digital

SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100

DIGITAL HiTest Notes

Part Number: EK-HSPXA-HN. C01

April 1998

Revision/Update Information: This is a revised document.

Digital Equipment Corporation
Maynard, Massachusetts

April 1998

Table of Contents

4 System Installation and Setup

5 Tests and Results

6 Problems and Solutions

7 Detailed Hardware Configuration

Figures

Tables

Preface

This document provides an overview of DIGITAL HiTest Suites and detailed technical information about the SAP Oracle UNIX AlphaServer 4100 HiTest Suite. This information includes the HiTest AppSet, the HiTest Foundation, configuration details, installation instructions, tuning parameters, problems encountered and their solutions, tests and test results, and system diagrams. Together, a HiTest Foundation and HiTest AppSet (Application Set) comprise all of the components in a HiTest Suite. The HiTest Foundation includes the hardware, operating system, middleware, and database software. The HiTest AppSet contains a collection of software specific to one class of customer solutions.

Audience

Primary users of this document are DIGITAL and Partners sales representatives and technical support personnel. Secondary audiences include product managers, customers, and the personnel responsible for installing, setting up, and operating a DIGITAL HiTest Suite.

Organization

This document is organized as follows:

Chapter Title	Description
Chapter 1 – Advantages of DIGITAL HiTest Suites	Provides a summary of the benefits of DIGITAL HiTest Suites and an overview of the Suite covered in this document.
Chapter 2 – About This DIGITAL HiTest Suite	Describes the specific characteristics of this HiTest Suite.
Chapter 3 – Configuration Data	Includes tables of configuration data about the hardware and software components that define the DIGITAL HiTest Template, and special configuration rules if any.
Chapter 4 – System Installation and Setup	Provides information for installing and setting up this DIGITAL HiTest Suite.
Chapter 5 – Tests and Results	Describes how the tests were set up including database organization, where data and programs were placed, and how the tests were run. It also describes system limits and characterization data.
Chapter 6 – Problems and Solutions	Discusses any problems and solutions that were discovered during testing.
Chapter 7 – Detailed Hardware Configuration	Contains more detailed information about the configuration of the hardware and software components listed in the Configuration Data chapter.
Appendix A – Postinstallation for ASE	Contains more detailed information about the Postinstallation for TruCluster ASE V1.5 in an SAP R/3 Environment.

Customer Feedback

What our readers think of this or any other DIGITAL documentation is important to us. If you have any comments, we would appreciate hearing from you. Send your comments to: *reader-comments@digital.com*.

Please reference the complete document title and part number (EK-HSPXA-HN. C01) in your correspondence about this document.

Ordering Information

Copies of this and other DIGITAL documents can be ordered by calling 1-800-DIGITAL.

This document and other HiTest documents can be downloaded from the DIGITAL HiTest web site, which also provides access to other HiTest information such as configuration tools and parts updates.

`http://cosmo.tay.dec.com/public/configsys/config_systems.htm`

You can also visit the Technical Support Center web page, which provides additional information such as pointers to benchmark centers and major technical training and events:

`http://cosmo.tay.dec.com` (Intranet)
`http://www.businesslink.digital.com` (Internet)

Related Documents

This document references the following manuals:

- *StorageWorks Array Controllers: HS Family of Array Controllers Users Guide* (EK-HSFAM-UG)
- *SAP R/3 DECsafe Package V2.2A*
- *DIGITAL UNIX Software Documentation* (is delivered with DIGITAL UNIX on CD-ROM). Logical Storage Manager documentation is included.
- *DIGITAL UNIX Installation Guide* V4.0D or higher (AA-QTLGB-TE)
- *SAP R/3 Installation on UNIX - ORACLE Database Guide*
- *SAP Installation Guide Operating Dependencies*
- *SAP Notes 74278* - R/3 Installation on UNIX
- *SAP Notes 74279* - R/3 Installation on UNIX - OS Dependencies
- *SAP Notes 74275* - R/3 Installation on UNIX - ORACLE Database

A copy of the *SAP R/3 DECsafe Package V2.2A* is available from the DIGITAL Intranet at: `http://www.fra.dec.com/SAP-Eng/available/cluster.html`.

The *SAP R/3 Installation on UNIX - ORACLE Database Guide* is delivered with the SAP R/3 software. For the *SAP Notes* send the FAX template, which is delivered with the software to SAP.

1
Advantages of DIGITAL HiTest Suites

This chapter describes what a HiTest Suite is, the suite components and advantages, and customer add-ons.

What Is a DIGITAL HiTest Suite?

DIGITAL HiTest Suites are guidelines for configuring a set of prequalified computer systems. A HiTest Suite often contains all the hardware and software needed for a complete customer solution. DIGITAL HiTest Suites can be used as a basis for configuring systems that satisfy a wide set of customer requirements. Typically, Suites target specific markets such as Data Warehousing or SAP Solution and Internet Servers.

In each HiTest Suite, the components are selected and the configurations designed to ensure system reliability, application performance, and ability to upgrade. The suite's hardware and software components have been successfully tested for interoperability.

The specifications for allowed ranges of hardware and software components, part numbers, description, and revision information are listed in the *DIGITAL HiTest Template* in Chapter 3.

DIGITAL HiTest Suite Components

The SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 HiTest Suite contains two groups of components: the *DIGITAL HiTest Foundation* and the *DIGITAL HiTest AppSet*.

The DIGITAL HiTest AppSet contains application software unique to the targeted market. The DIGITAL HiTest Foundation contains the operating system, middleware, database software, and hardware and can be used as a configuration guideline for the base platform for many applications and target markets. This Suite will meet the needs of low to medium SAP TruCluster ASE configurations. The System Management Station is an optional standalone personal computer system containing software used to manage the HiTest system.

Select components from the HiTest Template to configure a DIGITAL HiTest System. Any system configured as specified in the DIGITAL HiTest Template can be called a DIGITAL HiTest System.

Additional Hardware and Software

Besides the hardware and software specified in a DIGITAL HiTest Suite, additional hardware and software can be added to a HiTest System. Add-on hardware consists of accessory components such as printers, modems, and scanners that are supported by the operating system and other software. Adding these components should not affect interoperability and, therefore, the system can still be considered a DIGITAL HiTest System.

Customers who purchase a DIGITAL HiTest System that is configured below the maximum specified in the Template, can later add additional hardware up to the specified maximum range and still maintain the integrity of a DIGITAL HiTest System.

If additional hardware components beyond the maximum specified in the Template are configured into a system, you still have the assurance that the rest of the system has been thoroughly tested for component interoperability. Therefore, the risk of experiencing problems is greatly reduced.

2
About This DIGITAL HiTest Suite

This HiTest Suite satisfies the needs of customers who require high-performance computing capabilities on most reliable configurations. It is a customer demand to run the SAP business application 24/7/52, which includes online backup and a failover mechanism. The scalability of the AlphaServer 4100 ensures a broad range of configuration possibilities from small companies up to medium and large enterprises. The performance of this configuration is well known and the pretested systems guarantee a quick installation at a customer site.

The SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 HiTest Suite includes the following components:

- SAP R/3
- ORACLE7 for DIGITAL UNIX
- TruCluster Available Server for DIGITAL UNIX
- DIGITAL UNIX
- AlphaServer 4100
- StorageWorks ESA 10000

This chapter describes the following characteristics of the SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 HiTest Suite and evaluates the Suite in terms of each:

- Availability
- Installability
- Interoperability
- Price Range
- Scalability
- Services
- Year 2000 Compliance

Availability

Availability, which describes a computer system's ability to quickly recover from a failure, can be described in terms of the following:

- Data Protection – Ensures long-term data accessibility by providing the facility to do offline data backup.
- Data Availability – Stores redundant data on line for rapid, automatic data recovery in the event of a failure. Data availability is typically provided through the use of RAID technology.
- Platform Availability – Enables processing to continue during failure by using technologies that support failover to other components. Clustering, redundant power supplies, battery backup, and other components provide support for platform availability.
- Disaster Tolerance – Protects against computer room disasters such as fire, flood, and sabotage. Disaster Tolerant Systems require an additional system at a remote site and are more expensive than the previously defined alternatives. (The DIGITAL HiTest process does not test disaster tolerant configurations. If disaster tolerance is a requirement, your sales person can provide more information.)

Features of SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100

The DIGITAL HiTest process verified that each of the availability features provided by this Suite operate correctly and provide the protection required for all configurations.

Table 2-1 indicates availability features that are always included in this HiTest Suite when configured with the AppSet and those that are customer options.

Table 2-1: SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 Availability Features

Availability Feature	Enabling Technology	Always Included	Customer Optional
Data Protection	Backup and restore		Yes
Data Availability	Redundant disk storage (RAID)	Yes	
	Redundant disk controllers	Yes	
Platform Availability	System failover (clustering)	Yes	

Recommendations for SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 HiTest Suite

This HiTest Suite provides high availability while maintaining high performance because hardware with a high mean-time-between-failure (MTBF) is used. Use of redundant data storage ensures data availability. To reduce the risk of system failures, DIGITAL recommends that the following levels of availability features be considered for this HiTest Suite:

- Data Protection
  - Long term data accessibility is always provided with this HiTest Suite by the TZ-887 automated tape drive, which provides up to 280 GB of compressed data storage.
- Data Availability
  - High performance data access was maintained by using the non-redundant RAID 0 with large stripesets. Logical Storage Manager (LSM) was used for host based mirroring to provide data redundancy.
  - Redundant SCSI adapters were distributed across several PCI buses to maintain high performance while avoiding a single point of failure.
  - The use of shared SCSI bus storage is a requirement of the TruCluster Available Server configuration.
- Platform Availability
  - The TruCluster configuration provides redundancy at the system level. The available Server configuration of the TruCluster is used as this is the only configuration certified by the partner for use with their application.

It may be necessary to create failover scripts for other software applications that run on the clustered configurations of this DIGITAL HiTest Suite.

Understanding Availability Features

This section provides background information on the availability features included in this HiTest Suite.

Backup and Restore

Backup and restore ensures that data remains available from one day to the next. It is generally identified as a *data protection* technique because the stored information can also be removed to a remote, protected environment. DIGITAL offers a range of backup and restore capabilities from individual tape systems to automated tape libraries.

Disk Storage Technologies

This section describes the disk storage technologies used to provide availability for DIGITAL HiTest configurations.

Just a Bunch of Disks (JBOD)

Just a bunch of disks (JBOD) refers to a multiple disk drive configuration, internal or external to a host computer, in which there is no storage controller. Disk drives are managed by the host system. To increase availability, JBOD storage systems are often configured with hardware such as redundant power supplies and fans, or multiple SCSI buses.

Redundant Array of Independent Disks (RAID)

A Redundant Array of Independent Disks (RAID) is a collection of disks managed by specialized array management software. When using RAID, all disks in the RAIDset should be the same type. Array management software may be *host-based* (execute in the host computer) or *subsystem-based* (execute in an intelligent disk controller).[1]

Disk striping (RAID Level 0), is technically not RAID because it does not offer redundancy.

RAID Levels 2 and 3 are parallel access arrays (members are accessed concurrently). To ensure that all disks participate in every I/O request, the minimum chunk size is kept small (for example, a byte).

RAID Levels 4 and 5 are independent access arrays (members are not required to be accessed concurrently). By keeping the minimum chunk size at least as large as a disk sector (block), not all members have to participate in each I/O request.

RAID Levels 2 and 4 are not in general use.

Figure 2-1 summarizes the RAID levels. The shaded areas in Figure 2-1 refer to space used for redundancy features.

Figure 2-1: RAID Level Summary

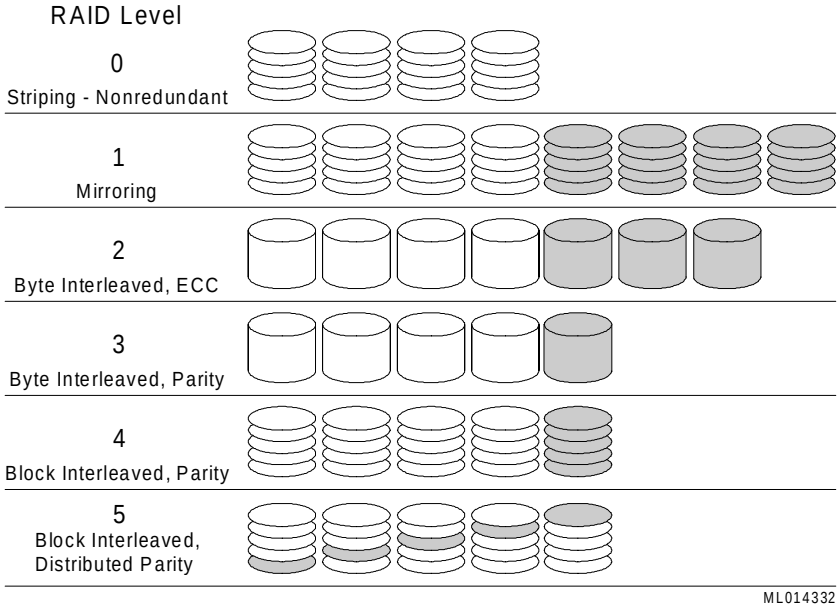

DIGITAL and other companies also use the terms RAID 0+1, RAID 1+5 and Adaptive 3/5 to refer to the combinations of these, and other, storage technologies. Table 2-2 describes the RAID types to consider when choosing a RAID configuration.

[1] The RAID technique was described by D. A. Patterson, G. Gibson, and R. H. Katz "A Case for Redundant Arrays of Inexpensive Disks (RAID)," Report No. UCB/CSD 87/391, University of California, Berkeley CA 1987.

Table 2-2: RAID Levels and Descriptions

RAID Level	Description	Advantages/Disadvantages
0	• Striping • Data segmented and distributed across several disks	+ increase in performance due to parallelism in read and write - no fault tolerance (*not* a high availability solution)
1	• Hardware Mirroring • Data written twice to different disk spindles within the disk array	+ good performance in read-intensive applications (data can be read in parallel from several disks) - slower in writes (multiple writes required) - spindle costs doubled
0+1	• Striped Mirroring • Combined level 0 and 1 • Data mirrored onto and striped across several disks • Best for performance-critical, fault-tolerant environments	+ good performance in reads (RAID 1) + write performance improved versus RAID 1 due to parallelism + adequate response maintained in event of disk failure - spindle costs doubled - recovery is I/O intensive
2	• Parallel access array • Striped • ECC on separate drives	+ high data transfer rate + ECC detects and corrects errors - low I/O request rate - not appropriate with modern drives
3	• Parallel access array • Small minimum chunk size • Check bit calculated from data • Parity bits on dedicated disk, data striped across remaining disks	+ good performance in reads due to parallelism (like RAID 0) + costs only slightly increased compared to disks without high availability solutions + good performance with long records (high data transfer rate) - write performance penalty due to check bit calculation - cannot overlap I/O (low I/O request rate)
4	• Independent access array • Parity disk	+ processes multiple requests simultaneously - parity disk is a bottleneck on writes
5	• Independent access array • Parity Bit • Check bit and data distributed (striped) across multiple disks • Best in environments that are mostly read and are not performance sensitive	+ good performance in reads due to parallelism (like RAID 0) + costs only slightly increased compared to disks without high availability solutions + overlapped I/O - write performance penalty due to check bit calculation

RAID Level	Description	Advantages/Disadvantages
1+5	• RAID 5 combined with mirroring • Mirroring provided by LSM or Volume Shadowing • Most reliable and highest performance solution	+ good performance in reads due to parallelism (like RAID 0) + double redundancy makes disk failure barely noticeable - spindle costs more than double - write performance penalty due to check bit calculation
Adaptive 3/5	• The best features of 3 and 5 • Adapts between Level 3 and Level 5 in response to changes in the application's workload	+ good performance in reads due to parallelism (like RAID 0) + costs only slightly increased compared to disks without high availability solutions + performs well with a wide variety of I/O loads even when load characteristics change minute by minute. - write performance penalty due to check bit calculation

Logical Storage Manager (LSM)

UNIX LSM provides low-overhead access to a logical view of storage devices. With LSM, logical disks can be created from any combination of physical disks of any size. LSM accomplishes this through the use of disk striping and mirroring. If a disk goes bad, the inherent redundancy keeps storage available.

LSM includes a visual management interface to simplify storage management. Disk hot-spots are highlighted so managers can see which parts of their storage configurations receive the most or least accesses. The graphical interface can be used to rearrange storage layout dynamically, either to grow or compress disk volumes as needs change, or to dynamically tune performance and system utilization.

Clustering

A cluster is a group of systems that works collectively to provide fast, uninterrupted computing service. Close cooperation can maximize performance and minimize down time. Within a cluster, individual systems and their components do not have to match the characteristics of mainframes, supercomputers, or fault-tolerant systems, yet can cooperate to achieve the same results. Most clusters provide a cost-effective solution to achieve optimal system availability and application performance, extensive scaling capability, and simplified system management.

DIGITAL UNIX TruCluster Solutions include:

TruCluster Available Server is an environment for providing a quick recovery, failover (high availability) environment for many UNIX applications, including standard database products, NFS, printers, network logins, and so forth. The functionality of the Available Server is included in the Production Server. TruCluster Available Server is designed for computing environments that can tolerate a short disruption, but need critical applications automatically restarted.

Redundant Components

Providing redundant components with a system increases that system's availability. If two identical components provide complete redundancy to each other, the availability factor for the set of components can increase by over 99% over a single component. Clustering and RAID are technologies that use redundancy of systems and of disks. Other components may be used in redundant configurations to further improve availability. Examples include disk controllers, power supplies, and network controllers. Many are hot swappable components that further increase availability by eliminating down time during replacement of the failed unit.

Installability

Installability is the ease with which hardware and software components can be installed and configured for use. Factors that are considered when evaluating installability include clarity of installation steps, number of steps and duration appropriate to the complexity of the product, and completeness of the installation and configuration information.

The DIGITAL HiTest process thoroughly examined all aspects of the installation of this HiTest Suite. The installation procedures that were used are documented in Chapter 4.

A few minor problems were detected that affect the software installation of more complex configurations. See Chapter 6 for more information.

DIGITAL Multivendor Computer Services (MCS) offers expert installation services.

Interoperability

Major components of this HiTest Suite have been tested for interoperability, including the application, database, operating system, hardware, firmware, and service packs and patches. Since interoperability problems are often related to inappropriate versions of components, the specific versions that are known to interoperate are documented. Minimum and maximum configurations for this Suite have been tested. The specific processes used for testing this Suite are described in Chapter 5.

The HiTest Notes provide solutions to interoperability problems in several ways. First, specific versions of all components are documented in Chapter 3. Second, installation and setup instructions in Chapter 4 are written so that many interoperability problems are avoided. Third, problems and solutions are documented in Chapter 6.

Some interoperability issues were discovered that related to reading large numbers of web pages, apparent TruCluster hangs, and online documentation. See Chapter 6 for more information and solutions.

Price Range

Figure 2-2 shows the approximate list price (U.S. dollars) for the minimum and maximum HiTest Systems that can be configured from the SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 HiTest Suite. These prices were effective as of March 1998. The price range can vary significantly over time and with the inclusion of service packages, consulting, country-specific prices, and other factors.

No matter what configuration is chosen, the resulting system is a leader in the industry due to the price and the power of the AlphaServer 4100 CPU.

Figure 2-2: SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 Price Range

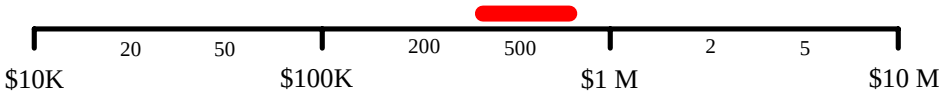

The purchase price of a system is only one factor affecting affordability. The cost of staff, space, maintenance, and upgrade also affect the total cost of ownership. The system value is determined by comparing these costs to the total benefit and deriving the return on investment (ROI). Because these costs and the benefits are unique to each customer, the ROI can best be determined by a joint customer and sales person team.

Scalability

For this HiTest Suite, scalability can be described in two ways. In terms of hardware, scalability refers to the additional hardware components that can be added to a system within and beyond the HiTest configuration. In terms of performance, scalability refers to the workload capability of the HiTest configuration.

Additional Hardware Components

Systems that are configured from this HiTest Suite can easily be upgraded both within and beyond the ranges specified in the Suite.

In Figure 2-3, hardware scalability for this Suite is illustrated in terms of memory, number of CPUs, and disk space. Within the limits set for the enclosures called for in this HiTest Suite, comparisons are shown for the minimum and maximum limits of the system configuration.

Note that within this DIGITAL HiTest Suite, the choice of system or cabinets can limit future expansion. If expansion is needed, select components that will allow future expansion.

Figure 2-3: SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 HiTest Suite Scalability

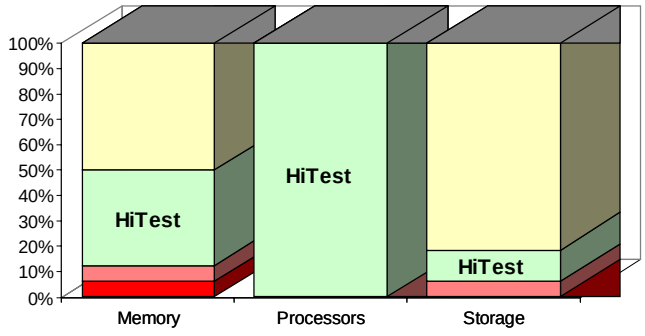

	Systems configured within this range involve tradeoffs between these, and sometimes other components.
	Systems configured within this range generally exceed the requirements of the market.
HiTest	Systems configured within this range meet the requirements of most customers.
	Systems configured within this range are generally insufficient to meet the requirements of the market.
	Systems cannot be configured in this range.

The AlphaServer 4x00 configurations of the SAP R/3 HiTest Suite meet the requirements for 80% of the SAP R/3 installations with 200 through 2000 users while providing significant expansion capability for situations that may reach beyond the scope of this HiTest Suite.

In general, systems can be configured beyond the limits illustrated in Figure 2-3 by adding additional storage cabinets, clusters, and other peripherals.

Workload Capability

Scalability also measures how performance is affected as additional resources and users are added. When scalability is measured by workload capability, the factors that are considered include the effectiveness of additional hardware; whether the system remains consistent as you add to it; and how expensive it is to add to it.

DIGITAL HiTest Suites are selected to provide an appropriate workload capability for the target application area. Often a choice of suites is available, each providing appropriate coverage for specific situations. HiTest works closely with other DIGITAL groups to ensure that a HiTest system will perform appropriately in a production environment. Many HiTest systems are tested and tuned for performance.

Characterization tests were performed for SD benchmark (SAP standard benchmark). These tests are described in Chapter 5.

Figure 2-4 shows the workload capability.

Figure 2-4: SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 Workload Capability

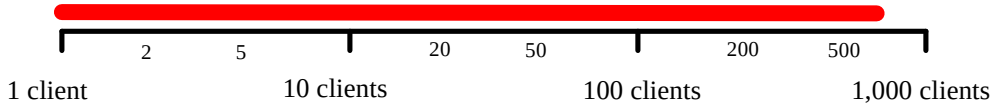

One SD user corresponds to about three users in the real application. Depending on which parts of the application are used, this number is higher or lower.

Benchmark data on other DIGITAL systems is available at the following URL:
`http://www.fra.dec.com/SAP-CC/Intranet/sizing/sizingliste.html`

Services

DIGITAL offers a range of service options. The following portfolio of Business Critical Services is available for HiTest Suites and is backed by the DIGITAL Uptime Guarantee.

Proof of Commitment: The DIGITAL Uptime Guarantee

The DIGITAL Uptime Guarantee is a formal contract that commits DIGITAL to keeping a client's business critical systems in operation at least 99.5% of the time, excluding outages beyond the control of DIGITAL, such as electrical shutdowns, environmental failures, and downtime caused by application failure. If uptime levels are lower than 99.5%, clients do not pay the full service charge.

Portfolio of Business Critical Services

The three vital elements of DIGITAL Business Critical Services are:

- Availability Review

  The first step in initiating a Business Critical engagement with DIGITAL is a customized, in-depth availability analysis of the computing environment, beginning with an overview of operating goals. This review identifies potential risks and trouble spots in hardware, software, operations, physical environment, and network. A comprehensive written report forms the basis for determining serviceability requirements.

- Business Critical Gold Support

  Clients who purchase Business Critical Gold Support work with a named technical account manager who serves as the single point of contact and ensures that problems are resolved quickly. A privileged hotline assures crisis response within 30 minutes. An assigned support team works with the account manager to apply continuous effort to critical problems. The on-site support agreement for Gold Support Customers provides coverage 24 hours a day and seven days a week. Additional benefits include:

  - Notification of software patches as soon as they become available
  - Notification of known problems and fixes
  - Monthly service activity review
  - Operating system upgrade impact planning
  - Bi-annual System Healthcheck assessments. These are conducted using advanced system-based tools to assess the performance and security of systems. The collected data is analyzed against accepted practices, and the findings, together with recommendations for corrective action, are documented in a summary report.

- Availability Partnership

  With Availability Partnership, system availability is maintained at the required level by measuring and analyzing actual system availability, and conducting regular updates to the original Availability Review. Particular focus is placed on:

  – Configuration and topology documentation
  – Availability status reporting
  – Change impact analysis
  – Proactive problem avoidance based on proactive patch/FCO/firmware management
  – Periodic detailed data collection and analysis
  – Availability model update
  – Contingency planning
  – Service planning and advising

Complementary Support Services

The three key Business Critical Services are augmented by:

- On-Site Parts Service

  DIGITAL works with the client to determine the appropriate inventory levels for their environment. A cost-effective *rental* parts solution is developed to maintain an on site inventory of spare parts.

- Installation and Startup

  DIGITAL offers rapid, worry-free implementation of new hardware and software – including systems, PCs, terminals, workstations, networking components, operating systems, layered products, applications, and software updates. Clients can choose hardware installation, software installation and startup, or both.

Meeting Client Needs Locally or Globally

With 450 service center locations in l00 countries, DIGITAL is prepared to deliver consistent and comprehensive service capabilities on a local or multinational basis. These services encompass:

- Total system support for servers, network operating system, applications, switching components, and PCs
- Multivendor support for a diverse range of products including networking equipment, applications, and peripherals
- Microsoft Solution Provider and Authorization Support Centers with the largest concentration of Microsoft certified engineers in the world

For More Information

To find out more about DIGITAL Business Critical Services, contact your local DIGITAL Multivendor Customer Services sales specialist or visit the Business Critical Services web site at `http://www.digital.com/services/mcs/mcs_critical.htm`.

Year 2000 Compliance

Year 2000 Compliance refers to whether computer systems will properly recognize the date change from December 31, 1999 to January 1, 2000. Current information on Year 2000 status of DIGITAL products can be obtained from the DIGITAL Year 2000 Program web site at `http://ww1.digital.com/year2000/`. Current information on the Year 2000 status of other vendor's products should be confirmed with those vendors.

While HiTest does not explicitly test for Year 2000 compliance in the components of this Suite, HiTest does check the published status of components where Year 2000 compliance would be a concern. The Year 2000 information presented here is accurate as of March 1998. Table 2-3 summarizes these findings.

The color codes used in the table represent the following categories of readiness:

- Blue – Version specified is Year 2000 ready today.
- Green – Currently not Year 2000 ready. Version to be Year 2000 ready specified with Year 2000 date noted.
- Yellow – Under evaluation.
- Red – Will not be made ready for Year 2000. Product will be removed from active status on or before 31 March 1998.
- N – Not Applicable - No Year 2000 implications exist for this component.

Table 2-3: SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 Year 2000 Compliance

Component	Year 2000 Status
AppSet software (SAP R/3)	Blue
Database Software (Oracle 7 for DIGITAL UNIX)	Blue
Middleware (TruCluster Available Server for DIGITAL UNIX)	Blue
Operating System (DIGITAL UNIX)	Blue
Platform (AlphaServer 4100)	Blue
Major Storage (if appropriate) (StorageWorks ESA 10000)	N

3
Configuration Data

This chapter describes the SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 HiTest Suite including the hardware, software, and firmware components and their revision levels. If required, special configuration rules are explained.

Hardware and Software Components

Table 3-1 and Table 3-2 identify the range of hardware and software components that can be configured using the HiTest Suite. These two tables form the DIGITAL HiTest Template. The ranges of hardware provided in this template include one through four GB of memory, eight 2.1 GB disks, six 9.1 GB disks and the ESA 10000 with different disk options (see the Special Configuration Rules section), and a 7-cartridge tape loader.

Table 3-3 lists the optional system management station hardware and software.

Table 3-4 lists the revision levels of the components.

The DIGITAL HiTest Template consists of three categories:

- AppSet Software – Includes software specific to one class of customer solutions, in this case SAP solution
- Foundation Hardware – Includes the base system, storage, and other hardware options
- Foundation Software – Includes the operating system, middleware, and database software

When ordering an item from a HiTest Template, select a quantity that is within the minimum and maximum range for the item. If the minimum quantity is zero (0), then the component is optional. If the minimum quantity is one or more, order at least the minimum quantity, but be cautious about exceeding the maximum quantity. The maximum quantity represents the greatest number of components that were tested for interoperability with all the other components in the Suite.

For more details on the HiTest Suite hardware configuration, see Chapter 7.

Table 3-1: DIGITAL HiTest Template – AppSet Software and Foundation Hardware

SAP HiTest AppSet Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 HiTest Foundation Hardware				
For documentation and updates:http://cosmo.tay.dec.com and http://www.businesslink.digital.com				
Line Item	**Description**	**Part Number**	**HiTest Range**	
			Min	**Max**
AppSet Software				
1	**SAP R/3 Version 3.1H** Contact SAP at http://www.sap.com	**SAP**	2	2
Foundation Hardware				
2	*Select two systems:* **AlphaServer 4100 5/533, 1 GB DIGITAL UNIX License** **AlphaServer 4100 5/533, 2 GB DIGITAL UNIX License** **AlphaServer 4000 5/533, 1 GB DIGITAL UNIX License** **AlphaServer 4000 5/533, 2 GB DIGITAL UNIX License** *Hardware includes*: • 5/533 MHz CPU with 4 MB cache • Memory • PB2GA-JB TRIO64 1 MB Graphics • DE500-AA 10/100 Mbit Fast Ethernet • KZPDA-AA FW SCSI and cable • SCSI CD-ROM drive • RX23L-AB 1.44 MB Floppy drive • LK47W-A2 PS/2 style keyboard • Three-button PS/2 compatible mouse *Software includes*: • DIGITAL UNIX Operating System and base license • Unlimited User license • DIGITAL UNIX Server Extensions • ServerWORKS • Internet AlphaServer Administration software • DECevent • BMC Patrol Agent	**DA-51KAC-FB** **DA-51KAC-GB** **DA-53KEC-FA** **DA-53KEC-GA**	2	2
3	*Select one enclosure:* **Pedestal with StorageWorks shelf** **Cabinet with StorageWorks shelf**	**BA30P-AB/BB** ❶ **H9A10-EL/EM**	2	2
4	*Select one enclosure:* **Single-bus UltraSCSI StorageWorks shelf for pedestal** **Single-bus UltraSCSI StorageWorks shelf for cabinet**	**BA36R-SD** **BA36R-RC**	2	2
5	**9.1 GB 7200 RPM UltraSCSI Disks** (internal)	**DS-RZ1DB-VW**	4	8
6	*For each system drawer installed in an H9A10-EL or H9A10-EM, order one:* **System Drawer Mounting Kit**	**CK-BA30A-BA/BB** ❶	2	2
7	**533 MHz CPU DIGITAL UNIX SMP UPG**	**KN305-DB**	0	6
8	**512 MB Memory Option** **1 GB Memory Option** **2 GB Memory Option** **Note:** This system supports a memory range from 1 to 8 GB per system. When selecting memory options, stay within the Template's 4 GB maximum.	**MS330-EA** **MS330-FA** **MS330-GA**	0	See Note
9	**PCI one-port FWD SCSI controller**	**KZPSA-BB**	5	5
10	**SCSI Y cable, 68 pin**	**BN21W-0B**	4	4
❶Indicates that geography-specific part number variants are available. Check the appropriate price book for details.				

SAP HiTest AppSet Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 HiTest Foundation Hardware				
For documentation and updates:http://cosmo.tay.dec.com and http://www.businesslink.digital.com				
Line Item	**Description**	**Part Number**	**HiTest Range**	
			Min	**Max**
11	**5 meter 16-bit SCSI cable**	**BN21K-05**	1	1
12	**10 meter VHDCI male to 68 HD male cable**	**BN38C-10**	4	4
13	**DIGITAL StorageWorks Enterprise Storage Array 10000** *Hardware includes:* • DS-HSZ70-AH (2) • BN37A-10 (2) • DS-BA370-AA (2)	**DS-SWXES-BA**	1	1
14	**Differential SCSI Terminator**	**H879-AA**	4	4
15	**GB 7200 RPM UltraSCSI Disks** (for Oracle and SAP executables)	**DS-RZ1DB-VW**	2	2
16	*Select additional one enclosure:* • Diskblock 1 with 4.3 GB Disks • Diskbblock 1, 2 with 4.3 GB Disks • Diskblock 1 with 9.1 GB Disks • Diskblock 1, 2 with 9.1 GB Disks • Diskblock 1, 2, 3 with 9.1 GB Disks • Diskblock 1, 2, 3, 4 with 9.1 GB Disks **Note:** For the database, you need a minimum of one additional diskblock of 4.3 or 9.1 GB disks. One diskblock equals 36 GB for the database.	 **DS-RZ1CB-VW** **DS-RZ1CB-VW** **DS-RZ1DB-VW** **DS-RZ1DB-VW** **DS-RZ1DB-VW** **DS-RZ1DB-VW**	 18 36 8 16 24 32	 18 36 8 16 24 32
17	**2.1 GB 7200 RPM UltraSCSI Disks** (for Oracle redologs)	**DS-RZ1BB-VW**	8	8
18	**FDDI-Controller Fiber - Dual Attach**	**DEFPA-DB**	2	2
19	**20-m SC to SC dual fiber cable**	**BN34B-20**	2	2
20	**140/280 GB 7-Cartridge Tape Loader System**	**TZ887-NE**	1	1
21	**SCSI Bus Extender and Signal Converter**	**DWZZA-AA**	1	1
22	**Single-ended SCSI cable**	**BC19J-IE**	1	1
23	*Select one high-resolution color monitor:* **15-in Flat-square with 0.28mm dot pitch** **17-in Trinitron with 0.28mm aperture grille pitch** **21-in Diamondtron with 0.28mm aperture grille pitch**	 **SN-VRCX5-WA** **SN-VRTX7-WA** **SN-VRCX1-WA**	2	2
24	**System Management Station**	**See Table 3-3**	0	1

Table 3-2: DIGITAL HiTest Template – Foundation Software

Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 HiTest Foundation Software						
For documentation and updates: http://cosmo.tay.dec.com and http://www.businesslink.digital.com						
Line Item	**Description**	**Part Number**	**HiTest Range**		**Required By**	
			Min	**Max**	**Fnd**[†]	**App**[†]
Foundation Software						
1	**DIGITAL UNIX for AlphaServer V4.0D**	**Included with item 2 of Table 3-1**	2	2	Yes	Yes
2	**DIGITAL UNIX V4.0D CD-ROM**	**QA-MT4AA-H8**	1	1	Yes	Yes
3	**TruCluster Available Server Software V1.5**	**QB-05SAG-AA**	2	2	Yes	Yes
4	**Oracle7 for DIGITAL UNIX V7.3.3**	**Included with item 1**	2	2	Yes	Yes
5	**HSZ70 Solutions Software for DIGITAL UNIX** *Includes:* • PC Card containing the storage controller software • StorageWorks Command Console • Licenses and documentation	**QB-5SBAB-SA**	4	4	Yes	Yes
6	**StorageWorks PLUS**, which includes: • Networker Save and Restore for DIGITAL UNIX V4.3 • DIGITAL UNIX Logical Storage Manager (LSM) • AdvFS Utilities	**QB-5RYAG-AA**	2	2	Yes	Yes
7	**Networker Save and Restore Application Interface for SAP R/3**	**QL-5JGAQ-AA**	2	2	Opt'l	Opt'l
8	**NSR Jukebox Tier 1 License**	**QL-04UAL-3B**	1	1	Opt'l	Opt'l
9	**Performance Advisor for DIGITAL UNIX, 3.0C CD-ROM**	**QA-054AA-H8**	2	2	Opt'l	Opt'l
10	**Performance Advisor for DIGITAL UNIX, 3.0C License**	**QL-0WFA9-AA**	2	2	Opt'l	Opt'l
11	**Hard copy of this Suite's HiTest Notes**	**EK-HSPXA-HN**	1	1	Yes	Yes
[†]Fnd = Foundation, App = AppSet						

Table 3-3: System Management Station Template

SAP HiTest AppSet System Management Station				
For documentation and updates: http://cosmo.tay.dec.com and http://www.businesslink.digital.com				
Line Item	**Description**	**Part Number**	**HiTest Range**	
			Min	**Max**
Note: This HiTest Suite supports the use of a system management station. When the management station option is included, this HiTest Template identifies the items required. When system management is provided through other means, this option may be omitted without invalidating the HiTest Suite.				
Management Station Hardware				
1	**DIGITAL PC 5100** *Hardware includes*: • 200 MHz Pentium CPU with MMX • 512 KB secondary cache • 32 MB memory • Integrated Fast Ethernet (10/100) • 16X CD-ROM • PCI 64-bit S3 ViRGE/GX graphics controller (with 2 MB) • 3.2 GB disk drive • 1.44 MB floppy *Software includes*: • Windows NT Workstation 4.0 (factory installed) **Note:** A functionally equivalent 80 x 86 system may be substituted without invalidating this HiTest Template.	**FR-DAB04-AF**	1	1
2	**Country Kit, North American**	**FR-PC94K-AA**	1	1
3	**32 MB SDRAM dual-bank DIMM Memory**	**FR-PCCAM-EC**	1	1
4	**Diamond 56.6 K Modem** **Note:** Used for page notification.	**FR-PCXFA-AA**	0	1
5	*Select one high-resolution monitor:* **21" (19.6" view) 1600 x 1200 @75Hz** **19" (18" view) 1600 x 1200 @75Hz** **17" (16" view) 1280 x 1024 @75Hz**	 **FR-PCXAV-WZ** **FR-PCXAV-TZ** **FR-PCXAV-YZ**	1	1
Software Installed on Managed System(s)				
6	**Windows NT Workstation 4.0** **Note:** Install Windows NT Service Pack 3 (available from `ftp://ftp.microsoft.com/bussys/winnt`)	**Included with item 1**	1	1
7	**Hummingbird Exceed, Version 6.0** Contact `http://www.hummingbird.com`	**Hummingbird**	1	1
8	**DIGITAL ServerWORKS Manager, Version 3.0**	**Included with the base system**	1	1
9	**StorageWorks Command Console, V1.1B**	**Included with the Storage-Works kit**		

SAP HiTest AppSet System Management Station				
For documentation and updates: http://cosmo.tay.dec.com and http://www.businesslink.digital.com				
Line Item	**Description**	**Part Number**	**HiTest Range**	
			Min	**Max**
10	*Choose one BMC product:* **BMC PATROLWATCH for ServerWORKS, V3.2**, **BMC PATROL Operator Console Windows NT** **BMC PATROL Developer Console Windows NT** **Note:** BMC products that are included with the AlphaServer are on the ServerWORKS Quick Launch CD. BMC PATROL Developer Console includes BMC PATROL Operator Console. Contact BMC at: `http://www.bmc.com`	**Included with base system** **BMC** **BMC**	0	1
11	**BMC PATROL Agent for Windows NT, lic. and media**	**QB-5KKAB-WA**	0	1
12	**BMC Operating System Knowledge Module for Windows NT, license and media**	**QB-5KLAB-WA**	0	1
Software Installed on Managed Systems				
13	**Base UNIX systems management tools**	**Included with DIGITAL UNIX**	1	1
14	**DIGITAL UNIX Management Agent**	**Included with item 8**	1	1
15	**StorageWorks Command Console Agent**	**Included with item 9**	1	1
16	**BMC PATROL Agent for UNIX**	**Included with the AlphaServer**	0	1
17	**BMC Operating System Knowledge Module** **Note: W*** refers to the class: WA - Desktop; WB - Workgroup; WC - Departmental; WD - Enterprise Server	**QB-5KLAA-W***	0	1

Table 3-4: Component Revision Levels

Hardware Component	Hardware	Firmware	Software
SRM console	–	5.0-2	–
AlphaBIOS	–	5.63-0	–
SCSI host adapter (KZPDA-AA)	Rev. B01	–	V1.19-5
FWD SCSI controller (KZPSA-BB)	Rev. P01	3.50-0-18	–
2.1 GB disks (DS-RZ1BB-VW)	–	DEC0656	–
4.3 GB disks (DS-RZ1CB-VW)	–	DEC0656	–
9.1 GB disk (DS-RZ1DB-VW)	–	LYJ0	–
533 MHz CPU (KN305-DB)	Rev. A01	3.0	–
1 GB Memory (MS330-FA)	Rev. B01	–	–
StorageWorks shelf power supply (DS-BA35X-HH)	Rev. A01	–	–
Software Component		**Version/Revision**	**Patch Level**
DIGITAL UNIX		4.0D	
SAP R/3		3.1H	
Networker Save & Restore (NSR) DIGITAL UNIX		4.3	
Networker Save & Restore Application Interface for SAP R/3		1.1	
HSZ70 Array Controller software		70Z-0	–

Special Configuration Rules

The special configuration rules for the Suite are as follows. For details, see the section Operating System Installation in Chapter 4.

- Root and Swap disk
  - The boot disk is a 9 GB disk.
  - For the root partition we used 256 MB. The swap1 will get 6,4 GB and /usr 2 GB.
  - The swap disk size depends on how much memory is used.
  - The complete boot disk is mirrored.

- Database disks
  The database consists of three major parts:
  - Oracle redo log files (origlogA, origlogB) and the mirror log files (mirrlogA, mirrlogB). The redo log files consists of two disks with a size of 2 GB each. The mirror log files also consist of two disks with a size of 2 GB each.

– SAP other (sapmnt, oracle, saparch, sapbackup, sapreorg and so on).
The SAP other disk is a 9 GB disk, to hand a lot of Archive files in saparch. On the other hand, all SAP and Oracle executables are installed on this disk. The remaining space is for all the directories and mountpoints in the SAP Environment.

– Database files (sapdata1 to sapdata6)
The database files are located on diskblocks. Each diskblock has a size of 36 GB. You can get a 36 GB diskblock with either 9 times 4 GB disks or with 4 times 9 GB disks. The 36 GB diskblock size should be the upper limit in case of data loss so the restore of the diskblock can be done in a useful timeframe. Storage is one ESA 10000 Enterprise Storage Array with four times HSZ70 each and additional disks. The four HSZ70 controllers in a single ESA 10000 are devided in two parts – two controllers and their failover. The disk configuration must be equal on both HSZ70s:

  Minimum Diskblock with 4 GB disks results 3 times 3 disk stripeset equals 36 GB.
  Minimum Diskblock with 9 GB disks results 2 times 2 disk stripeset equals 36 GB.
  Maximum Diskblock with 4 GB disks results 6 times 3 disk stripeset equals 72 GB.
  Maximum Diskblock with 9 GB disks results 8 times 2 disk stripeset equals 144 GB.

  The stripesets are mirrored with the Logical Storage Manager and put together with the AdvFS. For example:

  If you have 3 times a 3 disk stripe of 4 GB disks on the HSZ70, you will see three different volumes with the size of 12 GB after the configuration on the HSZ. Each volume is mirrored with the volumes on the other HSZ with the help of LSM. The AdvFS will put these 3 volumes together to a single volume with the size of 36 GB. So you use the striping of the HSZ, the mirroring of LSM, and the parallel mechanism of AdvFS. At the end, you will see one big volume of 36 GB. If you order more diskblocks you will have more AdvFS domains (72 GB, 108 GB or 144 GB).

R/3 Specific Configuration Rules

This section describes security and performance rules to follow for R/3-specific disk configuration.

Security Rules

These security rules are the most important and must be followed. Failure to do so could lead to loss of data.

The rules are:

- The Oracle log files and the Oracle archive files have to be on different disks.
- The archive files are not on a sapdata file system.
- The Oracle control files are on two different disks.

Performance Rules

The following rules are important to get the right performance:

- The Oracle log files are on different disks.
- The Oracle log files are not on sapdata file system.
- The archive files are on a single disk.

SAP did a performance layout using sapdata1 to sapdata6, this layout was not changed in our environment.

The following decisions were made:

- 36 GB diskblock contains — sapdata1 to sapdata6
- 72 GB diskblock contains — sapdata1 to sapdata3 on the first volume
sapdata4 to sapdata6 on the second volume
- 108 GB diskblock contains — sapdata1 and sapdata2 on the first volume
sapdata3 and sapdata4 on the second volume
sapdata5 and sapdata6 on the third volume
- 144 GB diskblock contains — sapdata1 and sapdata2 on the first volume
sapdata3 and sapdata4 on the second volume
sapdata5 on the third volume
sapdata6 on the fourth volume

For more information, see Installation of the ESA 10000 later in this document.

4
System Installation and Setup

This chapter describes how to install and set up a DIGITAL HiTest System configured from this DIGITAL HiTest Suite. System preparation includes installing hardware, operating system, and applications.

It is important to perform the operating system and application installations in the same sequence as described in the documentation:

1. Install UNIX on the DB-Server (hostname tcr001) and the Application-Server (hostname tcr002).
2. Install AdvFS utilities on both systems.
3. Install TruCluster ASE on both systems.
4. Install Logical Storage Manager (LSM) with the disk layout on the DB-Server, and only install LSM on the Application Server, which is done during encapsulation of the root disk (see Operating System Installation).
5. Install SAP R/3 on both systems.
6. Perform the postinstallation for TruCluster ASE in a SAP R/3 environment.

Hardware Installation

Install and interconnect the hardware as shown in Chapter 7.

The difference between the maximum and the minimum configuration is the amount of memory, CPU's and disks. The amount of CPU's does not influence the amount of disks and memory.

If more memory is added, the SAP kernel and profile parameters are modified and the disks for swapping have to be increased.

Amount of swapdisks (refer to *SAP Installation Guide Operating Dependencies*):

- Up to 2 GB RAM are cached by the swap1 (=6,4 GB) on the boot disk
- 3 – 5 GB RAM are cached by an additional 9 GB disk (=15,4 GB)
- 6 – 8 GB RAM are cached by another 9 GB disk (=24,4 GB)

The size of swapspace has to be identical on both systems. In case of failover, the SAP R/3 system needs the same available swapspace.

KZPSA Installation

If both systems are connected to the HSZ70s through shared-SCSI buses, be aware that the KZPSA-BB options are set to SCSI ID 7 by default. Set the SCSI ID to 6 on one of the two systems, as the following describes.

The following procedure shows how to change the SCSI IDs and check the system configuration:

1. Set the SCSI ID of the shared KZPSAs to 6, as follows:
```
P00>>>set kzpsa0_host_id 6
P00>>>set kzpsa1_host_id 6
```

2. Check the shared KZPSA option settings, as follows:
```
P00>>>show kzp*
kzpsa0_fast           1
kzpsa0_host_id        6     ←1st shared bus
kzpsa0_termpwr        1
kzpsa1_fast           1
kzpsa1_host_id        6     ← 2nd shared bus
kzpsa1_termpwr        1
kzpsa2_fast           1
kzpsa2_host_id        7     ←3rd for tape drive
kzpsa2_termpwr        1
```

3. Check all devices:
```
P00>>>show dev
polling ncr0 (NCR 53C810) slot 1,bus 0 PCI, hose 1 SCSI Bus ID 7
dka500.5.0.1.1     DKa500                   RRD46  0557
polling isp0 (QLogic ISP1020)slot 0,bus 2 PCI,hose 1 SCSI Bus ID 7
dkb0.0.0.2000.1    DKb0                     RZ1DB-CA  LYJ0
dkb600.6.0.2000.1  DKb600                   RZ1DB-CA  LYJ0
polling kzpsa0(DEC KZPSA)slot 3,bus 0 PCI,hose 1 TPwr 1 Fast 1 Bus
ID6
kzpsa0.6.0.3.1     dkc     TPwr 1 Fast 1 Bus ID 6   P01  A11
polling kzpsa1(DEC KZPSA)slot 4,bus 0 PCI,hose 1 TPwr 1 Fast 1 Bus
ID6
polling kzpsa2(DEC KZPSA)slot 5,bus 0 PCI,hose 1 TPwr 1 Fast 1 Bus
ID6
kzpsa2.6.0.5.1     dke     TPwr 1 Fast 1 Bus ID 6   P01  A11
polling floppy0 (FLOPPY) PCEB - XBUS hose 0
dva0.0.0.1000.0    DVA0                      RX23
polling pfi0(DEC PCI FDDI) slot 4,bus 0 PCI,hose 0
fwa0.0.0.4.0    08-00-2B-B9-DD-A0
polling tulip0 (DECchip 21140-AA) slot 5, bus 0 PCI, hose 0
ewa0.0.0.5.0       00-00-F8-04-99-19    Twisted-Pair
```

   Your output can differ, depending on what options you have ordered. The previous example shows the mandatory disks with diskblock 1 and 2 with 4 GB disks and no additional swap disk.

4. Check the CPUs as follows:

```
P00>>>show cpu
Primary CPU:      00
Active CPUs:      00    01
Configured CPUs:  00    01
SROM Revision:    V3.0  V3.0
```

   Your output can differ depending on which amount of CPU's you have ordered. In this example, two CPU's are shown.

5. Check the amount of Memory:

```
P00>>>show mem
Slot    Type       MB         Base Address
----   ------    ------      -------- -----------
 0       EDO      2048          0
Total             2048
```

Proceed to Installation of the ESA 10000 on both systems.

ESA 10000 Installation

The HSZ70 software (HSOF) is shipped separately. The HSZ will not function without the HSOF software.

Connect a serial terminal (vtxxx) to the first HSZ70 serial port. After the HSZ70 has been powered on, enter the following commands from the terminal to create and verify the controller configuration:

Note

Refer to the *StorageWorks Array Controllers: HS Family of Array Controllers Users Guide* for a complete description of the HSZ70 configuration procedures.

1. Set the prompt, date, failover, id's and preferred id's on both HSZ70:

```
HSZ> set this prompt="HSZ1_1> "
HSZ1_1> set failover copy=this
HSZ1_1> set other prompt="HSZ1_2> "
HSZ1_1> set this command_console_lun
HSZ1_1> set other command_console_lun
HSZ1_1> set this id=1,2,3,4
HSZ1_1> set this mirrored_cache
HSZ1_1> set other id=1,2,3,4
HSZ1_1> set this preferred_id = 1,3
HSZ1_1> set other preferred_id = 2,4
HSZ1_1> set this time=29-jan-1998:15:30:00
```

   The command command_console_lun creates a LUN for the hszterm or the SWCC. You will see one more disk in console mode on the UNIX system which is created automatically and looks like the following:

```
dkc103.1.0.3.1   Dkc103   HSZ70CCL    V70Z
```

2. Verify that the HSZ70 firmware (HSOF) is at revision V70Z-0, the cache size is 32 MB and in a GOOD state, and the battery state is GOOD:

```
HSZ_1> show this
Controller:
HSZ70-CX ZG74302938 Firmware V70Z-0, Hardware H01
Configured for dual-redundancy with ZG74202789
In dual-redundant configuration
Device Port SCSI address 6
Time:29-Jan-1998 15:30:06


Host port:
SCSI target(s) (1, 2, 3, 4),
Preferred target(s) (1, 3)
TRANSFER_RATE_REQUESTED = 20MHZ
Host Functionality Mode = A
Command Console LUN is target 1, lun 1

Cache:
32 megabyte write cache, version 4
Cache is GOOD
Battery is GOOD
No unflushed data in cache
CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
NOCACHE_UPS

Mirrored Cache:
32 megabyte write cache, version 4
Cache is GOOD
Battery is GOOD
No unflushed data in cache
HSZ_1>show other
Controller:
HSZ70-CX ZG74202789 Firmware V70Z-0, Hardware H01
Configured for dual-redundancy with ZG74302938
In dual-redundant configuration
Device Port SCSI address 7
Time:29-Jan-1998 15:30:53

Host port:
SCSI target(s) (1, 2, 3, 4),
Preferred target(s) (2, 4)
TRANSFER_RATE_REQUESTED = 20MHZ
Host Functionality Mode = A
Command Console LUN is target 1, lun 1

Cache:
32 megabyte write cache, version 4
Cache is GOOD
Battery is GOOD
No unflushed data in cache
CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
NOCACHE_UPS
```

```
Mirrored Cache:
32 megabyte write cache, version 4
Cache is GOOD
Battery is GOOD
No unflushed data in cache
```

3. Reboot the HSZ70 and check that the modifications are still valid:
   ```
   HSZ_1> reboot other
   HSZ_1> reboot this
   ```

4. Check the disk layout. If no disk is seen or disks are added after the config of the HSZ70, run either config or run cfmenu where cfmenu is menu driven.

Note

The disk layout depends on how many diskblocks for the database you have selected, 4 GB or 9 GB disks are possible in the database configuration. The disk configuration must be equal on both HSZs. Do not mix 4 GB and 9 GB disks for the diskblocks.

Depending on the diskblocks you have ordered, you have to setup the disks.
The following disks are mandatory regardless of how many diskblocks you have chosen for the database:

Figure 4-1: ESA 10000 Frontview - Mandatory Disks

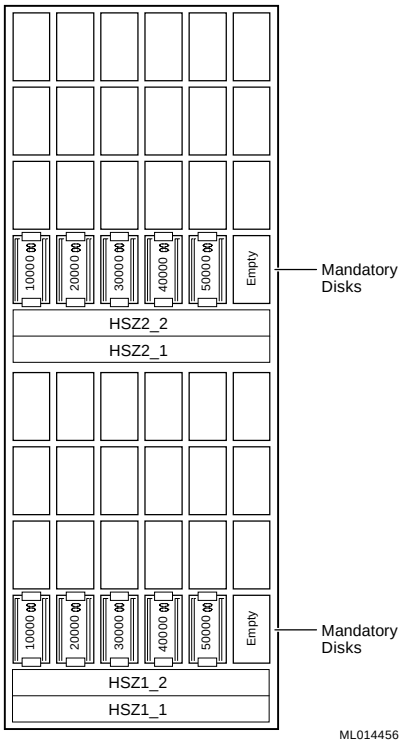

DISK10000 Redolog A 2GB
DISK20000 Redolog B 2GB
DISK30000 Mirrlog A 2GB
DISK40000 Mirrlog B 2GB
DISK50000 SAP_Rest 9GB

```
HSZ> init DISK10000 save_configuration
HSZ> init DISK20000 save_configuration
HSZ> init DISK30000 save_configuration
HSZ> init DISK40000 save_configuration
HSZ> init DISK50000 save_configuration
HSZ> add unit d100 DISK10000 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d200 DISK20000 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d300 DISK30000 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d400 DISK40000 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d403 DISK50000 writeback_cache
maximum_cached_transfer=1024
```

The disk DISK60000 is not used in both HSZ‘s.
This slot in the ESA 10000 should always be empty.
The following shows the setup for the other disks. Depending on the ordered diskblocks choose the correct part, which is valid.

Diskblock 1 with 4 GB Disks:

Figure 4-2: ESA 10000 - Diskblock 1 with 4 GB Disks

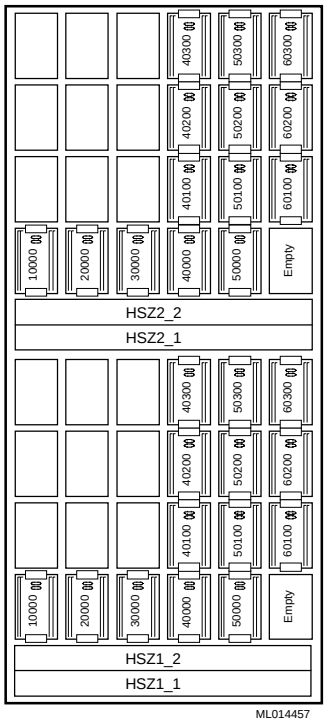

DISK40100	stripeset s1
DISK50100	stripeset s1
DISK60100	stripeset s1
DISK40200	stripeset s2
DISK50200	stripeset s2
DISK60200	stripeset s2
DISK40300	stripeset s3
DISK50300	stripeset s3
DISK60300	stripeset s3

```
HSZ> add stripe s1 DISK40100 DISK50100 DISK60100
HSZ> add stripe s2 DISK40200 DISK50200 DISK60200
HSZ> add stripe s3 DISK40300 DISK50300 DISK60300
HSZ> init s1 chunksize=256 save_configuration
HSZ> init s2 chunksize=256 save_configuration
HSZ> init s3 chunksize=256 save_configuration
HSZ> add unit d101 s1 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d201 s2 writeback_cache
maximum_cached_transfer=1024
```

```
HSZ> add unit d301 s3 writeback_cache
maximum_cached_transfer=1024
```

Diskblock 1 and 2 with 4 GB Disks:

Figure 4-3: ESA 10000 Frontview - Diskblock 1 and 2 with 4 GB Disks

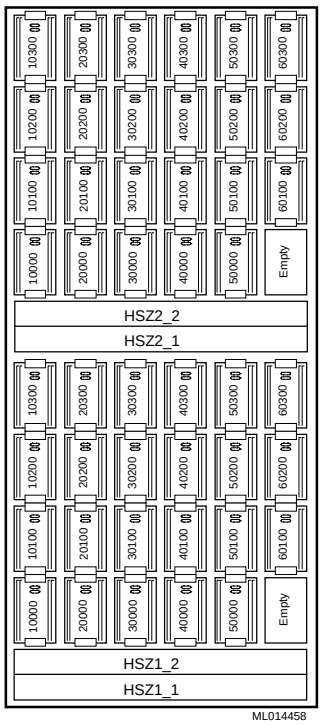

DISK40100	stripeset s1
DISK50100	stripeset s1
DISK60100	stripeset s1
DISK40200	stripeset s2
DISK50200	stripeset s2
DISK60200	stripeset s2
DISK40300	stripeset s3
DISK50300	stripeset s3
DISK60300	stripeset s3
DISK10100	stripeset s4
DISK20100	stripeset s4
DISK30100	stripeset s4
DISK10200	stripeset s5
DISK20200	stripeset s5
DISK30200	stripeset s5
DISK10300	stripeset s5
DISK20300	stripeset s5
DISK30300	stripeset s5

```
HSZ> add stripe s1 DISK40100 DISK50100 DISK60100
HSZ> add stripe s2 DISK40200 DISK50200 DISK60200
HSZ> add stripe s3 DISK40300 DISK50300 DISK60300
HSZ> add stripe s4 DISK10100 DISK20100 DISK30100
HSZ> add stripe s5 DISK10200 DISK20200 DISK30200
HSZ> add stripe s6 DISK10300 DISK20300 DISK30300
HSZ> init s1 chunksize=256 save_configuration
HSZ> init s2 chunksize=256 save_configuration
HSZ> init s3 chunksize=256 save_configuration
HSZ> init s4 chunksize=256 save_configuration
HSZ> init s5 chunksize=256 save_configuration
HSZ> init s6 chunksize=256 save_configuration
HSZ> add unit d101 s1 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d201 s2 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d301 s3 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d401 s4 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d102 s5 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d202 s6 writeback_cache
maximum_cached_transfer=1024
```

Diskblock 1 with 9 GB Disks:

Figure 4-4: ESA 10000 Frontview - Diskblock 1 with 9 GB Disks

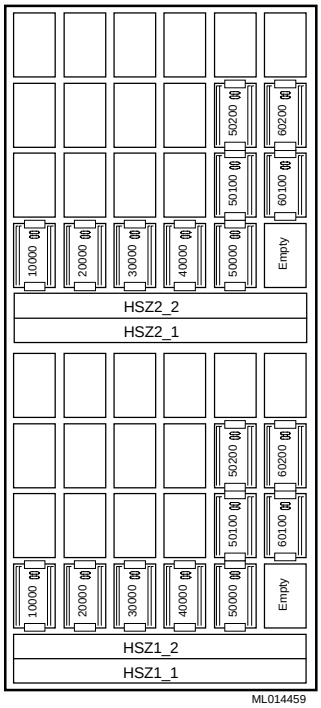

DISK60100 stripeset s1
DISK50100 stripeset s1
DISK60200 stripeset s2
DISK50200 stripeset s2

HSZ> **add stripe s1 DISK50100 DISK60100**
HSZ> **add stripe s2 DISK50200 DISK60200**
HSZ> **init s1 chunksize=256 save_configuration**
HSZ> **init s2 chunksize=256 save_configuration**
HSZ> **add unit d101 s1 writeback_cache maximum_cached_transfer=1024**
HSZ> **add unit d201 s2 writeback_cache maximum_cached_transfer=1024**

Diskblock 1 and 2 with 9 GB Disks:

Figure 4-5: ESA 10000 Frontview - Diskblock 1 and 2 with 9 GB Disks

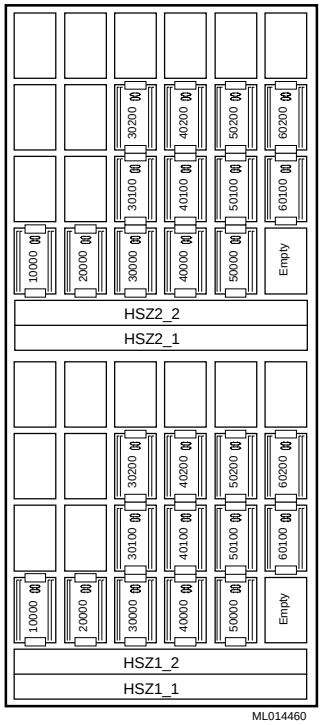

DISK60100	stripeset s1
DISK50100	stripeset s1
DISK60200	stripeset s2
DISK50200	stripeset s2
DISK40100	stripeset s3
DISK30100	stripeset s3
DISK40200	stripeset s4
DISK30200	stripeset s4

```
HSZ> add stripe s1 DISK50100 DISK60100
HSZ> add stripe s2 DISK50200 DISK60200
HSZ> add stripe s3 DISK40100 DISK30100
HSZ> add stripe s4 DISK40200 DISK30200
HSZ> init s1 chunksize=256 save_configuration
HSZ> init s2 chunksize=256 save_configuration
HSZ> init s3 chunksize=256 save_configuration
HSZ> init s4 chunksize=256 save_configuration
HSZ> add unit d101 s1 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d201 s2 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d301 s3 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d401 s4 writeback_cache
maximum_cached_transfer=1024
```

Diskblock 1 and 2 and 3 with 9 GB Disks:

Figure 4-6: ESA 10000 Frontview - Diskblock 1, 2, 3 with 9 GB Disks

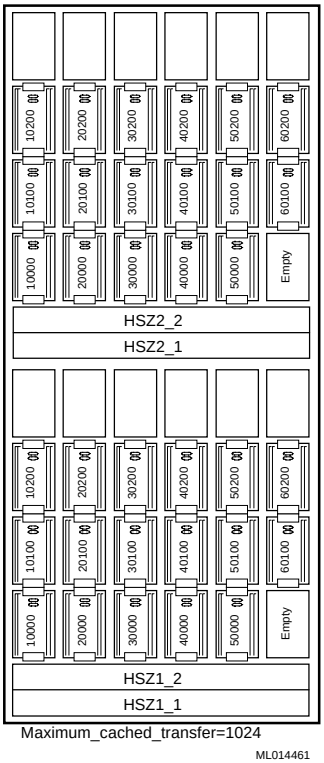

DISK60100 stripeset s1
DISK50100 stripeset s1
DISK60200 stripeset s2
DISK50200 stripeset s2
DISK40100 stripeset s3
DISK30100 stripeset s3
DISK40200 stripeset s4
DISK30200 stripeset s4
DISK20100 stripeset s5
DISK10100 stripeset s5
DISK20200 stripeset s6
DISK10200 stripeset s6

```
HSZ> add stripe s1 DISK50100 DISK60100
HSZ> add stripe s2 DISK50200 DISK60200
HSZ> add stripe s3 DISK30100 DISK40100
HSZ> add stripe s4 DISK30200 DISK40200
HSZ> add stripe s5 DISK10100 DISK20100
HSZ> add stripe s6 DISK10200 DISK20200
HSZ> init s1 chunksize=256 save_configuration
HSZ> init s2 chunksize=256 save_configuration
HSZ> init s3 chunksize=256 save_configuration
HSZ> init s4 chunksize=256 save_configuration
HSZ> init s5 chunksize=256 save_configuration
HSZ> init s6 chunksize=256 save_configuration
HSZ> add unit d101 s1 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d201 s2 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d301 s3 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d401 s4 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d202 s5 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d402 s6 writeback_cache
maximum_cached_transfer=1024
```

Diskblock 1 and 2 and 3 and 4 with 9 GB Disks:

Figure 4-7: ESA 10000 Frontview - Diskblock 1, 2, 3, 4 with 9 GB Disks

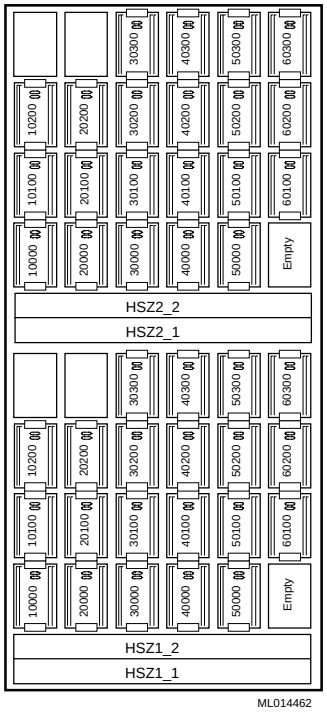

DISK60100 stripeset s1
DISK50100 stripeset s1
DISK60200 stripeset s2
DISK50200 stripeset s2
DISK40100 stripeset s3
DISK30100 stripeset s3
DISK40200 stripeset s4
DISK30200 stripeset s4
DISK20100 stripeset s5
DISK10100 stripeset s5
DISK20200 stripeset s6
DISK10200 stripeset s6
DISK60300 stripeset s7
DISK50300 stripeset s7
DISK40300 stripeset s8
DISK30300 stripeset s8

```
HSZ> add stripe s1 DISK50100 DISK60100
HSZ> add stripe s2 DISK50200 DISK60200
HSZ> add stripe s3 DISK30100 DISK40100
HSZ> add stripe s4 DISK30200 DISK40200
HSZ> add stripe s5 DISK10100 DISK20100
HSZ> add stripe s6 DISK10200 DISK20200
HSZ> add stripe s7 DISK50300 DISK60300
HSZ> add stripe s8 DISK30300 DISK40300
HSZ> init s1 chunksize=256 save_configuration
HSZ> init s2 chunksize=256 save_configuration
HSZ> init s3 chunksize=256 save_configuration
HSZ> init s4 chunksize=256 save_configuration
HSZ> init s5 chunksize=256 save_configuration
HSZ> init s6 chunksize=256 save_configuration
HSZ> init s7 chunksize=256 save_configuration
HSZ> init s8 chunksize=256 save_configuration
HSZ> add unit d101 s1 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d201 s2 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d301 s3 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d401 s4 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d202 s5 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d402 s6 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d102 s7 writeback_cache
maximum_cached_transfer=1024
HSZ> add unit d302 s8 writeback_cache
maximum_cached_transfer=1024
```

Operating System Installation

Install the DIGITAL UNIX 4.0D operating system with all kernel options.

1. Boot the CD containing the DIGITAL UNIX 4.0D distribution:
   P00>>>**boot device** (where device=CD)
2. Follow the steps in the *DIGITAL UNIX Installation Guide.*
3. When prompted to select the type of installation, select `UNIX Shell` and proceed to the Disk Initialization section to create disk device special files and check partitioning.

Disk Initialization

This section describes the steps required to create disk device special files, label the disks, and check the partitions. Use rz8 as root disk and rz14 as mirror of rz8.

1. Change to the /dev directory, as follows:
   # **cd /dev**
2. Use the MAKEDEV command to create disk device special files for rz8:
   # **./MAKEDEV rz8**
   ```
   MAKEDEV: special file(s) for rz8:
   rz8a rz8b rrz8b rz8c rrz8c rz8d rrz8d rz8e rrz8e rz8f rrz8f
   rz8g rrz8g rz8h rrz8h
   ```
3. Check the disk label information on rz8:
   # **disklabel -r rz8**
4. Start the disk label editor:
   # **disklabel -e rz8**
   ```
   /dev/rrz8a:
   type: SCSI
   disk: RZ1DB-CA
   label:
   flags: dynamic_geometry
   bytes/sector: 512
   sectors/track: 168
   tracks/cylinder: 20
   sectors/cylinder: 3360
   cylinders: 5273
   sectors/unit: 17773524
   rpm: 7200
   interleave: 1
   trackskew: 9
   cylinderskew: 9
   headswitch: 0           # milliseconds
   track-to-track seek: 0  # milliseconds
   drivedata: 0
   8 partitions:
   #    size     offset  fstype  [fsize bsize cpg]
   a:   524288        0  unused              # (Cyl.  0 - 156*)
   b: 13052884   524288  unused              # (Cyl.156*- 4040*)
   c: 17773524        0  unused    0     0   # (Cyl. 0 - 5289)
   d:  4194304 13577172  unused              # (Cyl.4040*- 5289*)
   e:        0        0  unused    0     0   # (Cyl.0- -1*)
   f:        0        0  unused    0     0   # (Cyl.0- -1)
   g:     1024 17771476  unused    0     0   # (Cyl.5289*- 5289*)
   ```

```
h:     1024 17772500  unused    0     0   # (Cyl.5289*- 5289*)
```

5. Save your edits and exit the editor:

```
:wq
write new label? [y]: y
#
# exit
INIT: SINGLE-USER MODE
Initializing system for DIGITAL UNIX installation. Please
wait...
*** Performing CD-ROM Installation
Loading installation process and scanning system hardware.
```

6. Go to the section, Operating System Custom Installation.

The setup must be equal on the database server and the Application Server except the hostnames and the internet addresses.

Operating System Custom Installation

This section describes the custom installation phase of the operating system installation.

1. Select `2) Custom Installation.`
2. Follow the display and enter data when prompted. The hostname for the first HiTest system was tcr001.
3. The following message is displayed when the disk scan occurs:

```
** Scanning system for disk devices. Please wait ...
Only one disk detected in this system (/dev/rz8, SCSI RZ29B
type).All file systems will be on that disk.
The rz1 disk has a non-default partition table.
Partition  Start      Size       End    Overlaps
Default
    a          0    131072     131071    c
    b     131072    401408     532479    c
    c          0   8380080    8380079    a b d e f g h
    d     532480   2623488    3155967    c g
    e    3155968   2623488    5779455    c g h
    f    5779456   2600624    8380079    c h
    g     532480   3936256    4468735    c d e
    h    4468736   3911344    8380079    c e f
Existing
    a          0   2000000    1999999    c
    b    2000000   2000000    3999999    c
    c          0   8380080    8380079    a b d e h
    d    4000000   2000000    5999999    c
    e    6000000   2000000    7999999    c
    f          0         0         -1
    g          0         0         -1
    h    8000000    380080    8380079    c
```

4. Select a disk for DIGITAL UNIX installation. The root file system will be placed on the "a" partition of this disk.

```
      Disk       Device     Device   Controller   Controller
      Type        Name      Number      Name        Number
1)    RZ1DB-CA    rz8          8        SCSI           1
2)    RZ1DB-CA    rz14        14        SCSI           1

Enter your choice: 1

The rz8 disk has a partition table that is not recommended.

Partition       Start       Size      End       Overlaps

Recommended
      a             0      262144     262143              c
      b        262144      262144     524287              c
      c             0    17773524   17773523   a b d e f g h
      d        524288     5749745    6274032         c g h
      e       6274033     5749745   12023777           c h
      f      12023778     5749746   17773523           c h
      g        524288     1433600    1957887           c d
      h       1957888    15815636   17773523       c d e f
Existing
      a             0      524288     524287              c
      b        524288    13052884   13577171              c
      c             0    17773524   17773523          a b d
      d      13577172     4194304   17771475              c
      e             0           0         -1
      f             0           0         -1
      g      17771476        1024   17772499
      h      17772500        1024   17773523
```

5. Select the existing table and do not use the default disk layout, as follows:

```
Choose which partition table to use.
1) Recommended table
2) Existing table
Enter your choice: 2
The default disk layout is:
  * root file system on the "a" partition, type UFS
  * /usr file system on the "g" partition, type UFS
  * /var as part of /usr
  * first swapping area (swap1) on the "b" partition
  * no second swapping area (swap2)
Use this default disk layout (y/n) ? n
```

6. Select the AdvFS file system for the root file system, as follows:

```
1) UFS -- UNIX file system
2) AdvFS -- advanced file system
Enter your choice: 2
```

7. Select the disk where the /usr file system will reside:

```
   Disk      Device    Device    Controller    Controller
   Type      Name      Number    Name          Number
1) RZ1DB-CA  rz8       8         SCSI          1
2) RZ1DB-CA  rz14      14        SCSI          1
Enter your choice: 1
```

8. Select the rz8 partition where the /usr file system will reside, as follows:

```
  Partition   Start         Size         End   Overlaps
1)    b      524288     13052884    13577171     c
2)    d    13577172      4194304    17771475     c
Enter your choice: 2
```

9. Select the file system type for the /usr file system:

```
1) UFS -- UNIX file system
2) AdvFS -- advanced file system
Enter your choice: 2
```

10. Select the disk where the first swapping area (swap1) will reside:

```
   Disk      Device    Device    Controller    Controller
   Type      Name      Number    Name          Number
1) RZ1DB-CA  rz8       8         SCSI          1
2) RZ1DB-CA  rz14      14        SCSI          1
Enter your choice: 1
```

    Swap1 will be on the b partition of the rz8 because all of its other partitions are already in use.

11. You may choose to have a second swapping area (swap2).

```
Do you want a second swapping area (y/n) ?n
```

12. You can make /var a separate file system, or you can have it share space on the /usr file system.

```
Should /var be a separate file system (y/n) ?n
```

13. Check the file system:

```
You have requested this file system layout:
  * root file system on rz8a, type AdvFS
  * /usr file system on rz8d, type AdvFS
  * /var will be on /usr
  * first swapping area (swap1) will be on rz8b
  * no second swapping area (swap2)
Is this the correct file system layout (y/n)? y
```

14. Install all mandatory and optional subsets.

15. Set the console boot variables, as follows, then boot:

```
>>> set boot_osflags A
>>> set bootdef_dev dkb0
>>> boot
```

16. When prompted, select all kernel options. At this point, the kernel is built and the system is rebooted.

17. Do not edit the configuration file.

18. When the system is up, start to modify the disklabel of disk rz14 to prepare it for mirroring the boot disk.

19. Write a disklabel to rz14:
    #disklabel -wr rz14 XYZ

    You can use XYZ as the disktype option because disklabel will check the disktype itself.

20. Edit the disklabel of disk rz14 and set the following partition values:
    #disklabel -e rz14

```
/dev/rrz14a:
type: SCSI
disk: RZ1DB-CA
label:
flags: dynamic_geometry
bytes/sector: 512
sectors/track: 168
tracks/cylinder: 20
sectors/cylinder: 3360
cylinders: 5273
sectors/unit: 17773524
rpm: 7200
interleave: 1
trackskew: 9
cylinderskew: 9
headswitch: 0           # milliseconds
track-to-track seek: 0  # milliseconds
drivedata: 0
8 partitions:

#    size     offset  fstype  [fsize bsize cpg]
a:   524288        0  unused              # (Cyl.  0 - 156*)
b: 13053908   524288  unused              # (Cyl.156*- 4041*)
c: 17773524        0  unused    0     0   # (Cyl. 0 - 5289)
d:  4194304 13578196  unused              # (Cyl.4041*- 5289*)
e:        0        0  unused    0     0   # (Cyl.0- -1*)
f:        0        0  unused    0     0   # (Cyl.0- -1)
g:        0        0  unused    0     0   # (Cyl.0*- -1*)
h:     1024 17772500  unused    0     0   # (Cyl.5289*- 5289*)
```

    This label differs from the label of the root disk because the volrootmir procedure needs 1024 blocks more of swapspace. The space of the g: partition is used to extend the swapspace by 1024 blocks. After volrootmir the disklabels are equal. The partition g: and h: can be used to enable Block Change Logging (BCL) or Dirty Region Log (DRL) in the future.

21. Encapsulate the boot disk:
    #volencap rz8

```
Setting up encapsulation for rz8.
- Disk rz8 is the system boot disk and LSM is not initialized.
Creating simple disk rz8e to initialize LSM and rootdg.
- Partition rz8a is the root partition which requires 2 passes to
encapsulate and the temporary use of a free partition.
Using partition rz8f for temporary root encapsulation.
- Creating nopriv disk for primary swap device rz8b.
- Creating nopriv disk for rz8d.

The following disks are queued up for encapsulation or use by LSM.
```

```
You must reboot the system to perform the actual encapsulations.
rz8e rz8a rz8f rz8b rz8d
```

The system will initialize LSM and reboot after the disk is prepared for encapsulation. See also the section Logical Storage Manager Documentation later in this Chapter.

Licenses

Licenses, also known as PAKs (Product Authorization Keys) are delivered with the system. The PAKs that come with this system include:

- OSF-BASE – Required for any system running the DIGITAL UNIX Operating System.
- StorageWorks PLUS – Required to use AdvFS Utilities, NSR and LSM product.
- TruCluster Available Server – Required to use TruCluster Available Server.

Register using the following steps:

1. **`#lmf register`** (or **`#lmfsetup`**)
   A template is displayed. At this moment, you are using the vi editor.
2. Add the information contained in the PAKs that came with the software.
3. Install all additional licenses.
4. Display the licenses with the command:
   # **`lmf list`**

Final Adjustments for the Operating System

1. Mirror the boot disk:
   #**`volrootmir -a rz14`**
   Mirroring system disk rz8 to disk rz14.

   ```
   This operation will destroy all contents on disk rz14.
   The partition map from rz8 will be copied to rz14 and all
   volumes associated with rz8 will be mirrored.

   Do you want to continue with this operation? (y or n) y
   Initializing rz14.

   Mirroring rootvol to rz14a.
   Mirroring swapvol to rz14b.
   Mirroring vol-rz8d to rz14d.
   ```

2. Now the disklabel of both disks - rz8 and rz14 should look like this:
   ```
   8 partitions:

   #    size     offset  fstype  [fsize bsize cpg]
   a:   524288        0  AdvFS              # (Cyl.  0 - 156*)
   b: 13051860   524288  swap              # (Cyl.156*- 4040*)
   c: 17773524        0  unused    0     0   # (Cyl. 0 - 5289)
   d:  4194304 13577172  AdvFS               # (Cyl.4040*- 5289*)
   e:     1024 13576148  LSMsimp   0     0   # (Cyl.4040- 4040*)
   f:        0        0  unused    0     0   # (Cyl.0- -1)
   g:     1024 17771476  unused    0     0   # (Cyl.5289*- 5389*)
   h:     1024 17772500  unused    0     0   # (Cyl.5289*- 5289*)
   ```

3. Now the boot disk rz8 is mirrored by disk rz14.
   Check that with the following command:
   #**volprint**
   The output should look like this:

```
TYPE NAME         ASSOC        KSTATE     LENGTH COMMENT
dg   rootdg       rootdg       -               -
dm   rz14a        rz14a        -          524288
dm   rz14b        rz14b        -        13051860
dm   rz14d        rz14d        -         4194304
dm   rz14e        rz14e        -               0
dm   rz8a         rz8a         -          524288
dm   rz8b         rz8b         -        13051860
dm   rz8d         rz8d         -         4194304
dm   rz8e         rz8e         -               0
sd   rz14a-01p    rootvol-02   -              16
sd   rz14a-01     rootvol-02   -          524272
sd   rz14b-01     swapvol-02   -        13051860
sd   rz14d-01     vol-rz8d-02  -         4194304
sd   rz8a-01p     rootvol-01   -              16
sd   rz8a-01      rootvol-01   -          524272
sd   rz8b-01      swapvol-01   -        13051860
sd   rz8d-01      vol-rz8d-01  -         4194304
plex rootvol-01   rootvol      ENABLED    524288
plex rootvol-02   rootvol      ENABLED    524288
plex swapvol-01   swapvol      ENABLED  13051860
plex swapvol-02   swapvol      ENABLED  13051860
plex vol-rz8d-01  vol-rz8d     ENABLED   4194304
plex vol-rz8d-02  vol-rz8d     ENABLED   4194304
vol  rootvol      root         ENABLED    524288
vol  swapvol      swap         ENABLED  13051860
vol  vol-rz8d     fsgen        ENABLED   4194304
```

4. Shutdown the system to set the bootdef_dev parameter to both disks:
   #**shutdown -h now**
   >>>**set bootdef_dev “dkb0 dkb600”**
   If the disk dkb0 fails the system tries to boot from dkb600.

5. Enable the Block Change Logging on the /usr disk vol -rz8d:

```
# voldisksetup -i rz8g publen=512 privlen=512
# voldisksetup -i rz8h publen=512 privlen=512
# voldisksetup -i rz14g publen=512 privlen=512
# voldisksetup -i rz14h publen=512 privlen=512
# voldg -g rootdg adddisk rz8g=rz8g
# voldg -g rootdg adddisk rz8h=rz8h
# voldg -g rootdg adddisk rz14g=rz14g
# voldg -g rootdg adddisk rz14h=rz14h
# volmake -g rootdg sd  rz8g-01 dm_name=rz8g dm_offset=0 len=1
# volmake -g rootdg sd  rz8h-01 dm_name=rz8h dm_offset=0 len=1
# volmake -g rootdg sd rz14g-01 dm_name=rz14g dm_offset=0 len=1
# volmake -g rootdg sd rz14h-01 dm_name=rz14h dm_offset=0 len=1
# volsd -g rootdg aslog vol-rz8d-01 rz8g-01
# volsd -g rootdg aslog vol-rz8d-02 rz14g-01
```

   The synchronization without BCL on the /usr partition takes about 20 minutes. With BCL enabled just a few seconds.

6. Check to see if the log disks are connected to the devices.
   Use either `volprint` or `dxlsm` to check that the BCL disk has joined the /usr disk vol -rz8d.
   #**`volprint -Ath`**
   or
   #**`dxlsm`**

   The devices rz8h-01 and rz14h-01 are reserved if you want to enable BCL on the /device rootvol. However there is performance disadvantage if you do enable it.

Configure Network Interfaces

After the kernel build and system reboot, login as root and start setup as follows:

1. # **`setup`**
2. From the setup menu, select:
   `1) Network Configuration Application`.
3. From the Network Configuration Application menu, select:
   `1  Configure Network Interfaces`

The environment in which a system is installed determines many of the network configuration settings. The following section list the settings used for this HiTest System as installed in the test laboratory. After the interfaces are configured, exit the menu and enter yes when prompted to have netsetup automatically restart the network services. Exit from the Setup menu and reboot the system. The DIGITAL UNIX installation is complete.

FDDI Controller (fta0)

Hostname for the system is tcr001.
IP address for interface fta0 is 1.0.0.1.
Subnet mask for fta0 is 255.255.255.0.
No additional ifconfig flags for this interface were set.

Fast Ethernet Controller (tu0)

Hostname for interface tu0 is tcr001e.
The IP address for interface tu0 is 155.56.201.101.
The subnet mask for tu0 is 255.255.255.0.
No additional ifconfig flags for this interface were set.

DIGITAL UNIX Patch Installation

This section describes how the DIGITAL UNIX patches were installed. Patches are dependent on the DIGITAL UNIX version (4.0D for this HiTest system). The patches are located at `ftp://ftp.service.digital.com/public/Digital_UNIX/v4.0d`.

- At the time the system was tested, no valid DIGITAL UNIX patches were available.
- If you are using Network Save and Restore, install the following Patch:
  `GURU://usr/guest/misc/osf/marshall/TCR150/libaseapi.tar`

Note

The contents of the patch directory are frequently updated. The patch file names may change as these updates occur.

Network Setup

This section describes setting up NFS and the hosts and rhosts files:

1. Before you start to setup nfs, create the following directories:
```
#mkdir/sammnt
#mkdir/sapmnt/TCR
#mkdir/usr/sap
#mkdir/usr/sap/trans
```

2. Start nfssetup as follows:
```
tcr001> nfssetup
```

3. Enter the following information when prompted by nfssetup:
```
NFS locking to be enabled [y]? y
Will you be exporting any directories [n] ? y
Do you want to allow non-root mounts [n] ? n
Enter the number of TCP daemons to run (0 to 128) [8] : 8
Enter the number of UDP daemons to run (0 to 120) [8] : 8
Would you like to run the property list daemon [n] ? n
Enter the number of block I/O daemons to run [7] : 7
Would you like to run the PC-NFS daemon [n] ? n
Would you like to run the automount daemon [n] ? n
Enter the directory pathname: /sapmnt/TCR
Netgroup/Machine name: tcr002
Enter the directory pathname: /usr/sap/trans
Netgroup/Machine name: tcr002
```

4. The NFS environment is confirmed by nfssetup, as follows:
```
    8 TCP server daemons, 8 UDP server daemons
    7 nfsiod daemons
    locking daemons installed
```

   Directory export list:
```
    /sapmnt/TCR exported to: tcr002
    /usr/sap/trans exported to: tcr002
```

Note

The /usr/sap/trans directory must also be on the shared devices. In a usual R/3 installation it is not the default. You must adjust for Availability Server purposes.

hosts and rhosts

Create the file .rhosts in the directory /. The HiTest system rhosts files contained the following lines:

```
tcr001 root
tcr001e root
tcr002 root
tcr002e root
<any other host> root
```

This allows rsh, rlogin and rcp between these systems.

/etc/hosts

Edit the file `hosts` in /etc to include the names of the hosts on the network.

Note

Your internet addresses and hostnames can differ from the following. If you use different names, you must adapt yours.

The lines added for the HiTest System environment are:

```
127.0.0.1        localhost
1.0.0.1          tcr001        ← FDDI host name DB Server
1.0.0.11         vtcr001       ← Virtual host name for TCR150 DB Server
155.56.201.101   tcr001        ← Ethernet name DB Server
155.56.201.111   vtcr001e      ← Virtual name for TCR150 PC Connection DB
                                 Server
1.0.0.2          tcr002        ← FDDI host name Appl. Server
1.0.0.12         vtcr002       ← Virtual host name for TCR150 Appl. Server
155.56.201.102   tcr002        ← Ethernet host name DB Server
155.56.201.112   vtcr002e      ← Virtual name for TCR150 PC Connection Appl.
                                 Server
```

The user's PC connects via the virtual name vtcr001e and vtcr002e to the SAP R/3 System. The communication between DB Server and Application Server is done via vtcr001 and vtcr002. The virtual addresses will move to the system where the specific instances run or fail over.

For example, if the DB Server fails, the virtual host name vtcr001 and vtcr001e will fail over to the Application Server and the PC's can connect to vtcr001e which is physically on tcr002. This is also valid for vtcr002 and vtcr002e on tcr001.

Postinstallation Procedures

The procedures in this section label the disks, configure additional swap space, and install additional software.

Labeling the Local Disks

The local disks cannot be added to swap unless they received a disklabel. Refer to the section Adding Swap space.

Labeling the Shared Disks

This section describes how the shared disks are labeled.

1. Check the HSZ disks with the following command:
   #file /dev/rrz*c|grep HSZ

   The output looks similar to the following:
   ```
   # /dev/rrzd20c:  character special (8/37058) SCSI
   #2 HSZ70 disk #163 (SCSI ID #4) (SCSI LUN #3)
   ```

2. Label all disks with the following command:
   #disklabel -wr <disk> XYZ
   You can use the disktype XYZ, because disklabel is checking the type of disk during the labeling.
   Do not try to disklabel the device which is called HSZ70CCL. This is a virtual SCSI device for the hszterm or SWCC. If you try to initialize it, you will get a write error.

Adding Swapspace

Swapspace is based on how much memory you use:

- If you are using up to 2 GB of Memory you do not need additional swapspace.
- If you are using 3 GB to 5 GB of Memory, add another 9 GB disk next to the root disk and modify the /etc/fstab file:
  /dev/rz9c swap2 ufs sw 0 2
- If you are using 6 GB to 8 GB of Memory, add another 9 GB disk next to swap2 and modify the /etc/fstab file:
  /dev/rz10c swap3 ufs sw 0 2

Do not forget to label the disks before you add them to swap:

```
# disklabel –wr rz9 XYZ
# disklabel –wr rz10 XYZ
```

The amount of swapspace has to be identical on both systems. In case of a failover, the R/3 system has to have the same available swapspace.

Installing the AdvFS Utilities

Install the AdvFS Utilities to take more than one device into one domain. Use the following steps:

1. Insert the CD DIGITAL UNIX 4.0D Associated Products Volume 2 into the CD-ROM drive and mount it as follows:
   ```
   # mount –r /dev/rz5c /mnt
   ```
2. Go to the specific directory and start the installation. Please refer to the documentation on the same CD-ROM in the directory: `/mnt/AdvFS_Advanced_Utilities/doc` The documentation is available in Text, HTML and Postscript format.
3. Change to the directory and install the software:
   ```
   # cd  /mnt/AdvFS_Advanced_Utilities/kit
   # setld –l .
   choose:     n) ALL of the above
   Is this correct ? (y/n):  yes
   ```
4. Check the screen for possible errors.

TruCluster Available Server

Install the TruCluster Available Server Software. The software is located on the DIGITAL UNIX V4.0D Associated Products Volume 2 CD-ROM.

1. Change to the directory /mnt/TruCluster/kit:
   ```
   #cd /mnt/TruCluster/kit
   ```
2. Start the installation. Please refer to the documentation on the same CD-ROM in the directory /mnt/TruCluster/doc. The documentation is available in text, HTML and Postscript format.
3. Start the installation and select: `ALL mandatory and optional subsets`
   ```
   # setld -l
   ```

4. Enter the following information (or information specific to your system), as prompted by the installation program:

```
Enter the IP name for the member network interface [tcr001]:tcr001
Is this correct? [n]: yes
Do you want to run the ASE logger on this node? [n]: yes
Enter the name of the kernel configuration file.[TCR001]: TCR001
Do you want to edit the configuration file ? (y/n) [n]: no
Select the controllers thet define the shared ASE I/O Buses.
Enter your choice (comma or space separated): 2 3
Are the above choices correct (y/n) [y]: yes
```

5. Exit the I/O Controller Name specification Menu:

```
All controllers connected to an I/O bus must be named the
same on all ASE members.  Enter the controller names for all
shared ASE I/O buses by assigning them one at a time or all
at once with the below options.
    Name     New Name   Controller  Slot    Bus    Slot
 1) scsi2    scsi2      pza0         0      pci1    3
 2) scsi3    scsi3      pza1         0      pci1    4
 f)  Assign buses starting at a given number
 p)  Assign buses as was done in pre-ASE V1.3
 v)  View non shared controllers
 s)  Show previous assignments
 r)  Reapply previous assignments
 q)  Quit without making any changes
 x)  Exit (done with modifications)
Enter your choice [f]: x

Your new scsi controller configuration is:
        Name    Controller      Slot    Bus     Slot
        scsi0   psiop0          0       pci1    1
        scsi1   isp0            0       pci2001 0
        scsi2   pza0            0       pci1    3
        scsi3   pza1            0       pci1    4
        scsi4   pza2            0       pci1    5
        scsi5   pza3            0       pci0    3

Is this ok? [y]: yes
```

6. The kernel is rebuilt after exiting from the menu. Check the kernel files, then move them and reboot as follows:

```
tcr001> ls -l /vm*
-rwxr-xr-x   1 root     system   12260144 May 26 10:02 /vmunix
-rwxr-xr-x   1 root     system   12240512 May 23 15:05 /vmunix.bef_patch
```

7. Move the kernel:

```
tcr001> mv /vmunix /vmunix.before.ase
tcr001> mv /sys/TCR001/vmunix /vmunix
tcr001>  reboot
```

8. Repeat the previous steps on the Application Server (tcr002 in this HiTest Suite), except labeling the shared disk. This has to be done only once.

Notes

If the configuration in /sys/conf/TCR001 is changed, especially SCSI setup, run /var/ase/sbin/ase_fix_config again to make the changes work with TruCluster V1.5. However, changing SCSI assignments will invalidate any volume definition that may exist then.

Refer to *SAP R/3 DECsafe Package V2.2A*, which is available from the DIGITAL Intranet at URL:
`http://www.fra.dec.com/SAP-Eng/available/cluster.html`

Read install.doc before you start to install TruCluster V1.5.

Logical Storage Manager (LSM)

Since the LSM Software is already installed and started (remember the mirroring of the boot disk), the next step is to create a new diskgroup (sapdg) and to add all the shared disks to that diskgroup.

From now on, the installation of the LSM Volumes can slightly differ depending of the diskblocks you have ordered (4 GB Disks or 9 GB Disks and the amount of diskblocks). The Mandatory Part is shown first and is valid for all installations. The diskblocks steps are shown step by step later in this chapter. This setup is done only on TCR001.

Mandatory Disksetup

The following chapter describes the disksetup, which is mandatory for all options:

1. Initial all mandatory disks which should be used by LSM:
   The first HSZ70 contains disk rz17 to disk rzd20.
   The second HSZ70 contains disk rz25 to disk rzd28.
   ```
   # voldisksetup -i rz17c nlog=1 nconfig=1
   # voldisksetup -i rz18c nlog=1 nconfig=1
   # voldisksetup -i rz19c nlog=1 nconfig=1
   # voldisksetup -i rz20c nlog=1 nconfig=1
   # voldisksetup -i rzd20c nlog=1 nconfig=1
   # voldisksetup -i rz25c nlog=1 nconfig=1
   # voldisksetup -i rz26c nlog=1 nconfig=1
   # voldisksetup -i rz27c nlog=1 nconfig=1
   # voldisksetup -i rz28c nlog=1 nconfig=1
   # voldisksetup -i rzd28c nlog=1 nconfig=1
   ```
2. Create the LSM disk group sapdg:
   ```
   # voldg init sapdg rz17=rz17c
   ```
3. Add all other mandatory disks to LSM group sapdg:
   ```
   # voldg -g sapdg adddisk  rz18=rz18c
   # voldg -g sapdg adddisk  rz19=rz19c
   # voldg -g sapdg adddisk  rz20=rz20c
   # voldg -g sapdg adddisk  rzd20=rzd20c
   # voldg -g sapdg adddisk  rz25=rz25c
   # voldg -g sapdg adddisk  rz26=rz26c
   # voldg -g sapdg adddisk  rz27=rz27c
   # voldg -g sapdg adddisk  rz28=rz28c
   # voldg -g sapdg adddisk  rzd28=rzd28c
   ```

4. Since Block Change Logging (BCL) will be used, create log subdisks to be associated with each Plex (see also LSM documentation).
The size of each log file is 1024 bytes.

```
# volmake -g sapdg sd rz17-01 dm_name=rz17 dm_offset=0 len=1k
# volmake -g sapdg sd rz18-01 dm_name=rz18 dm_offset=0 len=1k
# volmake -g sapdg sd rz19-01 dm_name=rz19 dm_offset=0 len=1k
# volmake -g sapdg sd rz20-01 dm_name=rz20 dm_offset=0 len=1k
# volmake -g sapdg sd rzd20-01 dm_name=rzd20 dm_offset=0 len=1k
# volmake -g sapdg sd rz25-01 dm_name=rz25 dm_offset=0 len=1k
# volmake -g sapdg sd rz26-01 dm_name=rz26 dm_offset=0 len=1k
# volmake -g sapdg sd rz27-01 dm_name=rz27 dm_offset=0 len=1k
# volmake -g sapdg sd rz28-01 dm_name=rz28 dm_offset=0 len=1k
# volmake -g sapdg sd rzd28-01 dm_name=rzd28 dm_offset=0 len=1k
```

5. Create the mandatory LSM data volumes:

   a) First you have to evaluate what the disksize is you want to add. This can easily be done with the help of volprint:

```
# volprint -th -g sapdg |grep dm
dm rz17 rz17c simple 1024 8376488 /dev/rrz17c
dm rz18 rz18c simple 1024 8376488 /dev/rrz18c
dm rz19 rz19c simple 1024 8376488 /dev/rrz19c
dm rz20 rz20c simple 1024 8376488 /dev/rrz20c
dm rzd20 rzd20c simple 1024 8376488 /dev/rrzd20c
```

   b) Next look for the disks rz17, rz18, rz19, rz20 and rzd20. The disks rz25, rz26, rz27, rz28 and rzd28 are created automatically during the mirroring.

   c) Evaluate the size of the subdisk. Use the values you got from the volprint command for the rz17. This example shows a 4 GB disk. In this configuration, the rz17 to the rz20 and the rz25 to rz28 are 2 GB disks. The rzd20 and rzd28 are 9 GB disks. Your display can differ from this. Subtract 2 from this value. You get 8376486 for the example of the rz17.

```
#valassist -g sapdg -U fsgen make redoA 8376486
alloc=0 align=0 rz17
```

   d) Create the subdisks for the mandatory disks for LSM. Do this for the following disks. Replace $SIZE with the values in volprint degraded by 2.

```
#volassist -g sapdg -U fsgen make redoA $SIZE
 alloc=0 align=0 rz17
#volassist -g sapdg -U fsgen make redoB $SIZE
 alloc=0 align=0 rz18
#volassist -g sapdg -U fsgen make mirrA $SIZE
 alloc=0 align=0 rz19
#volassist -g sapdg -U fsgen make mirrB $SIZE
 alloc=0 align=0 rz20
#volassist -g sapdg -U fsgen make SAPRest $SIZE
 alloc=0 align=0 rzd20
```

   e) Mirror the LSM disks in the background:

```
#volassist -g sapdg mirror redoA rz25 &
#volassist -g sapdg mirror redoB rz26 &
#volassist -g sapdg mirror mirrA rz27 &
#volassist -g sapdg mirror mirrB rz28 &
#volassist -g sapdg mirror SAPRest rz25 &
```

f) Associate the Log Plex with each Data Plex.
This step can only be done if the mirroring in the background is finished.

```
#volsd -g sapdg aslog redoA-01 rz17-01
#volsd -g sapdg aslog redoA-02 rz25-01
#volsd -g sapdg aslog redoB-01 rz18-01
#volsd -g sapdg aslog redoB-02 rz26-01
#volsd -g sapdg aslog mirrA-01 rz19-01
#volsd -g sapdg aslog mirrA-02 rz27-01
#volsd -g sapdg aslog mirrB-01 rz20-01
#volsd -g sapdg aslog mirrB-02 rz28-01
#volsd -g sapdg aslog SAPRest-01 rzd20-01
#volsd -g sapdg aslog SAPRest-02 rzd28-01
```

Diskblock 1 with 4 GB Disks

1. Initialize all diskblocks, which should be used by LSM:

```
# voldisksetup -i rzb17c nlog=1 nconfig=1
# voldisksetup -i rzb18c nlog=1 nconfig=1
# voldisksetup -i rzb19c nlog=1 nconfig=1
# voldisksetup -i rzb25c nlog=1 nconfig=1
# voldisksetup -i rzb26c nlog=1 nconfig=1
# voldisksetup -i rzb27c nlog=1 nconfig=1
```

2. Add all other diskblocks to diskgroup sapdg:

```
# voldg -g sapdg adddisk rzb17=rzb17c
# voldg -g sapdg adddisk rzb18=rzb18c
# voldg -g sapdg adddisk rzb19=rzb19c
# voldg -g sapdg adddisk rzb25=rzb25c
# voldg -g sapdg adddisk rzb26=rzb26c
# voldg -g sapdg adddisk rzb27=rzb27c
```

3. Since Block Change Logging will be used, create log subdisks to be associated with each Plex (see also LSM documentation about BCL).
The size of each log file is 1024 Bytes:

```
# volmake -g sapdg sd rzb17-01 dm_name=rzb17
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb18-01 dm_name=rzb18
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb19-01 dm_name=rzb19
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb25-01 dm_name=rzb25
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb26-01 dm_name=rzb26
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb27-01 dm_name=rzb27
  dm_offset=0 len=1k
```

4. Create the LSM data volumes. Get the size of the LSM subdisk with volprint as decribed before and create the disks:

```
# volassist -g sapdg -U fsgen make SAP1_S1  $SIZE
  alloc=0 align=0 rzb17
# volassist -g sapdg -U fsgen make SAP1_S2  $SIZE
  alloc=0 align=0 rzb18
# volassist -g sapdg -U fsgen make SAP1_S3  $SIZE
  alloc=0 align=0 rzb19
```

5. Mirror the LSM Disks in the background:
```
# volassist –g sapdg mirror SAP1_S1 rzb25 &
# volassist –g sapdg mirror SAP1_S2 rzb26 &
# volassist –g sapdg mirror SAP1_S3 rzb27 &
```

6. Associate the Log Plex with each Data Plex.
   This step can only be done if the mirroring in the background is finished:
```
# volsd –g sapdg aslog SAP1_S1-01 rzb17-01
# volsd –g sapdg aslog SAP1_S1-02 rzb25-01
# volsd –g sapdg aslog SAP1_S2-01 rzb18-01
# volsd –g sapdg aslog SAP1_S2-02 rzb26-01
# volsd –g sapdg aslog SAP1_S3-01 rzb19-01
# volsd –g sapdg aslog SAP1_S3-02 rzb27-01
```

Diskblock 1 and 2 with 4 GB Disks

1. Initialize all diskblocks, which should be used by LSM:
```
# voldisksetup -i rzb17c nlog=1 nconfig=1
# voldisksetup -i rzb18c nlog=1 nconfig=1
# voldisksetup -i rzb19c nlog=1 nconfig=1
# voldisksetup -i rzb20c nlog=1 nconfig=1
# voldisksetup -i rzc17c nlog=1 nconfig=1
# voldisksetup -i rzc18c nlog=1 nconfig=1
# voldisksetup -i rzb25c nlog=1 nconfig=1
# voldisksetup -i rzb26c nlog=1 nconfig=1
# voldisksetup -i rzb27c nlog=1 nconfig=1
# voldisksetup -i rzb28c nlog=1 nconfig=1
# voldisksetup -i rzc25c nlog=1 nconfig=1
# voldisksetup -i rzc26c nlog=1 nconfig=1
```

2. Add all other diskblocks to diskgroup sapdg:
```
# voldg -g sapdg adddisk rzb17=rzb17c
# voldg -g sapdg adddisk rzb18=rzb18c
# voldg -g sapdg adddisk rzb19=rzb19c
# voldg -g sapdg adddisk rzb20=rzb20c
# voldg -g sapdg adddisk rzc17=rzc17c
# voldg -g sapdg adddisk rzc19=rzc19c
# voldg -g sapdg adddisk rzb25=rzb25c
# voldg -g sapdg adddisk rzb26=rzb26c
# voldg -g sapdg adddisk rzb27=rzb27c
# voldg -g sapdg adddisk rzb28=rzb28c
# voldg -g sapdg adddisk rzc25=rzc25c
# voldg -g sapdg adddisk rzc26=rzc26c
```

3. Since Block Change Logging will be used, create Log subdisks to be associated with each Plex. (See also LSM documentation about BCL).
   The size of each log file is 1024 Bytes:
```
# volmake -g sapdg sd rzb17-01 dm_name=rzb17
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb18-01 dm_name=rzb18
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb19-01 dm_name=rzb19
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb20-01 dm_name=rzb20
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc17-01 dm_name=rzc17
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc18-01 dm_name=rzc18
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb25-01 dm_name=rzb25
```

```
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb26-01 dm_name=rzb26
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb27-01 dm_name=rzb27
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb28-01 dm_name=rzb28
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc25-01 dm_name=rzc25
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc26-01 dm_name=rzc26
  dm_offset=0 len=1k
```

4. Create the LSM data volumes. Get the size of the LSM subdisk with volprint as decribed before and create the disks:

```
# volassist -g sapdg -U fsgen make SAP1_S1  $SIZE
  alloc=0 align=0 rzb17
# volassist -g sapdg -U fsgen make SAP1_S2  $SIZE
  alloc=0 align=0 rzb18
# volassist -g sapdg -U fsgen make SAP1_S3  $SIZE
  alloc=0 align=0 rzb19
# volassist -g sapdg -U fsgen make SAP2_S4  $SIZE
  alloc=0 align=0 rzb20
# volassist -g sapdg -U fsgen make SAP2_S5  $SIZE
  alloc=0 align=0 rzc17
# volassist -g sapdg -U fsgen make SAP2_S6  $SIZE
  alloc=0 align=0 rzc18
```

5. Mirror the LSM Disks in the background:

```
# volassist –g sapdg mirror SAP1_S1 rzb25 &
# volassist –g sapdg mirror SAP1_S2 rzb26 &
# volassist –g sapdg mirror SAP1_S3 rzb27 &
# volassist –g sapdg mirror SAP2_S4 rzb28 &
# volassist –g sapdg mirror SAP2_S5 rzc25 &
# volassist –g sapdg mirror SAP2_S6 rzc26 &
```

6. Associate the Log Plex with each Data Plex. This step can only be done if the mirroring in the background is finished:

```
# volsd –g sapdg aslog SAP1_S1-01 rzb17-01
# volsd –g sapdg aslog SAP1_S1-02 rzb25-01
# volsd –g sapdg aslog SAP1_S2-01 rzb18-01
# volsd –g sapdg aslog SAP1_S2-02 rzb26-01
# volsd –g sapdg aslog SAP1_S3-01 rzb19-01
# volsd –g sapdg aslog SAP1_S3-02 rzb27-01
# volsd –g sapdg aslog SAP2_S4-01 rzb20-01
# volsd –g sapdg aslog SAP2_S4-02 rzb28-01
# volsd –g sapdg aslog SAP2_S5-01 rzc17-01
# volsd –g sapdg aslog SAP2_S5-02 rzc25-01
# volsd –g sapdg aslog SAP2_S6-01 rzc18-01
# volsd –g sapdg aslog SAP2_S6-02 rzc26-01
```

Diskblock 1 with 9 GB Disks

1. Initialize all diskblocks, which should be used by LSM:
```
# voldisksetup -i rzb17c nlog=1 nconfig=1
# voldisksetup -i rzb18c nlog=1 nconfig=1
# voldisksetup -i rzb25c nlog=1 nconfig=1
# voldisksetup -i rzb26c nlog=1 nconfig=1
```

2. Add all other diskblocks to diskgroup sapdg:
```
# voldg -g sapdg adddisk rzb17=rzb17c
# voldg -g sapdg adddisk rzb18=rzb18c
# voldg -g sapdg adddisk rzb25=rzb25c
# voldg -g sapdg adddisk rzb26=rzb26c
```

3. Since Block Change Logging will be used, create log subdisks to be associated with each Plex (see also LSM documentation about BCL). The size of each log file is 1024 Bytes:
```
# volmake -g sapdg sd rzb17-01 dm_name=rzb17
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb18-01 dm_name=rzb18
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb25-01 dm_name=rzb25
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb26-01 dm_name=rzb26
  dm_offset=0 len=1k
```

4. Create the LSM data volumes. Get the size of the LSM subdisk with volprint as decribed and create the disks:
```
# volassist -g sapdg -U fsgen make SAP1_S1  $SIZE
  alloc=0 align=0 rzb17
# volassist -g sapdg -U fsgen make SAP1_S2  $SIZE
  alloc=0 align=0 rzb18
```

5. Mirror the LSM Disks in the background:
```
# volassist –g sapdg mirror SAP1_S1 rzb25 &
# volassist –g sapdg mirror SAP1_S2 rzb26 &
```

6. Associate the Log Plex with each Data Plex. This step can only be done if the mirroring in the background is finished:
```
# volsd –g sapdg aslog SAP1_S1-01 rzb17-01
# volsd –g sapdg aslog SAP1_S1-02 rzb25-01
# volsd –g sapdg aslog SAP1_S2-01 rzb18-01
# volsd –g sapdg aslog SAP1_S2-02 rzb26-01
```

Diskblock 1 and 2 with 9 GB Disks

1. Initialize all diskblocks, which should be used by LSM:
```
# voldisksetup -i rzb17c nlog=1 nconfig=1
# voldisksetup -i rzb18c nlog=1 nconfig=1
# voldisksetup -i rzb19c nlog=1 nconfig=1
# voldisksetup -i rzb20c nlog=1 nconfig=1
# voldisksetup -i rzb25c nlog=1 nconfig=1
# voldisksetup -i rzb26c nlog=1 nconfig=1
# voldisksetup -i rzb27c nlog=1 nconfig=1
# voldisksetup -i rzb28c nlog=1 nconfig=1
```

2. Add all other diskblocks to diskgroup sapdg:

```
# voldg -g sapdg adddisk rzb17=rzb17c
# voldg -g sapdg adddisk rzb18=rzb18c
# voldg -g sapdg adddisk rzb19=rzb19c
# voldg -g sapdg adddisk rzb20=rzb20c
# voldg -g sapdg adddisk rzb25=rzb25c
# voldg -g sapdg adddisk rzb26=rzb26c
# voldg -g sapdg adddisk rzb27=rzb27c
# voldg -g sapdg adddisk rzb28=rzb28c
```

3. Since Block Change Logging will be used, create log subdisks to be associated with each Plex (see also LSM documentation about BCL).The size of each log file is 1024 Bytes:

```
# volmake -g sapdg sd rzb17-01 dm_name=rzb17
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb18-01 dm_name=rzb18
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb19-01 dm_name=rzb19
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb20-01 dm_name=rzb20
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb25-01 dm_name=rzb25
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb26-01 dm_name=rzb26
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb27-01 dm_name=rzb27
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb28-01 dm_name=rzb28
  dm_offset=0 len=1k
```

4. Create the LSM data volumes. Get the size of the LSM subdisk with volprint as decribed before and create the disks:

```
# volassist -g sapdg -U fsgen make SAP1_S1  $SIZE
  alloc=0 align=0 rzb17
# volassist -g sapdg -U fsgen make SAP1_S2  $SIZE
  alloc=0 align=0 rzb18
# volassist -g sapdg -U fsgen make SAP2_S3  $SIZE
  alloc=0 align=0 rzb19
# volassist -g sapdg -U fsgen make SAP2_S4  $SIZE
  alloc=0 align=0 rzb20
```

5. Mirror the LSM Disks in the background:

```
# volassist –g sapdg mirror SAP1_S1 rzb25 &
# volassist –g sapdg mirror SAP1_S2 rzb26 &
# volassist –g sapdg mirror SAP2_S3 rzb27 &
# volassist –g sapdg mirror SAP2_S4 rzb28 &
```

6. Associate the Log Plex with each Data Plex. This step can only be done if the mirroring in the background is finished:

```
# volsd –g sapdg aslog SAP1_S1-01 rzb17-01
# volsd –g sapdg aslog SAP1_S1-02 rzb25-01
# volsd –g sapdg aslog SAP1_S2-01 rzb18-01
# volsd –g sapdg aslog SAP1_S2-02 rzb26-01
# volsd –g sapdg aslog SAP2_S3-01 rzb19-01
# volsd –g sapdg aslog SAP2_S3-02 rzb27-01
# volsd –g sapdg aslog SAP2_S4-01 rzb20-01
# volsd –g sapdg aslog SAP2_S4-02 rzb28-01
```

Diskblock 1, 2 and 3 with 9 GB Disks

1. Initialize all diskblocks, which should be used by LSM:

```
# voldisksetup -i rzb17c nlog=1 nconfig=1
# voldisksetup -i rzb18c nlog=1 nconfig=1
# voldisksetup -i rzb19c nlog=1 nconfig=1
# voldisksetup -i rzb20c nlog=1 nconfig=1
# voldisksetup -i rzc17c nlog=1 nconfig=1
# voldisksetup -i rzc18c nlog=1 nconfig=1
# voldisksetup -i rzb25c nlog=1 nconfig=1
# voldisksetup -i rzb26c nlog=1 nconfig=1
# voldisksetup -i rzb27c nlog=1 nconfig=1
# voldisksetup -i rzb28c nlog=1 nconfig=1
# voldisksetup -i rzc25c nlog=1 nconfig=1
# voldisksetup -i rzc26c nlog=1 nconfig=1
```

2. Add all other diskblocks to diskgroup sapdg:

```
# voldg -g sapdg adddisk rzb17=rzb17c
# voldg -g sapdg adddisk rzb18=rzb18c
# voldg -g sapdg adddisk rzb19=rzb19c
# voldg -g sapdg adddisk rzb20=rzb20c
# voldg -g sapdg adddisk rzc17=rzc17c
# voldg -g sapdg adddisk rzc18=rzc18c
# voldg -g sapdg adddisk rzb25=rzb25c
# voldg -g sapdg adddisk rzb26=rzb26c
# voldg -g sapdg adddisk rzb27=rzb27c
# voldg -g sapdg adddisk rzb28=rzb28c
# voldg -g sapdg adddisk rzc25=rzc25c
# voldg -g sapdg adddisk rzc26=rzc26c
```

3. Since Block Change Logging will be used, create log subdisks to be associated with each Plex (see also LSM documentation about BCL). The size of each log file is 1024 Bytes:

```
# volmake -g sapdg sd rzb17-01 dm_name=rzb17
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb18-01 dm_name=rzb18
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb19-01 dm_name=rzb19
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb20-01 dm_name=rzb20
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc17-01 dm_name=rzc17
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc18-01 dm_name=rzc18
  dm_offset=0 len=1k
```

```
# volmake -g sapdg sd rzb25-01 dm_name=rzb25
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb26-01 dm_name=rzb26
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb27-01 dm_name=rzb27
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb28-01 dm_name=rzb28
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc25-01 dm_name=rzc25
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc26-01 dm_name=rzc26
  dm_offset=0 len=1k
```

4. Create the LSM data volumes. Get the size of the LSM subdisk with volprint as decribed before and create the disks:

```
# volassist -g sapdg -U fsgen make SAP1_S1  $SIZE
  alloc=0 align=0 rzb17
# volassist -g sapdg -U fsgen make SAP1_S2  $SIZE
  alloc=0 align=0 rzb18
# volassist -g sapdg -U fsgen make SAP2_S3  $SIZE
  alloc=0 align=0 rzb19
# volassist -g sapdg -U fsgen make SAP2_S4  $SIZE
  alloc=0 align=0 rzb20
# volassist -g sapdg -U fsgen make SAP3_S5  $SIZE
  alloc=0 align=0 rzc17
# volassist -g sapdg -U fsgen make SAP3_S6  $SIZE
  alloc=0 align=0 rzc18
```

5. Mirror the LSM Disks in the background:

```
# volassist –g sapdg mirror SAP1_S1 rzb25 &
# volassist –g sapdg mirror SAP1_S2 rzb26 &
# volassist –g sapdg mirror SAP2_S3 rzb27 &
# volassist –g sapdg mirror SAP2_S4 rzb28 &
# volassist –g sapdg mirror SAP3_S5 rzc25 &
# volassist –g sapdg mirror SAP3_S6 rzc26 &
```

6. Associate the Log Plex with each Data Plex. This step can only be done if the mirroring in the background is finished:

```
# volsd –g sapdg aslog SAP1_S1-01 rzb17-01
# volsd –g sapdg aslog SAP1_S1-02 rzb25-01
# volsd –g sapdg aslog SAP1_S2-01 rzb18-01
# volsd –g sapdg aslog SAP1_S2-02 rzb26-01
# volsd –g sapdg aslog SAP2_S3-01 rzb19-01
# volsd –g sapdg aslog SAP2_S3-02 rzb27-01
# volsd –g sapdg aslog SAP2_S4-01 rzb20-01
# volsd –g sapdg aslog SAP2_S4-02 rzb28-01
# volsd –g sapdg aslog SAP3_S5-01 rzc17-01
# volsd –g sapdg aslog SAP3_S5-02 rzc25-01
# volsd –g sapdg aslog SAP3_S6-01 rzc18-01
# volsd –g sapdg aslog SAP3_S6-02 rzc26-01
```

Diskblock 1, 2, 3 and 4 with 9 GB Disks

1. Initialize all diskblocks, which should be used by LSM:

```
# voldisksetup -i rzb17c nlog=1 nconfig=1
# voldisksetup -i rzb18c nlog=1 nconfig=1
# voldisksetup -i rzb19c nlog=1 nconfig=1
# voldisksetup -i rzb20c nlog=1 nconfig=1
# voldisksetup -i rzc17c nlog=1 nconfig=1
# voldisksetup -i rzc18c nlog=1 nconfig=1
# voldisksetup -i rzc19c nlog=1 nconfig=1
# voldisksetup -i rzc20c nlog=1 nconfig=1
# voldisksetup -i rzb25c nlog=1 nconfig=1
# voldisksetup -i rzb26c nlog=1 nconfig=1
# voldisksetup -i rzb27c nlog=1 nconfig=1
# voldisksetup -i rzb28c nlog=1 nconfig=1
# voldisksetup -i rzc25c nlog=1 nconfig=1
# voldisksetup -i rzc26c nlog=1 nconfig=1
# voldisksetup -i rzc27c nlog=1 nconfig=1
# voldisksetup -i rzc28c nlog=1 nconfig=1
```

2. Add all other diskblocks to diskgroup sapdg:

```
# voldg -g sapdg adddisk rzb17=rzb17c
# voldg -g sapdg adddisk rzb18=rzb18c
# voldg -g sapdg adddisk rzb19=rzb19c
# voldg -g sapdg adddisk rzb20=rzb20c
# voldg -g sapdg adddisk rzc17=rzc17c
# voldg -g sapdg adddisk rzc18=rzc18c
# voldg -g sapdg adddisk rzc19=rzc19c
# voldg -g sapdg adddisk rzc20=rzc20c
# voldg -g sapdg adddisk rzb25=rzb25c
# voldg -g sapdg adddisk rzb26=rzb26c
# voldg -g sapdg adddisk rzb27=rzb27c
# voldg -g sapdg adddisk rzb28=rzb28c
# voldg -g sapdg adddisk rzc25=rzc25c
# voldg -g sapdg adddisk rzc26=rzc26c
# voldg -g sapdg adddisk rzc27=rzc27c
# voldg -g sapdg adddisk rzc28=rzc28c
```

3. Since Block Change Logging will be used, create log subdisks to be associated with each Plex (see also LSM documentation about BCL). The size of each log file is 1024 Bytes:

```
# volmake -g sapdg sd rzb17-01 dm_name=rzb17
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb18-01 dm_name=rzb18
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb19-01 dm_name=rzb19
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb20-01 dm_name=rzb20
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc17-01 dm_name=rzc17
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc18-01 dm_name=rzc18
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc19-01 dm_name=rzc19
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc20-01 dm_name=rzc20
```

```
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb25-01 dm_name=rzb25
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb26-01 dm_name=rzb26
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb27-01 dm_name=rzb27
  dm_offset=0 len=1k
# volmake -g sapdg sd rzb28-01 dm_name=rzb28
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc25-01 dm_name=rzc25
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc26-01 dm_name=rzc26
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc27-01 dm_name=rzc27
  dm_offset=0 len=1k
# volmake -g sapdg sd rzc28-01 dm_name=rzc28
  dm_offset=0 len=1k
```

4. Create the LSM data volumes. Get the size of the LSM subdisk with volprint as decribed before and create the disks:

```
# volassist -g sapdg -U fsgen make SAP1_S1  $SIZE
  alloc=0 align=0 rzb17
# volassist -g sapdg -U fsgen make SAP1_S2  $SIZE
  alloc=0 align=0 rzb18
# volassist -g sapdg -U fsgen make SAP2_S3  $SIZE
  alloc=0 align=0 rzb19
# volassist -g sapdg -U fsgen make SAP2_S4  $SIZE
  alloc=0 align=0 rzb20
# volassist -g sapdg -U fsgen make SAP3_S5  $SIZE
  alloc=0 align=0 rzc17
# volassist -g sapdg -U fsgen make SAP3_S6  $SIZE
  alloc=0 align=0 rzc18
# volassist -g sapdg -U fsgen make SAP4_S7  $SIZE
  alloc=0 align=0 rzc19
# volassist -g sapdg -U fsgen make SAP4_S8  $SIZE
  alloc=0 align=0 rzc20
```

5. Mirror the LSM Disks in the background:

```
# volassist –g sapdg mirror SAP1_S1 rzb25 &
# volassist –g sapdg mirror SAP1_S2 rzb26 &
# volassist –g sapdg mirror SAP2_S3 rzb27 &
# volassist –g sapdg mirror SAP2_S4 rzb28 &
# volassist –g sapdg mirror SAP3_S5 rzc25 &
# volassist –g sapdg mirror SAP3_S6 rzc26 &
# volassist –g sapdg mirror SAP4_S7 rzc27 &
# volassist –g sapdg mirror SAP4_S8 rzc28 &
```

6. Associate the Log Plex with each Data Plex. This step can only be done if the mirroring in the background is finished:

```
# volsd -g sapdg aslog SAP1_S1-01 rzb17-01
# volsd -g sapdg aslog SAP1_S1-02 rzb25-01
# volsd -g sapdg aslog SAP1_S2-01 rzb18-01
# volsd -g sapdg aslog SAP1_S2-02 rzb26-01
# volsd -g sapdg aslog SAP2_S3-01 rzb19-01
# volsd -g sapdg aslog SAP2_S3-02 rzb27-01
# volsd -g sapdg aslog SAP2_S4-01 rzb20-01
# volsd -g sapdg aslog SAP2_S4-02 rzb28-01
# volsd -g sapdg aslog SAP3_S5-01 rzc17-01
# volsd -g sapdg aslog SAP3_S5-02 rzc25-01
# volsd -g sapdg aslog SAP3_S6-01 rzc18-01
# volsd -g sapdg aslog SAP3_S6-02 rzc26-01
# volsd -g sapdg aslog SAP4_S7-01 rzc19-01
# volsd -g sapdg aslog SAP4_S7-02 rzc27-01
# volsd -g sapdg aslog SAP4_S8-01 rzc20-01
# volsd -g sapdg aslog SAP4_S8-02 rzc28-01
```

A log subdisk (also called BCL subdisk) allows room on a disk for Logical Storage Manager support of Block Change Logging. When data is written to a volume, ranges of sector numbers are written to the log subdisk so that a record is kept of recent disk activity. When the system is restarted after a crash, these ranges of block numbers are used to limit the amount of data copying to recover plex consistency for the volume.

One log subdisk may be associated with each plex to greatly speed up recovery of the volume. However, the presence of a BCL subdisk degrades volume performance under normal conditions due to the writing of the log entries. Only one log subdisk may be associated with a plex at a time.

From this point, let the LSM GUI run in the background to monitor what happens. Start it from csh with the commands:

```
# setenv DISPLAY <node name>:0.0
# dxlsm &
```

where <*node name*> is a node with a graphical display.

The following figure is an example for the volume SAP01:

Figure 4-8: Example for Volume SAP01

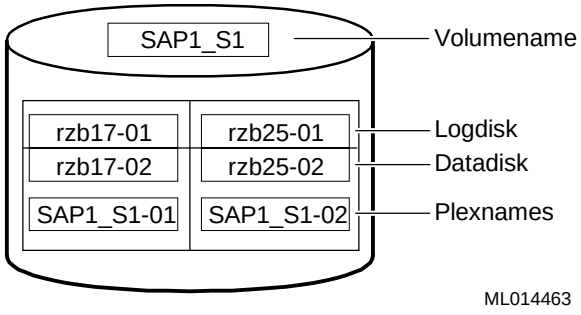

SAP R/3 Installation

This section describes how to install SAP R/3. The current release at installation time was 3.1H.

SAP System ID

The following are the specifics for the SAP R/3 installation:

- SID =TCR
  TCR will have the instance number 00 for the DIGITAL HiTest environment.

  The name of the SAP System is abbreviated to SID (SAP System ID).
  Since an R/3 System consists of exactly one database, the DB name and the SAP System ID can be identified. Contrary to that, one database consist of different DB Instances which are abbreviated to SID.
- Available RAM: Depends what is ordered. Use the exact value during the Installation.
- Installation directory: /sapmnt/TCR/install
- Modify SAPFS.TPL No

  Create the installation directory and set permission to 777 as follows:

1. #**mkdir /sapmnt/TCR/install**
2. #**chmod 777 /sapmnt/TCR/install**

SAP and Oracle Directory Structure

SAP enforces a rigid naming scheme concerning the entry points for the R/3 structures. The following entries must be present:

Table 4-1: Directory Structure

Entries	Description
/usr/sap/trans	Global directory for all SAP systems
/sapmnt/TCR	Systemwide data for one SAP system
/usr/sap/TCR	Instance-specific data
/oracle/stage	Installation and upgrade directory for the database software
/oracle/stage_733	ORACLE stage
/oracle/TCR	Directory for the TCR ORACLE instance
/oracle/TCR/sapdata1	SAP data
/oracle/TCR/sapdata2	SAP data
/oracle/TCR/sapdata3	SAP data
/oracle/TCR/sapdata4	SAP data
/oracle/TCR/sapdata5	SAP data
/oracle/TCR/sapdata6	SAP data
/oracle/TCR/origlogA	ORACLE redo logs 1 and 3
/oracle/TCR/origlogB	ORACLE redo logs 2 and 4
/oracle/TCR/mirrlogA	ORACLE mirror logs

Entries	Description
/oracle/TCR/mirrlogB	ORACLE mirror logs
/oracle/TCR/saparch	ORACLE archives of redologs
/oracle/TCR/sapreorg	Work directory for database administration
/oracle/TCR/sapbackup	Backup information

The directory /sapmnt/TCR and /usr/sap/trans should already exist. They are created during nfssetup. If they were not created, create them now.

Note

All other mountpoints which are on the shared buses are only mounted when the procedure for postinstallation of TruCluster V1.5 in a SAP R/3 environment is performed (detailed later in this chapter).

Map the SAP structures onto the available physical structure using the following commands:

1. The following directories are mandatory for all Database Options:

```
mkdir /oracle
mkdir /oracle/TCR
mkdir /oracle/stage
mkdir /oracle/stage/stage_733
mkdir /sapmnt
mkdir /sapmnt/TCR
mkdir /usr/users
mkdir /usr/sap
mkdir /usr/sap/trans
mkdir /usr/sap/TCR
mkfdmn /dev/vol/sapdg/SAPRest saprest
mkfset saprest oracle
mkfset saprest oraclestage
mkfset saprest sapmnt
mkfset saprest usrsap
mkfset saprest saptrans
mkfset saprest saparch
mkfset saprest sapreorg
mkfset saprest sapbackup

mount saprest#oracle /oracle/TCR
mount saprest#oraclestage /oracle/stage
mount saprest#sapmnt /sapmnt/TCR
mount saprest#usrsap /usr/sap/TCR
mount saprest#saptrans /usr/sap/trans

mkdir /oracle/TCR/saparch
mkdir /oracle/TCR/sapreorg
mkdir /oracle/TCR/sapbackup
mkdir /oracle/TCR/origlogA
mkdir /oracle/TCR/origlogB
mkdir /oracle/TCR/mirrlogA
mkdir /oracle/TCR/mirrlogB
mkdir /oracle/TCR/sapdata1
```

```
mkdir /oracle/TCR/sapdata2
mkdir /oracle/TCR/sapdata3
mkdir /oracle/TCR/sapdata4
mkdir /oracle/TCR/sapdata5
mkdir /oracle/TCR/sapdata6

mount saprest#saparch   /oracle/TCR/saparch
mount saprest#sapbackup /oracle/TCR/sapbackup
mount saprest#sapreorg  /oracle/TCR/sapreorg

mkfdmn /dev/vol/sapdg/redoA redologA
mkfset redologA logA
mkfdmn /dev/vol/sapdg/redoB redologB
mkfset redologB logB
mkfdmn /dev/vol/sapdg/mirrA mirrlogA
mkfset mirrlogA mlogA
mkfdmn /dev/vol/sapdg/mirrB mirrlogB
mkfset mirrlogB mlogB

mount redologA#logA   /oracle/TCR/origlogA
mount redologB#logB   /oracle/TCR/origlogB
mount mirrlogA#mlogA  /oracle/TCR/mirrlogA
mount mirrlogB#mlogB  /oracle/TCR/mirrlogB
```

2. Database disk with diskblock 1 with 4 GB disks (each diskblock of SAPn has 36 GB):

```
mkfdmn    /dev/vol/sapdg/SAP1_S1  sapdata
addvol    /dev/vol/sapdg/SAP1_S2  sapdata
addvol    /dev/vol/sapdg/SAP1_S3  sapdata

mkfset    sapdata     data1
mkfset    sapdata     data2
mkfset    sapdata     data3
mkfset    sapdata     data4
mkfset    sapdata     data5
mkfset    sapdata     data6

mount     sapdata#data1     /oracle/TCR/sapdata1
mount     sapdata#data2     /oracle/TCR/sapdata2
mount     sapdata#data3     /oracle/TCR/sapdata3
mount     sapdata#data4     /oracle/TCR/sapdata4
mount     sapdata#data5     /oracle/TCR/sapdata5
mount     sapdata#data6     /oracle/TCR/sapdata6
```

3. Database disk with diskblock 1 and 2 with 4 GB disks:

```
mkfdmn    /dev/vol/sapdg/SAP1_S1  sapdata
addvol    /dev/vol/sapdg/SAP1_S2  sapdata
addvol    /dev/vol/sapdg/SAP1_S3  sapdata

mkfdmn    /dev/vol/sapdg/SAP2_S4  sapdataA
addvol    /dev/vol/sapdg/SAP2_S5  sapdataA
addvol    /dev/vol/sapdg/SAP2_S6  sapdataA

mkfset    sapdata     data1
```

```
mkfset    sapdata     data2
mkfset    sapdata     data3
mkfset    sapdataA    data4
mkfset    sapdataA    data5
mkfset    sapdataA    data6

mount     sapdata#data1     /oracle/TCR/sapdata1
mount     sapdata#data2     /oracle/TCR/sapdata2
mount     sapdata#data3     /oracle/TCR/sapdata3
mount     sapdataA#data4    /oracle/TCR/sapdata4
mount     sapdataA#data5    /oracle/TCR/sapdata5
mount     sapdataA#data6    /oracle/TCR/sapdata6
```

4. Database disk with diskblock 1 with 9 GB disks:

```
mkfdmn    /dev/vol/sapdg/SAP1_S1  sapdata
addvol    /dev/vol/sapdg/SAP1_S2  sapdata

mkfset    sapdata     data1
mkfset    sapdata     data2
mkfset    sapdata     data3
mkfset    sapdata     data4
mkfset    sapdata     data5
mkfset    sapdata     data6

mount     sapdata#data1     /oracle/TCR/sapdata1
mount     sapdata#data2     /oracle/TCR/sapdata2
mount     sapdata#data3     /oracle/TCR/sapdata3
mount     sapdata#data4     /oracle/TCR/sapdata4
mount     sapdata#data5     /oracle/TCR/sapdata5
mount     sapdata#data6     /oracle/TCR/sapdata6
```

5. Database disk with diskblock 1 and 2 with 9 GB disks:

```
mkfdmn    /dev/vol/sapdg/SAP1_S1  sapdata
addvol    /dev/vol/sapdg/SAP1_S2  sapdata

mkfdmn    /dev/vol/sapdg/SAP2_S3  sapdataA
addvol    /dev/vol/sapdg/SAP2_S4  sapdataA

mkfset    sapdata     data1
mkfset    sapdata     data2
mkfset    sapdata     data3
mkfset    sapdataA    data4
mkfset    sapdataA    data5
mkfset    sapdataA    data6

mount     sapdata#data1     /oracle/TCR/sapdata1
mount     sapdata#data2     /oracle/TCR/sapdata2
mount     sapdata#data3     /oracle/TCR/sapdata3
mount     sapdataA#data4    /oracle/TCR/sapdata4
mount     sapdataA#data5    /oracle/TCR/sapdata5
mount     sapdataA#data6    /oracle/TCR/sapdata6
```

6. Database disk with diskblock 1, 2 and 3 with 9 GB disks:

```
mkfdmn    /dev/vol/sapdg/SAP1_S1  sapdata
addvol    /dev/vol/sapdg/SAP1_S2  sapdata

mkfdmn    /dev/vol/sapdg/SAP2_S3  sapdataA
addvol    /dev/vol/sapdg/SAP2_S4  sapdataA

mkfdmn    /dev/vol/sapdg/SAP3_S5  sapdataB
addvol    /dev/vol/sapdg/SAP3_S6  sapdataB

mkfset    sapdata     data1
mkfset    sapdata     data2
mkfset    sapdataA    data3
mkfset    sapdataA    data4
mkfset    sapdataB    data5
mkfset    sapdataB    data6

mount     sapdata#data1     /oracle/TCR/sapdata1
mount     sapdata#data2     /oracle/TCR/sapdata2
mount     sapdataA#data3    /oracle/TCR/sapdata3
mount     sapdataA#data4    /oracle/TCR/sapdata4
mount     sapdataB#data5    /oracle/TCR/sapdata5
mount     sapdataB#data6    /oracle/TCR/sapdata6
```

7. Database disk with diskblock 1, 2, 3 and 4 with 9 GB disks:

```
mkfdmn    /dev/vol/sapdg/SAP1_S1  sapdata
addvol    /dev/vol/sapdg/SAP1_S2  sapdata

mkfdmn    /dev/vol/sapdg/SAP2_S3  sapdataA
addvol    /dev/vol/sapdg/SAP2_S4  sapdataA

mkfdmn    /dev/vol/sapdg/SAP3_S5  sapdataB
addvol    /dev/vol/sapdg/SAP3_S6  sapdataB

mkfdmn    /dev/vol/sapdg/SAP4_S7  sapdataC
addvol    /dev/vol/sapdg/SAP4_S8  sapdataC

mkfset    sapdata     data1
mkfset    sapdata     data2
mkfset    sapdataA    data3
mkfset    sapdataA    data4
mkfset    sapdataB    data5
mkfset    sapdataC    data6
mount     sapdata#data1     /oracle/TCR/sapdata1
mount     sapdata#data2     /oracle/TCR/sapdata2
mount     sapdataA#data3    /oracle/TCR/sapdata3
mount     sapdataA#data4    /oracle/TCR/sapdata4
mount     sapdataB#data5    /oracle/TCR/sapdata5
mount     sapdataC#data6    /oracle/TCR/sapdata6
```

Starting the SAP R/3 Installation

SAP provides documentation to install their R/3 software. This section highlights the main SAP R/3 installation steps, and is intended to make the reader aware of the choices, and reasons for those choices, made during the SAP R/3 installation on this HiTest System.

Notes (Hinweise)

Every time a SAP installation or upgrade is performed, read the latest notes for information concerning your plans. To ensure that the notes are read, a password (included in the notes) is prompted by the SAP installation or upgrade procedure. The following is the list of notes for the installation of SAP R/3 3.1H with Oracle which are relevant to this HiTest Suite:

- 74278 R/3 Installation on UNIX
- 74279 R/3 Installation on UNIX - OS Dependencies
- 74275 R/3 Installation on UNIX - ORACLE Database

Check List

The document, *Check list - Installation Requirements: ORACLE*, is used to make sure that the system meets SAP requirements. This document is provided by SAP as part of the installation kit.

OS Dependencies

Complete the check list, then continue by using the *SAP R/3 Installation Guide Operating Dependencies* manual. The manual covers the following topics:

1. General Notes on NIS
2. Users and Groups
3. Services
4. Mounting a CD-ROM
5. Checking and Modifying the DIGITAL UNIX Kernel
6. File Systems/Raw Devices/Swap Space
7. Mounting Directories via NFS
8. Creating Groups and Users
9. SAP Tool Kinst
10. Troubleshooting

The following sections cover steps 4 and 5.

Mounting a CD-ROM

1. Create a mount point directory if it does not already exist:
   ```
   # mkdir /sapcd
   ```
2. Mount your CDs with the command:
   ```
   # mount -t cdfs -dr /dev/cd /sapcd
   ```

Checking and Modifying the DIGITAL UNIX Kernel

Since DIGITAL UNIX Version 3.0, a dynamic approach exists to change kernel parameters. Most of the system parameters can be specified in a file called `/etc/sysconfigtab`. Any modification in this file will be applied at the next system boot. A new kernel generation is not required.

The values for the system configuration file `/etc/sysconfigtab` are listed in *SAP R/3 Installation Guide Operating Dependencies* in the particular DIGITAL UNIX Chapter. Please apply these values as demanded.

Reboot the system:

```
# init 0
>>> boot
```

General Installation Preparations

Refer to the *R/3 Installation on UNIX-ORACLE Database Guide*.

Note

As described later in the postinstallation for TruCluster you need two SAP licenses. In case if the Central Instance fails and the ASE switches to the other System, there has to be a valid license.

Install a Dialog Instance

Install a dialog instance on the second system (tcr002) as described in *SAP R/3 Installation on UNIX-ORACLE Database* guide.

Create the following directories:

```
#mkdir /sapmnt/TCR/install
#mkdir /usr/users
```

Useful add ons:

1. Please do the following on both systems:
   Add this value to /etc/sysconfigtab in section vm:
   ubc-nfsloopback=1
2. Edit the following file and define the disks which are on the shared bus. This is done to prevent the advfsd daemon to poll thru this disks and create disk errors in the uerf file.
   ```
   # vi /var/opt/advfsd/disks.ignore
   rz18c
   rz19c
   ...
   rznnc
   ```
3. Edit the following file and define the disks which are on the shared bus. This is done to prevent the snmpd daemon to poll thru this disks and create Disk errors in the uerf file.
   ```
   # vi /etc/snmpd.conf
   hrDevice_ignore       /dev/rz18
   hrDevice_ignore       /dev/rz19
   ...
   hrDevice_ignore       /dev/rznm
   ```

Take a Full Backup

Use the DIGITAL UNIX command vdump to make backups of the disks. The backup will run for approximately 45 minutes.

When it is completed, do the following:
As root: # **<Ctrl>D**
As tcradm: **startsap**

Postinstallation for TruCluster V1.5 in a SAP R/3 Environment

To install the TruCluster V1.5 in a SAP R/3 environment, refer to the actual documentation at `http://www.fra.dec.com/SAP-Eng/available/cluster.html`

Note

These postinstallation adjustments are mandatory for the R/3 installation.

Differences to the Original Documentation

During the postinstallation tasks for TruCluster V1.5 in a SAP R/3 Environment on this DIGITAL HiTest System, parts of the procedure were performed differently than documented in the file `install.ps` (part of DECSAFE_V2_2.tar). In some cases, useful information was missing. The differences are described in Appendix A.

5
Tests and Results

The DIGITAL HiTest program tests for several types of problems that affect the system. The HiTest program works together with other organizations to obtain and share test information for other categories.

This chapter describes the overview of test results, how the tests were set up, and where the data and programs were placed.

Also covered in this chapter is the test environment, tools used for testing, test configuration, and the test process.

Overview of Results

Interoperability testing was performed successfully on the SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 HiTest Suite. Tests were performed to ensure the HiTest Suite met installability, interoperability, and availability criteria.

Test Environment

Figure 5-1 shows the SAP Oracle TruCluster ASE test environment.

Figure 5-1: Test Environment

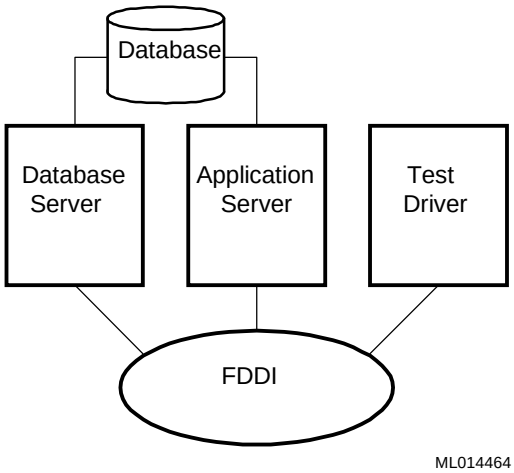

Test Tools

The following tools were used for interoperability testing:

- *ftp* and *tar* to move the client data file to the driver and expand it
- *benchinst* to create the structure of the simulation directory tree on the driver
- *cleandb* and *impsrc* to import the client data into the database on the HiTest system and change some source code
- *mmpv* (period shifter) to bring the booking period of the SAP transactions into the current month (must be rerun at import and at the start of every month)

The following test tools were used to create the load and measure the behavior of the system:

- *mkapl* to define the load parameters (number of users, number of loops)
- *mksim* to create all scripts and additional directories for a load
- *benchrun* to start one load
- *vmubc* to watch the overall CPU and memory behavior of the HiTest system
- *iostat* to watch the disk behavior
- R/3 transaction ST02 to watch the memory behavior of R/3
- R/3 transaction SM50 to watch the behavior of the various R/3 processes

Test Configuration

To stress test the HiTest configuration and to prove its viability, a standardized SAP benchmark method is being used. To put a meaningful load onto the HiTest System, the following conditions must be met:

- A third system (called driver) is connected to the DB and Application Server through a FDDI connection that is able to connect to the virtual hostnames.
- The driver simulates the user load with the help of SAP-written scripts and executables. This benchmark environment is available for all customers if a person from the specific Competence Center is available and runs the tests.

Note

Do not use this benchmark software in Production Environments. You will get no support from SAP.

- Get the benchmark software from SAP network together with the newest VERY_IMPORTANT.doc. All Competence Centers know the location of this Kit. Together with the kit you get three descriptions:
  1. Installation of the SAP R/3 benchmark
  2. Hints for the SAP R/3 benchmark
  3. Executing the SAP R/3 benchmark
- Create a user on the driver who will drive the benchmark and modify the environment. Check and modify the network so that all systems can connect to each other.
- Unpack the benchmark tar file.

System Limits and Characterization Data

It was not in the scope of our testing to specifically determine system limitations or provide comprehensive performance characterization. The focus was a functional testing in a typical client situation.

Sizing information can be found:
`http://www.fra.dec.com/SAP-Cc/Intranet/sizing/sizingliste.html`

Test Process and Results

The following information describes the test results:

1. A short 10-User benchmark was run as a sanity check. No failures should appear. Please check Syslog with Transaction SM21 for all Instances.
2. Run a 100 SD User Benchmark to the normal distributed Systems. Fifty to the DB Server with the Central System and 50 to the Application Server.
3. Relocate the service Application Server from vtcr002 to vtcr001 to check if both Instances can run on the DB Server in case the Application Server dies.
4. Run 50 Users to the DB Server with the Central System and 50 to the Application Server which was failed over to the DB Server System.
5. Relocate the Service Application Server vtcr002 back to the Application Server System and rerun 100 SD User (50/50).
6. Relocate the Service DB Server from vtcr001 to vtcr002 to check if both instances can run on the Application Server in the case the DB Server dies.
7. Run 50 Users to the Application Server and 50 to the DB Server which was failed over to the Application Server System.
8. Relocate the service DB Server vtcr001 to the Application Server system and rerun 100 SD User (50/50).
9. Switch off one HSZ70 to simulate a powerfail and check the mirrorsets.
   The System will wait until the HSZ70 is back again.

   This is a feature, the system manager does not see that the HSZ is gone. If the second HSZ70 faults, the System will be unavailable and a restore has to be done because the data on the HSZ70s is gone.

   If you run 100 SD User to that scenario, the user will be stuck until the HSZ70 is back online.
10. Remove a disk out of a working stripeset to simulate a power fail of a disk. This is done during a 100 SD benchmark.

    LSM finds a faulty stripeset and marks the faulty disk “disabled.”
11. Move the faulty plex out of the LSM-Volume and recreate it. After that, bring the plex back to the volume and check the shadowcopy.
12. Rerun 100 SD User to check everything works fine.

In the failover situations tested, there were no problems with our SD-benchmark. The system reacts as expected and can handle the failover situations.

6
Problems and Solutions

This chapter describes problems encountered during the testing. Where appropriate, a solution for each problem is given which provides a fix or workaround. An impact statement is also provided.

Foundation Hardware

The following problems were identified:

KZPSA Misc Errors

Problem The system crashes with KZPSA misc errors.

Impact System crashes.

Solution A single KZPSA was not used during one test (the one for the tapedrive). This causes a crash. Make sure that there are devices connected to the KZPSA. If you do not want to use a KZPSA, you have to remove it from the system. A single Terminator does not fix that problem.

Foundation Operating System

Error in UERF

Problem Disk errors in the uerf (UNIX Error Formatter).

Impact /usr partition is getting full.

Solution These errors are from the dtadvfs daemon, which is frequently (5 min.) checking the disks. Also the disks on the shared buses are checked and found reserved. This will flood the uerf with harmless disk errors. To disable the disk spanning it is not enough to define these disksin the /usr/var/opt/advfsd/disks.ignore file as decribed in the advfs_support notesfile. The only solution is to stop the advfsd with the following command:

```
/sbin/rc3.d/S53advfsd stop
```

If you want to use the dtadvfs again you can restart it:

```
/sbin/rc3.d/S53advfsd start
```

To prevent the system starting the daemon during system startup, remove the link to the original file or move it to a different name.

Directory /usr/users not found

Problem During the installation of the Central Instance, the R3INST will fail with the following error during the phase “Creating UNIX Users”:

Directory /usr/users not found.

Impact SAP R/3 installation is interrupted.

Solution Create the directory and use it as the parent directory for R/3 Administrator.

Cannot Change /sapcd/DEC/INSTALL

Problem During R/3 Installation in the phase "copy RDBMS SW" the procedure prints:
`Cannot change to /sapcd/DEC/INSTALL`

Impact SAP R/3 installation is interrupted.

Solution Check log file R3INST.log, exit R3INST and restart again.
The fault came from the function “change cd-mountpoint with R3INST-assistence.”

Short printout of R3INST.log:

```
1997-May-28 10:11:11 I exit_on_label 4 Continuing with
incorrect CD-LABEL
1997-May-28 10:11:11 E ik011_cd_to_exe 3 Cannot change to
/sapcd/DEC/INSTALL
1997-May-28 10:11:14 E ik011_check_instdir 3 Copying templates
from CD failed !
1997-May-28 10:11:15 E ik011_adapt_user 3 Installation
templates from Kernel CD are missing !
```

end unsuccessfully

Problem During R/3 Installation in the phase "DB Load" the procedure prints:
`Out of 1 started processes did 1 end unsuccessfully`

Impact SAP R/3 installation is interrupted.

Solution You forgot to install Oracle NETV2.
See *SAP R/3 Installation on UNIX - ORACLE Database guide*.

entries in TCPDB

Problem If you check the SAP R/3 installation and you take a look at the syslog with transaction SM21, you will see the following error:
In table TCPDB are 0 entries.

Impact SAP R/3 installation has an error message.

Solution Apply SAP note 15023.

APPL-SERVER not found

Problem During the installation of the Dialog Instance on the Application Server the R3INST will fail with the following error:

Expected line [APPL-SERVER] not found.

Impact SAP R/3 installation is interrupted.

Solution Do not use the installation directory, which you have used to install Central Instance.

Foundation Software

The following problems were identified:

Oracle Installation

Not able to run root.sh

Problem User ora<sid> is not able to run root.sh as user root.

Impact Oracle installation is interrupted.

Solution Add user ora<sid> to group system in /etc/group and restart.

Failure during check of directories

Problem During SAP R/3 installation in phase “R/3 Installation on UNIX Oracle Database” the R3INST will fail during check of directories.

Impact Oracle installation is interrupted.

Solution Create directory /oracle/stage/stage_733 and do not use the default value oracle/stage/stage_723.

local bin Directory

Problem During run of root.sh the system is asking for the path of the “local bin directory.” The default is /usr/local/bin but it has to exist.

Impact Oracle installation is interrupted.

Solution If you answer with default, you have to create the directory bin manually.

```
# mkdir /usr/local/bin
```

AppSet Software

No problems were encountered.

Testphase

No connect to the database

Problem During the installation of the benchmark environment you cannot connect the database from the driver system. You can test the connection with:

R3trans –d

You will see the following error in the local directory in file trans.log, which is created during connection:

2EETW169 no connect possible: “DBMS = ORACLE --- ORACLE_SID = ‘SDR’ ”

Impact SD benchmark is not running.

Solution The variable dbs_ora_tnsname is not set. Do that in the $HOME/.cshrc file of the benchmark user.

Setenv dbs_ora_tnsname = SDR

Unable to Obtain Requested Swapspace

Problem During start up of R/3 the following error appears on the console terminal:

"`unable to obtain requested swapspace`"

This error can occur after a memory upgrade.

Impact Performance problems.

Solution Add more swap space; add a disk.

Benchinst fail

Problem The benchinst during installation of the benchmark environment will fail when compiling the file benchrun.c. This is an error in the SAP CAR tool, which is nearly similar to the UNIX tar.

Impact SD benchmark is not running.

Solution Cd /$SIMDIR/src
vi benchrun.c
goto line 374
remove the * at the end of the line
save the file
restart the benchinst

Perl is needed

Problem Since R/3 3.1H and the equivalent benchmark software 3.1H, the command language Perl is used to run the benchmark software either on UNIX or on NT.

Perl is not available on standard UNIX and NT systems.

Impact SD benchmark is not running.

Solution Get Perl (Perl15003setld.tar) from the following web site:

`ftp://ftp.digital.com/pub/Digital`

If you have untared and installed it with setld, you can check the version.

perl –v

This is perl, version 5.003 with EMBED
build under dec_osf at Sep 20 1996 13:47:02
+ suidperl security patch

Copyright…

Cleandb fail

Problem The cleandb during installation of the benchmark environment will fail with the following error:

/$SIMDIR/mandt/mandt_exp not found.

Impact SD benchmark is not running.

Solution Go to this directory and create a softlink from mandt_exp.31H to mandt_exp and restart cleandb.

ln –s mandt_exp.31H mandt_exp

Unable to extend table SAPR3.MDKP

Problem Error during cleandb in the $SIMDIR/tmp/import900.log

Unable to extend table SAPR3.MDKP by 1280 in tablespace PSAPBTABD.

Impact SD benchmark fails.

Solution Add new datafile for tablespace PSAPBTABD with sapdba.

Impsrc fail

Problem The impsrc fail during installation of the benchmark environment with the following error:

/$SIMDIR/mandt/mandt_exp_source not found

Impact SD benchmark is not running.

Solution Go to this directory and create a softlink from mandt_exp_source.31H to mandt_exp_source and restart impsrc.

ln –s mandt_exp_source.31H mandt_exp_source

Mmpv fail

Problem The mmpv during installation of the benchmark software will fail with the following error in a SAP icon.

SAPGUI Icon: This failure should be handled by the caller of DPTM-layer.

Impact SD benchmark is not running.

Solution The variable SAPRELEASE is not set to 31H.

Set this variable in $HOME/.cshrc of the benchmark user and restart mmpv.

Mess-tools not started

Problem The benchrun during the run of benchmark will not start the mess-tools.

The name ‘mess’ is coming from the German word ‘Messung’, that means measurement.

Impact SD benchmark is not running.

Solution Do *not* use the hosttype PR in the apl file of the local simulation directory.

This is not supported.

7
Detailed Hardware Configuration

This chapter describes the minimum and maximum hardware configuration for the SAP Oracle TruCluster ASE DIGITAL UNIX AlphaServer 4100 HiTest Suite by providing the following:

- System Diagram
- HiTest System Slot Configurations
- Input/Output Slot Usage

System Diagram

Figure 7-1 shows a diagram of the maximum configuration of this HiTest Suite and Table 7-1 lists the major cables.

Figure 7-1: System Diagram

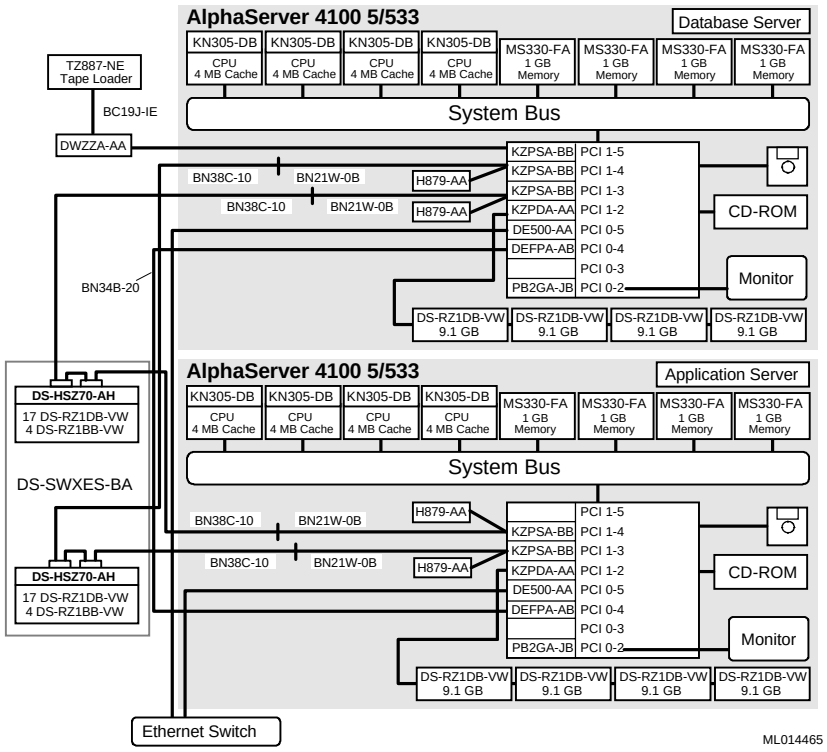

The root disk must be in the first slot of the BA356 shelf. The mirror disk is in the last slot. These two disks should be seen as rz8 and rz14 during the UNIX Installation.

Table 7-1: Configuration Cabling

Part Number	Qty	Description	From	To
BN21K-05	1	SCSI bus	KZPSA-BB (DB Server)	DWZZA-AA
BN21W-0B	4	SCSI bus	KZPSA-BB	BN38C-10
BN38C-10	4	SCSI bus	BN21W-0B	ESA 10000
BN34B-20	2	SCSI bus	DEFPA	DEFPA
BC19J-IE	1	SCSI bus	DWZZA-AA	TZ887-NE

HiTest System Slot Configuration

Figure 7-2 shows the HiTest System Slot Usage and describes the minimum and maximum hardware configurations used in this HiTest Template.

Figure 7-2: HiTest System Slot Usage

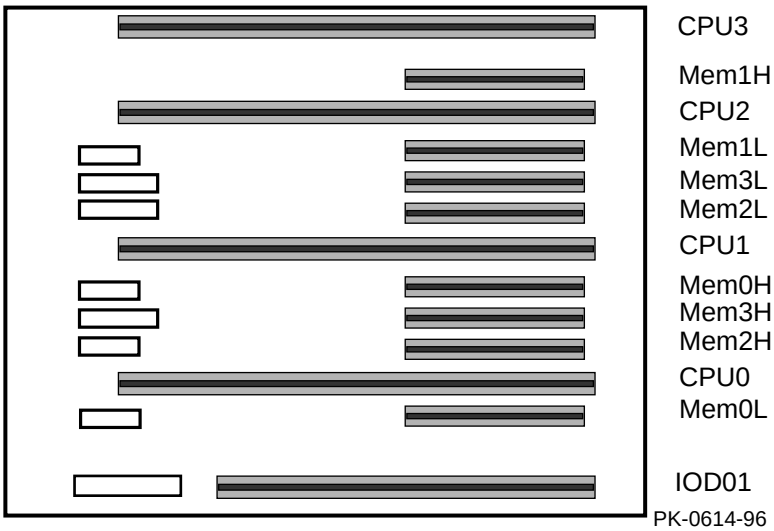

Table 7-2: System Slot Usage (Minimum and Maximum Configurations)

Slot	Minimum Options	Maximum Options	Description
CPU3	open	KN305-DB	533 MHz CPU 4 MB cache
Mem1H	open	MS330-FA	Memory pair 1 (2 of 2)
CPU2	open	KN305-DB	533 MHz CPU 4 MB cache
Mem1L	open	MS330-FA	Memory pair 1 (1 of 2)
Mem3L	open	open	
Mem2L	open	MS330-FA	Memory pair 2 (1 of 2)
CPU1	open	KN305-DB	533 MHz CPU 4 MB cache
Mem0H	MS320-FA	MS330-FA	Memory pair 0 (2 of 2)
Mem3H	open	open	
Mem2H	open	MS330-FA	Memory pair 2 (2 of 2)
CPU0	KN304-DB	KN305-DB	533 MHz CPU 4 MB cache
Mem0L	MS320-FA	MS330-FA	Memory pair 0 (1 of 2)
IOD01	Bridge	Bridge	System bus to PCI bus bridge module

Input/Output Slot Usage

Figure 7-3 and Table 7-3 show the input/output (I/O) slot usage for the minimum and maximum configurations of this HiTest Template.

Figure 7-3: I/O Slot Usage

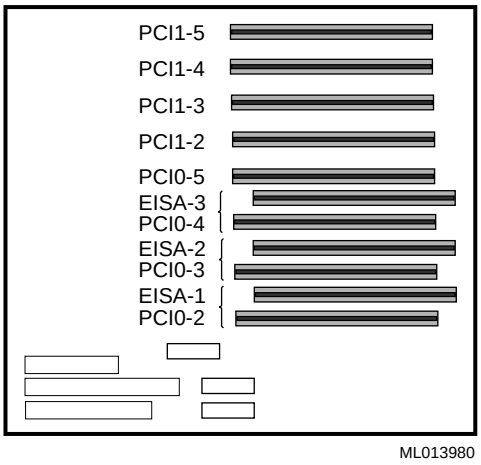

Table 7-3: I/O Slot Usage (Minimum and Maximum Configurations)

Slots	Minimum Configuration Options	Maximum Configuration Options	Description
PCI1-5	KZPSA-BB	KZPSA-BB	FWD SCSI controller
PCI1-4	KZPSA-BB	KZPSA-BB	FWD SCSI controller
PCI1-3	KZPSA-BB	KZPSA-BB	FWD SCSI controller
PCI1-2	KZPDA-AA	KZPDA-AA	FWSE SCSI controller
PCI0-5	DE500-AA	DE500-AA	Ethernet controller
EISA-3/ PCI0-4	DEFPA-BA	DEFPA-BA	FDDI controller
EISA-2/ PCI0-3			Not used
EISA-1/ PCI0-2	PB2GA-JB	PB2GA-JB	TRIO64 Graphics

A

Postinstallation for TruCluster V1.5 in a SAP R/3 Environment

This appendix contains the `rc_serv.conf` file which was modified for SAP R/3 environments on both systems equally. The position MOUNT_FS_POINTS can slightly differ.

```
#########################################################################
#########
#
# PROGRAM:          /usr/local/ase/rc_serv.conf
# LAST.MOD:             Nov/25/1997
# PURPOSE:          Site specific definitions used by all ASE scripts
# ENVIRONMENT:      DIGITAL UNIX V3.2, ASE V1.3 or higher, SAP R/3 30E
# AUTHOR:                 DEC/SAP CC, April 1997
# SUBSTITUTIONS:    (required) : pick your database - Oracle,
Informix, Adabas
#                   and change the following parameters
#                   for your environment.
#
#########################################################################
#######
#########################################################################
#########
#
# the following substitutions are site specific - change where
necessary
#
#########################################################################
########
#########################################################################
########
#
# Global paramenters for the failover database environment
# Insert the values for your environment.
# If the SAP frontends connect this database R/3 instances via a
second network,
# insert the second alias for the specified instance to the parameter
DB1_ALIAS2_HOST
# and the netmask to NETMASK_NETADAPTER_2. if not, don't insert any
value.
#
#########################################################################
#########
ASEDIR="/usr/local/ase"          # directory where the scripts reside
SAPSYSTEMNAME=TCR         # SID of the database
DB1_SERV=vtcr001          # name of the disk-database service in the
asemgr
DB1_ALIAS_HOST=vtcr001           # virtuel name of the db and central
instance for net-adapter 1
DB1_INSTANCE=DVEBMGS00           # instance name of the central
```

```
instance
DB1_ALIAS2_HOST=vtcr001e  # virtual name of the central instance for
net-adapter 2
NETMASK_NETADAPTER_2=255.255.255.0             # netmask for the
second network adapter
NETADAPTER_2=tu0               # name of the second network adapter
(ex. tu0)
#########################################################################
########
DB1_INSTANCE_ID=`print - ${DB1_INSTANCE} | awk '{n=length($0); printf
"%s", substr($0,n-1)}'`
DB1_STARTSAP="startsap_vtcr001_00"           # start script for the
central instance
DB1_STOPSAP="stopsap_vtcr001_00"       # stop script for the central
instance
TIMEOUT=200                              # timeframe in sec. to stop
the R/3 system + database
DB1_START_LOG="${ASEDIR}/log.start_db1_service"     # log file for
the start action
DB1_STOP_LOG="${ASEDIR}/log.stop_db1_service" # log file for the stop
action
#########################################################################
#########
#
# Global paramenters for the application server environment
# Insert the values for your environment.
# If the SAP frontends connect this R/3 application instances via a
second network
# insert the second alias for the specified instance at the parameter
AS1_ALIAS2_HOST
# if not, don't insert any value.
#
#########################################################################
########
AS1_SERV=vtcr002         # name of the user-defined-application
service in the asemgr
AS1_ALIAS_HOST=vtcr002          # virtuel name of the appl. server
instance for net-adapter 1
AS1_ALIAS2_HOST=vtcr002e      # virtual name of the appl. server
instance for net-adapter 2
AS1_INSTANCE=D01         # instance nam eof the application server
instance
AS1_SCRIPT=as_serv.ksh        # script name to start the application
service
#########################################################################
########
AS1_INSTANCE_ID=`print - ${AS1_INSTANCE} | awk '{n=length($0); printf
"%s", substr($0,n-1)}'`
AS1_STARTSAP="startsap_vtcr002_01"     # start script for the appl.
server instance
AS1_STOPSAP="stopsap_vtcr002_01" # stop script for the appl. server
instance
AS1_START_LOG="${ASEDIR}/log.start_as1_service"           # log file
for the start action
AS1_STOP_LOG="${ASEDIR}/log.stop_as1_service"       # log file for
the stop action
#########################################################################
########
# uncomment following lines if you have external instances
# Here are the global paramenters for one external application server
# If you have more than one external application server this part has
# to be multiplied.
# Insert the values for your environment.
# if not, don't insert any value.
#
#########################################################################
#########
#
```

```
#APP_INSTANCE=D02                # Instance name of the external
application
#ASEDIR="/usr/local/ase"         # directory where the scripts reside
#SAPSYSTEMNAME=ASE               # SID of the database
#START_SAP="startsap"            # startsap command for the external
application
#STOP_SAP="stopsap"              # stopsap command for the external
application
#SAP_RESTART_LOG="${ASEDIR}/log.app_restart"     # log file for the
restart action
#
#APP_INSTANCE_ID=`print - ${APP_INSTANCE} | awk '{n=length($0);
printf "%s", substr($0,n-1)}'`
#
#########################################################################
#########
#
# uncomment following lines if you have external instances and modify
# the server name and the external instance names in the
EXTERNAL_INSTANCES
# parameter:
# e.g. alpha10 D02
#      alpha11 D03 etc.
#
#########################################################################
#########
#
#RESTART_SCRIPT=${ASEDIR}/app_restart.ksh
#EXTERNAL_INSTANCES=" \
#alpha10 D02 \
#alpha11 D03 \
#"
#
#########################################################################
#########
#
# Global paramenters for the second database running on the failover
server
# Insert the values for your environment, if you use a second
database system.
#
#########################################################################
#########
DB2_HOSTNAME=                          # "hostname -s" of second
database server
DB2_STARTSAP=                          # start script for the second
instance
DB2_STOPSAP=                       # stop script for the second
instance
DB2_INSTANCE=                          # instance name of the central
instance
DB2_SAPSYSTEMNAME=                 # SID of the second database
DB2_START_LOG="${ASEDIR}/log.start_db2_service"       # log file for
the start action
DB2_STOP_LOG="${ASEDIR}/log.stop_db2_service"     # log file for the
stop action
TRUL='tr "[:upper:]" "[:lower:]"'
DB2USER="`echo $DB2_SAPSYSTEMNAME | $TRUL`adm"
#########################################################################
#########
TRUL='tr "[:upper:]" "[:lower:]"'
ADMUSER="`echo $SAPSYSTEMNAME | $TRUL`adm"
ASE_ADMIN="root"
ASETMP="/var/ase/tmp"
SCRIPT="$0"
HOSTNAME=`/usr/bin/hostname -s`
NFS_LOCK_INFO="/sapmnt/${SAPSYSTEMNAME}/.ase/nfs_lock"
NFS_EXPORT="/usr/sap/trans /sapmnt/${SAPSYSTEMNAME}"
```

```
NFS_SAPMNT="${DB1_ALIAS_HOST}:/sapmnt/${SAPSYSTEMNAME}"
NFS_TRANS="${DB1_ALIAS_HOST}:/usr/sap/trans"
NFS_OPTS="-o bg,soft,timeo=10,retrans=10,retry=10"
#########################################################################
#########
#
# uncomment following lines if database is ORACLE
# and modify the parameter, domain- and directory structure for your
# environment
#
#########################################################################
#########
DB_TYPE=ORACLE
ORAUSER="ora`echo $SAPSYSTEMNAME | $TRUL`"
SGADEF=/oracle/${SAPSYSTEMNAME}/dbs/sgadef${SAPSYSTEMNAME}
MOUNT_FS_POINTS=" \
saprest#oracle /oracle/TCR \
saprest#oraclestage /oracle/stage \
saprest#sapmnt /sapmnt/TCR \
saprest#usrsap /usr/sap/TCR \
saprest#saptrans /usr/sap/trans \
saprest#saparch /oracle/TCR/saparch \
saprest#sapbackup /oracle/TCR/sapbackup \
saprest#sapreorg /oracle/TCR/sapreorg \
redologA#logA /oracle/TCR/origlogA \
redologB#logB /oracle/TCR/origlogB \
mirrlogA#mlogA /oracle/TCR/mirrlogA \
mirrlogB#mlogB /oracle/TCR/mirrlogB \
sapdata#data1 /oracle/TCR/sapdata1 \
sapdata#data2 /oracle/TCR/sapdata2 \
sapdata#data3 /oracle/TCR/sapdata3 \
sapdataA#data4 /oracle/TCR/sapdata4 \
sapdataA#data5 /oracle/TCR/sapdata5 \
sapdataA#data6 /oracle/TCR/sapdata6 \
"
#########################################################################
#########
#
# uncomment following lines if database is ADABAS
# and modify the domain and directory structure for your
# environment
#
#########################################################################
########
#
#DB_TYPE=ADABAS
#MOUNT_FS_POINTS=" \
#sap#adabas  /adabas/${SAPSYSTEMNAME} \
#sap#sapmnt  /sapmnt/${SAPSYSTEMNAME} \
#sap#usr_sap /usr/sap/${SAPSYSTEMNAME} \
#sap#trans      /usr/sap/trans \
#"
#########################################################################
#########
#
# uncomment following lines if database is INFORMIX
# and modify the domain and directory structure for your
# environment
#
#########################################################################
#########
#
#DB_TYPE=INFORMIX
#MOUNT_FS_POINTS=" \
#sap#informix       /adabas/${SAPSYSTEMNAME} \
#sap#sapmnt  /sapmnt/${SAPSYSTEMNAME} \
#sap#usr_sap /usr/sap/${SAPSYSTEMNAME} \
#sap#trans   /usr/sap/trans \
```

```
#"
#
#########################################################################
#########
#
# SUBROUTINE DEFINITION
# do not modify this section.
#
#########################################################################
#########
MOUNT_POINTS=""
set $MOUNT_FS_POINTS
while [ "X$1" != "X" ]
do
  MOUNT_POINTS="${MOUNT_POINTS} $2"
  shift;shift
done
UMOUNT_POINTS=`for i in ${MOUNT_POINTS};do print - "$i";done|sort -r`
#########################################################################
#########
MAIL_ADMIN () {
       if [ -n "${ASE_ADMIN}" ]; then
       mailx -s "Critical ASE/R3-Error:" ${ASE_ADMIN} < $1
       fi
}

kill_proc_on () {
  # stop processes with open files.
  ${ASEDIR}/bin/fuser -ck $*
}

proc_timeout () {
  # start this function in the background before the proc you want to
timeout
  # eg. # proc_timeout sleep 20 5 & sleep 300; echo "exit status: $?"
  proc_string=$1; wait_seconds=$2; interval=$3; i=0

  # wait interval seconds to let the process start
  sleep $interval

  while [ $i -lt $wait_seconds ]
  do
    is_running=`ps axww | grep -w "$proc_string" | grep -v grep | wc
-l`
    if [ $is_running -eq 0 ]
    then
      break
    else
      i=`expr $i + $interval`; sleep $interval
    fi
  done
  if [ $i -ge $wait_seconds ]
  then
    pids=`ps axww | grep -w "$proc_string" | grep -v grep | awk
'{print $1}'`
    if [ "$pids" != "" ]; then
      echo "\ntimeout reached, stopping process: $proc_string
processid: $pids"
      kill -KILL $pids
    fi
   fi
}

nfs_wait_loop () {
# function, to wait for a clean umount of mount points
# parametrs proc_string, interval
```

```
  proc_string=$1; interval=$2; is_running=0
  while [ $is_running -eq 0 ]
  do
    mount | grep -w "$proc_string" | grep -v grep
    is_running=$?
    if [ $is_running -eq 1 ]
    then
      break
    else
      /sbin/umount -f $proc_string
      kill_proc_on  $proc_string
      sleep $interval
      is_running=0
    fi
  done
}

db_action () {
# function for specific database actions (ORACLE, INFORMIX, ADABAS)
# during start and stop actions
# parameters $1=DB_TYPE, $2=start,stop
       case $1 in
       ORACLE)
              case $2 in
              start)
                  TNS_ADMIN=`su - ${ORAUSER} -c 'echo ${TNS_ADMIN}'`
                  print - "ASE-Info: starting the oracle listener on
{HOSTNAME}"
                     su - ${ORAUSER} -c "lsnrctl status" > /dev/null
                     if [ $? -eq 0 ]
                     then
                         print - "ASE-Info: listener is already
running"
                     else
                         su - ${ORAUSER} -c "umask 0;lsnrctl start" >
/dev/null
                         case $returncode in
                            0) print - "ASE-Success: startup of oracle
listener k"
                            ;;
                            *) print - "ASE-Error: startup of oracle
listener faild"
                           /var/ase/sbin/nfs_ifconfig ${HOSTNAME} stop
${DB1_ALIAS_HOST}
                               exit 1
                            ;;
                         esac
                     fi
                  if [ -f ${SGADEF}.dbf -o -f ${SGADEF}.ora ]
                  then
                  print - "ASE-Info: ${SGADEF} exists - start DB with
startup recover"
                  if [ -x /oracle/${SAPSYSTEMNAME}/bin/svrmgrl ]
                  then
                     print - "connect internal;\nstartup;\nexit;\n"| \
                         su - ${ADMUSER} -c "svrmgrl"
                     returncode=$?
                  Else
                    print - "connect internal;\nstartup
recover;\nexit;\n"| \
                     su - ${ADMUSER} -c "sqldba lmode=y"
                     returncode=$?
                  Fi
                  case $returncode in
                  0) print - "ASE-Success: startup recover succeeded"
                  ;;
                  *) print - "ASE-Error: startup recover failed"
                  MAIL_ADMIN ${DB1_START_LOG}
```

```
                /var/ase/sbin/nfs_ifconfig ${HOSTNAME} stop
${DB1_ALIAS_HOST}
                 exit 1
                 ;;
                 esac
                   fi
             ;;
             stop)
                /usr/bin/ipcs -m | grep ${ORAUSER} | grep -v grep
             if [ $? -ne 0 ]
             then
               print - "ASE-Info: no more shared memory for
${ORAUSER} present"
             else
                  /usr/bin/ipcs -m | grep ${ORAUSER} | grep -v grep
|\
                  while read m id rest
                  do
                   echo $id
                      /usr/bin/ipcrm -m $id
                      case $? In
                      0)   print - "ASE-Warning: not deleted, shared
memory still set ";;
                      *)   print - "ASE-Success: shared memory
succesfully deleted";;
                      esac
                    done
             fi
             ;;
             esac
      ;;
      ADABAS)
             case $2 in
                start)
                i=0
                while [ $i -ne $2 ]
                do
                      su - ${ADMUSER} -c "x_server stop"   >
/dev/null
                      sleep $2
                      su - ${ADMUSER} -c "x_server start"  >
/dev/null
                      if [ $? -eq 0 ]
                         then
                                print - "ASE-Action: x_server
successfully started"
                                break
                         else
                                print - "ASE-Warning: x_server start
failed, try it again"
                                i=`expr $i + 1`
                      fi
                done
             ;;
             stop)
             ;;
             esac
      ;;
      INFORMIX)
      ;;
esac
}
restart_of_appl(){
      if [ ${AS1_SERV} != " " ]
      then
             print - " "
             print - "ASE-Action: Execute sleep 50;/usr/sbin/asemgr -
r ${AS1_SERV}"
```

```
            ${ASEDIR}/bin/execwrap "sleep 50;/usr/sbin/asemgr -r
${AS1_SERV}"
            print - "ASE-Info: done"
      fi
}

restart_of_external_appl(){
    set $EXTERNAL_INSTANCES
    while [ "X$1" != "X" ]
    do
      print - " "
      print - "ASE-Action: ping/check for hostname ${1} if alive"
      /usr/sbin/ping -c 1 ${1} >/dev/null
      returncode=$?
      case $returncode in
      0) print - "ASE-Info : ${1} is responding over the net"
         print - "ASE-Action: Trigger the restart of application
instance ${2} "
         print - "            on server $1 using
${ASEDIR}/${RESTART_SCRIPT}"
         ${ASEDIR}/bin/execwrap "${ASEDIR}/bin/tcp_client_ipc ${1}
${RESTART_SCRIPT}"
         returncode=$?
          case $returncode in
          0) print - "ASE-Success: restart of instance ${2} on $1
succesfully triggered"
             ;;
          *) print - "ASE-Error: restart of instance ${2} on $1
failed"
             MAIL_ADMIN ${DB1_START_LOG}
             ;;
          esac
        ;;
      *) print - "ASE-Info : ${1} is not respondig over the net"
         print - "ASE-Warning: restart of instance $2 not possible"
        ;;
      esac
      shift;shift
    done
}

stop_second_database () {

        print - " "
        print - "ASE-Action: check if instance ${DB2_INSTANCE} is
running"
        PID=`ps -o pid,command -A|grep "${DB2_SAPSYSTEMNAME}" | \
        grep -v grep|awk '{print $1}'`
        if [ "X${PID}" != "X" ]
        then
                print - "ASE-Action: stop local database server
${DB2_INSTANCE}"
                print - "            with su - ${DB2USER} -c
${DB2_STOPSAP}"
                proc_timeout ${DB2_STOPSAP} ${TIMEOUT} 10 &
                su - ${DB2USER} -c "${DB2_STOPSAP}" > /dev/null
                returncode=$?
                case $returncode in
                0) print - "ASE-Success: Database stop command
${DB2_STOPSAP} succeeded"
                   print - "ASE-Action: Check if
${DB2_SAPSYSTEMNAME}_${DB2_INSTANCE} is really down"
                   PID=`ps -o pid,command -A|grep
"${DB2_SAPSYSTEMNAME}_${DB2_INSTANCE}" | \
grep -v grep|awk '{print $1}'`
                   if [ "X${PID}" != "X" ]
                   then
                        print - "ASE-Action: Kill SAP Instance
```

```
${DB2_SAPSYSTEMNAME}_${DB2_INSTANCE}"
                        /bin/kill -KILL  ${PID}
                   fi
                   print - "ASE-Success:
${DB2_SAPSYSTEMNAME}_${DB2_INSTANCE} is really down"
                ;;
                *) print - "ASE-Error: ${DB2_STOPSAP} failed"
                   MAIL_ADMIN ${DB2_STOP_LOG}
                ;;
                esac
        else
                print - " "
                print - "ASE-Info: second database server is not
running"
        fi
}
#############################################################################
```