and check the logfile in /etc/cmcluster/<SID> for consistency

HP Somersault Integration with SGeSAP

HP Somersault was introduced with SAP R/3 Release 4.6B to protect the SPOF enqueue server. HP Somersault is only supported in an MC/ServiceGuard cluster environment. The installation steps described in this section must be followed when enabling HP Somersault with SGeSAP. Refer to Chapter 3, “SGeSAP Administration,” for additional information about HP Somersault in SGeSAP.

OS559 Optional Step:

Configure HPSOM to identify HP Somersault use.

Set `HPSOM=1` if HP Somersault is enabled and the SAP R/3 enqueue server is protected by HP Somersault. This is the only parameter to set within SGeSAP that enables HP Somersault. The default value is `HPSOM=0`.

The environment variable `HPSOM_HOME` sets the launching or shutdown of HP Somersault. The standard path `/var/opt/hpsom/<SID>` is automatically set in the SGeSAP environment. Refer to the hidden options section in `sap.conf` to set a non standard path for `HPSOM_HOME`.

When `HPSOM` is set to 1 there are additional parameters to be set in the SAP R/3 profile and standard HP Somersault must be successfully configured. Refer to your HP Somersault documentation.

MC/ServiceGuard Configuration

Logon as root on the primary host.

OS560 Optional Step:

Recommendation

Set `AUTO_VG_ACTIVATE=0` in `/etc/lvmrc`. Edit the `custom_vg_activation()` function if needed. Distribute the file to all cluster nodes.

OS570 Optional Step:

Recommendation

Set `AUTOSTART_CMCLD=1` in `/etc/rc.config.d/cmcluster`. Distribute the file to all cluster nodes.

IS580 Installation Step:

Create cluster configuration file in `/etc/cmcluster`:

```
cmquerycl -v -C /etc/cmcluster/sapr3.conf [-n <nodename> ...].
```

IS590 Installation Step:

Customize the cluster configuration file. Specify the volume groups listed in Step IS040. You might need to increase the timing parameters.

IS600 Installation Step:

Obsolete.

IS610 Installation Step:

Specify package name, node names and subnets in the package configuration files `db.conf` and `ci.conf` or in `dbci.conf`.

IS620 Installation Step:

If you implement the one package concept:

The standard package name is `dbci<SID>`. If you want to use different variables follow OS530.

If you implement the two package concept:

The standard package names are `db<SID>` and `ci<SID>`. If you want to use different names here, make sure that you also specified the

appropriate variables in Step OS530.

IS630 Installation Step:

Specify the package control script names in the package configuration files:

For one package concept specify in `/etc/cmcluster/<SID>/dbci.conf`:

```
RUN_SCRIPT /etc/cmcluster/<SID>/dbci.cnt1
HALT_SCRIPT /etc/cmcluster/<SID>/dbci.cnt1 \
```

For two package concept specify in `/etc/cmcluster/<SID>/db.conf`:

```
RUN_SCRIPT /etc/cmcluster/<SID>/db.cnt1
HALT_SCRIPT /etc/cmcluster/<SID>/db.cnt1
```

For two package concept specify in `/etc/cmcluster/<SID>/ci.conf`:

```
RUN_SCRIPT /etc/cmcluster/<SID>/ci.cnt1
HALT_SCRIPT /etc/cmcluster/<SID>/ci.cnt1
```

For each Application Server package specify in
`/etc/cmcluster/<SID>/ap<SID><ASNR>.conf`:

```
RUN_SCRIPT /etc/cmcluster/PRC/ap<SID><ASNR>.cnt1
HALT_SCRIPT /etc/cmcluster/PRC/ap<SID><ASNR>.cnt1
```

IS650 Installation Step:

Follow the directions in Steps IS040 and IS310 to define volume groups, logical volumes, IP addresses and subnets in the package control scripts `dbci.cnt1`, `db.cnt1`, `ci.cnt1` and `ap<SID><ASNR>.cnt1` as appropriate.

Fill the `IP[*]` and `SUBNET[*]` array with the IP addresses and subnet addresses that the package is attached to. For example a 1-pkg control script `dbci.cnt1`:

```
IP[0]="15.27.218.238"
SUBNET[0]="15.27.216.0"
```

The filesystems you specify in the `LV[*]`, `FS[*]` and `FS_MOUNT_OPT[*]` array are not identical to the filesystems that are exported. For example:

```
LV[0]="/dev/vgDBPRD/lvoracle"; FS[0]="/oracle/PRD";
FS_MOUNT_OPT[0]="-o rw"
LV[1]="/dev/vgDBPRD/lvora805"; FS[1]="/oracle/805_64";
FS_MOUNT_OPT[1]="-o rw"
```

```
LV[2]="/dev/vgDBPRD/lvsapmnt"; FS[2]="/export/sapmnt/PRD";  
FS_MOUNT_OPT[2]="-o rw"  
LV[3]="/dev/vgPRC/lvusrsap"; FS[3]="/usr/sap/PRD/DVEBMGS20";  
FS_MOUNT_OPT[3]="-o rw"  
LV[4]="/dev/vgTRANS/lvtrans"; FS[4]="/export/usr/sap/trans";  
FS_MOUNT_OPT[4]="-o rw"
```

The directories that are handled by the automounter must be exported by the scripts if they are part of the packages, for example, the `/export/sapmnt/<SID>` directory is usually exported. Exported directories can usually be found beneath the special export directory `/export`. The directories to be exported are specified including their export options, using the `XFS[*]` array of `db.cnt1` or `dbci.cnt1`.

The transport directory is also part of the package in standard installations. In those cases it also has to be mentioned here. On INFORMIX installations without local database executables, access must be granted to the database filesystem too.

Only allow access from machines inside of the cluster that are configured as backup hosts and additional application servers of this specific SAP R/3 system. Make sure to allow access for all addresses that the servers use so they can reach the directories. Allowing the right addresses is essential if your hosts use more than one LAN card.

For example:

```
XFS[0]="-o access=hpcc006:hpcc008:hpcc071:hpccdemo  
/export/sapmnt/<SID>"  
XFS[1]="-o access=hpcc006:hpcc008:hpcc071:hpccdemo  
/export/usr/sap/trans"  
XFS[2]="-o access=hpcc006:hpcc008:hpcc071:hpccdemo  
/export/informix/<SID>"  
"
```

IS660 Installation Step:

To enable the SAP R/3 specific scripts change the `customer_defined_commands` sections of the package control script(s):

For the one package concept:

```
function customer_defined_run_cmds  
{  
/etc/cmcluster/<SID>/sapdbci.cnt1 startDBCI <SID>  
test_return 51  
}
```

Step by Step Installation Guide

MC/ServiceGuard Configuration

```
function customer_defined_halt_cmds
{
/etc/cmcluster/<SID>/sapdbci.cnt1 stopDBCI <SID>
test_return 52
}
```

For a two package concept:

Use the commands stopDB, startDB in file db.cnt1 or stopCI, startCI in ci.cnt1 respectively.

For an Application Server package in ap<SID><ASNR>.cnt1:

```
{
function customer_defined_run_cmds
# ADD customer defined run commands.
/etc/cmcluster/<SID>/sapap.cnt1 startAP <ASNR> <SID> <SIDADM>
test_return 51
}
function customer_defined_halt_cmds
{
# ADD customer defined halt commands.
/etc/cmcluster/<SID>/sapap.cnt1 startAP <ASNR> <SID> <SIDADM>
test_return 52
}
```

OR680 Oracle Database Step:

If you use ORACLE database:

Remove INFORMIX specific files from the cluster directory:

```
rm /etc/cmcluster/customer.sh /etc/cmcluster/customer.csh
```

IR690 Informix Database Step:

If you use INFORMIX database:

Move INFORMIX specific files into the INFORMIX home directory. Read the command example carefully. Make sure you move the 'dot files'. For example:

```
mv /etc/cmcluster/customer.sh /home/informix/.customer.sh
mv /etc/cmcluster/customer.csh /home/informix/.customer.csh
chown informix:informix /home/informix/.customer.sh
chown informix:informix /home/informix/.customer.csh
```

IS700 Installation Step:

Copy all integration files below `/etc/cmcluster` to the other cluster hosts using `rcp(1)`.

IS710 Installation Step:

Create the binary cluster configuration file and distribute it using `cmapplyconf(1m)`.

IS720 Installation Step:

If you plan to use an additional package for the transport directory:

Specify this package in the last position of the `cmapplyconf` command. Later, if you force a shutdown of the whole cluster with `cmhaltcl -f`, the package with the transport directory is the last one stopped. This prevents the transport directory from disappearing before all SAP R/3 systems in the cluster have completed their shutdown.

Automounter Configuration

Repeat the steps in this section for each node of the cluster and for each external application server host.

Logon as root.

IS730 Installation Step:

Check that the Automounter is active. In `/etc/rc.config.d/nfsconf`, the section for the `autofs` configuration must look like:

```
AUTOMOUNT=1
AUTO_MASTER="/etc/auto_master"
AUTOMOUNT_OPTIONS="-f $AUTO_MASTER"
AUTOMOUNTD_OPTIONS=
AUTOFS=1
```

Older installations on HP-UX 10.x and installations without `autofs` require a slightly different syntax:

```
AUTOMOUNT=1
AUTO_MASTER="/etc/auto_master"
AUTO_OPTIONS="-f $AUTO_MASTER"
```

IS740 Installation Step:

Make sure that at least one NFS client daemon and one NFS server daemon is configured to run. This is required for the automounter to work. Check the listed variables in `/etc/rc.config.d/nfsconf`. They should be specified as greater or equal to one. For example:

```
NFS_CLIENT=1
NFS_SERVER=1
NUM_NFSD=4
NUM_NFSIOD=4
```

IS750 Installation Step:

Add the following line to your `/etc/auto_master` file:

```
/- /etc/auto.direct
```

IS760 Installation Step:

Create a file called `/etc/auto.direct`.

For each directory configured to be mounted below `/export` in Step IS040, add a line to this file. For example:

```
/usr/sap/trans <relocdbci_s>:/export/usr/sap/trans  
/sapmnt/<SID> <relocdbci_s>:/export/sapmnt/<SID>
```

For INFORMIX databases, add another entry. For example:

```
/informix/<SID> <relocdbci_s>:/export/informix/<SID>
```

For the two package concept, typically the relocatable IP address of the database package `<relocdb_s>` is used.

When configuring AUTOFS the automounter map file `/etc/auto.direct` must be executable. Make sure to set the appropriate permissions of `/etc/auto.direct` to 644.

IS770 Installation Step:

Restart the automounter with:

```
/sbin/init.d/nfs.client stop
```

and

```
/sbin/init.d/nfs.client start
```

NOTE

Never kill the automount process. Always use `nfs.client` to stop or start it.

Never stop the NFS client while the automounter directories are still in use by some processes. If `nfs.client stop` reports that some filesystems could not be unmounted, the automounter may refuse to handle them after `nfs.client start`.

After configuring all automounters by performing the above mentioned steps on all hosts, verify that the setup works correctly to this point. Do not continue with the following application dependent sections without doing this! Perform the tests using the debug-switch

OS780 Optional Step:

Create a file called `/etc/cmcluster/debug` on all nodes in the cluster.

Step by Step Installation Guide

Automounter Configuration

This enables you to start the MC/ServiceGuard packages without running the SAP R/3 specific steps.

Reboot all machines.

Start the cluster and run all packages. This might be done automatically. If everything is working correctly, the log-file(s) in `/etc/cmcluster/<SID>` will indicate successful package startup.

On the primary host(s) it should be possible to run the database and the Central Instance. Be sure to start the listener as `ora<sid>` before trying to bring up an ORACLE database. All application servers should work after starting them manually.

Now, that everything works stop all application servers. Stop the Central Instance and the database. Stop the listener if needed. Do not stop the packages yet.

Database Configuration

This section describes the following:

- Additional Steps for Informix
- Additional Steps for Oracle

Additional Steps for Informix

Logon as root to the primary host of the database where the package is running in debug mode.

IR790 Informix Database Step:

Perform the following steps as an INFORMIX user:

```
su - informix
```

Comment out the `remsh` sections in the files called `.dbenv.csh` and `.dbenv.sh` in the home directory. If they are missing, check for alternative files with hostnames in them:

```
.dbenv_<local>.csh and .dbenv_<local>.sh.
```

The `remsh` section looks similar to:

```
# remsh <local> date >& /dev/null
# if ( $status <= 0 ) then
#   echo Remote login check to dbserver <local> failed.
#   echo Check if <local> is alive and file ...
#   echo on <local> contains an entry for ...
# endif
```

Copy the two `.dbenv` files over to the INFORMIX home directory on all cluster nodes, all external application server hosts, and MC/ServiceGuard by using `ftp`.

If you do not do this, switching time increases dramatically; each `remsh` command to the original database host causes delay if the node is not available.

IR800 Informix Database Step:

In Step IR690 you copied two files to the INFORMIX home directory of the primary node. At this time, still as an INFORMIX user, customize

Step by Step Installation Guide

Database Configuration

these files by replacing the string `relocdb` with your individual `<relocdbci>` (or `<relocdb>` in case of two package concept). For example:

`.customer.sh`:

```
##### .customer.sh #####
DBRELOC=relocdb; export DBRELOC
ONCONFIG=onconfig.relocdb.${INFORMIX_DBID}; export ONCONFIG
TEMP=`netstat -i | awk ' {print $4 }' | sort -u | \
  awk -F. '$1~/'$DBRELOC'/ {if (length($1) == length("'"$DBRELOC'"")) {\
    {print 1;exit}}'`
if [ ${TEMP:=0} -eq 1 ]; then
    INFORMIXSERVER=${DBRELOC}${INFORMIX_DBID}shm
else
    INFORMIXSERVER=${DBRELOC}${INFORMIX_DBID}tcp
fi
export INFORMIXSERVER
#####
```

`.customer.csh`:

```
#####.customer.csh #####
setenv DBRELOC relocdb
setenv ONCONFIG onconfig.${DBRELOC}.${INFORMIX_DBID}
setenv TEMP `netstat -i | awk ' {print $4 }' | \
  sort -u | awk -F. '$1~/'$DBRELOC'/ \
  {if (length ($1) == length("'"$DBRELOC'"")) {print 1;exit}}'`
if ( $TEMP == 1 ) then
setenv INFORMIXSERVER ${DBRELOC}${INFORMIX_DBID}shm
else
setenv INFORMIXSERVER ${DBRELOC}${INFORMIX_DBID}tcp
endif
#####
```

Copy the `.customer.sh` and `.customer.csh` to all INFORMIX home directories on all nodes including the application servers outside of the cluster.

IR810 Informix Database Step:

Perform the following steps as `<sid>adm` user:

```
su - <sid>adm
```

Copy the files that were manipulated in Steps IS760 and IS770 to the home directory:

```
cp /home/informix/.dbenv* ~
cp /home/informix/.customer.* ~
```

Copy the files over to the home directories of <sid>adm on all cluster nodes and all external application server hosts.

IR820 Informix Database Step:

Perform the following steps as an INFORMIX user.

Modify the parameters DBSERVERNAME and DBSERVERALIAS in the ONCONFIG file:

```
/informix/<SID>/etc/onconfig.<primary>.<DBID>.
```

The parameter default to:

```
<local><INFORMIX_DBID><shm|tcp>
```

Set them to:

```
<relocdb><INFORMIX_DBID><shm|tcp>
```

For example:

```
DBSERVERNAME relocdb<sid>shm  
DBSERVERALIASES relocdb<sid>tcp
```

IR830 Informix Database Step:

Rename the ONCONFIG file to:

```
/informix/<SID>/etc/onconfig.<relocdb>.<DBID>
```

IR840 Informix Database Step:

Add a line with rel-IP-name of the database package to the file:

```
/informix/<SID>/etc/sqlhosts.soc.
```

After SAP R/3 installation this file will be similar to:

```
demo_on onipcshm on_hostname on_servername  
demo_se seipcpip se_hostname sqlexec  
<local><sid>shm onipcshm <local> sapinf<SID>  
<local><sid>tcp onsoctcp <local> sapinf<SID>
```

Change the <local> entries to <relocdb> of the database package. For example:

```
<relocdb><sid>shm onipcshm <relocdb> sapinf<SID>  
<relocdb><sid>tcp onsoctcp <relocdb> sapinf<SID>
```

IR845 Informix Database Step:

If you are using SAP R/3 version 4.0x:

Create additional entries in the `INSTVERS` table of SAP R/3. Print out the content of the table. Note the ID column. This is the primary key of the table. All new entries need a unique key.

Review the `STATUS` column. There should be one row containing a 0 as status. This indicates the current, successful installation.

Copy this row using a unique key as well as secondary hostnames for `HOSTNAME` and `DBHOSTNAME`.

For the two package concept or more than two cluster nodes include entries for all possible combinations. For example, you need two additional entries for a two-package concept with two cluster nodes to reflect that the two packages can run on different nodes. Refer to Step OR915 for more information.

Additional Steps for Oracle

Logon as root to the primary host of the database where the package is running in debug mode.

OR850 Oracle Database Step:

Perform the following step as `<sid>adm`.

To ensure that the database recovers itself after a crash, change the `startdb` script in the `/sapmnt/<SID>/exe` directory.

Change the lines:

```
#
# Startup the database without changing the ARCHIVELOG state
#
echo "connect internal;" > $SRVMGRDBA_CMD_FILE
echo "startup;" >> $SRVMGRDBA_CMD_FILE
echo "exit;" >> $SRVMGRDBA_CMD_FILE
eval $SRVMGRDBA command=@$SRVMGRDBA_CMD_FILE >> $LOG 2>&1
```

to:

```
#
# Startup the database without changing the ARCHIVELOG state
#
echo "connect internal;" > $SRVMGRDBA_CMD_FILE
```



```
echo "startup mount;" >> $SRVMGRDBA_CMD_FILE
echo "recover database;" >> $SRVMGRDBA_CMD_FILE
echo "alter database open;" >> $SRVMGRDBA_CMD_FILE
echo "exit;" >> $SRVMGRDBA_CMD_FILE
eval $SRVMGRDBA command=@$SRVMGRDBA_CMD_FILE >> $LOG 2>&1
```

OR860 Oracle Database Step:

If you are using SAP R/3 version 3.x:

Perform the following steps as ora<sid>.

Configure the listener to listen on the relocatable name of the database package. To do this, change all references from <local> to the relocatable name <relocdb> in the files. For example:

```
/usr/sap/trans/listener.ora
/usr/sap/trans/tnsnames.ora
```

OR870 Oracle Database Step:

If you are using SAP R/3 version 4.0A or higher:

Perform the following steps as ora<sid>.

Configure the listener to listen on the relocatable name of the database package. To do this, change all references from <local> to the relocatable name <relocdb> in the files on the shared volume group. Be careful if these files were customized after SAP R/3 installation. For example:

```
/oracle/<SID>/network/admin/listener.ora
/oracle/<SID>/network/admin/tnsnames.ora
```

OR880 Oracle Database Step:

If you are using SAP R/3 version 4.0A or higher:

Copy /oracle/<SID>/network/admin/tnsnames.ora to all additional application server hosts. Be careful if these files were customized after SAP R/3 installation.

OR885 Oracle Database Step:

If you are using SAP R/3 version 4.0A or higher and you are using the two package concept:

As with all application servers, the Central Instance needs a subset of

Step by Step Installation Guide

Database Configuration

the files in `/oracle/<SID>` to work properly. Verify if your Central Instance can run on any host without an application server installed. These hosts need a local copy of the filetree.

If you have access to the SAP R/3 install disks you can get the files from the SAP R/3 archive `ORACLI.CAR` or `OCL<oclrrel>.CAR` on the kernel CD. The extractor program `CAR` functions the same as the `tar (1m)` command. It is shipped with SAP R/3.

If the Central Instance was installed on a node different from the database node, you can get the files there. Or you can use any additional application server host as source. Copy the filetree, for example, by using the procedure explained in Step IS260. The filetree should contain the following subdirectories: `lib`, `network`, `ocommon`, and `rdbms`.

Do not copy `/oracle/<SID>` from the database host or you will copy much more than you want. Refer to SAP R/3 OSS Note 180430 for more details.

OR890 Oracle Database Step:

If you are using SAP R/3 version 3.x:

If existing locally, copy `/etc/listener.ora` to all adoptive nodes in the cluster.

OS900 Optional Step:

If you are using SAP R/3 version 3.x and plan to use more than one SAP R/3 system inside of your cluster:

It is possible that more than one database is running on the same node. Even though one listener is capable of serving many database instances problems can occur in switchover environments because needed filesystems may not be available at the startup time of the listener. Use a dedicated listener process for each database.

You can use the standard `listener.ora` file that is created during the installation of *one* SAP R/3 system `<SID1>` as a template. Double its contents. For example:

```
cat listener.ora listener.ora >listener.ora.new
mv listener.ora.new listener.ora
```

Now the file consists of two identical parts.

- In the first part of the file:

Replace each occurrence of the word `LISTENER` by a new listener name. You can choose what suits your needs, but it is recommended to use the syntax `LISTENER<SID1>`:

```
( host = <relocdb_1> )
```

Change nothing.

- In the second part of the file: Replace each occurrence of the word `LISTENER` by a new listener name different from the one chosen above. For example, use `LISTENER<SID2>` if `<SID2>` is the SID of the second SAP R/3 system. Replace any other occurrence of `<SID1>` by `<SID2>`. The line should be modified to contain the appropriate relocatable address belonging to the database package (`db` or `dbci`) of the second system. For example:

```
( host = <relocdb_2> )
```

In the line:

```
( port = 1527 )
```

a new previously unused port should be placed. For example:

```
( port = 1528 )
```

Adapt the `(host=...)` and `(port=...)` lines corresponding to the values you have chosen in the `listener.ora` file.

Test your setup by starting the listeners as `ora<sid1/2>`:

```
lsnrctl start LISTENER<SID1/2>
```

Create an `/etc/services` entry for the new port you specified above. Use `tlisrv<SID2>` as service name. The name is not needed anyhow. This entry has to be made on *all* hosts that run an instance that belongs to the system. This includes all external application server hosts outside of the cluster.

OS910 Optional Step:

If you use the two package concept:

Set the optional parameter `SQLNET.EXPIRE_TIME` in `sqlnet.ora` to a reasonable value in order to take advantage of the Dead Connection Detection feature of ORACLE. The parameter file `sqlnet.ora` resides either in `/usr/sap/trans` or in `/oracle/<SID>/network/admin`.

The value of `SQLNET.EXPIRE_TIME` determines how often (in seconds) SQL*Net sends a probe to verify that a client-server connection is still active. If the Central Instance switches, the application servers may

Step by Step Installation Guide

Database Configuration

crash, thereby leaving shadow processes running on the database host. While the CI package cleans up the application server hosts, it does not touch the ORACLE shadow processes running on the database host. Remove them, because their number increases with every CI package switch. After an application server crash, a connection to the database shadow process may be left open indefinitely. If the `SQLNET.EXPIRE_TIME` parameter is specified, SQL*Net sends a probe periodically to determine whether there is an invalid connection that should be terminated. It finds the dead connections and returns an error, causing the server process to exit.

OR915 Oracle Step:

If you are using SAP R/3 version 4.0A or higher:

Perform the following step as `<ora>sid`. Start the listener if it is not already running:

```
lsnrctl start <listenername>
```

Start `svrmgrl` and type:

```
> connect internal;
> startup;
> select * from sapr3.instvers;
```

The output is similar to:

ID	R3RELEASE	SAPSID	SY OPSYS	HOSTNAME	DBHOSTNA	DBTYPE	DATUM			
TIME	A	STATUS	TOOL	INFO						
N	1	40B	<SID>	<INSTNR>	HP-UX	<pri_host>	<pri_host>	ORACLE	19980527	094221
	2	R3SETUP	[RFCRADDBDIF_IND_IND]	CRFCJOB-RADDBDIF						
N	0	40B	<SID>	<INSTNR>	HP-UX	<pri_host>	<pri_host>	ORACLE	19980527	094219
	1	R3SETUP	[RFCRADDBDIF_IND_IND]	CRFCJOB-RADDBDIF						
N	2	40B	<SID>	<INSTNR>	HP-UX	<pri_host>	<pri_host>	ORACLE	19980527	094250
	3	R3SETUP	[RFCRADDBDIF_IND_IND]	CRFCJOB-RADDBDIF						
N	3	40B	<SID>	<INSTNR>	HP-UX	<pri_host>	<pri_host>	ORACLE	19980527	094252
	0	R3SETUP	[RFCRADDBDIF_IND_IND]	CRFCJOB-RADDBDIF						

- Review the `ID` column. This is the primary key of the table. All new entries need a unique key.
- Review the `STATUS` column. There should be one row containing a 0 as status. This indicates the current, successful installation. Copy this row using a unique key as well as secondary hostnames for

HOSTNAME and DBHOSTNAME.

If you use one package concept with a two node cluster, enter the following line in the svrmgr1 session and you are done:

```
> insert into sapr3.instvers values
(<key>,'40B',<SID>,<INSTNR>,'HP-UX',<sec_host>,
<sec_host>,'ORACLE','19980527','094252','N','0','MCSG',
'[RFCRADDBDIF_IND_IND]CRFCJOB-RADDBDIF');
```

For two package concept or more cluster nodes add entries for all possible combinations. For example, you need two additional entries for a two-package concept with two cluster nodes to reflect that the packages can run on different nodes:

```
> insert into sapr3.instvers values
(<key>+1,'40B',<SID>,<INSTNR>,'HP-UX',<pri_host>,
<sec_host>,'ORACLE','19980527','094252','N','0','MCSG',
'[RFCRADDBDIF_IND_IND]CRFCJOB-RADDBDIF');
> insert into sapr3.instvers values
(<key>+2,'40B',<SID>,<INSTNR>,'HP-UX',<sec_host>,
<pri_host>,'ORACLE','19980527','094252','N','0','MCSG',
'[RFCRADDBDIF_IND_IND]CRFCJOB-RADDBDIF');
```

Stop the database before continuing.

OR916 Oracle Step:

Additional steps for Oracle 8.1.x DB:

For the SGeSAP <=V3.02 the sapbackup directory needs to be in \$ORACLE_HOME/sapbackup. If you are using an Oracle 8.1.x DB the \$ORACLE_HOME directory is set to /oracle/<SID>/8.1.x whereas the sapbackup directory is still placed /oracle/<SID>/sapbackup. It's best practice to create the symbolic link \$ORACLE_HOME/sapbackup -> /oracle/<SID>/sapbackup.

OR917 Oracle Step:

additional steps for Application server and Oracle 8.1.x DB

If you run an Oracle DB >= 8.1.x and install additional R/3 Application Server follow the procedure described in OSS303238.

SAP R/3 System Configuration

Logon as <sid>adm on the primary node on which the Central Instance has been installed. The appropriate MC/ServiceGuard package should still run on this host in debug mode.

IS920 Installation Step:

Change into the profile directory by typing the alias:

```
cdpro
```

In the DEFAULT.PFL change the following entries and replace the hostname with the relocatable name. For example:

```
SAPDBHOST = <relocdb>  
rdisp/mshost = <relocci>  
rdisp/sna_gateway = <relocci>  
rdisp/vbname = <relocci>_<SID>_<instnr>  
rdisp/enqname = <relocci>_<SID>_<instnr>  
rdisp/btcname = <relocci>_<SID>_<instnr>  
rslg/collect_daemon/host = <relocci>
```

The following parameters are only necessary if an application server is installed on the adoptive node. For example:

```
rslg/send_daemon/listen port  
rslg/collect_daemon/listen port  
rslg/collect_daemon/talk port
```

NOTE

If you are using the one package concept use <relocdbci> instead of <relocdb> or <relocci>. Never use server LAN addresses for any of these parameters.

IS930 Installation Step:

In the <SID>_DVEBMGS<INSTNR> profile add or modify the following entries:

```
SAPLOCALHOST = <relocci>  
SAPLOCALHOSTFULL = <relocci>.<domain>
```

Because we are using relocatable IP addresses for SAP R/3 services, the SAP R/3 profile parameters `SAPLOCALHOST` and `SAPLOCALHOSTFULL` are very important for SAP R/3 operations.

NOTE

The profile name is sometimes extended by the hostname. You do not need to change this filename to include the relocatable hostname.

`SAPLOCALHOST` is set to the hostname per default at startup time and is used to build the SAP R/3 application server name:

```
<SAPLOCALHOST>_<SID>_<INSTNR>
```

This parameter represents the communication path inside an SAP R/3 system and between different SAP R/3 systems. `SAPLOCALHOSTFULL` is used for rfc-connections. Set it to the fully qualified hostname.

The application server name appears in the server list held by the message server, which contains all instances, hosts and services of the instances. The application server name or the hostname is also stored in some system tables on the database.

When using the default addressing scheme for the application server name, this name changes during a failover to another node because the backup node has a different hostname than the primary node. This creates a need to cleanup ABAP code similar to the processes required in version 1.0 of the SGeSAP Integration scripts.

When setting the `SAPLOCALHOST` parameter to the name associated with the relocatable IP address - which moves with the package to the other node - the application server name stays the same after a switchover. This enables all services which make use of the application server name to continue working after a switchover without any change.

The name of a relocatable IP address is also used to specify the `SAPDBHOST`. This is different from the standard setup with previous SGeSAP Integration versions.

IS940 Installation Step:

The parameter `SAPLOCALHOSTFULL` must be set even if you do not use DNS. In this case you should set it to the name without the domain name:

```
SAPLOCALHOSTFULL=<relocci>
```

IS950 Installation Step:

If you use a SAP R/3 release prior to 3.1H:

Add the following entry to DEFAULT.PFL:

```
gw/netstat =/usr/bin/netstat -in | cut -c 29-
```

IS960 Installation Step:

Starting with SAP R/3 releases 4.5 TPPARAM has been replaced by 2 new transport configuration files: DOMAIN.CFG and TP_DOMAIN_<SID>.PFL

Modify the dbhost entries appropriately.

In the file /usr/sap/trans/bin/TPPARAM, modify the dbhost entry as follows:

```
<SID>/dbhost = <relocdb>
```

IS970 Optional Step:

If you have already received your licenses from SAP R/3 install them on all the nodes where the Central Instance can start. Refer to the *MC/ServiceGuard Extension for SAP R/3 Release Notes* for further information on how to do this. The package comes up without a license, too. But certain restrictions apply to the SAP R/3 application. A warning is printed into the package log-file.

SAP R/3 will now be ready to run on all cluster nodes. Test the manual startup on all adoptive nodes.

IS980 Installation Step:

Switch the packages to the adoptive nodes with debug mode still enabled.

Start SAP R/3 as <sid>adm on the adoptive nodes manually. You should end up with SAP R/3 running on any cluster node.

IS982 Installation Step:

On all cluster nodes, remove the file /etc/cmcluster/debug.

IS984 Installation Step:

Connect with a SAPGUI. Import the changed SAP R/3 profiles within SAP R/3. The transaction is called RZ10.

After importing the profiles, check with rsparam in SE38 if the parameters SAPLOCALHOST and SAPLOCALHOSTFULL are correct. If you do

not import the profiles, the profiles within SAP R/3 can be edited by the SAP R/3 Administrator. The values listed from SAP R/3 will be wrong and, when saved, will overwrite the values which edited on the HP-UX level.

IS986 Installation Step:

The destination for print formatting, which is done by a Spool Work process, uses the application server name.

If the application server name stays the same because `SAPLOCALHOST` has been set to the relocatable name, after a switch no changes need to be done.

Printing works consistently. A print job in process at the time of the failure is canceled and needs to be reissued manually after the failover. To make a spooler highly available on the Central Instance server, set the destination of the printer to `<reloci>_<SID>_<INSTNR>` using transaction `SPAD`.

IS988 Installation Step:

Batch jobs can be scheduled to run on a particular instance.

You select a particular instance by its hostname at the time of job scheduling. The application server name and the hostname retrieved from the Message Server are stored in the Batch control tables `TBTCO`, `TBTCS...`

When the batch job is ready to run, the application server name is used to start it on the appropriate instance. When using `<reloci>` to build the application server name for the SAP R/3 instance, you do not need to change batch jobs, which are tied to the locality, after a switchover, even if the hostname which is also stored in the above tables differs.

IS990 Installation Step:

Within the SAP R/3 Computing Center Management System (CCMS) you can define operation modes for SAP R/3 instances.

An operation mode defines a resource configuration for the instances in your SAP R/3 system. It can be used to determine which instances are started and stopped, and how the individual services are allocated for each instance in the configuration.

An instance definition for a particular operation mode consist of the number and types of work processes as well as start and instance profiles (starting with version 3.0 the CCMS allows profile maintenance from

within SAP R/3).

When defining an instance for an operation mode, enter the hostname and the system number of the application server. By using <reloci> to fill in the hostname field, the instance is working under control of the CCMS after a failover without any change.

If an instance is running on the standby node in normal operation (and is stopped when switching over), the control console shows this instance to be down (for example, you will get a red node on a graphical display) after the switchover.

IS995 Installation Step:

Configure the frontend PCs to attach to <reloci>. Most of the time, this can be achieved by distributing a new `saplogon.ini` file to the windows directory.

Installation is complete.

The next step is to do some comprehensive switchover testing covering all possible failure scenarios. It is important that all relevant SAP R/3 application functionalities are tested on the switchover nodes. There exist several documents provided by HP or SAP R/3 that can guide you through this process.

3 SGeSAP Administration

This chapter describes administration of the special features of the MC/ServiceGuard Extension for SAP R/3 (SGeSAP). Topics presented in this chapter include:

- Switching SGeSAP Off and On
- SGeSAP Administration Issues
- SGeSAP Administration With HP Somersault Enabled

Refer to your SAP R/3 documentation for more information about managing SAP R/3 environments. For HP-UX system administration, refer to the HP-UX manual *Managing Systems and Workgroups* (Part Number B2355-90157), and for MC/ServiceGuard issues, refer to *Managing MC/ServiceGuard* (Part Number B3936-90026 for HP-UX 10.20 and B3936-90045 for HP-UX 11.0 and 11i).

Switching SGeSAP Off and On

This section provides a brief description of how to switch off SGeSAP. Your individual setup may require additional steps that are not included in this document. Contact your HP consultant for additional information.

Switching off the SGeSAP Integration means that the SAP R/3 system will not run on the relocatable IP address. If you are performing an SAP R/3 upgrade, make sure that your SAP R/3 does not require any relocatable IP addresses. Switching off the SGeSAP Integration is the first step in SAP R/3 upgrade related work. This ensures full support from SAP R/3 during the upgrade procedure.

Switch Off SGeSAP

Refer to the following sections for descriptions of how to switch of SGeSAP.

Switching off SGeSAP starts with the following processes. Perform the steps described no matter which database you are using:

- Halt Packages
- Turn on Debug Mode
- Start Packages

Perform the steps in one of these two sections, selecting the additional steps for ORACLE or INFORMIX as appropriate to your database.

- Additional steps for ORACLE
- Additional steps for INFORMIX

Follow the steps in these sections to set up SAP R/3 internally. Follow these steps no matter which database your system has. They are always the same, independent from the underlying database:

- Check SAP R/3
- Backup Files Again

Halt Packages

If the SGeSAP packages are still running, halt them using:

```
cmhaltpkg <package>
```

If you use the two package concept, stopping the database package also brings down the Central Instance package. If you use an additional package that contains the transport directory, you do not need to stop that package. This allows any other SAP R/3 system that shares the transport directory to not be affected and continue running.

Turn on Debug Mode

Turn on the debug mode of the SGeSAP Integration. Then MC/ServiceGuard can provide all the filesystems and relocatable IP addresses, but the scripts do not start the database or SAP R/3. If a file named `debug` exists in `/etc/cmcluster` SGeSAP Integration packages start and no SAP R/3 specific script starts. For example:

You must create a debug-file on all nodes of the cluster:

```
touch /etc/cmcluster/debug
```

Using this command prevents the SAP R/3 system from starting up on the other host as the result of a failure. This ensures that SAP R/3 does not run with incorrect profiles.

If you have more than one SAP R/3 system within the cluster you can also turn on debug mode for specific systems. If the file `/etc/cmcluster/<SID>/debug` exists, the additional scripts (start and stop SAP R/3) for this system are not started. Other systems are controlled as usual by MC/ServiceGuard.

Start Packages

Start the packages with:

```
cmrunpkg <package>
```

Remember that with the debug mode turned on, the Central Instance package of the two package concept is not triggered by the database package. You have to start it manually.

Additional steps for ORACLE

Additional steps for switching off SAP R/3 with an ORACLE database include:

- Backup Files

Save the configuration files so you can switch on the SGeSAP

SGeSAP Administration

Switching SGeSAP Off and On

Integration again. Save the files as user <sid>adm. This list of files you have to save include:

```
listener.ora
/usr/sap/trans/bin/TPPARAM
tnsnames.ora
/usr/sap/<SID>/SYS/profile/DEFAULT.PFL
/usr/sap/<SID>/SYS/profile/<SID>_DVEBMGS<INSTNR>
```

- **Stop Listener**

If you have an oracle database check to see if the listener is running. If it is, stop it as user ora<SID> with:

```
lsnrctl stop
```

If you have configured names for the listeners, which allows you to have more than one listener running, make sure to specify the correct name:

```
lsnrctl start <LISTENERNAME>
```

- **Change Profiles**

The references of the relocatable name must be changed to the real hostname in the following files:

```
listener.ora
tnsnames.ora
/usr/sap/<SID>/SYS/profile/DEFAULT.PFL
/usr/sap/<SID>/SYS/profile/<SID>_DVEBMGS<INSTNR>
/usr/sap/trans/bin/TPPARAM
```

NOTE

/usr/sap/trans/listener.ora is a link to /etc/listener.ora.

- **Start Listener**

Start oracle listener as oracle user with:

```
lsnrctl start <LISTENERNAME>
```

Additional steps for INFORMIX

Additional steps for switching off SAP R/3 with an INFORMIX database include:

- **Rename Customer Shell Scripts**

In the home directories of the users Informix and <sid>adm you can find two files called `.customer.sh` and `.customer.csh`. Rename them to prevent them from being executed the next time the database starts. For example:

```
mv .customer.csh .customer.sh.off
mv .customer.csh .customer.csh.off
```

Execute these lines in both directories. Make sure not to change owner or permissions of the files. Even though the home directory of <sid>adm is usually local, it is enough to do this on only one machine of the cluster. Choose the machine where you want to run the database later.

- **Backup Files**

Save the configuration files so you can switch the SGeSAP Integration on again. The list of files you have to save includes:

```
/informix/<SID>/etc/onconfig.<rel-IP_name_of_db>.<DBID>
/informix/<SID>/etc/sqlhosts.soc
/usr/sap/trans/bin/TPPARAM
/usr/sap/<SID>/SYS/profile/DEFAULT.PFL
/usr/sap/<SID>/SYS/profile/<SID>_DVEBMGS<INSTNR>
```

- **Change Profiles**

Replace any reference to the relocatable IP address(es) in the files you previously saved. Insert the local hostname of your machine. This is the name you get when you use the hostname command.

As informix user go to `/informix/<SID>/etc`. In the `onconfig` file replace the:

```
<rel-IP_name_of_DB><INFORMIX_DBID><shm|tcp>
```

entries with:

```
<local_hostname><INFORMIX_DBID><shm|tcp>
```

entries, for example:

```
DBSERVERNAME <local_hostname><sid>shm
             # Name of default database server
DBSERVERALIASES <local_hostname><sid>tcp
             # List of alternate dbservernames
```

SGeSAP Administration

Switching SGeSAP Off and On

Afterwards rename the file:

```
mv onconfig.<rel-IP_name_db>.<DBID> \  
onconfig.<local_hostname>.<DBID>
```

In the same way replace the appropriate entries:

```
<rel-IP_name_of_db><DBID><shm|tcp> on.. <rel-IP_name_db>  
<socket>
```

in `sqlhosts.soc` with:

```
<local_hostname><DBID><shm|tcp> on.. <local_hostname>  
<socket>
```

As user `<sid>adm` replace any occurrence of:

```
<rel-IP_name_of_db> or <rel-IP_name_of_ci>
```

with:

```
<local_hostname>
```

in:

```
/usr/sap/trans/bin/TPPARAM  
/usr/sap/<SID>/SYS/profile/DEFAULT.PFL  
/usr/sap/<SID>/SYS/profile/<SID>_DVEBMGS<INSTNR>
```

Check SAP R/3

To check SAP R/3:

Start all SAP R/3 Instances as `<SID>adm`.

NOTE

The automounter is still using the relocatable IP-addresses. You do not need to turn off the automounter. MC/ServiceGuard and the packages remain running and provide all necessary resources.

- Import the profiles within SAP R/3. Transaction code: RZ10.
- Check SAP R/3 setup. `Ping(1M)` to all application servers from the Central Instance and the other way round. Transaction code: SM59
- Check that batch jobs in SAP R/3 are not scheduled to run on the

relocatable IP address. Transaction code: SM37

- Relocate printers to the real hostname. Transaction code: SPAD
- Check operational modes within SAP R/3. You must setup new operation modes for the new hostname. Transaction code: RZ04
- Do all testing described in the document *SAP BC High Availability*.

Backup Files Again

After successfully finishing all tests, take another copy of the files you saved in the Sections “Additional steps for ORACLE” or “Additional steps for INFORMIX”. Refer to this step for a list of the files.

You previously renamed the `onconfig` file if you have an INFORMIX database. You might need this second backup for reference. If you plan to switch on the SGeSAP Integration again at a later point in time, make sure that no additional changes in the files occurred in the between times. Any changes would be lost if you replaced the altered files with the backup of the originals.

Use the backup created in this step, the backup of the altered files, later on for comparison with current active versions.

Switch On SGeSAP

Refer to the following sections for descriptions of how to switch on SGeSAP.

- Stop SAP R/3
- For ORACLE Database
- For INFORMIX Database
- Start RSAP R/3
- Turn Off Debugging Mode
- Reconfigure SAP R/3

Stop SAP R/3

Stop the database and SAP R/3 manually on all hosts in and outside of the cluster with:

```
stopsap all
```

For ORACLE Database

Restore Profiles:

- Check to see if the listener is running. Stop the listener process as ora<SID>:

```
lsnrctl stop <LISTENERNAME>
```

- Restore all profiles with the versions that use relocatable addresses. You saved these versions when you switched off the SGeSAP. If appropriate, incorporate any changes that occurred to these files during the switch off process. You can use `diff` and the second backup you took to easily find out if anything changed.
- Restore the following files together with their destination paths:

```
listener.ora  
/usr/sap/trans/bin/TPPARAM  
tnsnames.ora  
/usr/sap/<SID>/SYS/profile/DEFAULT.PFL  
/usr/sap/<SID>/SYS/profile/<SID>_DVEBMGS<INSTNR>
```

- Restart the listener as oracle user:

```
lsnrctl start <LISTENERNAME>
```

For INFORMIX Database

Restore Profiles

- Restore all profiles with the versions that use relocatable addresses. You saved these versions when you switched off the SGeSAP. If appropriate, incorporate any changes that occurred to these files during the switch off process. You can use `diff` and the second backup you took to easily find out if anything changed.
- Restore the following files together with their destination paths:

```
/informix/<SID>/etc/onconfig.<rel-IP_name_of_db>.<DBID>  
/informix/<SID>/etc/sqlhosts.soc/usr/sap/trans/bin/TPPAR  
AM  
/usr/sap/<SID>/SYS/profile/DEFAULT.PFL  
/usr/sap/<SID>/SYS/profile/<SID>_DVEBMGS<INSTNR>
```

- Rename Customized Files

As informix user perform the following from your home directory:

```
rm /informix/<SID>/etc/onconfig.<local_hostname>.<DBID>  
mv .customer.csh.off .customer.sh  
mv .customer.csh.off .customer.csh
```

As <sid>adm in home directory:

```
mv .customer.csh.off .customer.sh  
mv .customer.csh.off .customer.csh
```

Start RSAP R/3

Start SAP R/3 as <SID>adm using:

```
startsap all
```

Turn Off Debugging Mode

Delete /etc/cmcluster/debug and/or /etc/cmcluster/<SID>/debug if they exist.

Reconfigure SAP R/3

To reconfigure SAP R/3:

- Import the profiles within SAP R/3. Transaction code: RZ10.
- Check SAP R/3 setup. Try to ping (1M) to all application servers from the Central Instance and the other way round. Transaction code: SM59.
- Check that the batch jobs in SAP R/3 are scheduled to run on the relocatable IP address not on the local hostname. Transaction code: SM36
- Relocate printers to the relocatable name. Transaction code: SPAD
- Check operational modes within SAP R/3. You must setup operation modes for the relocatable name.
- Do all the testing as described in the document *SAP High Availability*.

SGeSAP Administration Issues

This section describes the new aspects that a System Administrator of any SGeSAP cluster should always be aware of.

Installation of the SGeSAP Integration Scripts significantly changes the hardware and software setup of your system. This affects the way you administer your SAP R/3. To get more detailed information on your specific setup, please refer to the documentation you receive from the person that installs the SGeSAP Integration on your system.

You no longer need to treat the Central Instance of your SAP R/3 system and its accompanying database Instance as though it runs on a dedicated host. With SAP R/3 they are wrapped up inside one or more MC/ServiceGuard packages and packages can be sent to any of the hosts that are inside of your MC/ServiceGuard cluster.

This provides not only a mechanism to cope with hardware failures but also a new amount of flexibility and opportunities. If you have to maintain the host machine on which your SAP R/3 is running, you can “send the running SAP R/3 over” to another host. This causes a few minutes of downtime only, in contrast to the significantly higher amount of downtime if you have to wait until the maintenance is completed before you can restart your SAP R/3 Instances.

On the other hand you have to be more careful in changing the setup. This applies to hardware changes as well as software issues.

The balance of this section describes the following SGeSAP administration aspects:

- Hardware Aspects
- HP-UX Software Aspects
- SAP R/3 Administration Aspects
- MC/ServiceGuard Administration Aspects
- SGeSAP Administration Aspects

Hardware Aspects

If you add new hardware and SAP R/3 needs access to it to work properly, make sure to allow this access from any host of the cluster by

appropriately planning the connectivity. For example:

It is possible to increase database disk space by adding a new RAID-Array to the primary host on which your Database Instance normally runs.

Setup with shared disks, so the new RAID-Array is visible on the new host and does not create an error.

You can add the RAID-Array using the conventional method you are used to. But remember: The fact that your SAP R/3 system runs correctly after the changes does *not* imply that it will work after a switchover to a different host as well.

For the same reason a high available SAP R/3 Spooling Service needs a connection to the printers from any host in the cluster.

If you do not feel comfortable in changing your hardware setup, please contact your HP-consultant.

HP-UX Software Aspects

Depending on your particular setup, you will have to deal with three important groups of directories that need special treatment:

- *common directories that are kept local on any node*

Using a standard setup the following directories and their files are kept locally on any host of the cluster:

`/etc/cmcluster`—the directory in which MC/ServiceGuard keeps its configuration files

`/home/<SID>adm`—the home directory of the SAP R/3 system Administrator.

It is your responsibility to synchronize the contents of these directories on all hosts of the cluster. `/home/<SID>adm` does not need to be the same on all of the hosts. For example:

It is possible to install an additional application server on a host of the cluster. The application server will not be part of any package. It is local to this host. The SAP startup scripts are only needed on this dedicated host. You do not need to distribute them to other hosts.

The standard HP-UX configuration files are local, too. Never delete the mutual `.rhosts` entries of the root user and `<SID>adm` on any of the nodes. Never change entries in `/etc/hosts`, `/etc/services`,

`/etc/passwd` or `/etc/group` on only some of the nodes. Keep them unified.

If you use an ORACLE database, be aware that the listener configuration file of SQL*Net V2 is kept locally as `/etc/listener.ora` by default, too.

- *directories that reside on shared disks*

Changing these files on any host of the cluster applies the change to the whole cluster.

Files in the following directories and all subdirectories are typically shared:

```
/usr/sap/<SID>/DVEBMGS<instance_id>  
/export/usr/sap/trans  
/export/sapmt/<SID>  
/export/informix or /oracle/<SID>
```

They are only available on a host if the package they belong to is running on it. MC/ServiceGuard switches them to another node with the package. If you use a two package concept all directories belong to the database package (db) apart from the first which belongs to the Central Instance package (ci). Please refer to the description of your particular setup to obtain a list of your shared directories.

- *directories that are treated by the automounter*

These directories are mounted automatically as needed. This is true not only for the nodes of the cluster. If you use external application servers, they also use them.

Automounter directories are:

```
/sapmnt/<SID>  
/usr/sap/trans  
/informix
```

They are NFS-mounted from their equivalents in the `/export` directory of the node(s) which run(s) the package(s). The automounter setup uses the relocatable IP addresses. The directories are soon available again after a switchover has taken place.

There are two important issues concerning these directories:

— The directories below `/export` are exported without root

permissions.

This happens according to the recommendations of SAP R/3 and enhances the security of the installation. The effect is, that the root user cannot modify these directories or their contents. With standard permissions set, the root user cannot even see the files. This is not an error and the system runs without problems. If you want to modify anything as root, please use the equivalent directory below /export on the host the package runs on.

- If the database package is halted, you cannot log in as <SID>adm unless you keep the binaries local.

The reason for this is, that /usr/sap/<SID>/SYS/exe is part of the path of <SID>adm. Without local binaries this directory links to /sapmnt/<SID> which in fact is handled by the automounter. The automounter cannot contact the host belonging to the relocatable address that is configured because the package is down. The system hangs. To avoid this, always keep local copies of the executables.

SAP R/3 Administration Aspects

As far as SAP R/3 is concerned, nearly everything remains the same. The only difference is the use of the relocatable IP-addresses. During installation of the SGeSAP Integration, the profiles are changed to contain only relocatable IP-addresses for the database as well as the Central Instance. You can check this using transaction RZ10. In the DEFAULT.PFL the entries are altered:

```
SAPDBHOST= <relocatable_db_name>
rdisp/mshost= <relocatable_ci_name>
rdisp/vbname= <relocatable_ci_name>_<SID>_<nr>
rdisp/enqname= <relocatable_ci_name>_<SID>_<nr>
rdisp/btcname= <relocatable_ci_name>_<SID>_<nr>
rslg/collect_daemon/host = <relocatable_ci_name>
```

There are also two additional profile parameters SAPLOCALHOST and SAPLOCALHOSTFULL included as part of the Instance Profile of the Central Instance. Anywhere SAP R/3 uses the local hostname internally, this name is replaced by the relocatable values <relocatable_ci_name> or <relocatable_ci_name>.domain.organization of these parameters. Make sure that they are always defined and set to the correct value. This is vital for the system to function correctly.

Normally use the relocatable name of your Central Instance package if SAP R/3 asks for the hostname of the machine the Central Instance runs on. Unfortunately, SAP R/3 sometimes uses local addresses by itself. If the system is set up correctly as described in Chapter 2, “Step by Step Installation Guide,” you do not have to worry about this. The values change if you switch over to another host.

The destination for print formatting, which is done by a Spool Work process, uses the application server name. Use the relocatable name if you plan to use Spool Work processes with your Central Instance. In these cases no changes need to be done in case of a failover - printing will work persistently. Note that a print job in process at the time of the failure will be canceled and needs to be reissued manually after the failover. To make a spooler highly available on the Central Instance, set the destination of the printer to `<relocatable_ci_name>_<SID>_<nr>` using the transaction SPAD. Print all *time critical* documents via the high available spool server of the Central Instance. Print requests to other spool servers stay in the system after failure until the host is available again and the spool server has been restarted. These requests can be moved manually to other spool servers if the failed server is unavailable for a longer period of time.

Batch jobs can be scheduled to run on a particular instance. Generally speaking, it is better not to specify a destination host at all. Sticking to this rule, the batch scheduler chooses a batch server which is available at the start time of the batch job. However, if you wish to specify a destination host, specify the batch server running on the highly available Central Instance. The application server name and the hostname (which is retrieved from the Message Server) are stored in the batch control tables TBTCO, TBTCS,.... In case the batch job is ready to run, the application server name is used to start it. Therefore, when using the relocatable name to build the application server name for the SAP R/3 Instance, you do not need to change batch jobs which are tied to it after a switchover. This is true even if the hostname which is also stored in the above tables, differs.

Plan to use saplogon to application server groups instead of saptemu/sapgui to individual application servers. When logging on to an application server group with two or more application servers, the SAP R/3 user does not need a different login procedure if one of the application servers of the group fails. Also, using login groups, provides workload balancing between application servers, too.

Within the CCMS you can define operation modes for SAP R/3 Instances.

An operation mode defines a resource configuration. It can be used to determine which instances are started and stopped and how the individual services are allocated for each instance in the configuration. An instance definition for a particular operation mode consists of the number and types of Work processes as well as Start and Instance Profiles. When defining an instance for an operation mode you need to enter the hostname and the system number of the application server. By using relocatable names to fill in the hostname field, the instance will be working under control of the CCMS after a failover without a change. Note however, that if an instance is running on the standby node in normal operation and is stopped during the switchover, the control console will show the instance to be down afterwards.

Only configure the update service on a node for Application Services running on the same node. As a result, the remaining SAP R/3 servers, running on different nodes, are not affected by the outage of the update server. However, if the update server is configured to be responsible for application servers running on different nodes, any failure of the update server leads to subsequent outages at these nodes. Configure the update server on the highly available Central Instance. Using local update servers should only be considered, if performance issues require it.

MC/ServiceGuard Administration Aspects

MC/ServiceGuard keeps information about the cluster configuration, it especially needs to know the relocatable IP addresses and its subnets, your Volume Groups, the Logical Volumes and their mountpoints. Check with your HP consultant for information about the way MC/ServiceGuard is configured to suite your SAP R/3 system. If you touch this configuration, you have to reconfigure your cluster.

SGeSAP Administration Aspects

The SGeSAP needs some additional information about the configuration of your SAP R/3 system. It gathers this information from the file `/etc/cmcluster/sap.conf`. It is very important that the information in this file is consistent with the way your system is configured. This file must be available on all nodes of the cluster. Under normal operation there is no need to touch this file. But it is a good idea for you as an administrator to have a look at it to prevent you from doing things that can cause the setup to be no longer able to switch. Comments are provided within the file and most things are self-explanatory if you are

SGeSAP Administration

SGeSAP Administration Issues

trained to understand shell scripts. For example, the following administration activities are possible but they must be accompanied by an adaptation of the `sap.conf` file on all cluster nodes:

- changing the SAP System ID
- changing the name of the SAP System Administrator
- migrating to another database vendor
- adding/deleting additional application servers
- changing Instance Numbers
- changing the name belonging to a relocatable address
- changing the name of a SGeSAP-Integration package
- changing hostnames of hosts inside the cluster
- changing hostnames of hosts that run additional application servers
- changing the location of any SAP-specific directory in the filesystem
- changing the name of the ORACLE listener

After performing any of the above mentioned activities the ability to failover correctly has to be tested again. Be aware that changing the configuration of your SAP R/3 cluster in any way can lead to the loss of warranty. Always make sure to plan those steps together with your HP consultant.

SGeSAP Administration With HP Somersault Enabled

The administration of the SAP R/3 cluster is slightly different when HP Somersault is configured. In a failover scenario, for example, if the Central Instance (CI) host crashes and the CI switches over to the backup node, the HP Somersault protected enqueue service remains on the server set nodes. When the CI restarts on the backup node the work processes reconnect to the enqueue service.

For maintenance reasons it might be necessary to shutdown the complete SAP R/3 system.

- To shutdown the SAP R/3 system, including the complete CI and the HP Somersault unit, type:

```
cmhaltpkg dbciPKG/ciPKG
```

This is the only time when the HP Somersault unit dumps its in-memory Enqueue table to the backup file ENQBCK in `/usr/sap/<SID>/SYS/global`.

- To startup the SAP R/3 system, including the complete CI and the HP Somersault Enqueue unit, type:

```
cmrunpkg dbciPKG/ciPKG
```

It is essential to have the Enqueue backup file ENQBCK located in `/usr/sap/<SID>/SYS/global`.

To enable HP Somersault:

1. Place the Enqueue backup file where the HP Somersault system can find it: `/usr/sap/<SID>/SYS/global/ENQBCK`.

`/usr/sap/<SID>/SYS/global` is usually linked to `/sapmnt/<SID>/global`. All valid backup files must be moved there.

2. Follow the procedures in Step OS559.

3. Edit the SAP R/3 profile:

`/usr/sap/<SID>/SYS/profile/DEFAULT.PFL`.

Add the following two lines:

```
ENHA_Identifier           =$(SAPSYSTEMNAME)_HPSOM_ENQUEUEE  
enqueue/process_location =REMOTEA
```

To disable HP Somersault:

1. Place the Enqueue backup file to the SAP R/3 specified location.
2. Follow the procedures in Step OS559.

3. Edit the SAP R/3 profile:

```
/usr/sap/<SID>/SYS/profile/DEFAULT.PFL.
```

Remove the following two lines:

```
ENHA_Identifier          =$(SAPSYSTEMNAME)_HPSOM_ENQUEUE  
enqueue/process_location =REMOTEGA
```

For additional information on the administration of HP Somersault refer to your HP Somersault documentation.

Symbols

.ora file, 98
.profile
 modifying, 62
 set -u option, 61
.rhosts file, 59, 62

A

additional Application Server
 ASHOST array, 65
 ASNR array, 65
 ASTREAT array, 66
 sap.conf arrays, 65
addresses
 host machine, 20
 IP, 20
 relocatable, 60
administrating
 MC/ServiceGuard, 121
 SAP R/3, 116, 119
 SGeSAP, 122
agent files
 backing up, 113
Application Server
 automounter, 23
 deactivate, 72
 directory handling, 22
 external, 62
 internal, 51
 previously installed, 48
 reactivate, 72
 stopping, 75
application server, 15
ASHOST array, 71
 additional Application Server,
 65
ASNR array
 additional Application Server,
 65
ASTREAT array, 71
 additional Application Server,
 66

auto.direct file, 91
auto_master file, 90
AUTO_VG_ACTIVATE function,
 84
automounter, 87
 configuration, 90
 directories, 23
 NFS client daemon, 90
 relocatable IP addresses, 23
 restarting, 91
 SAP R/3, 118
AUTOSTART_CMCLD function,
 84

B

backing up
 agent files, 113
 Informix files, 111
 Oracle files, 109
batch jobs, 120
 SAP R/3 configuration, 105
binaries
 local, 24
binary cluster configuration file
 distributing, 89

C

CCMS, 105
Central Instance
 CINR parameter, 64
 directory, 59
 multiple configurations, 57
 node, 12, 13
 relocatable hostname, 66
 restart, 67
 starting host, 74
 stopping, 75
changing
 profiles, 110, 111
checking
 SAP R/3, 112
CINR parameter, 64

CIRELOC parameter, 66
cleanipc tool, 69
cleanup policies, 34
CLEANUP_POLICY parameter,
 68
cmapplyconf, 89
cmquerycl
 configuring MC/ServiceGuard,
 84
Computing Center Management
 System, 105
configuration
 directories, 21
 distributing, 52
 files, 63
 MC/ServiceGuard, 36
 MC/ServiceGuard files, 26
 one package, 12
 package models, 12
 run-time, 35
 two package, 13
configuring
 automounter, 90
 cluster node, 59
 database, 93
 DNS, 61
 external Application Server, 62
 Informix, 93
 integration, 63
 LANs, 61, 62
 MC/ServiceGuard, 84
 multiple Central Instance
 nodes, 57
 Oracle, 96
 SAP R/3, 102
 SAP R/3 batch jobs, 105
connectivity
 SAP R/3, 117
consolidating system
 environment, 70
control files
 MC/ServiceGuard, 28
 MC/ServiceGuard files, 37

control logic, 35
crossmounting, 22
customized files
 renaming, 114

D

database
 ASHOST array reconnect
 functions, 71
 ASTREAT array
 reconnect functions, 71
 configuration, 93
 DB parameter, 64
 Informix, 93
 node, 12, 13
 reconnect, 70
 relocatable hostnames, 66
 specifying, 76
 DB parameter, 64
 DB_RECONNECT parameter,
 64
 DBHOSTNAME
 Informix, 96
 DBRELOC parameter, 66
 DBSERVERALIAS
 Informix, 95
 DBSERVERNAME
 Informix, 95
deactivate
 Application Server, 72
debug
 file, 92
 stopping, 115
 turn mode off, 109
defaults
 redefining, 76
DELAY_INTERVALS
 parameter, 67
depot
 installing, 63
device
 minor numbers, 43

dialog instance handling
 SAP R/3, 33
Dialog Instances, 75
directories
 application server, 22
 automounter, 23
 Central Instance, 59
 common, 21
disable
 HP Somersault, 124
distributing, 59
 configurations, 52
 local directory tree, 58
DNS
 configuring, 61

E

enable
 HP Somersault, 123
Enqueue Server
 HP Somersault, 31, 83
executables
 local, 25
 sapcpe mechanism, 25
exported directories, 87
extension files
 one package model, 28
 two package model, 29
external Application Server
 configuration, 62

F

failover
 Central Instance restart, 67

G

groups
 primary node, 53
 propagating, 49

H

halt
 packages, 108
hardware
 SAP R/3, 117
High Availability
 HP Somersault, 31
host machine addresses, 20
hostname
 additional Application Server,
 65
 Informix, 96
 minor numbers, 43
HP Somersault, 31, 32
 administration, 123
 disable, 124
 enable, 123
 installing, 83
 maintenance, 123
 shutdown, 123
 startup, 123

I

Informix
 backing up files, 111
 hostname, 96
 Oracle cleanup, 88
 switching off, 110
installing
 depot, 63
 HP Somersault, 83
 one package configuration, 63
 SAP R/3, 49
 SAP R/3 on primary node, 50
Instance ID, 75
 additional Application Server,
 65
integration
 cluster host files, 89
 configuration, 63
internal Application Server, 51
internal application servers, 15

-
- IP addresses, 20
 - MC/ServiceGuard configuration, 86
 - IPC resources, 69
 - K**
 - kernel
 - synchronizing, 55
 - L**
 - LAN
 - configuring, 61, 62
 - IP addresses, 20
 - listener
 - starting, 110
 - stopping, 110
 - local
 - binaries, 24
 - executables, 25, 48, 51
 - local directory
 - distributing, 58
 - log on
 - Application Server group, 120
 - logical volumes
 - MC/ServiceGuard configuration, 86
 - LVM
 - setup, 43
 - M**
 - MC/ServiceGuard
 - administrating, 121
 - configuration, 36
 - configuration files, 26
 - configuring, 84
 - control files, 28, 37
 - directories, 21
 - IP addresses, 86
 - logical volumes, 86
 - SAP R/3 benefits, 11
 - subnets, 86
 - volume groups, 86
 - minor numbers, 43
 - mountpoints, 22
 - creating, 50
 - Oracle, 57
 - SAP R/3 installed, 46
 - N**
 - network
 - IP addresses, 20
 - NFS client
 - automounter configuration, 90
 - node
 - central application servers, 15
 - Central Instance, 12, 13
 - configuring cluster, 59
 - database, 12, 13, 15
 - groups, 53
 - names, 84
 - primary, 12, 13
 - primary installation, 50
 - services, 54
 - synchronization, 53
 - users, 53
 - O**
 - onconfig file
 - Informix, 95
 - one package
 - CIRESTART parameter, 65
 - configuration, 12
 - extension files, 28
 - installation files, 63
 - script files, 28
 - Oracle
 - backing up files, 109
 - configuration, 96
 - directories, 21
 - Informix cleanup, 88
 - mountpoints, 22, 57
 - relocatable, 97
 - switching off, 109
 - P**
 - package control files, 63
 - package control scripts
 - package configuration, 86
 - SAP R/3 scripts, 87, 88
 - packages
 - additional Application Server, 66
 - configuration files, 26, 84
 - control files, 28
 - halt, 108
 - IP addresses, 20
 - name, 84
 - package control script, 86
 - starting, 109
 - volume groups, 44
 - primary node, 12, 13
 - Central Instance installed, 55
 - Informix installed, 56
 - Oracle installed, 57
 - print
 - Spool Work process, 105
 - product files
 - SGeSAP, 35
 - profiles
 - changing, 110, 111
 - restoring, 114
 - R**
 - reactivate
 - Application Server, 72
 - reconfiguring
 - SAP R/3, 115
 - relocatable
 - external Application Servers, 62
 - hostnames, 66
 - IP addresses, 20, 60
 - Oracle, 97
 - relocatable IP addresses
 - automounter, 23
 - remsh

-
- Informix, 93
 - renaming
 - customized files, 114
 - shell scripts, 111
 - restarting
 - automounter, 91
 - restoring
 - profiles, 114
 - RMADM array, 75
 - RMDEP array, 75
 - RMDEPHOST array, 75
 - RMDEPNR array, 75
 - RMNR array, 75
 - root
 - directory permissions, 24
 - run-time configuration, 35
 - run-time functions, 35
- S**
- SAP R/3
 - administrating, 116, 119
 - automounter directories, 118
 - checking, 112
 - common directories, 117
 - configuration, 102
 - connectivity, 117
 - dialog instance handling, 33
 - directories, 21
 - hardware administration, 117
 - installing, 49
 - MC/ServiceGuard benefits, 11
 - monitoring processes, 16
 - mountpoints, 46
 - package control scripts, 87, 88
 - pre-installed, 46
 - propagating users, groups, 49
 - reconfiguring, 115
 - shared directories, 118
 - shared memory resource
 - handling, 34
 - single server, 33
 - starting, 115
 - stopping, 113
 - stopping instances, 75
 - UNIX administration, 117
 - volume groups, 46
 - sap.conf
 - arrays, 65
 - Central Instance, 64
 - customizing, 64
 - database, 64
 - editing file, 122
 - file, 63
 - one package, 65
 - two package concept, 64
 - sapcpe, 25, 48, 51
 - SAPDBHOST, 119
 - SAPDBHOSTFULL, 119
 - SAPLOCALHOST parameter, 103
 - SAPLOCALHOSTFULL parameter, 103
 - saprouter, 74
 - scripts
 - one package structure, 28
 - two package structure, 29
 - server machine set
 - HP Somersault, 32
 - server set machines, 32
 - SERVER_CONSOLIDATION parameter, 70
 - services
 - primary node, 54
 - setup
 - LVM, 43
 - SGeSAP
 - administrating, 122
 - product files, 35
 - switching off, 108
 - switching on, 113
 - shared disks
 - SAP R/3, 118
 - shared memory, 68
 - resource handling, 34
 - SAP R/3, 34
 - shared volume groups, 59
 - shell scripts
 - renaming, 111
 - shutdown
 - HP Somersault, 123
 - WAIT_OWN_AS parameter, 68
 - single point of failure, 15
 - SPACECHECK entries, 50
 - SPOF, 15, 31
 - Spool Work, 120
 - print configuration, 105
 - SQLNET.EXPIRE_TIME parameter, 99
 - starting
 - listener, 110
 - packages, 109
 - SAP R/3, 115
 - startup
 - HP Somersault, 123
 - steps
 - coding, 41
 - stopping
 - debug mode, 115
 - listener, 110
 - SAP R/3, 113
 - SAP R/3 instances, 75
 - subnets, 84
 - MC/ServiceGuard configuration, 86
 - svrmgrl
 - Oracle configuration, 100
 - switching off
 - Informix, 110
 - Oracle, 109
 - SGeSAP, 108
 - switching on
 - SGeSAP, 113
 - synchronization
 - nodes, 53

T

tasks

- coding, 41

transaction reset, 70

transport directory, 87

- additional packages, 89

TRANSRELOC parameter, 66

turning off

- debug mode, 109

two package

- configuration, 13

- DB_RECONNECT parameter,
64

- extension files, 29

- MC/ServiceGuard, 84

- script files, 29

U

UNIX

- SAP R/3, 117

- SAP R/3 directories, 117

update service, 121

users

- primary node, 53

- propagating, 49

V

volume groups, 49

- cluster configuration file, 84

- distributing, 59

- MC/ServiceGuard

 - configuration, 86

- MC/ServiceGuard packages,
44

- required, 43

- SAP R/3 installed, 46

W

WAIT_OWN_AS parameter

- shutdown, 68

witness set machines, 32

Index
