SAP Software on UNIX: OS Dependencies

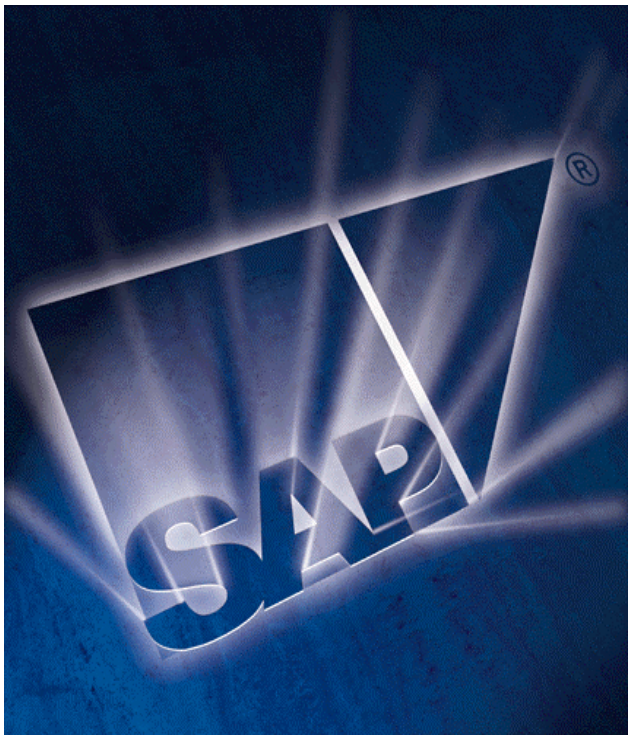

4.6C SR2

SAP® AG • Neurottstr. 16 • D-69190 Walldorf

51 013 679

Copyright

©Copyright 2001 SAP AG. All rights reserved.

No part of this brochure may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.

Microsoft®, WINDOWS®, NT®, EXCEL®, Word® and SQL Server® are registered trademarks of Microsoft Corporation.

IBM®, DB2®, OS/2®, DB2/6000®, Parallel Sysplex®, MVS/ESA®, RS/6000®, AIX®, S/390®, AS/400®, OS/390®, and OS/400® are registered trademarks of IBM Corporation.

ORACLE® is a registered trademark of ORACLE Corporation, California, USA.

INFORMIX®-OnLine *for SAP* and Informix® Dynamic Server™ are registered trademarks of Informix Software Incorporated.

UNIX®, X/Open®, OSF/1®, and Motif® are registered trademarks of The Open Group.

HTML, DHTML, XML, XHTML are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Laboratory for Computer Science NE43-358, Massachusetts Institute of Technology, 545 Technology Square, Cambridge, MA 02139.

JAVA® is a registered trademark of Sun Microssystems, Inc. , 901 San Antonio Road, Palo Alto, CA 94303 USA.

JAVASCRIPT® is a registered trademark of Sun Microsystems, Inc., used under license for technology invented and implemented by Netscape.

ADABAS® is a registered trademark of Software AG.

SAP, SAP Logo, mySAP.com, mySAP.com Marketplace, mySAP.com Workplace, mySAP.com Business Scenarios, mySAP.com Application Hosting, WebFlow, R/2, R/3, RIVA, ABAP, SAP Business Workflow, SAP EarlyWatch, SAP ArchiveLink, BAPI, SAPPHIRE, Management Cockpit, SEM, SAP DB™ are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world. All other products mentioned are trademarks or registered trademarks of their respective companies.

Documentation in the SAP Service Marketplace

You can find this documentation at the following address:

`http://service.sap.com/instguides`

Typographic Conventions

Type Style	Represents
Interface Text	Words or characters that appear on the screen. This includes system messages, field names, screen titles, pushbuttons, menu names, and menu options.
Document Title	Cross-references to other documentation
`User Entry`	Exact user entry. These are words and characters that you enter exactly as they appear in the documentation.
`File Name`	File names, batch files, paths or directories, and screen messages
`<Variable User Entry>`	Variable user entry. Pointed brackets indicate that you replace these words and characters with appropriate entries.
NAME	Names of elements in the SAP System. These include report names, program names, transaction codes, table names, and ABAP language elements.
`KEY`	Keys on your keyboard. These include function keys (for example, `F2`) and the `ENTER` key.

Icons

Icon	Meaning
	Caution
	Example
	Note
	Recommendation
	Syntax
	Tip

Contents

SAP Software on UNIX: OS Dependencies 7

1 Network Information Service Environment 8

2 AIX: OS-Dependent Installation Steps 10
2.1 AIX: Preparing the Installation 10
2.2 AIX: Mounting a CD-ROM 11
2.3 AIX: Volume Groups, File Systems, Raw Devices, Swap Space 12
2.3.1 AIX: Size of a Logical Partition 12
2.3.2 AIX: Setting up Swap Space 12
2.3.3 AIX: Creating Volume Groups 13
2.3.4 AIX: Setting up File Systems 14
2.3.5 AIX: Setting up Raw Devices 14
2.4 AIX: Mounting Directories via NFS 16
2.5 AIX: Creating UNIX Groups and Users 18
2.6 AIX: Troubleshooting 19

3 Compaq Tru64 UNIX: OS-Dependent Installation Steps 20
3.1 Compaq: Mounting a CD-ROM 20
3.2 Compaq: Checking and Modifying the UNIX Kernel 21
3.3 Compaq: File Systems, Raw Devices and Swap Space 23
3.3.1 Compaq: Preparing Hard Disks 23
3.3.2 Compaq: Setting up Swap Space 26
3.3.3 Compaq: Setting up Standard File Systems 27
3.3.4 Compaq: Setting up Advanced File Systems 28
3.3.5 Compaq: Setting up Raw Devices 29
3.4 Compaq: Mounting Directories via NFS 30
3.5 Compaq: Creating UNIX Groups and Users 31
3.6 Compaq: Troubleshooting 31

4 HP-UX: OS-Dependent Installation Steps 32
4.1 HP-UX: Mounting a CD-ROM 32
4.2 HP-UX: Checking and Modifying the UNIX Kernel 33
4.3 HP-UX: File Systems, Raw Devices and Swap Space 36
4.3.1 HP-UX: Setting up Swap Space 36
4.3.2 HP-UX: Setting up File Systems using LVM 38
4.3.3 HP-UX: Setting up File Systems using SAM 40
4.3.4 HP-UX: Accessing Raw Devices 41
4.4 HP-UX: Mounting Directories via NFS 42
4.5 HP-UX: Creating UNIX Groups and Users 44
4.6 HP-UX: Troubleshooting 44

5 IBM NUMA-Q: OS-Dependent Installation Steps 45
5.1 IBM NUMA-Q: Preparing the Installation 45
5.2 IBM NUMA-Q: Mounting a CD-ROM 45
5.3 IBM NUMA-Q: Checking System and Kernel Parameters 46
5.4 IBM NUMA-Q: File Systems, Raw Devices, Swap Space 48
5.4.1 IBM NUMA-Q: Preparing Hard Disks 48
5.4.2 IBM NUMA-Q: Setting up Swap Space 49
5.5 IBM NUMA-Q: Creating UNIX Groups and Users 50
5.6 IBM NUMA-Q: Troubleshooting 50

6 Linux: OS-Dependent Installation Steps 51
6.1 Linux: Mounting a CD-ROM 51
6.2 Linux: Checking and Modifying the Linux Kernel 51
6.3 Linux: File Systems and Swap Space 52
6.3.1 Linux: Preparing Hard Disks 52
6.3.2 Linux: Setting up Swap Space 53
6.3.3 Linux: Setting up Standard File Systems 53
6.4 Linux: Exporting Directories via NFS 54
6.5 Linux: Creating Linux Groups and Users 55

7 ReliantUNIX: OS-Dependent Installation Steps 56
7.1 ReliantUNIX: Preparing the Installation 56
7.2 ReliantUNIX: Mounting a CD-ROM 56
7.3 ReliantUNIX: Checking and Modifying the UNIX Kernel 57
7.4 ReliantUNIX: File Systems, Raw Devices, Swap Space 60
7.4.1 ReliantUNIX: Setting up Swap Space 60
7.4.2 ReliantUNIX: File System 62
7.4.3 ReliantUNIX: Accessing Raw Devices 65
7.5 ReliantUNIX: Mounting Directories via NFS 66
7.6 ReliantUNIX: Creating UNIX Groups and Users 67
7.7 ReliantUNIX: Troubleshooting 68

8 Solaris: OS-Dependent Installation Steps 69
8.1 Solaris: Preparing the Installation 69
8.2 Solaris: Mounting a CD-ROM 69
8.3 Solaris: Checking and Modifying the UNIX Kernel 70
8.4 Solaris: File Systems, Raw Devices, Swap Space 72
8.4.1 Solaris: Preparing Hard Disks 73
8.4.2 Solaris: Setting up Swap Space 74
8.4.3 Solaris: Creating File Systems 77
8.4.4 Solaris: Accessing Raw Devices 78
8.5 Solaris: Mounting Directories via NFS 79
8.6 Solaris: Creating UNIX Groups and Users 80
8.7 Solaris: Troubleshooting 81

9 Heterogeneous SAP System Installation 82

SAP Software on UNIX: OS Dependencies

Purpose

During an installation or upgrade, some of the tasks that you need to perform are operating system-dependent.

This documentation contains information on the following operating systems:

- AIX [page 10]
- Compaq Tru64 UNIX [page 20]
- HP-UX [page 32]
- IBM NUMA-Q [page 45]
- Linux [page 51]
- ReliantUNIX [page 56]
- Solaris [page 69]

In addition, this documentation contains information on the following:

- Network Information Service Environment [page 8]
- Heterogeneous SAP System Installation [page 82]

Integration

Some operating system-dependent tasks may be described in other documentation, for example, in the installation or upgrade documentation.

Read **SAP Note 387078** (SAP Software on UNIX - OS Dependencies). This note contains the most recent information regarding OS Dependencies, as well as corrections to this documentation.
Make sure that you have the most recent version of the note. You can find SAP Notes in the SAP Service Marketplace (`http://service.sap.com/notes`).

1 Network Information Service Environment

General Notes on NIS

In environments where Network Information Service (NIS) is used, the person in charge of the installation is responsible for distributing users over the network.

All users must have identical environment settings.

If you change the environment delivered by SAP, such as variables, paths, and so on, SAP will not assume responsibility.

The SAP installation tool R3SETUP checks all required users, groups, and services on the local machine.

If you manage users, groups and/or services network-wide in your company, you should consider the following:

- R3SETUP uses NIS commands to check the users, groups and services. However, R3SETUP does not change NIS configurations.
- R3SETUP creates the necessary users, groups and services locally. Use these entries to adjust your NIS Server. Examine the R3SETUP messages carefully.

It is also possible to generate the network-wide entries prior to the R3SETUP run. See the section for your operating system.

Users and Groups

The following tables give an overview of the users and their primary groups:

Users and their Primary Groups

Database System	User	Primary Group
SAP DB	<sapsid>adm	sapsys
	sqd<sapsid>	sapsys
DB2 UDB for UNIX and Windows	<sapsid>adm	sapsys, db<sapsid>ctl
	sapr3	sapsys
	db2as	db2asgrp
	db2<sapsid>	db<sapsid>adm
DB2 UDB for OS/390	<sapsid>adm	sapsys
Informix	<sapsid>adm	sapsys
	sapr3	sapsys
	informix	informix
Oracle	<sapsid>adm	sapsys, oper, dba
	ora<sapsid>	dba, oper

Groups and Members

Database System	Group	Members
SAP DB	sapsys	sqd<sapsid>, <sapsid>adm
DB2 UDB for UNIX and Windows	sapsys	<sapsid>adm, sapr3
	db<sapsid>ctl	<sapsid>adm
	db2asgrp	db2as
	db<sapsid>adm	db2<sapsid>
DB2 UDB for OS/390	sapsys	<sapsid>adm
	sapctrl	<sapsid>adm
Informix	sapsys	sapr3, <sapsid>adm
	informix	informix
	super_archive	informix, <sapsid>adm
Oracle	sapsys	<sapsid>adm
	oper	<sapsid>adm, ora<sapsid>
	dba	<sapsid>adm, ora<sapsid>

Compaq Tru64 UNIX:
The user <sapsid>adm must be a member of the group mem.

The user ID (UID) and group ID (GID) of SAP users and groups must be identical for all servers belonging to any SAP System.
This does **not** mean that all users and groups have to be installed on all SAP servers.
There are no other requirements concerning certain numbers for UIDs or GIDs.

R3SETUP checks if the users and groups already exist and creates them if necessary.
R3SETUP chooses free user and group IDs unless you are installing a dialog instance. In this case the IDs must be the same as on the central instance host.

Services

R3SETUP checks if the required services are available on the host and creates them if necessary.

See the log messages about the service entries and adapt the network-wide (NIS) entries accordingly.

2 AIX: OS-Dependent Installation Steps

This documentation has been written with the support of the IBM-Bull Competence Center. It is not intended to replace the documentation of the AIX operating system.

All OS-dependent steps can be done in two ways:

- Manually

  This means that AIX commands with the appropriate options have to be keyed in.

- Using SMIT

  This is a menu-driven system administration tool.

The following text describes the use of the AIX tool SMIT. If problems arise with the function keys, you can also use `ESC` and the corresponding number to simulate the function key (for example, `F4=ESC 4`).

2.1 AIX: Preparing the Installation

Portable Stream Environment (PSE)

SAP requires the Portable Stream Environment. The command **`strload`** enables you to load PSE, PSE drivers and PSE modules. This command must run at boot time from the `inittab`.

PSE is not needed in a **DB2 for OS/390** installation.

To enable the PSE, enter the command:
`mkitab strload:2:once:/usr/sbin/strload`

To check the PSE, as user **`root`**, enter the command:
`lsitab strload`

This has to return the `inittab` entry specified above with the **`mkitab`** command.

2.2 AIX: Mounting a CD-ROM

1. Log on as user **`root`**.
2. Add a CD-ROM file system.
   a. Enter the command **`smitty`**.
   b. Select:

      Physical & Logical Storage → *File Systems* →
      Add / Change / Show / Delete File Systems →
      CDROM File Systems → *Add a CDROM File System*

   c. Press `F4` to get a list of device names.

      Select a device name.

   d. Enter the mount point **`<CD-mountdir>`**.

      (For example, `<CD-mountdir>` is `/sapcd`)

   e. Press `F10` to exit `smitty`.
3. Mount the CD-ROM.
   a. Enter the command **`smitty`**.
   b. Select:

      Physical & Logical Storage →
      File Systems → *Mount a File System*

   c. Place the cursor on *File System Name* and press `F4`.
   d. Select `/sapcd`.
   e. Place the cursor on field *Type of File System*. Press `F4`.
   f. Select `cdrfs`.
   g. Change *Mount as Read Only system* to `YES`.
   h. Press `ENTER`.
   i. Press `F10` to exit `smitty`.

2.3 AIX: Volume Groups, File Systems, Raw Devices, Swap Space

The following describes how to create the volume groups, file systems, raw devices and swap space for your installation.

The structure of the SAP system depends on the disk configuration of the customer system. For this reason, there are no general procedures for setting up the file systems.

Consider the following when you specify the sizes of file systems (not applicable for **DB2 for OS/390**):
Your operating system needs up to 10% of disk space for administration.

AIX does not allow block 0 to block 15 on the disk to be part of the raw device used by the database system. When you reboot the machine, the operating system writes system information to these blocks. Data from the database will be overwritten and a restore of the database becomes necessary.

See the following sections for more information:

- Size of a Logical Partition [page 12]
- Setting up Swap Space [page 12]
- Creating Volume Groups [page 13]
- Setting up File Systems [page 14]
- Setting up Raw Devices [page 14]

2.3.1 AIX: Size of a Logical Partition

In the Logical Volume Manager, the size of a logical partition is equal to the size of the physical partition of the volume group to which the partition belongs.

At the time a volume group is created, the size can be chosen from the range 1, 2, ...256 MB. For example, when a volume group is created with the parameter `Physical Partition SIZE` of 128 MB, each logical partition will be 128 MB.

2.3.2 AIX: Setting up Swap Space

Use

You can set up swap space now if you have a suitable volume group. You also can perform this step at the end of the OS preparation.

Prerequisites

You can find the minimum swap space requirement in the installation documentation in Hardware and Software Requirements Check. If possible, you should use an entire physical disk as a swap space partition.

Procedure

1. Determine the size of the installed RAM using the command:
   `lscfg | grep mem`
2. Determine the allocated swap space:
   a. Enter the command **`smitty`**.
   b. Select:

      Physical & Logical Storage → Logical Volume Mgr → Paging Space → List All Paging Spaces

      or enter the command
      `lsps -a`
3. Add another paging space if necessary using `smitty`:
   a. Select:

      Physical & Logical Storage → Logical Volume Mgr → Paging Space → Add Another Paging Space

      A list of volume group names is displayed.
   b. Select a volume group.
   c. Enter the size of paging space in logical partitions.
   d. Set *Start using this paging space NOW ?* to **`YES`**.
   e. Set *Use this paging space each time the system is RESTARTED* to **`YES`**.
   f. Press `F10` to exit `smitty`.
   g. To check the results, follow the procedure described in step 2.

2.3.3 AIX: Creating Volume Groups

Prerequisites

Before you start to set up the necessary file systems for SAP, carefully plan their distribution to available disks. Decide whether one or more volume groups (VGs) should be used.

Procedure

1. Enter the command **`smitty`**.
2. Select:

   Physical & Logical Storage → Logical Volume Manager → Volume Groups → Add a Volume Group
3. Enter a volume group name, for example, **`sapr3vg`**.

   Press `F4` to get a list of disks, which have not yet been assigned to a volume group.
4. With `F7`, mark the disks you want to be included in the volume group and press `ENTER`.
5. Press `F10` to exit.

Repeat these steps for each volume group that has to be created.

2.3.4 AIX: Setting up File Systems

1. Create one logical volume for each file system listed in the appropriate SAP profile.
   a. Select:

      Physical & Logical Storage → *Logical Volume Manager* →
      Logical Volumes → *Add a Logical Volume*

   b. Enter a volume group name, for example, **`sapr3vg`**.
   c. Enter a logical volume name, for example, **`lvsap01`**.
   d. Enter the number of logical partitions (see section *Size of a Logical Partition*).
   e. Press `F3` until the *Physical & Logical Storage* menu appears.
2. Create the file systems.
   a. Select:

      Physical & Logical Storage → *File Systems* →
      Add/Change/Show/Delete File Systems → *Journaled File Systems* →
      Add a Journaled File System on a previously defined Logical Volume

   b. Press `F4` to get a list of logical volumes.
   c. Select one logical volume.

      Select the logical volume with the desired size.

      Enter mount point as given in the file system list files.

      DB2 UDB for UNIX and Windows: **`/db2/<SAPSDID>`**
      Informix: **`/informix/<SAPSID>`**
      Oracle: /**`oracle/<SAPSID>/sapdata1`**
      SAP DB: **`/sapdb/<SAPSID>/sapdata`**

      Set *mount automatically* to **`yes`**.

   d. Press `F10` to exit `smitty`.

2.3.5 AIX: Setting up Raw Devices

1. Create volume group.

   See the section *Creating Volume Groups*.
2. Create logical volume:
   a. Enter the command **`smitty`**.
   b. Select:

      Physical & Logical Storage → *Logical Volume Manager* →
      Logical Volumes → *Add a Logical Volume*

c. Enter volume group name.

d. Enter logical volume name (for example: **`lvsap02`**).

e. Enter the number of logical partitions.

f. Set the logical volume TYPE to `r3data` and press `ENTER`.

g. Press `F10` to exit SMIT.

The logical volume can be accessed through `/dev/r<name of raw log.vol>`.

Name of logical volume: `lvsap02` leads to the corresponding device name `/dev/rlvsap02`.

3. Accessing raw devices

For each logical volume, create a link to access the volume.

DB2 UDB for UNIX and Windows: Not necessary
Informix:
```
ln -s /dev/rlvsap01    \
      /informix/<SAPSID>/sapdata/physdev1/data1
```
Oracle: Not necessary
SAP DB: Not necessary (R3SETUP creates the links)

2.4 AIX: Mounting Directories via NFS

Use

There are two ways of mounting directories via NFS:

- Manually
- Using SMIT

Procedure

Mounting Directories via NFS using SMIT

To mount directories via NFS from the host on which the central instance runs:

1. Log on as user **root**.
2. Use SMIT to start NFS services at the host on which the central instance runs:
   a. Enter the command **smitty**.
   b. Select:

      Communications Applications and Services → *NFS* →
      Network File System (NFS) → *Configure NFS on this System* → *Start NFS*

   c. In the line *Start NFS now*, *on system restart or both*, enter:
      both.
   d. Press ENTER.
3. Export the directory (for example /sapmnt/<SAPSID>/exe) with read or read-write access for the host on which the additional instance runs:
   a. Enter the command **smitty**.
   b. Select:

      Communications Applications and Services → *NFS* →
      Network File System (NFS) → *Add a Directory to Exports List*

   c. Enter the path of the directory which you want to export (for example, **/sapmnt/<SAPSID>/exe**).
   d. Choose export mode (use read-write or read-only as required by SAP).

      In the line *HOSTS allowed root access,* enter the name of the host on which the additional instance runs. For security reasons, this root access should be disabled after the installation.

   e. In the line *Export directory now, system restart or both*, enter:
      both
   f. Press ENTER.
4. Create the mount point at the host on which the additional instance runs:
   /usr/bin/mkdir /sapmnt/<SAPSID>/exe
5. Mount the directory on the host on which the additional instance runs
   a. Enter the command **smitty**.

b. Select:

Communications Applications and Services → *NFS* →
Network File System (NFS) → *Add a File System for Mounting*

c. Enter the path name of the mount point.

d. Enter the path name of the remote directory (the directory of the central instance).

e. Enter the host where the remote directory resides.

f. Set *Mount now, add entry to /etc/filesystems or both* to **both**.

g. Set */etc/filesystems entry will mount the directory on system RESTART* to **yes**.

h. Change *Mount file system soft or hard* to **soft** and press ENTER.

6. If you exported the directory with read-write access, check whether the host on which the future additional instance will run has write access to the directory `/sapmnt/<SAPSID>/exe`:

cd /sapmnt/<SAPSID>/exe
touch test (create a file called test)
ls -l test (check if file test is created)
rm test (remove the file test)

Mounting Directories via NFS manually

To mount directories via NFS from the host on which the central instance runs:

1. Log on as user **root**.
2. Enter the directories you wish to export into the file `/etc/exports`.
3. Enter:
exportfs -a

2.5 AIX: Creating UNIX Groups and Users

Prerequisites

Informix :
Due to operating system functions, it is possible that the group `super_archive` is not inserted correctly in the `/etc/group` file. Check whether the group name was truncated to `super_` here and change it to `super_archive`.

Procedure

For the users and groups created by R3SETUP, see the appropriate section.

Creating UNIX Groups and Users

Use SMIT to create UNIX groups and users:

1. Create groups:
   a. Enter the command **`smitty`**.
   b. Select:

      Security & Users → *Groups* → *Add a Group*

   c. Enter a group name, for example, **`sapsys`**.

      Set *administration group* to `true`.

   d. Press `F3` until the *Security & Users* menu appears.
2. Create users:
   a. Select *Security & Users* → *Users* → *Add a User*
   b. Enter a user name, for example **`<sapsid>adm`**.
   c. Enter all necessary values.
3. Set the initial password using the following command:
   `passwd <user>`

```
passwd <sapsid>adm
```

Checking Created Users

Check, as **`root`**, all existing users To do this:

1. Enter the command **`smitty`**.
2. Select:

   Security & Users → *Users* → *Change/Show Characteristics of a User*

3. Press F4 to get a list of users.

4. For user `root` and each created user `<user>`:
   a. Select `<user>`.
   b. Change field *Soft CPU time* to **`-1`** (this is the default value).
   c. Change field *Soft CORE file size* to **`2097151`** (this is the default value).
   d. Change field *Soft FILE size* to **`4194302`**.
   e. Change field *Soft DATA segment* to **`-1`**.
   f. Change field *Soft STACK size* to **`-1`**.

   You must make sure that the system-wide default *HARD* values are not explicitly defined to be lower than the number indicated above. Check the file `/etc/security/limits` under the *default:* stanza. If they are not explicitly set, then the values are as shown in the table at the top of the file.

Checking the Operating System

1. Enter the command **`smitty`**.
2. Select:

   System Environments → *Change/Show Characteristics of Operating System*
3. Change *Maximum number of PROCESSES allowed per user* to **`500`**.
4. Press `F10` to exit SMIT.

2.6 AIX: Troubleshooting

If you suspect that hardware problems are the reason for errors in your installation, analyze the system log. To do this, enter the command:

```
errpt -a | more
```

3 Compaq Tru64 UNIX: OS-Dependent Installation Steps

This documentation has been written with the support of the Compaq Competence Center. It is not intended to replace the documentation of the Compaq Tru64 UNIX (OSF/1) operating system.

3.1 Compaq: Mounting a CD-ROM

1. Log on as user **root**.
2. Create a mount point for the CD-ROM with the command:
   mkdir <CD-mountdir> (for example, <CD-mountdir> is /sapcd)
3. Mount the CD-ROM with the command:
   mount -dr -t cdfs /dev/rz<n>c <CD-mountdir>

   <n> is the hardware or SCSI address of the CD-ROM (0-6).

   If the command fails with the message Invalid argument, the kernel is not configured for the use of CD-ROMs. In this case, proceed as follows:

   a. Edit the file /usr/sys/conf/<HOST>
      Insert the following line in the appropriate section:
      options CDFS

   b. Generate a new kernel with
      /usr/sbin/doconfig -c <HOST>
      <HOST> is the hostname of the computer in uppercase letters.

   c. Copy the new kernel to the **root** directory
      cp /usr/sys/<HOST>/vmunix /.

   d. Reboot the system.

3.2 Compaq: Checking and Modifying the UNIX Kernel

Use

To run an SAP System, you must check and, if necessary, modify the Compaq Tru64 UNIX operating system kernel.

Procedure

Checking the UNIX Kernel

You can check the UNIX kernel with the following UNIX commands:

```
sysconfig -q ipc
sysconfig -q proc
sysconfig -q vm
```

Compare the output of these commands with the values in the list below.

Corrections of the kernel parameters are published in the installation notes. These notes contain the most recent information regarding the installation.

We recommend that all Compaq Tru64 UNIX kernel modifications be performed by your system administrator.

Changing the UNIX Kernel

To change the system configuration defined in `/etc/sysconfigtab`:

1. Create a file named `sapr3.stanza` in any directory `<path>` (for example, `/etc/sapr3.stanza`). It must contain the following lines for the subsystem entries `ipc` and `proc`:

```
ipc:
        sem-mni                     =          512
        sem-msl                     =          512
        sem-opm                     =          512
        sem-ume                     =          512
        shm-max                     =   2139095040
        shm-mni                     =          512
        shm-seg                     =          512

proc:
        maxusers                    =          512
        max-per-proc-address-space  =   8589934592
        max-per-proc-data-size      =   8589934592
        max-per-proc-stack-size     =    268435456
        max-proc-per-user           =          512
        per-proc-address-space      =   8589934592
        per-proc-data-size          =   8589934592
        per-proc-stack-size         =    268435456
```

2. Insert the following lines for the subsystem entry `vm` into the file `sapr3.stanza`.

```
vm:
        ubc-minpercent              =           3
        vm-ubcseqstartpercent       =           3
        ubc-borrowpercent           =           3
        ubc-maxpercent              =           7
        vm-maxvas                   = 17179869184
        vm-vpagemax                 =     1048576
```

3. Use the following commands to newly add the three subsystem entries to `/etc/sysconfigtab`:

   sysconfigdb -a -f <path>/sapr3.stanza ipc
   sysconfigdb -a -f <path>/sapr3.stanza proc
   sysconfigdb -a -f <path>/sapr3.stanza vm

   or if the subsystem entries are already present in `/etc/sysconfigtab` use:
   sysconfigdb -u -f <path>/sapr3.stanza
   to update the complete subsystem entries with your new settings in `sapr3.stanza`.

4. Reboot the system after changing the kernel parameters. Check the kernel parameter settings after rebooting the system.

3.3 Compaq: File Systems, Raw Devices and Swap Space

The following describes how to create the volume groups, file systems, raw devices and swap space for your installation.

The structure of the SAP system depends on the disk configuration of the customer system. For this reason, there are no general procedures for setting up the file systems.

The installation documentation (section Requirements Checklist for Compaq Tru64 UNIX) lists the space requirements for the swap space.

Consider the following when you specify the sizes of file systems:
- Your operating system needs up to 10% of disk space for administration.
- 1 MB is 1024 KB

See the installation documentation for additional planning information.

For new installations, you are recommended to use the standard file system or the Advanced File System (ADVFS, license required).

In both cases, you can also use LSM (logical storage manager) volumes (license required). To do this, replace `<diskname>` by the name of the created logical volumes in the following sections.

See the following sections for more information:

- Preparing Hard Disks [page 23]
- Setting up Swap Space [page 26]
- Setting up Standard File Systems [page 27]
- Setting up Advanced File Systems [page 28]
- Setting up Raw Devices [page 29]

3.3.1 Compaq: Preparing Hard Disks

`disklabel` Command

The command **`disklabel`** displays the start/stop cylinder, the size of the partition, the label and the disk type of a disk device.

As user **`root`**, enter the command:
`disklabel -r /dev/<rawdisk device>`

```
disklabel -r /dev/rrz3c
```

3.3 Compaq: File Systems, Raw Devices and Swap Space

Compaq Tru64 UNIX does not allow block 0 to block 15 on the disk to be part of the raw device used by the database system. When the machine is rebooted, the operating system writes a label here. Consequently, data from the database is overwritten and a restore of the database becomes necessary.
If you use a RAID System or Advanced File System (advfs) then the device driver will take care of this restriction.
Otherwise, when you initially setup a raw device, the standard 'c' partition (`/dev/rrz?c`) may **not** be used since it normally includes track zero on the disk. A disk that is intended for use as a raw device, must be reformatted appropriately to ensure that the raw device begins at track two.
Check your disk configuration as soon as possible to make sure that your raw devices meet the above requirements. To do this:

a. Log on as user **`root`** and change to the directory `/<db-system>/<SAPSID>/sapdata`
b. List the links to the raw devices with the command: **`ls -lR`**
c. Display all partitions on the disks with the command: **`disklabel -r /dev/<raw_device>`**

Assume **`ls -lR`** produces the output
`lrwxrwxrwx 1 root 18 Nov 29 data3 -> /dev/rrz3h`
then the partition `h` of this disk is used as a raw device.

Always use partition `c` of the raw device in the **`disklabel`** command to display the content of the entire disk.

The output of **`disklabel -r /dev/rrz3c`** looks similar to:

```
# /dev/rrz3c:
type: SCSI
disk: rz26
label:
flags:
bytes/sector: 512
sectors/track: 57
tracks/cylinder: 14
sectors/cylinder: 798
cylinders: 2570
sectors/unit: 2050860
rpm: 3600
interleave: 1
trackskew: 0
cylinderskew: 0
headswitch: 0              # milliseconds
track-to-track seek: 0     # milliseconds
drivedata: 0
```

```
8 partitions:
#       size   offset    fstype   [fsize bsize   cpg]
 a:  131072        0    unused     1024  8192     # (Cyl.     0 - 164*)
 b:  262144   131072    unused     1024  8192     # (Cyl.  164*- 492*)
 c: 2050860        0    unused     1024  8192     # (Cyl.     0 - 2569)
 d:  552548   393216    unused     1024  8192     # (Cyl.  492*-1185*)
 e:  552548   945764    unused     1024  8192     # (Cyl. 1185*-1877*)
 f:  552548  1498312    unused     1024  8192     # (Cyl. 1877*-2569*)
 g: 1657644   393216    4.2BSD     1024  8192 16  # (Cyl.  492*-2569*)
 h:  838444  1212416    unused     1024  8192     # (Cyl. 1519*-2569*)
```

Since partition `a` contains track zero, it is necessary to choose an other partition as raw device. To edit and change this output for (re-)partitioning the disk, it should be directed in a file using the command:
`disklabel -r /dev/rrz3c > <tmp_file>`

Partitioning Disks

1. Edit the output file `<tmp_file>` you generated in the previous step and change the boldface lines accordingly:

```
# /dev/rrz3c:
type: SCSI
disk: rz26
label:
flags:
bytes/sector: 512
sectors/track: 57
tracks/cylinder: 14
sectors/cylinder: 798
cylinders: 2570
sectors/unit: 2050860
rpm: 3600
interleave: 1
trackskew: 0
cylinderskew: 0
headswitch: 0                   # milliseconds
track-to-track seek: 0  # milliseconds
drivedata: 0

8 partitions:
#     size  offset  fstype [fsize bsize      cpg]
a:  ....16       0  unused   1024  8192  16       # (Cyl.    0 - 164*)
b:  262144  131072  unused   1024  8192           # (Cyl.  164*- 492*)
c: 2050860       0  unused   1024  8192           # (Cyl.    0 - 2569)
d:  552548  393216  unused   1024  8192           # (Cyl.  492*- 1185*)
e:  552548  945764  unused   1024  8192           # (Cyl. 1185*- 1877*)
f:  552548 1498312  unused   1024  8192           # (Cyl. 1877*- 2569*)
g: 1657644  393216  4.2BSD   1024  8192  16       # (Cyl.  492*- 2569*)
h:  131056      16  unused   1024  8192  16       # (Cyl. 1519*- 2569*)
```

- Only partitions a and h were modified !
- Never change partition c because it always represents the entire disk.
- Partition a should always span block 0 and block 1.
- Partitions a and c are the only ones containing track zero.
- After activating this partition all partitions but a and c can be used as raw devices.

2. Activate the disk using the command:
   `disklabel -R -r <raw_device> <tmp_file> <disktype>`

   where `<disktype>` is the value of `disk:` in the above `<tmp_file>`.

   `disklabel -R -r /dev/rrz3c <tmp_file> rz26`

   Always use partition `c` of the raw device in the above command.

3.3.2 Compaq: Setting up Swap Space

1. Determine the size of the installed RAM with the command:
   `vmstat -P | head -2`
2. Determine the allocated swap space with the command:
   `/usr/sbin/swapon -s`
3. Determine the required size of the swap space.

   You can find the minimum swap space requirement in the installation documentation in section Hardware and Software Requirements Check.

   If possible, you should use an entire physical disk as swap space partition.
4. Increase the swap space if necessary.

   To define additional swap partitions, insert the entry:
   `/dev/<disk partition> swap<n> ufs sw 0 2` in the file `/etc/fstab`.

   Insert entry `/dev/rz1b swap1 ufs sw 0 2`
5. To activate new swap partitions, use the command:
   `/usr/sbin/swapon -a`
6. Check the results by using the command:
   `/usr/sbin/swapon -s`
7. Execute `memlimits`, to verify paging space size and kernel settings.

   Do not execute this step if you are installing a standalone DB server.

- Make sure that the UNIX kernel, paging space and user limits are already configured for the SAP System as described in the installation documentation in section Hardware and Software Requirements Check and in section Compaq: Checking and Modifying the UNIX Kernel [page 21].
- Unpack the file memlimits:

  ```
  cd <INSTDIR>

  /<CD-DIR>/UNIX/<OS>/SAPCAR -xgvf     \
     /<CD-DIR>/UNIX/<OS>/SAPEXE.SAR memlimits
  ```

- Start `memlimits`. Enter:
  `./memlimits -l 20000`

  If error messages occur, increase your paging space and rerun `memlimits` until it is error free.

3.3.3 Compaq: Setting up Standard File Systems

The following table shows the variables and their corresponding values.

Variables in Standard File Systems

Variable	Description
`<diskname>`	Full path of a disk, for example, `/dev/rz11c`
`<rawdiskname>`	Full path of a raw disk, for example, `/dev/rrz11c`
`<disktype>`	Disk type, for example, `rz28`
`<mountpointname>`	Full path of a mount point, for example, SAP DB: `/sapdb/<SAPSID>/sapdata1` Informix: Not applicable Oracle: `/oracle/<SAPSID>/sapdata1`

To set up a standard file system:

1. Ask the customer which disks are available and make sure that they do **not** belong to one of the categories listed in the following table:

   Avoidable Disk Categories

Disk Category	How to Check
Mounted disks	`grep <diskname> /etc/fstab`
Swapdevices	`swapon -s | grep <diskname>`
Advanced FS	`cd /etc/fdmns; showfdmn * | grep <diskname>`
Volumes within LSM (Logical Storage manager)	`volprint -Ath | grep <diskname>`

2. Create disk label, if necessary:
   `disklabel -w -r <rawdiskname> <disktype>`
3. Create empty file system:
   `newfs -m 0 <rawdiskname>`

4. Create mount point:
   mkdir <mountpointname>
5. Add line to /etc/fstab:
   <diskname> <mountpointname> ufs rw 1 0
6. Mount file system:
   mount <mountpointname>

3.3.4 Compaq: Setting up Advanced File Systems

The following table shows the variables and their values.

Variables in Advanced File Systems

Variable	Description
<diskname>	Full path of a disk, for example, /dev/rz11c
<disktype>	Disk type, for example, rz28
<mountpointname>	Full path of a mount point, for example, SAP DB: /sapdb/<SAPSID>/sapdata1 Informix: Not applicable Oracle: /oracle/<SAPSID>/sapdata1
<domainname>	ame of a domain, for example, sapdomain1
<filestename>	Name of a file set, for example, sapfileset1

To set up an advanced file system:

1. Ask the customer which disks are available and make sure that they do **not** belong to one of the categories listed in the following table:

   Avoidable Disk Categories

Disk Category	How to Check
Mounted disks	grep <diskname> /etc/fstab
Swapdevices	swapon -s | grep <diskname>
Advanced FS	cd /etc/fdmns; showfdmn * | grep <diskname>
Volumes within LSM (Logical Storage manager)	volprint -Ath | grep <diskname>

2. Check required licenses:
   lmf list | grep ADVFS-UTILITIES
3. Create file domain:
   mkfdmn -t <disktype> <diskname> <domainname>
4. To add more disks to the domain, specify:
   addvol -t <disktype> <diskname> <domainname>
5. Install one or more file sets in the domain:
   mkfset <domainname> <filesetname>
6. Create mount point:
   mkdir <mountpointname>

7. Add line to `/etc/fstab`:
   `<domainname>#<filesetname> <mountpointname> advfs rw 1 0`
8. Mount file system:
   `mount <mountpointname>`

3.3.5 Compaq: Setting up Raw Devices

Use

Some DBMSs prefer raw devices. Therefore, you have to maintain your disks for raw device access. Keep the following points in mind:

- The installation tool can handle disk partitions up to 2 GB
- The first few disk blocks should not be contained in an active raw device partition
- You are recommended **not** to use a disk with a raw device partition for anything other than raw devices (that is, if one partition of a disk is used as raw device, no other partition of this disk should be used for file system, swap space, or anything else except raw devices.)
- You should use a disk for raw device **only** when you are absolutely sure that this disk is not used in any other way.

Procedure

1. Ask the customer which disks are available and make sure that they do **not** belong to one of the categories listed in the following table:

   Avoidable Disk Categories

Disk Category	How to Check
Mounted disks	`grep <diskname> /etc/fstab`
Swap devices	`swapon -s | grep <diskname>`
Advanced FS	`cd /etc/fdmns; showfdmn * | grep <diskname>`
Volumes within LSM (Logical Storage manager)	`volprint -Ath | grep <diskname>`

2. Create disk label, if necessary:
   `disklabel -w -r <rawdiskname> <disktype>`
3. Modify the partition table of a particular raw device `<rawdiskname>`.

   `prompt>` **`disklabel -r -e <rawdiskname>`**

If you perform the above command with an existing device name, you will be put into an editor and you can edit the partition lines as required. For example, if you want to use partition b to hold the whole disk except the first 16 blocks as your raw partition, your partition table should look as follows:

```
#      size    offset     fstype    [fsize bsize    cpg]
a:    131072         0      ADVfs                      # (Cyl.     0 - 8
b:   4110464        16     unused      1024   8192      # (Cyl.    82*- 3
c:   4110480         0     unused      1024   8192      # (Cyl.     0 - 2
d:    442343    131072     unused      1024   8192      # (Cyl.    82*- 3
e:    442343    131072     unused      1024   8192      # (Cyl.    82*- 3
f:    442343    131072     unused      1024   8192      # (Cyl.    82*- 3
g:   3537065    573415      ADVfs                      # (Cyl.   362*- 2
h:    442343    131072     unused      1024   8192      # (Cyl.    82*- 3
```

Do not use partitions `a` and `c` because they contain the first 16 blocks of the disk.

3.4 Compaq: Mounting Directories via NFS

1. Log on as user **`root`** to the NFS server and proceed as follows (The following assumes that the central instance host is the NFS server):
   a. Make sure that your host is configured as NFS server. You can set up your host as NFS server with the command:
   `/usr/sbin/nfssetup`
   b. For each directory (being in a local file system) you want to export add a line to the local file `/etc/exports`.

```
#export /sapcd to host osfsap2
/sapcd  -ro  osfsap2
#export this directory read and write to all hosts
/usr/sap/trans  -root=  <nfs_cli_host_1> : ... :
<nfs_cli_host_n>,
   access=      <nfs_cli_host_1> : ... : <nfs_cli_host_n>
```

(nfs_cli_host_n: hostname of the NFS client where you want to install the SAP instance).

```
#export the directories (set remote root uid = 0)
/sapmnt/<SAPSID>/profile  -root= <nfs_cli_host_1>: ...
/sapmnt/<SAPSID>/global   -root= <nfs_cli_host_1>: ...
```

For security reasons, the option
`-root= <nfs_cli_hostname_1>: ... :<nfs_cli_host_n>`
should be used during installation only.

2. Log on as user **`root`** to the host where the file system should be imported:
   a. For each remote file system, add a line to `/etc/fstab`

```
<nfs_server>:/usr/sap/trans   /usr/sap/trans  nfs rw 1 0
```

   b. Mount the file systems with:
   `mount  <nfs_server>:<file_system>  <mount_point>`

```
mount   <nfs_server>:/usr/sap/trans   /usr/sap/trans
```

3.5 Compaq: Creating UNIX Groups and Users

The following table contains commands for creating and administering UNIX groups and users.

Commands for Creating and Administering UNIX Groups and Users

Action	Command
Creating UNIX groups	`addgroup`
Creating UNIX users	`adduser`
Modifying UNIX users	`vipw`
Deleting UNIX users	`userdel`

If home directory `/usr/sap/C11` is chosen, **`adduser`** creates the home directory `/usr/sap/C11/c11adm`
You cannot force adduser to use `/usr/sap/C11` as the home directory for the user **`c11adm`**. You must edit `/etc/passwd` and correct the entry of the home directory to `/usr/sap/C11`. See the man pages of `passwd` to find out the structure of `passwd`-entries.

3.6 Compaq: Troubleshooting

If you suspect that hardware problems are the reason for errors in your installation, analyze the system log. To do this, enter the command:

```
/usr/sbin/uerf  -R  |  more
```

4 HP-UX: OS-Dependent Installation Steps

This documentation describes the OS-dependent installation steps for HP-UX. It has been written with the support of the HP Competence Center. It is not intended to replace the documentation of the HP-UX operating system.

4.1 HP-UX: Mounting a CD-ROM

Mounting a CD-ROM Manually

1. Log on as user **root**.
2. Create a mount point for CD-ROM with the command:
   mkdir <CD-mountdir>

   (usually <CD-mountdir> is /sapcd).
3. Make sure that the driver is part of the kernel (skip this step if the CD drive is already working):
   grep cdfs /stand/system

   If the driver is not configured, you have to add the string cdfs to the file /stand/system and rebuild the kernel. The procedure to build a new kernel is described in Checking and Modifying the UNIX Kernel [page 33] (section Manual Configuration of the Kernel). Reboot the system after rebuilding the kernel.
4. Mount the CD-ROM with the command:
   mount -r -F cdfs /dev/dsk/<diskdevice> <CD-mountdir>

   <diskdevice> is c0t4d0, a CD drive with hardware address 4.

Mounting a CD-ROM Using SAM

1. Enter the command
   /usr/sbin/sam
2. Select:

   Disks and Filesystems → *Disk Devices* → *Actions* → *Mount*
3. Enter the mount directory
   <CD-mountdir>

   (for example, <CD-mountdir> is /sapcd).
4. Perform task.
5. Exit SAM.

4.2 HP-UX: Checking and Modifying the UNIX Kernel

Use

To run an SAP System, you must check and, if necessary, modify the UNIX operating system kernel.

We recommend that all UNIX kernel modifications be performed by your UNIX system administrator.

There are two ways to perform the necessary changes to the UNIX operating system:

- Manually
- Using SAM

Reboot the system after changing the kernel parameters.

Procedure

1. Check the following table for the recommended kernel parameters for HP-UX.

Corrections of the kernel parameters are published in the installation notes. These notes contain the most recent information regarding the installation.

Only HP-UX 11.x
If a default kernel value is greater than the ones suggested under SAP Requirements, do not change the value.

Recommended Kernel Parameters for HP-UX

Parameter Group	Parameter	Description	32-bit SAP Kernel Requirements	64-bit SAP Kernel Requirements
Memory Paging	maxswapchunks	Maximal number of swap space chunks	≥ 1200	≥ 2048
Message Parameters	msgseg	Message segments	32767	32767
	msgssz	Message segment size	32	32
	msgmnb	Max. sum of messages in a queue	65535	65535
	msgtql	Max. number of message headers	2046	2046
	msgmap	msgtql + 2	2048	2048
Semaphores	semaem	'adjust on exit' -Max. value	16384	16384
	semmnu	No.of semaphores UNDO	≥ 256	≥ 256
	semmns	No.of semaphores	1024	1024
	semmni	Semaphores keys	520	520
	semume	UNDO keys	100	100
	semvmx	Max.value semaphore	32767	32767
Shared Memory	shmmax	Max Shared memory	1.073.741.824	≥ 17.179.869.184
	shmmni	Max. Shared Memory keys	≥ 256	≥ 256
	shmseg	Shared Memory segments	≥ 100	200
File System	maxfiles	soft-limit opened files	512	1024
	maxfiles_lim	hard-limit opened Files	1024	2048
	nflocks	No. of file locks	200	800
	nfile	No. of files	8192	8192
	ninode	No. of open inodes	8192	8192

Process Management	maxuprc	No. of processes per user	100	400
	maxdsiz	Max. data segment size	990.056.448	990.056.448
	maxdsiz_64bit	Max. data segment size 64-bit	N/A	4.294.967.296
	nproc	No. of processes system-wide		10 * maxusers
Miscellaneous Parameters	maxusers	Max. "USER" (system resource allocation)	128	128
	dbc_min_pct	Min. percentage of memory to be used by dynamic buffer cache	5	5
	dbc_max_pct	Max. percentage of memory to be used by dynamic buffer cache	8	8

The values listed in the table above are required by the SAP System.
If the currently configured kernel parameters have already been set by another application to approximately the values recommended by SAP, the values shown in the table may not be high enough. The increase could be insufficient. As a result, the start-up of the SAP System or of other applications may fail.

2. Adapt the UNIX kernel values manually or using SAM if necessary.

Manual Configuration of the Kernel

1. Change the kernel parameters according to the table *Recommended Kernel Parameters for HP-UX* in the file
   /stand/system
2. Generate a new kernel after making the changes using the following command:
   mk_kernel -o /stand/vmunix -s /stand/system
3. Reboot your system.

Configuration of the Kernel Using SAM

1. Enter the command
   `/usr/sbin/sam`
2. Select:

   Kernel Configuration → *Configurable Parameters*
3. Choose the parameter you want to modify and select:

   Actions → *Modify Configurable Parameter*
4. Modify all kernel parameters according to the table *Recommended Kernel Parameters for HP-UX*.
5. Select *Process New Kernel* from the *Actions* menu.
6. Exit SAM.
7. Reboot your system.

4.3 HP-UX: File Systems, Raw Devices and Swap Space

The following describes how to create the volume groups, file systems, raw devices and swap space for your installation.

The structure of the SAP system depends on the disk configuration of the customer system. For this reason, there are no general procedures for setting up the file systems.

Consider the following when you specify the sizes of file systems:
- Your operating system needs up to 10% of disk space for administration.
- 1 MB is 1024 KB.

See the following sections for more information:

- Setting up Swap Space [page 36]
- Setting up File Systems using LVM [page 38]
- Setting up File Systems using SAM [page 40]
- Accessing Raw Devices [page 41]

4.3.1 HP-UX: Setting up Swap Space

1. Find the minimum swap space requirements in the installation documentation in section Hardware and Software Requirements Check.
2. Find out whether you have to increase the swap space:
   a. Determine the size of the installed RAM: Run SAM and choose:

      Performance Monitors → *System Properties* → *Memory*
   b. Find out how much swap is currently configured on your system:
      `/usr/sbin/swapinfo`

3. Increase the swap space if necessary. You have two options:
   - Manually
   - Using SAM
4. Verify the paging size and the kernel settings if you are not installing a standalone DB server.

Setting up Swap Space Manually

1. Create a logical volume:
   `lvcreate -n <LVName> /dev/<VGName>`
2. Define size and allocate the logical volume to a disk:

   ```
   lvextend -L <size in MB> /dev/<VGName>/<LVName>  \
   /dev/dsk/<diskdevice>
   ```
3. Add the following entry to the file `/etc/fstab`:
   `/dev/<VGName>/<LVName>  /swap swap  defaults 0 0`
4. Activate the swap space with the command:
   `/usr/sbin/swapon -a`
5. Check activation with the command:
   `/usr/sbin/swapinfo -tm`

Setting up Swap Space Using SAM

1. Enter the command
   `/usr/sbin/sam`
2. Select

   Disks and Filesystems → *Swap* → *Actions* →
   → *Add Device Swap* → *Using the LVM*
3. Choose a partition for swap and choose *OK*.
4. Exit SAM.

Verify paging space size and kernel settings

Do not execute this step if you are installing a standalone DB server.

1. Make sure that the UNIX kernel, paging space and user limits are already configured for the SAP System as described in the installation documentation in section Hardware and Software Requirements Check and the above text.
2. Execute `memlimits`, to verify paging space size and kernel settings.
   a. Unpack the file `memlimits`:

   ```
   cd <INSTDIR>

   /<CD-DIR>/UNIX/<OS>/SAPCAR
   -xgvf /<CD-DIR>/UNIX/<OS>/SAPEXE.SAR memlimits
   ```

b. Start `memlimits`.

For a 32-bit SAP kernel, enter: **`./memlimits`**

For a 64-bit SAP kernel, enter: **`./memlimits -l 20000`**

3. If error messages occur, increase your paging space and rerun `memlimits` until it is error free.

4.3.2 HP-UX: Setting up File Systems using LVM

Use

Using a logical volume manager allows you to distribute partitions (logical volumes) across several disks (physical volumes). The individual logical volumes are grouped together into volume groups.

File systems can be larger than physical disks, but not larger than the volume group.

Procedure

1. Examine device configuration

   Enter the command:
   `ioscan -f -C disk`

   This command provides the logical unit (LU) number and the hardware addresses of all available devices, using the device class disk.

   The following command scans all disks for logical volumes:
   `vgscan -pv`

   Make sure you use option **`-p`** (preview), otherwise `/etc/lvmtab` will be updated. LVM can coexist in a system that uses fixed partitions.

2. Prepare disks

   To assign an unused disk to a physical volume, enter:
   `pvcreate /dev/rdsk/<diskdevice>`

3. Create volume group directory `<VG Name>` and group device file (for example `SAPR3`). For each volume group in the system, there must be a volume group directory that has a character device file named group in it:

   `mkdir /dev/<VG Name>`
   `mknod /dev/<VG Name>/group c 64 0x<nn>0000`

4. Create the volume group

   To create a volume group, you specify which physical volumes (disks) belong to the group:
   `vgcreate /dev/<VG Name> /dev/dsk/<diskdevice>`

   To add another disk to an existing volume group, enter:
   `vgextend /dev/<VG Name> /dev/dsk/<diskdevice>`

5. Examine the size of volume group

   To see how many physical disks you have in a volume group, enter:
   `vgdisplay /dev/<VG Name>`

6. Calculate the free space in the volume group:
   `FREE_Space = Free physical extents (PEs) * PE_Size`

7. Create one logical volume for each file system listed in SAPFS.PAR:
   `lvcreate  /dev/<VG Name>`

   Allocate the logical volume to a disk with the command:

   ```
   lvextend -L <size in MB> /dev/<VGName>/<LVName>  \
   /dev/dsk/<diskdevice>
   ```

   SAPFS.PAR lists the required size for each file system. The size <in MB> should be a multiple of PE_Size, or the size will be rounded up.

   You can find out the size of the logical volume with either of these commands:

   ```
   vgdisplay  -v  /dev/<VG Name>
   lvdisplay  /dev/<VG Name>/<LV Name>
   ```

   Write down the device names of the logical volumes (for example, `lv12`). You will need the device names when creating and mounting the file systems.

   The following steps are only needed for file systems, not for raw devices. If you set up raw devices, see section *Accessing Raw Devices* for more information.

8. Determine the disk type with the command:
   `diskinfo /dev/rdsk/<diskdevice>`
9. Create the file systems required by SAP.

   For `sapdata1` to `sapdata<n>` enter:

   - **HFS:**
     ```
     newfs -L -i 20000 -F hfs -b 8192 -f 8192 -m 1  \
     /dev/<VG Name>/r<LV Name>
     ```
   - **JFS:**
     ```
     newfs -F vxfs -b 8192   /dev/<VG Name>/r<LV Name>
     ```

   For all others, enter:

   - **HFS:**
     ```
     newfs -L -F hfs /dev/<VG Name>/r<LV Name>
     ```
   - **JFS:**
     ```
     newfs -F vxfs /dev/<VG Name>/r<LV Name>
     ```

   The logical volume is identified by the device file that you defined when you created the logical volume.
10. Create mount directories.
11. Add the new file system to the `/etc/fstab`.

    HFS:
    ```
    /dev/<VG Name>/<LV Name> /<mountdir> hfs defaults 0 2
    ```
    JFS:
    ```
    /dev/<VG Name>/<LV Name> /<mountdir> vxfs delaylog, \
     nodatainlog 0 2
    ```

12. Mount the file systems using the command:
    `mount -a`

The mount sequence is determined via the file `/etc/fstab`.

4.3.3 HP-UX: Setting up File Systems using SAM

Use

SAM is not able to build file systems with 8 K fragment size.

Procedure

1. Enter the command
   `/usr/sbin/sam`
2. Select

   Disks and Filesystems → *File Systems* → *Actions* → *Add Local File System* → *Using the LVM*
3. Add all disks.
4. Select

   Disks and Filesystems → *Volume Groups* → *Actions* → *Create*
5. Create all volume groups.
6. Select

   Disks and Filesystems → *Logical Volumes* → *Actions* → *Create*
7. Create all logical volumes.
8. Exit SAM.

SAM already creates filesystems with the `newfs` command, but the filesystems for `sapdata1` to `sapdata<n>` have to be rebuild with special options.

9. Perform the steps 9 to 12 of Setting up File Systems Using LVM [page 38]:
   - Create the file system required by SAP
   - Create mount Directories
   - Add the new file system to the `/etc/fstab`
   - Mount the file system using the command

4.3.4 HP-UX: Accessing Raw Devices

File systems and raw devices differ in the way that data is written to and read from disk:

- **Buffering:** Reads and writes to a file system are buffered in a UNIX system. To be absolutely sure that all data is physically present on a disk, the buffers and files must be synched.

  Writes to a raw device are unbuffered: the system writes directly to the disk. These writes are faster and more secure.

- **File access:** Accessing files on a UNIX file system is transparent. Accessing data on a raw device is only possible with a special application.

Some databases prefer raw devices. To access these raw devices you can create symbolic links. The following examples show symbolic link commands:

SAP DB: not necessary, because R3SETUP creates the links
Informix:

```
ln -s /dev/rdsk/<diskdevice>    \
   /informix/<SAPSID>/sapdata/physdev1/data1
ln -s /dev/rdsk/<diskdevice>    \
   /informix/<SAPSID>/sapdata/physdev<m>/data<n>
```

Oracle: not necessary

4.4 HP-UX: Mounting Directories via NFS

Use

There are two ways of mounting directories via NFS:

- Manually
- Using SAM

Procedure

Mounting Directories via NFS using SAM

On the host on which the main instance runs, proceed as follows:

1. Enter the command
   `/usr/sbin/sam`
2. Select

   Networking and Communications → *Networked File Systems* →
   → *Exported Local File Systems* → *Actions* → *Add*
3. Enter the *Local Directory Name*

   ```
   /sapmnt/CUS
   ```
4. Select *Specify UID for unknown user* and enter at *User ID* the value `0`.
5. Select *Specify Root User Access* and add the *Remote System Names*.
6. Type `OK`.
7. Exit SAM.

On the host on which the additional instance runs:

1. Enter the command
   `/usr/sbin/sam`
2. Select

   Networking and Communications → *Networked File Systems* →
   → *Mounted Remote File Systems* → *Actions* → *Add Remote File Systems* →
   → *Using NFS*
3. Enter:
   - The *Local Directory Name*
   - The *Remote Server Name* of the host exporting the file system
   - The *Remote Directory Name*
4. Enable the mount *Now* and *On boot.*
5. Specify *Read-Write Access.*

6. Type ok.
7. Exit SAM.

Mounting Directories via NFS manually

On the host on which the main instance runs, proceed as follows:

1. Add the file system you want to export to the file `/etc/exports` using the option

```
-root=   <nfs_cli_hostname_1>:...:< nfs_cli_hostname_n>,>   \
   access=     <nfs_cli_hostname_1>:...:< nfs_cli_hostname_n>
```

```
/sapmnt/C11/exe   -root=hw5111:hw5115,
access=hw5111:hw5115
```

When you encounter problems with your input similar to the example above, try something like
`/sapmnt/C11/exe   -anon=y, access=hw5111:hw5115`

For security reason the option
`-root= <nfs_cli__hostname_1>: ... :<nfs_cli_hostname_n>`
should be used during installation only.

2. Make the file system available to NFS clients:
`/usr/sbin/exportfs -a`

On the host on which the additional instance runs:

1. Add the remote file system to `/etc/fstab`

```
hwi173:/sapmnt/CUS  /sapmnt/CUS nfs defaults 0 0
```

2. Mount the file system

```
mount hwi173:/sapmnt/CUS
```

4.5 HP-UX: Creating UNIX Groups and Users

For the users and groups created by R3SETUP, see Network Information Service Environment [page 8].

Use SAM to create UNIX groups and users:

1. Enter the command
   `/usr/sbin/sam`
2. Select:

   Accounts for Users and Groups → *Local Users* → *Actions* → *Add*
3. Enter the required users.
4. Select:

   Accounts for Users and Groups → *Groups* → *Actions* → *Add*
5. Enter the required groups.
6. Exit SAM.
7. Verify that TZ settings in the following are consistent:

   ```
   /etc/TIMEZONE
   /etc/profile
   /etc/csh.login
   ```

4.6 HP-UX: Troubleshooting

If you suspect that hardware problems are the reason for errors in your installation, analyze the system log and perform the following steps:

1. Check your console system messages are printed here.
2. Check the output of the command
   `/usr/sbin/dmesg`

   The output is listed in `/var/adm/messages`.

5 IBM NUMA-Q: OS-Dependent Installation Steps

This documentation has been written with the support of the IBM NUMA-Q Competence Centers. It is not intended to replace the documentation of the IBM NUMA-Q operating system.

Other documentation is needed to configure the IBM NUMA-Q system for use by an SAP application.

5.1 IBM NUMA-Q: Preparing the Installation

The following table lists the directories that are required on the system before the installation can start.

Directories that need to be created before the Installation

Directory	Description
`/usr/sap/trans`	Directory where the configuration file `.sapconf` will reside, but it is not needed as the NFS mount point of a file system.
`/usr/sap/exe`	Directory that is used to store the local database binaries for the SAP instance.
`/oracle/stage`	Directory used by Oracle for a permanent staging area.
`/temp/install`	Installation directory referred to as `<INSTDIR>`. Do **not** use the `/tmp` directory since it is cleaned out upon rebooting of the system.

These directories need to be created by user `root` before the installation tool is started. The first three directories should have the file access modes of `0775`, while `/temp/install` should have the modes of `0777`.

In addition, any file systems that are to be mounted for the SAP instance and Oracle tablespaces need to be created, initialized and mounted before the installation tool is started.

5.2 IBM NUMA-Q: Mounting a CD-ROM

1. Log on as user `root`.
2. Insert the CD into the CD-ROM drive.
3. Create the mount point for the CD-ROM:
   `mkdir /sapcd`
4. Mount the CD-ROM:
   `/etc/mount -r -f cdfs -o showdot, toupper /dev/dsk/cd1  /sapcd`
5. Before exchanging the CD-ROM, the CD file system must be unmounted first, using the command:
   `/etc/umount /sapcd`

5.3 IBM NUMA-Q: Checking System and Kernel Parameters

To assure that the system and kernel parameters are set correctly, perform the following steps:

1. Check and modify the kernel.

   The following table shows the recommended kernel parameters for high performance SAP System installation on IBM NUMA-Q.

Recommended Kernel Parameters for IBM NUMA-Q

ptx/Admin Adjustment	Parameter	Installation Req.
set	BUFPCT	2
set	FDIV_BUG	0
set	NTLI	8192
set	NTRW	4160
set	MAXAIO	4096
set	NABUF	8192
set	NPROC	2048
set	MAXUP	2048
set	SHMMAX	2147483647
set	SHMGAP	2097152
set	SEMMNS	3100
set	SEMMNU	3100
set	SEMMSL	350
set	SEMUME	50
set	NFIFO	3000
set	NFILE	24000
set	NINODE	2048
set	NOFILETAB	12000
set	N_UDP_PCB_FREE	2048
set	TCP_SENDSPACE	32
set	MAXUSERS	1024
set	SHM_LOCK_UID	-1
set	SHM_LOCK_OK	1
set	TCP_DELAYACK	20
set	NABUF_DYN_DIV	1
set	NABUF_DYN_MULT	2
set	N_TCP_PCB_FREE	2048
set	N_TCP_PCB_HDRS	64

set	MSGMAX	16384
set	MSGMNB	65535
set	MSGSEG	4096
set	MSGSSZ	64
set	NTICLTS	2048
set	MSGTQL	90
set	SHMMNI	1000
set	SHM_LOCKDF_OK	1
set	SHM_LOCKDF_UID	-1

If a kernel parameter is not set within the current kernel or the current value of a kernel parameter is too low, the system configuration file needs to be updated and a new kernel built. Use the ptx/Admin tool **`menu`** to change the kernel parameters, to build a new kernel and to schedule the installation of the new kernel upon reboot of the system. For more information on how to use ptx/Admin, see the IBM NUMA-Q documentation *Dynix/ptx System Administration Guide Volume 1*, chapter 3.

2. Setting up VMTUNE parameters

   To better utilize the virtual memory system, the adjustments of the VMTUNE environment listed in the following table should be made.

   Recommended Virtual Memory Parameters for IBM NUMA-Q

Parameter	IBM NUMA-Q Default	Installation Requirement
maxdirty	400	10240
dirtyhigh	200	9216
dirtylow	100	8192

   This can be done by creating a file `/etc/rc2.d/S99vmtune` that contains:
   `/etc/vmtune -f -maxdirty 10240 -dirtyhigh 9216 -dirtylow 8192`.

   For more information, see the IBM NUMA-Q documentation *Dynix/ptx System Configuration and Performance Guide.*

3. Rebooting the system

   If a new kernel has been built, the system has to be booted.

   If kernel parameters were manually inserted into the system configuration file of the previous kernel, these parameters have to be reinserted into the system configuration file of the new kernel, and the kernel build has to be restarted.

   The new kernel is booted by entering:
   `init 6`

5.4 IBM NUMA-Q: File Systems, Raw Devices, Swap Space

The following describes how to create the volume groups, file systems, raw devices and swap space for your installation.

The structure of the SAP system depends on the disk configuration of the customer system. For this reason, there are no general procedures for setting up the file systems.

Consider the following when you specify the sizes of file systems:
- Your operating system needs up to 10% of disk space for administration.
- See *Installation Planning* in your installation documentation for additional disk planning information.

For new installations, you are recommended to use the standard file system partitions or file systems generated by the IBM NUMA-Q Volume Manager (ptx/SVM, license required).

See the following sections for more information:

- Preparing Hard Disks [page 48]
- Setting up Swap Space [page 49]

5.4.1 IBM NUMA-Q: Preparing Hard Disks

Use

Additional file systems are required for the database.

These file systems are created using either command line utilities or ptx/Admin (`menu`) as user `root`.

Procedure

1. Plan your total disk configuration
2. Use ptx/Admin to create custom Volume Table of Contents (VTOC) on each of the physical disks using the tool **`menu`**. Choose *System Administration* → *Disk Drive Management* → *Create a Custom VTOC File*
3. Format the partitions using the tool **`menu`**. Choose *System Administration* → *File System Management* → *Make a New Filesystem*
4. Add these file systems to the file system table `/etc/vfstab` using the tool **`menu`**. Choose *System Administration* → *File System Management* → *Add Filesystem Table Entry*
5. Label the partitions using the tool **`menu`**. Choose *System Administration* → *File System Management* → *Create-Change Disk Labels*
6. Mount these file systems using the tool **`menu`**. Choose *System Administration* → *File System Operations* → *Mount a Filesystem*
7. Check the free space available of the new files system using the tool **`menu`**. Choose *System Administration* → *File System Operations* → *Report Free Disk Space*.

Interpret the available free space by taking the value in the *Avail* column, which is reported in 512byte blocks, and divide by 2048 to determine the available free space in megabytes.

8. Check the mode of the mount points, which must be `0775`.

5.4.2 IBM NUMA-Q: Setting up Swap Space

1. Determine the size of the installed RAM by using the command:
   `/etc/showcfg | grep MEM`

   Sum up the memory entries `size=xxx.0mb (xxx.0mb)` of each MEM line to get the total memory of the system.

2. Determine the allocated swap space by using the command:
   `/etc/swap -l`

3. Determine the required size of the swap space.

   The minimum swap space requirement depending on installed RAM is displayed in the following table.

   Swap Space Requirements

Installed RAM	Required Swap Space
≤ 192 MB	1 GB
> 192 MB and ≤ 512 MB	1 GB + 2 * RAM
> 512 MB	3 * RAM

   If possible, you should use an entire physical disk as swap space partition.

   For a 64-bit SAP kernel, at least 20 gigabyte of swap space is recommended. For more information, see **SAP Note 153641**.

4. Increase the swap space if necessary.

   To define additional swap partitions, use the following command:
   `/etc/swap -a <partition> <length>`

   where `<partition>` is a device that is not being used by a file system.

   `/etc/swap -a /dev/vx/dsk/rootdg/sd1s2 0`

5. To activate new swap partitions, reboot the system using the command:
   `init 6`

6. Upon reboot of the system, check the results by using the command:
   `/etc/swap -l`

5.5 IBM NUMA-Q: Creating UNIX Groups and Users

Groups and users can be created manually using the ptx/Admin tool **`menu`**.

For the users and groups created by R3SETUP see Network Information Service Environment [page 8].

5.6 IBM NUMA-Q: Troubleshooting

If you suspect that hardware problems are the reason for errors in your installation, analyze the system log and perform the following steps:

1. Check your console. System messages are printed here.
2. Check the last lines of the system messages file:
   ```
   tail /usr/adm/messages
   ```

6 Linux: OS-Dependent Installation Steps

This documentation is not intended to replace the documentation of the Linux operating system.

6.1 Linux: Mounting a CD-ROM

1. Log on as user **root**.
2. Create a mount point for the CD-ROM with the command:
   mkdir <CD-mountdir> (for example, <CD-mountdir> is /sapcd)
3. Mount the first CD-ROM device with the command:
   mount -t iso9660 -r <device> <CD-mountdir>

   where **<device>** is **/dev/cdrom** for non-SCSI CD-ROM devices and **/dev/scd<n>** for SCSI drives with the device number **<n>**.

   If the file names on the mounted CD-ROM are written in lowercase letters, remount the CD-ROM with the following commands:
   umount <device>
   mount -t iso9660 -r -omap=off <device> <CD-mountdir>

6.2 Linux: Checking and Modifying the Linux Kernel

Use

Normally, checking and modifying the Linux Kernel is not necessary, as the kernel has been adjusted by SAP. Nevertheless, it is recommended to run a special test tool that analyzes several system parameters.

Prerequisites

Make sure that the SAP Kernel was booted. The following command should report a Linux kernel version adapted by SAP:
uname -a

Procedure

1. Install the test tool saposcheck provided on the kernel CD with the following command:
   rpm -i saposcheck.i386.rpm

   Check sapserv<x> for the newest version of saposcheck.

   If other RPM packages are missing on your system for saposcheck, you will be asked to install them. The according installation procedure is similiar to that of saposcheck. These supplementary RPMs can be found either on the kernel CD or on Linux Operating System CDs.

2. After the installation of saposcheck:
   a. As user **root**, enter the command:
      saposcheck
   b. Check the output for errors.

The saposcheck package contains a readme file that contains information on both the activities of saposcheck and the possibilities of setting some system parameters manually.

6.3 Linux: File Systems and Swap Space

The following describes how to create the volume groups, file systems and swap space for your installation.

The structure of the SAP system depends on the disk configuration of the customer system. For this reason, there are no general procedures for setting up the file systems.

The installation documentation (section Hardware and Software Requirements Check) lists the space requirements for the swap space.

Consider the following when you specify the sizes of file systems:
- Your operating system needs up to 10% of disk space for administration.
- 1 MB is 1024 KB

See the installation documentation for additional planning information.

See the following sections for more information:

- Preparing Hard Disks [page 52]
- Setting up Swap Space [page 53]
- Setting up Standard File Systems [page 53]

6.3.1 Linux: Preparing Hard Disks

Use

The commands **fdisk** or **cfdisk** display the start/stop cylinder, the size of the partition, the label and the disk type of a disk device.

Procedure

Execute one of the following commands as user **root**:
fdisk /dev/<device> or **cfdisk /dev/<device>**

6.3.2 Linux: Setting up Swap Space

Prerequisites

Make sure that the UNIX kernel, paging space and user limits are already configured as described in the installation documentation (section Hardware and Software Requirements Check) and in section Checking and Modifying the Linux Kernel [page 51].

Procedure

1. Determine the allocated swap space by using the command:
   `cat /proc/swaps`
2. Determine the required size of the swap space.

   You can find the minimum swap space requirement in the installation documentation (section Hardware and Software Requirements Check).

   If possible, you should use an entire physical disk as swap space partition.
3. To configure swap space, use the following command:
   `linuxconf`

   Maximum size for swap partitions is 2GB. If more swap space than 2 GB is required, use multiple partitions.
4. Unpack the file `memlimits`:
   ```
   cd <INSTDIR>
   /<CD-DIR>/UNIX/<OS>/SAPCAR -xgvf /<CD-DIR>/UNIX/<OS>/SAPEXE.SAR \
      memlimits
   ```
5. Start `memlimits`:
   `./memlimits -a yes_w -s mf`
6. If error messages occur, increase your paging space and rerun `memlimits` until it is error free.

   For more information, see **SAP Note 165892**.

6.3.3 Linux: Setting up Standard File Systems

Prerequisites

The following table shows the variables and their corresponding values.

Variables in Standard File Systems

Variable	Description
`<diskname>`	Full path of a disk, for example, `/dev/sda`
`<mountpointname>`	Full path of a mount point, for example, SAP DB: `/sapdb/<SAPSID>/sapdata1` Informix: `/informix/<SAPSID>/sapdata` Oracle: `/oracle/<SAPSID>/sapdata1` DB2 UDB: `/db2/<SAPSID>/sapdata1`

Procedure

1. Create empty file system:
   mke2fs -b 4096 /dev/<device>
2. Create mount point:
   mkdir <mountpointname>
3. Add line to `/etc/fstab`:
   /dev/<device> <mountpointname> ext2 defaults 1 0
4. Mount file system:
   mount <mountpointname>

6.4 Linux: Exporting Directories via NFS

To export directories via NFS, perform the following steps (the following assumes that the central instance host is the NFS server):

1. Log on as user **root** to the NFS server.
2. Make sure that your host is configured as NFS server by making sure that the output of the command

   chkconfig --list nfs

   looks like:
   ```
   nfs    0:off    1:off    2:off    3:on    4:on    5:on    6:off
   ```
   You can set up your host as NFS server using the command:
   /etc/rc.d/init.d/nfs start
3. To export a directory from a local file system you can
   - Use the tool **linuxconf**
   - Perform the configuration manually

   To perform the configuration manually:

   a. Add a line to the local file `/etc/exports:`

   ```
   #/etc/exports
   <directory>        <hostname>(<options>)
   ```

   To export the directory `/usr/sap/trans` in read-only mode to the NFS client host.`wdf.sap-ag.de`:
   ```
   #/etc/exports
   /usr/sap/trans      host.wdf.sap-ag.de(ro)
   ```
   To export the directory in read-write mode:
   ```
   #/etc/exports
   /usr/sap/trans      host.wdf.sap-ag.de(rw)
   ```
   To export the directory to all NFS clients of the domain using a wildcard (`*`):
   ```
   #/etc/exports
   /usr/sap/trans      *.wdf.sap-ag.de(rw)
   ```

There must **not** be a blank between hostname and options. Otherwise, the directory will be exported both read-only to hostname **and** in consideration of the options to everyone.

b. To activate the changes (that is, inform the NFS daemon about the changes performed in `/etc/exports`), enter:
`exportfs -r`

To get a list of all currently exported directories, enter
`exportfs -v`

For further details please consult the man page by entering `man exports`.

4. Log on as user **`root`** to the host where the file system should be imported.
5. Mount the file systems with:
`mount <nfs_server>:<file_system> <mount_point>`

```
mount   <nfs_server>:/usr/sap/trans   /usr/sap/trans
```

6.5 Linux: Creating Linux Groups and Users

For the users and groups created by R3SETUP see Network Information Service Environment [page 8].

You can administer UNIX users with the graphical tool **`userconf`** or with the following commands:

- Create UNIX groups with **`groupadd`**.
- Create UNIX users with **`useradd`**.
- Modify UNIX users with **`usermod`**.
- Delete UNIX users with **`userdel`**.

7 ReliantUNIX: OS-Dependent Installation Steps

This documentation has been written with the support of the SIEMENS Competence Center. It is not intended to replace the documentation of the ReliantUNIX operating system.

ReliantUNIX: Throughout this documentation, *RM600* refers to the *RM600* as well as to the *RM300* and *RM400* machines.

7.1 ReliantUNIX: Preparing the Installation

Installation Directory

In a standard ReliantUNIX installation, the contents of the `/tmp` directory is removed while rebooting the system. We recommend that you create another directory like `/temp` for the installation.

7.2 ReliantUNIX: Mounting a CD-ROM

1. Log on as user **`root`**.
2. Create a mount point for the CD-ROM with:
   `mkdir <mountpoint>`
3. Determine the device name of the CD-ROM with:
   `autoconf -l | grep CD`

   which produces output like
   `ios0/sdisk00?  OS02  (CD-ROM)`
4. Mount the CD-ROM with:
   `mount -F hs /dev/<device_name>s0  <mountpoint>`
5. Set the x-bits for the executables:
   `/usr/sbin/cdmntsuppl -F 0555 <mountpoint>`

```
mkdir /sapcd
mount -F hs /dev/ios0/sdisk005s0  /sapcd
cdmntsuppl -F 0555 /sapcd
```

7.3 ReliantUNIX: Checking and Modifying the UNIX Kernel

Use

To run an SAP System, you must check and, if necessary, modify the UNIX operating system kernel.

Procedure

The recommended values of the kernel parameters may change. Refer to **SAP Note** *SAP Software on UNIX - OS Dependencies*.

Kernel parameters should not be decreased when adapting the following requirements.

1. Check the following table for the recommended kernel parameters for ReliantUNIX.

Recommended Kernel Parameters for ReliantUNIX

Parameter Group	Parameter	Description	32-bit SAP Kernel Req.	64-bit SAP Kernel Req.
Semaphores	SEMMAP	Semaphores map	512 (=SEMMNS/2)	512 (=SEMMNS/2)
	SEMMNI	Semaphores keys	1024	1024
	SEMMNS	No.of semaphores	1024	1024
	SEMMNU	No.of semaphores UNDO	700	700
	SEMUME	UNDO keys	40	40
	SEMVMX	Max. value semaphores	32767	32767
	SEMMSL	Max. number of sem./ key	250 (if more than one SAP System)	250 (if more than one SAP System)
File System	NBUF	Buffer header file system code	512	512
	HFNOLIM	Hard limit opened files	2048	2048
	SFNOLIM	Soft limit opened files	2048	2048
	HFSZLIM	Max. file size (hard limit)	0x7FFFFFFF	0x7FFFFFFFFF
	SFSZLIM	Max. file size (soft limit)	0x7FFFFFFF	0x7FFFFFFFFF
	HCORLIM	Max. core size (hard limit)	0x7FFFFFFF	0x7FFFFFFF
	SCORLIM	Max. core size (soft limit)	0x7FFFFFFF	0x7FFFFFFF

Process	HDATLIM	Max. bytes proc. data segm.	0x7FFFFFFF	0x7FFFFFFFF
	SDATLIM	Softlimit bytes proc. data segm.	0x7FFFFFFF	0x7FFFFFFFF
	HVMMLIM	Max. bytes proc. address space	0x7FFFFFFF	0x7FFFFFFFF
	SVMMLIM	Softlimit bytes proc. address space	0x7FFFFFFF	0x7FFFFFFFF
	SSTKLIM	Softlimit max. proc. stack size	0x7FFFFFFF	0x7FFFFFFF
	HSTKLIM	max. proc. stack size	0x7FFFFFFF	0x7FFFFFFF
	MAXUP	No. of proc. per user	500	500
	GPGSLO	Page daemon start value	600	600
	SM_SIZE	Perc. of phys. mem. for buffer cache	10 (if RAM ≥ 512:5)	10 (if RAM ≥ 512:5)
Shared Memory	SHMMAX	Max. shared memory	1073741824	4294967296
	SHMMNI	Max. shared memory keys	256	256
	SHMSEG	Max. shared memory segments	100	100
Other	AIO_HIWAT	No. reserved async. I/O structures (only for raw device installations)	4096	4096
	AIO_MAX	Max. No. of async I/O structures (only for raw device installations)	6144	6144
	MAXUSERS	Max. "USER" (system resource allocation)	512	512
	HRTIME	No. of timer	NPROC/4 (minimum 1000)	NPROC/4 (minimum 1000)
	HRVTIME	No. of timer (SW)	NPROC/4 (minimum 1000)	NPROC/4 (minimum 1000)
	IDLE_MIGRATE	Migration to an idle CPU (RM600 only)	1	1
	IPPORT_USER_START	Start of area for dynamical allocation of ports	60000	60000
	IPPORT_USER_END	End of area for dynamical allocation of ports	65535	65535
	KV_SIZE	Kernel virtual segment size	0x6000000 (if RAM ≥ 1.5GB: 0xC000000)	0x6000000 (if RAM ≥ 1.5GB: 0xC000000)

After creation of a new kernel, the machine must be rebooted.

2. Adapt the UNIX kernel parameters using the shell command **`idtune`**. To do this:

   a. Check the actual setting of the parameters using the command:
   `strings /unix|grep <Parameter>`

   The default kernel settings are defined in the file `/etc/conf/cf.d/mtune`.

   b. Adapt all parameters listed in the table *Recommended Kernel Parameters for ReliantUNIX* separately.

   Parameters are changed in the file `/etc/conf/cf.d/stune` by setting the value to the requirements as described below.

It is not recommended to change any parameter settings in the file `/etc/conf/cf.d/mtune`.

It is not possible to define a parameter in `/etc/conf/cf.d/stune` with a higher value than the defined maximum in `/etc/conf/cf.d/mtune`.

Adjust the parameters with the command:
`/etc/conf/bin/idtune <parameter> <parametervalue>`

```
idtune  SEMMNS  1024
```

The command idtune checks the minimum/maximum values listed in `/etc/conf/cf.d/mtune`

and writes the changes to the file `/etc/conf/cf.d/stune`.

If a range error occurs, the minimum/maximum value can be modified (through editing `mtune`), if reasonable.

If a value for the 64-bit kernel requirements is higher than the maximum setting defined in `/etc/conf/cf.d/mtune`, set the parameter to this maximum value and ask your Technical Service to check whether it is possible to increase the maximum value in `/etc/conf/cf.d/mtune`.

c. Create a new kernel.

- First check if you are on a multi or on a single processor machine. Enter the command:
  `uname -M`

  This shows the number of CPUs and the RAM size:

  `1/256` : single processor system and 256 MB RAM

  `2/*` : multi processor system

- Create the kernel with the commands

  `/etc/conf/bin/idbuild -M` (multi-processor system)

  `/etc/conf/bin/idbuild -S` (single processor system)

d. To activate the new operating system kernel, the system must be rebooted with:
`shutdown -y -g<grace period> -i6`

7.4 ReliantUNIX: File Systems, Raw Devices, Swap Space

The following describes how to create the volume groups, file systems, raw devices and swap space for your installation.

The structure of the SAP system depends on the disk configuration of the customer system. For this reason, there are no general procedures for setting up the file systems.

Consider the following when you specify the sizes of file systems:
– Your operating system needs up to 10% of disk space for administration.
– 1 MB is 1024 KB.

See the following sections for more information:

- Setting up Swap Space [page 60]
- File Systems [page 62]
- Accessing Raw Devices [page 65]

7.4.1 ReliantUNIX: Setting up Swap Space

1. Find the minimum swap space requirements in the installation documentation in section Hardware and Software Requirements Check.
2. Determine the size of the installed RAM.

   Use the following command to display the number of CPUs and megabytes of memory:
   `uname -M`

The following **sample** output describes two CPUs and 512 MB of main memory:

```
2/512
```

3. Determine the allocated swap space and define added swap space if necessary.

   The following command displays size and location of currently defined swap space:

   `swap -l`

Sample output:

```
path                        dev   swaplo   blocks   free
/dev/ios0/sdisk000s1        4,1   0        261960   234352
```

   Additionally, the swap space can be looked up in the file `/etc/vfstab`, where disk slices used for swap space are marked with `'swap'`. Slices reserved for swap space cannot be used for a file system.

   Additional swap space can be defined using the command:

   `swap -a <diskdevice> <start block> <number of blocks>`

   The **`<number of blocks>`** defines the number of 512-byte blocks to be added.

The following command creates a new swap area on disk 1, slice 3 starting at block 0. The size of the area is 50000 blocks (25 MB).

```
swap -a /dev/ios0/sdisk001s3 0 50000
```

   Edit the file /etc/vfstab and add the new swap are.

```
/dev/ios0/sdisk001s3  /dev/ios0/rsdisk001s3 -  swap  - -
rw
```

4. Verify the paging space size and the kernel settings.

Do not execute this step, if you are installing a standalone DB server.

   – Make sure that the UNIX kernel paging space and user limits are already configured for the SAP System as described in the installation documentation in section Hardware and Software Requirements Check and the information above.

   – Execute `memlimits` to verify paging space size and kernel settings.

     Unpack the file `memlimits`:

```
cd <INSTDIR>

/<CD-DIR>/UNIX/<OS>/SAPCAR -xgvf    \
   /<CD-DIR>/UNIX/<OS>/SAPEXE.SAR memlimits
```

Start memlimits:

- For a 32-bit SAP kernel, enter: **./memlimits**
- For a 64-bit SAP kernel, enter: **./memlimits -l 20000**

If error messages occur, increase your paging space and rerun `memlimits` until it runs without errors.

7.4.2 ReliantUNIX: File System

Use

The configuration of virtual disks and the creation of file systems can be done manually using shell commands or using of **sysadm**.

Procedure

Creating File Systems on Hard Disk Partitions

1. Create new file systems using **newfs**. (File system type: ufs)

   The file system partition is created with the command:

   SAP DB: newfs -m 0 <device name>

   Informix: newfs -m 0 <device name>

   Oracle:
   ```
   newfs -i <15000 per GB> -o t -m 0 -b 8192   \
         -f 8192 <device name>
   ```

   The option '**-m <number>**' describes the minimum free space threshold as a percentage.

   Better performance can be achieved using the Veritas File System (vxfs):

   ```
   mkfs -F vxfs <device name> -
   ```
   resp.
   ```
   mkfs -F vxfs -o ninode=<15000 per GB>,bsize=8192 <device name> -
   ```

2. Create directories and mount the file systems.

   The directories to hold the file systems must be created using **mkdir**. The newly created file systems can then be defined in `/etc/vfstab` and mounted into the directories.

Configuring of Virtual Disks Manually

Make sure you follow the recommendations in the installation documentation in section *Installation Planning* concerning the distribution of file systems on distinct physical disks when configuring virtual disks.

To configure virtual disks manually, perform the following steps:

1. Define the structure of the virtual disks to be used so that they contain the file systems by editing the file `/etc/dktab`.

   The file `/etc/dktab` describes the configuration of all virtual disks. It defines the disk partitions used to create a virtual disk. To establish a virtual disk, this file has to be edited.

   There are two types of lines in `dktab` used to configure virtual disks:

a. Virtual disk declaration lines

b. Physical disk declaration lines

The new virtual disk has to be defined with one virtual disk declaration line followed by a number (equal to the number of slices) of physical disk declaration lines:

- Syntax of virtual disk declaration line:
  `<virtual disk name> <type> <number of slices> <cluster size>`
- Syntax of physical disk declaration line:
  `<slice name> [<start block> <number of blocks>]`

Explanation of parameters:

- The `<virtual disk name>` has to follow the pattern '`/dev/vd/vdisk<number>`'.
- The value of `<type>` is one of

  `'concat'` to combine partitions of different sizes

  `'stripe'` to combine partitions of the same size, situated on different disks (allowing for an even workload distribution)

  `'simple'` to make a certain number of blocks available as a partition
- `<number of slices>` counts the number of pieces needed to form the virtual disk. Virtual disks of type `'simple'` do not need this parameter.
- `<cluster size>` is only needed for virtual disks of type `'striped'` and should be defined as '`128`'
- `<slice name>` is the device name of the piece to be used.
- `[<start block> <number of blocks>]` have to be defined if you are using only parts of slices.

```
A concatenated virtual partition
/dev/vd/vdisk0  concat    4
/dev/ios0/sdisk001s4
/dev/ios0/sdisk001s5
/dev/ios0/sdisk002s4
/dev/ios0/sdisk002s5

A striped virtual partition with a cluster size of 128 blocks
/dev/vd/vdisk1  stripe    2    128
/dev/ios0/sdisk001s4
/dev/ios0/sdisk002s4

A tiny 2400 block simple virtual partition
/dev/vd/vdisk2 simple
/dev/ios0/sdisk010s7      0    2400
```

2. Configure the virtual disks with the **`dkconfig`** utility.

   After being defined by editing `/etc/dktab`, the new virtual disk systems have to be configured.

   Use the following command to configure separate entries:
   `dkconfig -c <virtual disk name>`

or the following command to configure all virtual disks described in `/etc/dktab`:
`dkconfig -vac`

3. Create new file systems using newfs.

   After defining the virtual disks in `/etc/dktab` and executing `dkconfig`, file systems can be generated on them (similar to using a raw device).

   The file system partition is created with the command:

   SAP DB: **`newfs -i <15000 per GB> -m 1 -o s <virtual disk name>`**

   Informix: **`newfs -i <15000 per GB> -m 1 -o s <virtual disk name>`**

   Oracle:
   ```
   newfs -i <15000 per GB> -m 1 -o s -b 8192 -f 8192   \
     <virtual disk name>
   ```

   The option '**`-m <number>`**' describes the minimum free space threshold as a percentage.

   Better performance can be achieved using the Veritas File System (vxfs):

   ```
   mkfs -F vxfs <device name> -
   ```
   resp.
   ```
   mkfs -F vxfs -o ninode=<15000 per GB>,bsize=8192 <device name>
   -
   ```

4. Create directories and mount the file systems.

   The directories to hold the file systems must be created using **`mkdir`**.

   The newly created file systems may then be defined in `/etc/vfstab` and mounted into the directories.

Preparing Virtual Disks With sysadm

1. Create the virtual disk nodes using **`mknod`** if they are not already defined.
2. Define the structure of the virtual disks to be used to contain the file systems.

   Configure the virtual disks:

   – Enter the command **`sysadm`**.
   – For each file system listed in the appropriate SAP template or documentation, select *file_systems* → *vdisk* → *<type of disk to be configured>*
   – Fill in the type dependent form and press *Save* to trigger update of `/etc/dktab` and execution of **`dkconfig`**.

3. Create new file systems.

   Create directories and mount the file systems:

   – Enter the command **`sysadm`**.
   – For each file system listed in the appropriate SAP template or documentation, select *file_systems* → *make*
   – Fill in the form and save your changes.

7.4.3 ReliantUNIX: Accessing Raw Devices

File systems and raw devices differ in the way data is written to and read from disk:

- **Buffering:** Reads and writes to a file system are buffered in a UNIX system. To be absolutely sure that all data is physically present on a disk, the buffers and files must be synched.

  Writes to a raw device are **unbuffered**: the system writes directly to the disk. These writes are faster and more secure.

- **File access:** Accessing files on a UNIX file system is transparent. Accessing data on a raw device is only possible with a special application.

Some databases prefer raw devices. To access these raw devices you can create symbolic links. The following examples show symbolic link commands:

SAP DB: Not necessary because R3SETUP creates the links
Informix:

```
ln -s /dev/ios0/<diskdevice>    \
   /informix/<SAPSID>/sapdata/physdev1/data1
   .
   .
ln -s /dev/ios0/<diskdevice>    \
   /informix/<SAPSID>/sapdata/physdev<m>/data<n>
```

Oracle: Not necessary

It is important to also list the used raw devices in `/etc/vfstab`. Otherwise, system tools like `dkmap` do not recognize that the disks are already in use.

Entry in `/etc/vfstab`:

```
/dev/ios0/<diskdevice> /dev/ios0/r<diskdevice> - raw 0 no
rw
```

7.5 ReliantUNIX: Mounting Directories via NFS

Use

If directories are to be accessed from several machines in parallel, mount directories via NFS.

Procedure

To mount directories via NFS, perform the following steps on the server side and on the client side.

Server Side

The machine that actually contains the directories is the server. Local NFS resources you want to make available for mounting by remote systems must be known to the server. Perform the following steps on the server side:

1. Enter the following command to display the local resources already accessible to be mounted remotely:
   `/usr/sbin/share`
2. A local directory can be shared by entering the command:

   ```
   /usr/sbin/share -F nfs -o root=<application server>,     \
      rw=<client side> -d <description> <file system>
   ```

   where **`<client side>`** is the host name of the machine the directory is to be mounted via NFS, **`<description>`** is a user-defined text to describe the directory and **`<file system>`** is the name of the directory to be shared.
3. If the directories are to be shared permanently, the share commands can be placed in the description file `/etc/dfs/dfstab` using the syntax described above. The share commands are executed automatically when the system is entering init state 3.

Client Side

The machine where the remote resources are mounted via NFS is the client. Perform the following steps on the client side:

1. Enter the following command to make sure the respective directory is shared by the server side:
   `/usr/sbin/showmount -e <server side>`

   All directories already shared by host **`<server side>`** are displayed.
2. Mount the directory from the server side via NFS:
   `/sbin/mount -F nfs <server side>:<pathname>  <newpathname>`

   where **`<server side>`** is the host name of the machine that actually contains the resource, **`<pathname>`** is the path name of the resource on the server side, and **`<newpathname>`** is the path name of the mount on the client side.
3. If the directories are to be mounted permanently, the mount descriptions can be placed in the file `/etc/vfstab` in the following form:
   `<server side>:<pathname> - <newpathname> nfs - yes -`

The directory `/usr/sap/trans` residing on host `host1` is to be mounted on `host2` via NFS.
Server Side (host1):
Enter:

```
/usr/sbin/share -F nfs -o root=host2,rw=host2 -d"info"  \
   /usr/sap/trans
```

Add the same command line to `/etc/dfs/dfstab`.
Client Side (host2):
Enter:

```
/usr/sbin/mount -F nfs host1:/usr/sap/trans  \
   /usr/sap/trans
```

Add the following line to `/etc/vfstab`:

```
host1:/usr/sap/trans - /usr/sap/trans nfs - yes -
```

7.6 ReliantUNIX: Creating UNIX Groups and Users

Use

Groups and users can either be created manually using the shell commands `useradd` and `groupadd` or using the tool `sysadm`.

Procedure

For the users and groups created by R3SETUP see Network Information Service Environment [page 8].

Informix only:
Due to operating system functions, it is possible that the group `super_archive` is not inserted correctly in the `/etc/group` file.
Please check whether the group name was truncated to `super_` here and change it to `super_archive`.

Add the user **`root`** to group `dba` manually (entry in `/etc/group`).

Manual Creation of UNIX Groups and Users

Use the following commands to create groups and users:

```
groupadd  <new_group>
useradd -g <login_group> -d <home_directory> \
            -s <login_shell>  <new_user>
```

To delete a user, enter the command:

```
userdel -r <user>
```

Creation of UNIX Groups and Users using sysadm

1. Enter the command **sysadm**.
2. Select *users* → *add user/group*.
3. Fill in the form and save your changes.

7.7 ReliantUNIX: Troubleshooting

If you suspect that hardware problems are the reason for errors in your installation, analyze the system.

To do this, enter the command:

```
more  /var/adm/log/messages
```

8 Solaris: OS-Dependent Installation Steps

This documentation has been written with the support of the SUN Competence Center. It is not intended to replace the documentation of the Solaris operating system.

8.1 Solaris: Preparing the Installation

C++ Runtime Environment

Make sure that the Shared Library Patch for C++ is installed.

Installation Directory

In a standard Solaris installation `/tmp` is a tmpfs file system.

tmpfs is a memory-based file system that uses kernel resources relating to the VM system and page cache as a file system.

Once mounted, a tmpfs file system provides standard file operations and semantics. tmpfs is so named because all files and directories are lost after reboot or unmounts. For this reason, we recommend that you create a directory `/temp` on an ufs file system.

8.2 Solaris: Mounting a CD-ROM

Use

In the Solaris operating system, a layer of software manages CD-ROM and other media. This is the volume manager, which automates the interaction between you and your CD-ROMs. Do **not** use this volume manager to mount CD-ROMs needed for the installation, as the manager does not use the *nomaplcase* option. Without this option, some files on CDs can not be found. Mount CD-ROMs for the installation manually as described in the following procedure.

For more information on configuring the volume management, see the *SunOS Routine System Administration Guide*.

Procedure

1. Log on as user **`root`**.
2. Check if `vold` (volume daemon) is running:
   `ps -ef | grep vold`
3. Kill `vold` if it is running.
4. Insert the CD-ROM into your disk drive.
5. Create a mountpoint:
   `mkdir /sapcd`
6. Mount the CD-ROM:
   `mount -F hsfs -o nomaplcase,ro /dev/dsk/c0t6d0s2 /sapcd`

After the installation, you can start `vold` with the command `/usr/sbin/vold`.

8.3 Solaris: Checking and Modifying the UNIX Kernel

Use

Before installing an SAP System, you must check and, if necessary, modify the UNIX operating system kernel.

Procedure

The Solaris kernel modules are automatically loaded when needed. This makes rebuilding of the kernel unnecessary.

1. Edit the file `/etc/system`. It must contain the following lines:
   - For 32-bit SAP Kernel:

```
*Shared memory parameters set for SAP
*
set shmsys:shminfo_shmmin=1
set shmsys:shminfo_shmmax=<physical memory, min. 536870912,
                                          max. 2147483648>
set shmsys:shminfo_shmmni=256
set shmsys:shminfo_shmseg=100
*
* End of shared memory parameters
*
* Semaphores parameters set for SAP
*
set semsys:seminfo_semmap=256
set semsys:seminfo_semmni=4096
set semsys:seminfo_semmns=4096
set semsys:seminfo_semmnu=4096
set semsys:seminfo_semume=64
set semsys:seminfo_semmsl=2048
set semsys:seminfo_semopm=50
*
* End of semaphores parameters
*
* IPC Message queues Parameters set for SAP
*
set msgsys:msginfo_msgmap=2048
set msgsys:msginfo_msgmax=8192
set msgsys:msginfo_msgmnb=16384
set msgsys:msginfo_msgmni=50
set msgsys:msginfo_msgssz=32
set msgsys:msginfo_msgtql=2048
set msgsys:msginfo_msgseg=32767
*
* End of message queues parameters
*
* File descriptor parameters set for SAP
*
set rlim_fd_cur=1024
set rlim_fd_max=1536
```

```
*
* End of file descriptor parameters
*
* Other parameters for SAP:
*
* End of other parameters
```

– For 64-bit SAP Kernel:

```
* Shared memory parameters set for SAP
*
set shmsys:shminfo_shmmin=1
set shmsys:shminfo_shmmax=1073741824
set shmsys:shminfo_shmmni=256
set shmsys:shminfo_shmseg=200
*
* End of shared memory parameters
*
* Semaphores parameters set for SAP
*
set semsys:seminfo_semmap=256
set semsys:seminfo_semmni=4096
set semsys:seminfo_semmns=4096
set semsys:seminfo_semmnu=4096
set semsys:seminfo_semume=64
set semsys:seminfo_semmsl=2048
set semsys:seminfo_semopm=50
*
* End of semaphores parameters
*
* IPC Message queues Parameters set for SAP
*
set msgsys:msginfo_msgmap=2048
set msgsys:msginfo_msgmax=8192
set msgsys:msginfo_msgmnb=16384
set msgsys:msginfo_msgmni=50
set msgsys:msginfo_msgssz=32
set msgsys:msginfo_msgtql=2048
set msgsys:msginfo_msgseg=32767
*
* End of message queues parameters
*
* File descriptor parameters set for SAP
*
set rlim_fd_cur=1024
set rlim_fd_max=1536
*
* End of file descriptor parameters
*
* Other parameters for SAP:
*
* End of other parameters
```

DB2 UDB for UNIX and Windows:
Use the following parameter values:

```
set shmsys:shminfo_shmmax: Choose a max. value of 90% of main
                                memory (value must be less than 4 GB).
set shmsys:shminfo_shmmni=300
set semsys:seminfo_semmap=4098
set msgsys:msginfo_msgmax=65535
set msgsys:msginfo_msgmnb=65535
set msgsys:msginfo_msgmap=258
set msgsys:msginfo_msgmni=2046
set msgsys:msginfo_msgssz=32
set msgsys:msginfo_msgtql=1024
set msgsys:msginfo_msgseg=32768
```

2. Reboot the system with the command:
   `shutdown -i6 -y -g0`

8.4 Solaris: File Systems, Raw Devices, Swap Space

The following describes how to create the volume groups, file systems, raw devices and swap space for your installation.

The structure of the SAP system depends on the disk configuration of the customer system. For this reason, there are no general procedures for setting up the file systems.

Consider the following when you specify the sizes of file systems:
- Your operating system needs up to 10% of disk space for administration.
- 1 MB is 1024 KB.

Solaris with Veritas File System:
If you are installing Oracle 8 on Solaris with Veritas file system, make sure you have Veritas Version 3.2.6 or higher.

See the following sections for more information:

- Preparing Hard Disks [page 73]
- Setting up Swap Space [page 74]
- Creating File Systems [page 77]
- Accessing Raw Devices [page 78]

8.4.1 Solaris: Preparing Hard Disks

prtvtoc Utility

prtvtoc prints the content of the VTOC (volume table of contents). It displays the start/stop cylinder, the size of the partition and the mounted directory. Execute **prtvtoc** as follows:

1. Log on as user **root**.
2. Enter:
   /usr/sbin/prtvtoc -s <rawdevice>

Solaris does not allow track zero on the disk to be part of the raw device used by the database system. When the machine is rebooted, the operating system writes a label here. Consequently, data from the database is overwritten and a restore of the database becomes necessary.
If you use a RAID System or a Volume Management Software from Sun (Solstice DiskSuite or Veritas Volume Manager) then the device driver will take care of this restriction.
Otherwise, when you initially setup a raw device, the standard second partition (/dev/rdsk/c?t?d?s2) may **not** be used since it normally includes track zero on the disk. A disk that is intended for use as a raw device, must be reformatted appropriately to ensure that the raw device begins at track one.
Check your disk configuration as soon as possible to make sure that your raw devices meet the above requirements. To do this:
Log on as root and change to the directory
/<db-system>/<SAPSID>/sapdata
List the links to the raw devices by entering the command:
ls -lR
Display all partitions on a disk (substitute "?"):
/usr/sbin/prtvtoc <rawdevice>

Assume **ls -lR** produces the output
lrwxrwxrwx 1 root 18 Nov 29 data3 -> /dev/rdsk/c0t3d0s6
then the partition 6 of this disk is used as a raw device.
Assume **/usr/sbin/prtvtoc /dev/rdsk/c0t3d0s6** produces the output:

```
                   First    Sector    Last
Part.  Tag  Flags  Sector   Count     Sector   Mount Directory
0      0     00     0        7050      7049
2      5     01     0        3929670   3929669
6      4     00     7050     3922620   3929669
```

Then the first sector of partition 6 is **not** equal zero.
If the partition you want to use begins with sector zero (as shown for partition 2 in the above example) it cannot be used as a raw device partition.

Partitioning Disks

1. Look for a mounted partition with the format utility:

```
format
format> partition
partition> print
```

2. Define partition size:

```
partition> label
partition> quit
format> disk
```

3. Chose next disk:

```
format> quit
```

For more information, see the Solaris documentation.

8.4.2 Solaris: Setting up Swap Space

1. Determine the size of the installed RAM with the command:
   prtconf | grep "Memory"
2. Determine the allocated swap space:
   swap -s
3. Increase the swap space if necessary. You can find the minimum swap space requirement in the installation documentation in section Hardware and Software Requirements Check.

   To increase swap space, you have the following options:

   – Create a swap file

   – Use an entire disk or partition.
4. Verify paging space size and kernel settings.

Creating a Swap File

To create additional swap space without reformatting a disk, first create a swap file using the `mkfile` command. `mkfile` creates a file that is suitable for use either as an NFS mounted or local area. You can specify the size in bytes (default), kilobytes, blocks or megabytes using the **k**, **b** or **m** suffixes.

1. Make sure that you have enough free space on your system for the additional swap:
   df -k
2. Create the swap file:
   mkfile <size>[k | b | m] <filename>

```
mkfile 500m /sap/newswap
```

3. Once the swap file is created you make it accessible using the swap command:
   `swap -a <filename>`

```
swap -a /sap/newswap
```

4. To make the swap file available automatically each time the system boots, add an entry to the `/etc/vfstab` file.

```
/sap/newswap - - swap - no -
```

   Since the file system must be mounted before the swap file, make sure that the entry that mounts the file system comes before the swap entry.

Using an Entire Disk or Partition

Understanding disk device names:
The disk naming convention is based on logical (not physical) device names. To specify a slice (partition) on a disk with a bus controller (SCSI), use a device name with these conventions: **`c<w>t<x>d<y>s<z>`**

Key:

`<w>`	logical controller number
`<x>`	physical bus target number (for example, SCSI-ID)
`<y>`	drive number
`<z>`	slice (partition) 0-7

To use an entire disk or partition:

1. Look for an unused partition or disk. Verify that the partition is unused by looking at `/etc/vfstab`, `/etc/opt/SUNWmd/md.tab` or using the **`vxprint`** utility.
2. Make your swap accessible.
   `swap -a <blockdevice>`
3. To make the swap partition available automatically each time the system boots, add an entry to the `/etc/vfstab` file.
   `<blockdevice> - - swap - no -`

Setting up Metadevices

For advanced disk management, use the **Solstice DiskSuite** package or the **Veritas Volume Management Software.** They provide performance (striping), high availability (mirroring, RAID 5) and maintenance (hotspares, growing file systems) using **logical metadevices**.

- Solstice DiskSuite
  - All information about the logical devices is kept in several replicas of a metadatabase. To access a stripe, for example, use its nickname.
  - A `<metadevice>` is defined in the `/etc/opt/SUNWmd/md.tab`. The format is: `/dev/md/dsk/d[0..n]`.
  - To access, for example, the raw partition of the metadevice `d10`, use: `/dev/md/rdsk/d10`.

    For striping and mirroring with **Solstice DiskSuite** equal physical partitions are required.

  - To set up metadevices:

    Check the metadatabase:
    `/usr/opt/SUNWmd/sbin/metadb`
    Define the metadevices:
    `vi /etc/opt/SUNWmd/md.tab.`
    Initialize the metadevices with:
    `/usr/opt/SUNWmd/sbin/metainit -a`
    Check metadevices with:
    `/usr/opt/SUNWmd/sbin/metastat`
    Use raw metadevices:
    Instead of `/dev/rdsk/<diskdevice>` you should use `/dev/md/rdsk/d[0..n]`.

  For more information, see the **Solstice DiskSuite** documentation.
- Veritas Volume Management Software

  For more information, see the **Veritas Volume Management Software** documentation.

Verify Paging Space Size and Kernel Settings

Do not execute this step if you are installing a standalone DB server.

- Make sure that the UNIX kernel, paging space and user limits are already configured for the SAP System as described in the installation documentation in section *Hardware and Software Requirements* and in the previous section.
- Execute `memlimits`, to verify paging space size and kernel settings:

  a. Unpack the file `memlimits`:

  ```
  cd <INSTDIR>

  /<CD-DIR>/UNIX/<OS>/SAPCAR -xgvf
     /<CD-DIR>/UNIX/<OS>/SAPEXE.SAR memlimits
  ```

b. Start `memlimits`.

   For a 32-bit SAP Kernel, enter: **`./memlimits`**

   For a 64-bit SAP Kernel, enter: **`./memlimits -l 20000`**

c. If error messages occur, increase your paging space and rerun `memlimits` until it is error free.

8.4.3 Solaris: Creating File Systems

Creating File Systems on Hard Disk Partitions

1. Look for already configured and mounted file systems using the command:
   `df -k`
2. Look for the physical available disks and partitions using the **`format`** or **`/usr/sbin/prtvtoc <rawdevice>`** utility.
3. Verify that the partition is unused.
4. Create new file systems using the command **`newfs`**:

   Oracle (sapdata file systems):
   : **`newfs -i 20000 -m 1 -b 8192 -f 8192 <rawdevice>`**

   Standard SAP file systems:
   : **`newfs <rawdevice>`**

`newfs /dev/rdsk/<diskdevice>`
The disk naming convention is based on logical (not physical) device names. To specify a slice (partition) on a disk with a bus controller (SCSI), use a device name `<diskdevice>` with these conventions: **`c<w>t<x>d<y>s<z>`**
Key:

`<w>`	logical controller number
`<x>`	physical bus target number (for example, SCSI-ID)
`<y>`	drive number
`<z>`	slice (partition) 0-7

Mounting File Systems

Create directories (**`mkdir`**) and mount the file systems using the command:
`mount <blockdevice> <mountdir>`

To make the file systems available automatically each time the system boots, add an entry to the `/etc/vfstab` file:

- If you do **not** use **`Solstice DiskSuite`**, add the following entry:

  ```
  /dev/dsk/<diskdevice> /dev/rdsk/<diskdevice> \
     <mountdir> ufs <fsck pass#> yes -
  ```

  where `<fsck pass#>` is a number greater than 6.

- If you use **`Solstice DiskSuite`**, add the following entry:

  ```
  /dev/md/dsk/<metadevice> /dev/md/rdsk/<metadevice> \
     <mountdir> ufs - no -
  ```

8.4.4 Solaris: Accessing Raw Devices

File systems and raw devices differ in the way that data is written to and read from disk:

- **Buffering:** Reads and writes to a file system are buffered in a UNIX system. To be absolutely sure that all data is physically present on a disk, the buffers and files must be synchronized.

  Writes to a raw device are **unbuffered**: the system writes directly to the disk. These writes are faster and more secure.

- **File access:** Accessing files on a UNIX file system is transparent. Accessing data on a raw device is only possible with a special application.

Some databases prefer raw devices. To access these raw devices, you can create symbolic links. The following examples show symbolic link commands:

SAP DB: ln -s <rawdevice> /sapdb/<SAPSID>/sapdata/DISK01
Informix: `ln -s <rawdevice> \`
`/informix/<SAPSID>/sapdata/physdev1/data1`
Oracle: Not necessary

With Solstice DiskSuite :
SAP DB: Not necessary because R3SETUP creates the links
Informix: ln -s /dev/md/rdsk/<metadevice> \
/informix/<SAPSID>/sapdata/physdev1/data1
Oracle: Not necessary

With Veritas Volume Manager:
SAP DB: Not necessary because R3SETUP creates the links
Informix: ln -s /dev/vx/rdsk/<volumegroup>/<volumename> \
/informix/<SAPSID>/sapdata/physdev1/data1
Oracle: Not necessary

8.5 Solaris: Mounting Directories via NFS

Procedure

To mount directories via NFS from the host on which the central instance runs, log on as user `root` and proceed as follows:

On the host on which the main instance runs:

1. Enter the command:
   `/usr/sbin/share`
2. Edit the file `/etc/dfs/dfstab` to add file systems shared via NFS:
   `vi /etc/dfs/dfstab`

   Add the following line for each file system, for example:

   ```
   share -F nfs -o root=<nfsclient1>:<nfsclient2>   \
      -d "description" /usr/sap/trans
   ```

   After your SAP System has been installed successfully, in the above line you have to change `-o root` to `-o rw` (or remove `anon=0`, respectively) for all exported directories, for example:

   ```
   share -F nfs -o rw=<nfsclient1>:<nfsclient2>   \
      -d "description" /usr/sap/trans
   ```

3. If the `/etc/dfs/dfstab` was empty, the NFS server is not active.

   Start the NFS server with the command:
   `/etc/init.d/nfs.server start`
4. To see if the NFS server is active and which partitions are mountable, enter the command:
   `showmount -e <NFS-server>`

On the host on which the additional instance runs:

1. If you are mounting NFS disks for the first time, the NFS client software is not active.

   Start the NFS client software with the command:
   `/etc/init.d/nfs.client start`
2. Edit the file `/etc/vfstab` to mount the directory:
   `vi /etc/vfstab`

   Add the following line for each file system, for example:
   `db_server:/usr/sap/trans - /usr/sap/trans nfs - yes -`

   If the mount point exists, mount `/usr/sap/trans`, for example, with the command:
   `mount /usr/sap/trans`

8.6 Solaris: Creating UNIX Groups and Users

Use

Groups and users can be created manually using the shell commands **useradd** and **groupadd**. If you run a graphical user interface, you can use the **admintool**.

Procedure

For the users and groups created by R3SETUP, see Network Information Service Environment [page 8].

Informix only:
Due to operating system functions, it is possible that the group `super_archive` is not inserted correctly in the `/etc/group` file.
Check whether the group name was truncated to `super_` here and change it to `super_archive` after creation of all the SAP System users.

User **root** cannot be added to group `dba` using **groupadd**. You have to use the **admintool** for this entry.

1. Use `groupadd` and `useradd` according to the following syntax:

   ```
   groupadd <new group>

   useradd -g <login group> -d <homedirectory> \
      -s <login shell>  <new user>
   ```

2. Check the following values for the users:

   – In `csh` call **limit**

   Output:

   ```
   cputime                 unlimited
   filesize                unlimited
   datasize                2097148 kbytes    (<-- important)
   stacksize               8192 kbytes
   coredumpsize            unlimited
   descriptors             1024
   memorysize              unlimited         (<-- important)
   ```

– In `sh` or `ksh` call **`ulimit -a`**

Output:

```
time(seconds)                   unlimited
file(blocks)                    unlimited
data(kbytes)                    2097148         (<-- important)
stack(kbytes)                   8192
coredump(blocks)                unlimited
nofiles(descriptors)            1024
memory(KBytes)                  unlimited       (<-- important)
```

3. If your parameter settings differ from the settings above, change these values accordingly.

If you have to change the value for `descriptors` to `1024`, proceed as follows:

1. Add the following line to the `.cshrc` file after the line containing `case Sun*:`
   `limit descriptors 1024`
2. Add the following line to the `.profile` file after the line containing `Sun*)`
   :
   `ulimit -n 1024`

8.7 Solaris: Troubleshooting

If you suspect that hardware problems are the reason for errors in your installation, analyze the system. To do this, use the following commands:

```
dmesg   |   more
more  /usr/adm/messages
```

9 Heterogeneous SAP System Installation

Use

This documentation describes how to install SAP instances in a heterogeneous UNIX environment, that is:

- You run instances on different UNIX platforms.
- You mix 32-bit instances and 64-bit instances in one landscape.

This section does **not** apply to installations of SAP instances in a mixed UNIX / Windows NT environment. See the documentation *R/3 Installation on Windows NT* for more information.

Prerequisites

If your database is **Informix**, you need to perform some additional steps as well as the procedure below. Refer to Heterogeneous SAP System Installation for Informix.

Procedure

Local Software Installation

Proceed as follows to store the SAP executables locally:

1. Install the central instance and the database as described in the documentation *R/3 Installation on UNIX*.
2. Log on as user **`root`** on the dialog instance host.
3. Create `<sapmnt>/<SAPSID>/exe` as a local directory on the dialog instance.
4. On the central instance host, export the directories `<sapmnt>/<SAPSID>/global` and `<sapmnt>/<SAPSID>/profile` via NFS and mount them on each dialog instance host as described in section *Mounting directories via NFS*.
5. Mount the CD drive (via NFS) and insert the CD with the SAP executables on the dialog instance host.
6. Install the dialog instance.

Shared Software Installation

It is also possible to perform a shared software installation. The executables are installed in a differently named directory on the server and mounted via NFS. Shared software installation makes upgrading and maintenance of SAP software easier while producing a greater load on the network.

To perform a shared software installation, install the executables of dialog instance hosts on different platforms to the central instance host on the central instance host in the directories `<sapmnt>/<SAPSID>/<platform>/exe` (`/sapmnt/BZ6/AIX_32/exe` for example) and NFS export the directories in read-write mode. On the dialog instance host, create `<sapmnt>/<SAPSID>/exe`, `<sapmnt>/profile` and `<sapmnt>/global` as a NFS mount to the corresponding directories on the central instance host (`<sapmnt>/<SAPSID>/exe` must be a link to `<sapmnt>/<SAPSID>/<platform>/exe`

created above on the central instance host). Now, install the dialog instance. Take account of the following aspects during the dialog instance installation:

- During the installation of the first dialog instance of a platform, R3SETUP must be forced to copy the correct executables from CD-ROM to the directory `<sapmnt>/<SAPSID>/exe` (which is `<sapmnt>/<SAPSID>/<platform>/exe` on the central instance host).
- In the command file for R3SETUP, specify the key `EXTRACT_ON_NFS=YES` in the steps that extract `SAPEXE.SAR` and `SAPEXEDB.SAR`. (When installing additional dialog instances of this platform, do not specify this key.)