



SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400

DIGITAL HiTest Notes

Part Number: EK-HSPXB-HN. C01

July 1998

Revision/Update Information: This is a revised document.

**Digital Equipment Corporation
Maynard, Massachusetts**

July 1998

Digital Equipment Corporation makes no representations that the use of its products in the manner described in this publication will not infringe on existing or future patent rights, nor do the descriptions contained in this publication imply the granting of licenses to make, use, or sell equipment or software in accordance with the description.

Possession, use, or copying of the software described in this publication is authorized only pursuant to a valid written license from DIGITAL or an authorized sublicensor.

© Digital Equipment Corporation 1998. All rights reserved.

The following are trademarks of Digital Equipment Corporation: DIGITAL, the DIGITAL logo, Alpha, AlphaServer, TruCluster, ServerWORKS, and StorageWorks.

Third-party trademarks:

Oracle is a registered trademark and Oracle8 is a trademark of Oracle Corporation.

SAP is a registered trademark of SAP Aktiengesellschaft.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd.

Windows is a registered trademark and Windows NT is a trademark of Microsoft Corporation.

All other trademarks are the property of their respective owners.

Table of Contents

1 Advantages of DIGITAL HiTest Suites

- What Is a DIGITAL HiTest Suite? 1-1
- DIGITAL HiTest Suite Components..... 1-1
- Additional Hardware and Software..... 1-2

2 About This DIGITAL HiTest Suite

- Availability2-2
 - Features of SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 84002-2
 - Recommendations for SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400
 - HiTest Suite2-2
 - Understanding Availability Features2-3
 - Backup and Restore2-3
 - Disk Storage Technologies.....2-3
 - Just a Bunch of Disks (JBOD).....2-3
 - Redundant Array of Independent Disks (RAID)2-3
 - Logical Storage Manager (LSM).....2-5
 - Clustering2-6
 - Redundant Components2-6
- Installability2-6
- Interoperability2-6
- Price Range2-7
- Scalability2-7
 - Additional Hardware Components2-7
 - Workload Capability.....2-8
- Services2-9
 - Proof of Commitment: The DIGITAL Uptime Guarantee2-9
 - Portfolio of Business Critical Services2-9
 - Complementary Support Services2-10
 - Meeting Client Needs Locally or Globally2-10
 - For More Information2-10
- Year 2000 Compliance2-11

3 Configuration Data

- Hardware and Software Components3-1
- Special Configuration Rules3-7
- R/3 Specific Configuration Rules3-8
 - Security Rules3-8
 - Performance Rules.....3-8

4 System Installation and Setup

Hardware Installation	4-1
KZPSA Installation.....	4-2
ESA 10000 Installation	4-3
Database Disks.....	4-6
Example with Eighteen Disks for the Database	4-8
Setup the Second ESA 10000.....	4-9
Setup the Database Disks on HSZ3 and HSZ4 (Optional).....	4-9
Performance Hints	4-9
Setup with Three Stripsets and Two Disks Each on the HSZ3 and HSZ4.....	4-10
Setup with Three Stripsets and Six Disks Each on the HSZ3 and HSZ4.....	4-11
Setup with Four Stripsets and Six Disks Each on the HSZ3 and HSZ4	4-12
Operating System Installation.....	4-13
Disk Initialization	4-13
Operating System Custom Installation	4-14
Licenses.....	4-18
Final Adjustments for the Operating System	4-18
Configure Network Interfaces	4-20
FDDI Controller (fta0)	4-20
Fast Ethernet Controller (tu0).....	4-20
DIGITAL UNIX Patch Installation	4-20
Network Setup.....	4-21
Hosts and Rhosts.....	4-22
/etc/hosts.....	4-22
Postinstallation Procedures	4-22
Labeling the Local Disks	4-22
Labeling the Shared Disks.....	4-23
Adding Swap space.....	4-23
Installing the AdvFS Utilities	4-23
TruCluster Available Server	4-24
Logical Storage Manager (LSM).....	4-25
Mandatory Disksetup	4-25
First ESA 10000 Installation	4-27
Second ESA 10000 Installation (Optional if the Database is Bigger than 164 GB)	4-28
SAP R/3 Installation	4-29
SAP System ID.....	4-30
SAP and Oracle Directory Structure.....	4-30
Mandatory Directories for all Database Options.....	4-30
Database with Disks with One ESA 10000 (includes HSZ1 and HSZ2):.....	4-31
Database Disks with Two ESA 10000, Six Stripsets	4-32
Database Disks with Two ESA 10000, Seven Stripsets.....	4-32
Starting the SAP R/3 Installation	4-33
Useful Add Ons	4-33
Postinstallation for TruCluster V1.5 in a SAP R/3 Environment	4-33
Differences to the Original Documentation	4-34

5 Tests and Results

Overview of Results	5-1
Test Environment	5-1
Test Tools.....	5-2
Test Configuration.....	5-2
System Limits and Characterization Data	5-3

Test Process and Results.....	5-3
6 Problems and Solutions	
Foundation Operating System.....	6-1
Error in UERF	6-1
AppSet Software.....	6-1
Local bin Directory.....	6-1
SM51 Failed	6-2
Entries in TCPDB.....	6-2
Import of Online Documentation Failed.....	6-3
Errors in dev_w files in the Instance Work Directory.....	6-3
APPL-SERVER not found	6-4
Error in dev_w.....	6-4
Error During Making Oracle Intelligent Agent.....	6-4
Testphase	6-4
No Connect to the Database.....	6-4
Unable to Obtain Requested Swapspace	6-4
Perl Is Needed	6-5
Mess-tools not Started	6-5
Error After Relocating with TruCluster AS.....	6-5
Error in Deassign TCP/IP alias	6-5
Error for Each Simulated User.....	6-6
Problems with Sapgui	6-6
7 Detailed Hardware Configuration	
System Diagram	7-2
HiTest System Slot Configuration	7-3
A Instructions to Modify the TCR Available Server SAP R/3 Installation	
Stop R/3	A-1
Setting up the XNTPD Time Daemon.....	A-1
Modifications to Provide a Failover without Database Dependencies	A-1
General Modifications	A-2
Modify the /etc/exports-files.....	A-5
Modifications for the Oracle Database.....	A-6
Create a Disk Service	A-7
Create a User-Defined Service.....	A-8
Task to Do after Relocation of the Database.....	A-8
Procedure How to Put (SAP R/3 upgrade) a R/3 System in a Failover Environment	A-9
TruCluster AS Manager Utility.....	A-9
How to Set Up a Disk Service.....	A-10
How to Set Up a User-Defined Service.....	A-14
How to Relocate the Services	A-16
Transactions and Tests to Check the R/3 System	A-16
Modifications in R/3 to Prevent the Database Running on the Application Server to Pop-Up a Sapgui Error	A-17
Testing R/3 Basic Functionality.....	A-17
Testing Message and Gateway	A-18
Testing Enqueue and Update Services	A-18
Batch Services.....	A-19
Testing Spool Service.....	A-20

Contents

Application Tests	A-20
Additional Tests	A-21
Test the Transport System.....	A-21
Testing the Installed DECsafe Configuration	A-22
Debugging and Diagnosing Tool sys_check	A-22
Tasks for Setting the ASE Off-line.....	A-22
Restrictions and Limitations	A-23
Parameter in the rc_serv.conf file	A-24

Figures

Figure 2-1: RAID Level Summary.....	2-4
Figure 2-2: SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 Price Range	2-7
Figure 2-3: SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 HiTest Suite Scalability	2-8
Figure 2-4: SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 Workload Capability.....	2-9
Figure 4-1: ESA 10000 Frontview - Mandatory Disks.....	4-5
Figure 4-2: ESA 10000 - Database with Nine Times 4 GB Disks	4-7
Figure 4-3: ESA 10000 Frontview - Database with Eighteen Disks with 4 GB Disks	4-8
Figure 4-4: ESA 10000 Frontview HSZ3 and HSZ4.....	4-10
Figure 4-5: ESA 10000 Frontview HSZ3 and HSZ4 with 3 Stripsets (6 Disks).....	4-11
Figure 4-6: ESA 10000 Frontview HSZ3 and HSZ4 with 4 Stripsets (6 Disks).....	4-12
Figure 4-7: Example for Volume SAP01	4-29
Figure 5-1: Test Environment.....	5-1
Figure 7-1: System Diagram	7-2
Figure 7-2: HiTest System Slot Usage	7-3
Figure 7-3: I/O Slot Usage.....	7-4

Tables

Table 2-1: SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 Availability Features	2-2
Table 2-2: RAID Levels and Descriptions	2-4
Table 2-3: SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 Year 2000 Compliance	2-11
Table 3-1: DIGITAL HiTest Template – AppSet Software and Foundation Hardware	3-2
Table 3-2: DIGITAL HiTest Template – Foundation Software	3-4
Table 3-3: System Management Station Template	3-5
Table 3-4: Component Revision Levels	3-7
Table 7-1: Configuration Cabling	7-3
Table 7-2: System Slot Usage (Minimum and Maximum Configurations)	7-3
Table 7-3: I/O Slot Usage (Minimum and Maximum Configurations).....	7-4

Preface

This document provides an overview of DIGITAL HiTest Suites and detailed technical information about the SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 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 – Instructions to modify the TCR Available Server SAP R/3 Installation	Contains more detailed information about the TCR Available Server SAP R/3 Installation.

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-HSPXB-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* (including *Guide to the DECsafe Available Server*)
- *DIGITAL UNIX Software Documentation* (is delivered with DIGITAL UNIX on CD-ROM).
Logical Storage Manager documentation is included in the *DIGITAL UNIX Software Documentation*
- *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 86856 - R/3 Installation on UNIX*
- *SAP Notes 86855 - R/3 Installation on UNIX - OS Dependencies*
- *SAP Notes 86857 - 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.

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 AS DIGITAL UNIX AlphaServer 8400 HiTest Suite contains three groups of components: the *DIGITAL HiTest Foundation*, and the *DIGITAL HiTest AppSet*, and *System Management Station*.

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 medium- to high-end SAP TruCluster AS 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.

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 8400 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 AS DIGITAL UNIX AlphaServer 8400 HiTest Suite includes the following components:

- SAP R/3
- Oracle8 for DIGITAL UNIX
- TruCluster Available Server for DIGITAL UNIX
- DIGITAL UNIX
- AlphaServer 8400
- StorageWorks ESA 10000

This chapter describes the following characteristics of the SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 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 off-line 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 AS DIGITAL UNIX AlphaServer 8400

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 AS DIGITAL UNIX AlphaServer 8400 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 AS DIGITAL UNIX AlphaServer 8400 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).

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.

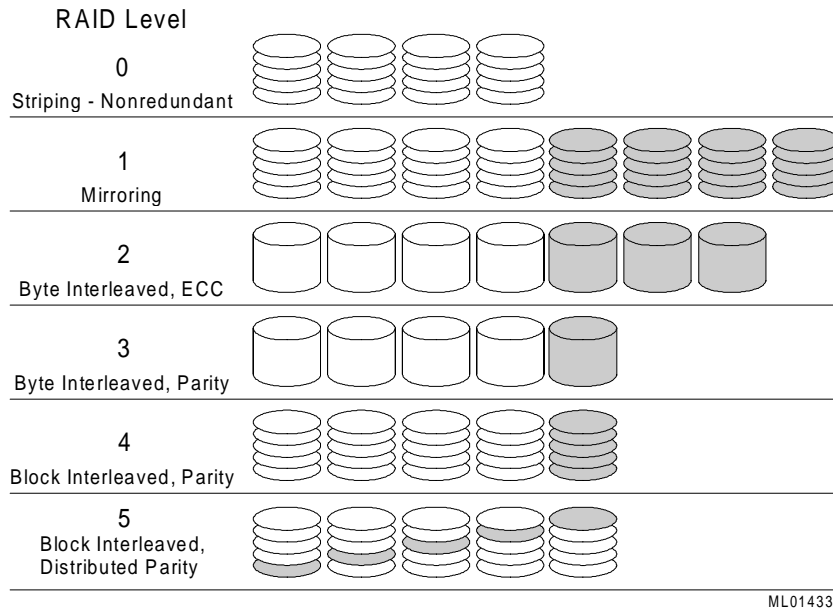
About This DIGITAL HiTest Suite

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.

Table 2-2: RAID Levels and Descriptions

RAID Level	Description	Advantages/Disadvantages
0	<ul style="list-style-type: none"> Striping Data segmented and distributed across several disks 	<ul style="list-style-type: none"> + increase in performance due to parallelism in read and write - no fault tolerance (<i>not</i> a high availability solution)
1	<ul style="list-style-type: none"> Hardware Mirroring Data written twice to different disk spindles within the disk array 	<ul style="list-style-type: none"> + 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	<ul style="list-style-type: none"> Striped Mirroring Combined level 0 and 1 Data mirrored onto and striped across several disks Best for performance-critical, fault-tolerant environments 	<ul style="list-style-type: none"> + 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

RAID Level	Description	Advantages/Disadvantages
2	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • Independent access array • Parity disk 	+ processes multiple requests simultaneously - parity disk is a bottleneck on writes
5	<ul style="list-style-type: none"> • 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
1+5	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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 AS DIGITAL UNIX AlphaServer 8400 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 8400 CPU.

Figure 2-2: SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 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.

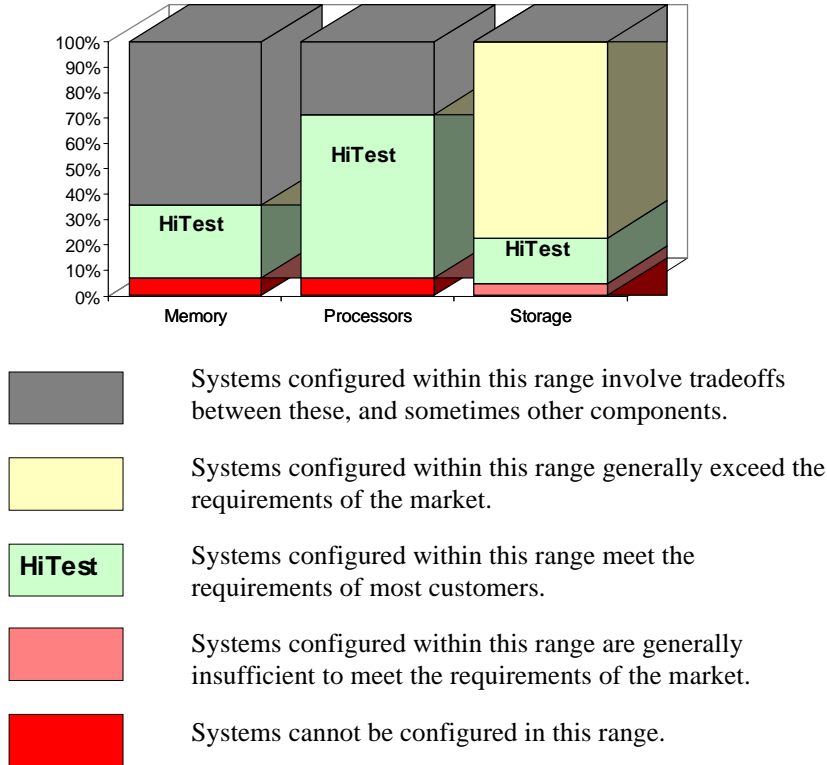
Additional Hardware Components

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 AS DIGITAL UNIX AlphaServer 8400 HiTest Suite Scalability



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 AS DIGITAL UNIX AlphaServer 8400 Workload Capability



One SD user corresponds to about three users in 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.

About This DIGITAL HiTest Suite

- 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 100 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 AS DIGITAL UNIX AlphaServer 8400 Year 2000 Compliance

Component	Year 2000 Status
AppSet software (SAP R/3)	Blue
Database Software (Oracle8 for DIGITAL UNIX)	Blue
Middleware (TruCluster Available Server for DIGITAL UNIX)	Blue
Operating System (DIGITAL UNIX)	Blue
Platform (AlphaServer 8400)	Blue
Major Storage (if appropriate) (StorageWorks ESA 10000)	N

Configuration Data

This chapter describes the SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 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 ten GB of memory, eight 2.1 GB disks, six 9.1 GB disks and the ESA 10000 with different disk options (see 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 AS DIGITAL UNIX AlphaServer 8400 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 4.0B Contact SAP at http://www.sap.com	SAP	2	2
Foundation Hardware				
2	<i>Select two systems:</i> AlphaServer 8400 5/625 System, 2 GB AlphaServer 8400 5/625 System, 4 GB <i>Hardware includes:</i> <ul style="list-style-type: none"> • Two 5/625 MHz CPUs; 4 MB cache • I/O module with four I/O channels (KFTHA-AA) • Memory (2 GB or 4 GB) • 4.3 GB 3.5" SCSI disk • One BN38C-02 PCI FWD SCSI cable • 600 MB CD-ROM drive • One KZPBA-CA, PCI UltraSCSI Adapter • KZPAA-AA, PCI FNSE SCSI ctr, for CD-ROM only • BN21H-02 2-meter SCSI cable • DE500-AA 10/100 Mbit EtherWORKS 32-bit NIC • Shielded console cable • BA670-AA UltraSCSI StorageWorks Plug-in-unit • One DWLPB-AA, PCI 12 slot Plug-in-unit • Two H7263-Ax non-BBU capable 48 VDC power regulators; includes power cord <i>Software includes:</i> <ul style="list-style-type: none"> • DIGITAL UNIX Operating System and base license • Unlimited User license • DIGITAL UNIX Server Extensions • ServerWORKS • Internet AlphaServer System Software • DECEvent • BMC PATROL Agent 	DA-292GF-CA ⓘ DA-292GG-CA ⓘ	2	2
3	Dual 5/625 CPU module	758P2-AX	0	8
4	512 GB Memory Module 1 GB Memory Module 2 GB Memory Module 4 GB Memory Module Note: This HiTest Template supports a memory range from 2 GB to 10 GB per system. When selecting memory options, stay within the Template's 10 GB maximum.	MS7CC-DA MS7CC-EA MS7CC-FA MS7CC-GA	0	See Note
5	Console Terminal	VT510-AG	2	2
6	Console Keyboard	LK46W-A2	2	2
7	Second PCI Expansion Box (mounts in DWLPB-AA)	DWLPB-BA	0	2
8	PCI plug-in unit with one PCI box	DWLPB-AA	0	2
9	SCSI Differential terminator	H879-AA	4	8
10	SCSI Y cable, 68-pin	BN21W-0B	4	8
ⓘ Indicates that geography-specific part number variants are available. Check the appropriate price book for details.				

SAP HiTest AppSet Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 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	SCSI Adapter PCI host (FWD), HD68 connector	KZPSA-BB	4	8
12	VHDCI female to HD68 male connector	BN38C-10	4	8
13	StorageWorks ESA 10000 <i>Hardware includes:</i> <ul style="list-style-type: none"> • DS-HSZ70-AH (2) • BN37A-10 (2) • DS-BA370-AA (2) 	DS-SWXES-BA	1	2
14	4.3 GB UltraSCSI disks with snap-in carriers	DS-RZ1CB-VW	10	82
15	9.1 GB UltraSCSI disks with snap-in carriers (internal disks)	DS-RZ1DB-VW	4	10
16	9.1 GB UltraSCSI disks with snap-in carriers (archive logs)	DS-RZ1DB-VW	2	2
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	176-Slot Tape library unit w/6 TZ89 drives	DS-TL896-BA	1	1
20	SCSI Adapter PCI host (FWD), HD68 connector	KZPSA-BB	3	3
21	15-meter SCSI cable	BN21K-15	3	3
22	20-meter SC to SC dual fiber cable	BN34B-20	2	2
23	Order items 21, 22, and 23 to create FDDI environment: DECconcentrator 900MX module with no modular PMDs factory installed Note: If the customer doesn't have a FDDI environment you have to connect the FDDI interfaces of the two systems to a FDDI Hub. A two-node FDDI connection does not work in an TruCluster AS environment, if one interface fails.	DEF6X-MA	0	1
24	DEChub ONE Ethernet single-slot hub chassis Includes power supply	DEHUA-CX	0	1
25	Multimode Fiber MMF SC connector	DEFXM-SC	0	1
26	System Management Station	See Table 3-3	0	1

Table 3-2: DIGITAL HiTest Template – Foundation Software

Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 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-05SAQ-AA	2	2	Yes	Yes
4	Oracle8 for DIGITAL UNIX	Included with item 1	2	2	Yes	Yes
5	HSZ70 Solutions Software for DIGITAL UNIX <i>Includes:</i> • PC Card containing the storage controller software • StorageWorks Command Console • Licenses and documentation	QB-5SBAB-SA	1	1	Yes	Yes
6	HSZ70 Solutions Software for DIGITAL UNIX including everything listed in item 5, except documentation.	QB-5SBAB-SB	3	7	Yes	Yes
7	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
8	Networker Save and Restore Application Interface for SAP R/3	QL-5JGAQ-AA	2	2	Opt'l	Opt'l
9	NSR Jukebox Tier 1 License	QL-04WAL-3B	1	1	Opt'l	Opt'l
10	Performance Advisor for DIGITAL UNIX, 3.0C CD-ROM	QA-054AA-H8	2	2	Opt'l	Opt'l
11	Performance Advisor for DIGITAL UNIX, 3.0C License	QL-0WFA9-AA	2	2	Opt'l	Opt'l
12	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 5510 ST System <i>Hardware includes:</i> <ul style="list-style-type: none"> • 233 MHz Pentium CPU with 32 KB cache • 512 KB secondary cache • 32 MB SDRAM memory • Integrated Fast Ethernet (10/100) • 24X CD-ROM • Integrated 3D Graphics adapter • 3.2 GB Ultra DMA disk drive • 1.44 MB floppy • Integrated Audio • Country Kit, North American <i>Software includes:</i> <ul style="list-style-type: none"> • Windows NT Workstation 4.0 (factory installed) Note: A functionally equivalent 80 x 86 system may be substituted without invalidating this HiTest Template.	FR-G2B2A-VB	1	1
2	32 MB SDRAM dual-bank DIMM Memory	FR-PCCAM-EC	1	1
3	Diamond 56.6K Modem Note: Used for page notification.	FR-PCXFA-AA	0	1
4	<i>Select one high-resolution monitor:</i> 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				
5	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
6	Windows NT Service Pack 3 Contact Microsoft at http://www.microsoft.com or ftp://ftp.microsoft.com/bussys/winnt or call the Microsoft Order Desk in the United States at (800) 360-7561 between 6:30 A.M. and 5:30 P.M., Pacific time.	Microsoft	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	1	1

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	<p><i>Choose one BMC product:</i> BMC PATROLWATCH for ServerWORKS, V3.2,</p> <p>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</p>	<p>Included with base system</p> <p>BMC BMC</p>	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	<p>BMC Operating System Knowledge Module Note: W* refers to the class: WA – Desktop; WB – Workgroup; WC – Departmental; WD – Enterprise Server</p>	QB-5KLAA-W*	0	1

Table 3-4: Component Revision Levels

Hardware Component	Hardware	Firmware	Software
SRM console	–	5.1-30	–
AlphaBIOS	–	5.35-0	–
SCSI host adapter (KZPAA-AA)	Rev. B01	–	
FWD SCSI controller (KZPSA-BB)	Rev. P01	A11	–
2.1 GB disks (DS-RZ1BB-VW)	–	DEC0656	–
4.3 GB disks (DS-RZ1CB-VW)	–	DEC0656	–
9.1 GB disk (DS-RZ1DB-VW)	–	LYJ0	–
625 MHz CPU (758P2-AX)	C02	5.1-30	–
2 GB Memory (MS7CC-FA)	B01	–	–
DEFPA	Rev D1	3.10	
PCI Ultra SCSI Adapter (KZPBA-CA)	A01	5.53	
StorageWorks shelf power supply (DS-BA35X-HH)	Rev. A01	–	–
Software Component		Version/Revision	Patch Level
DIGITAL UNIX		4.0D Rev.878	Patch Kit 001
SAP R/3		4.0B	
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 9.1 GB disk. Remove the 4.3 GB disks which are delivered with the 8400's and use them for the database.
 - 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.

Configuration Data

- SAP other (sapmnt, oracle, saparch, sapbackup, sapreorg and so on). The SAP other disk is a 9.1 GB disk, to handle 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 stripesets. The size of a stripeset depends on how much disks are in it. One stripeset can have from two to six disks (4.3 or 9.1GB disks). 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 divided in two parts - two controllers and their failover.
An optional ESA 10000 can be ordered to extend the database from 164 GB with 9.1 GB disks to 370 GB. Again the two HSZs mirror each other in this ESA 10000.
4.3 GB disks and 9.1 GB disks can be used for the database.
The disk configuration must be equal on each pair of HSZ.

For example:

You have 3 times 3 disks with a size of 4.3 GB each on both HSZs in the first ESA 10000. Build three stripesets with a size of 12.9 GB each and enable them to the operating system. Six devices on two controllers will be seen. Three volumes are the mirror part of the other three volumes on base of LSM.

Then you have three big volumes on which you will create AdvFS domains, because we need six file systems (sapdata1 to sapdata6) and we have to place two sapdata file systems on one volume.

Use the striping of the HSZ, The mirroring of LSM and the Advanced File System.

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.

If you are using one ESA 10000:

- 1st stripeset on HSZ1 in ESA 10000 contains sapdata1 and sapdata2
- 2nd stripeset on HSZ1 in ESA 10000 contains sapdata3 and sapdata4
- 3rd stripeset on HSZ1 in ESA 10000 contains sapdata5 and sapdata6

Do not forget the mirror disks on the 2nd HSZ (HSZ2) in the same ESA 10000.

If you are using two ESA 10000:

- 1st stripeset on HSZ1 in the 1st ESA 10000 contains sapdata1
- 2nd stripeset on HSZ1 in the 1st ESA 10000 contains sapdata2
- 3rd stripeset on HSZ1 in the 1st ESA 10000 contains sapdata3
- 1st stripeset on HSZ1 in the 2nd ESA 10000 contains sapdata4
- 2nd stripeset on HSZ1 in the 2nd ESA 10000 contains sapdata5
- 3rd stripeset on HSZ1 in the 2nd ESA 10000 contains sapdata6

Do not forget the mirror disks on the 2nd HSZ (HSZ4) in the same ESA 10000.

If your database extends 324 GB you have to create a 4th stripeset in the 2nd ESA 10000. You have to create a directory sapdata7 and mount this volume to this directory.

During SAP R/3 installation you have to spread the Data across 7 file systems.

SAP R/3 basic knowledge is mandatory for this installation.

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

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 tcr003) and the Application Server (hostname tcr004).
2. Install AdvFS utilities on both systems.
3. Install TruCluster AS 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 AS 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)
- 8 – 10 GB RAM are cached by another 9 GB disk (=33.4 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 how to check the system configuration:

1. Set the SCSI ID of the shared KZPSAs to 6, as follows:

```
P00>>>set kzpsa1_host_id 6
P00>>>set kzpsa2_host_id 6
P00>>>set kzpsa3_host_id 6
P00>>>set kzpsa4_host_id 6
```

2. Check the shared KZPSA option settings, as follows:

```
P00>>>show kzp*
kzpsa0_fast           1
kzpsa0_host_id       6      ←system
kzpsa0_termprwr      1
kzpsa1_fast           1
kzpsa1_host_id       6      ←1st shared bus
kzpsa1_termprwr      1
kzpsa2_fast           1
kzpsa2_host_id       6      ←2nd shared bus
kzpsa2_termprwr      1
kzpsa3_fast           1
kzpsa3_host_id       6      ←3rd shared bus
kzpsa3_termprwr      1
kzpsa4_fast           1
kzpsa4_host_id       6      ←4th shared bus
kzpsa4_termprwr      1
kzpsa5_fast           1
kzpsa5_host_id       6      ←for tape drive
kzpsa5_termprwr      1
```

3. Check all devices:

```
P00>>>show dev
```

4. Check the CPUs as follows:

```
P00>>>show cpu
Primary CPU:         00
Active CPUs:         00   01
Configured CPUs:    00   01
SRROM 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. Each HSZ70 contains two dual-redundant controllers (for example HSZ1_1 and HSZ1_2). Each ESA 10000 contains two HSZs. HSZ1 and HSZ3 are the bottom HSZs in the ESA 10000. HSZ2 and HSZ4 are the upper HSZs in the ESA 10000. HSZ1_1, HSZ2_1, HSZ3_1 and HSZ4_1 are the bottom controllers.

Set the prompt, date, failover, id's and preferred id's on HSZ1 and HSZ2:

```
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
```

System Installation and Setup

```
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 cfmnu where cfmnu is menu driven.

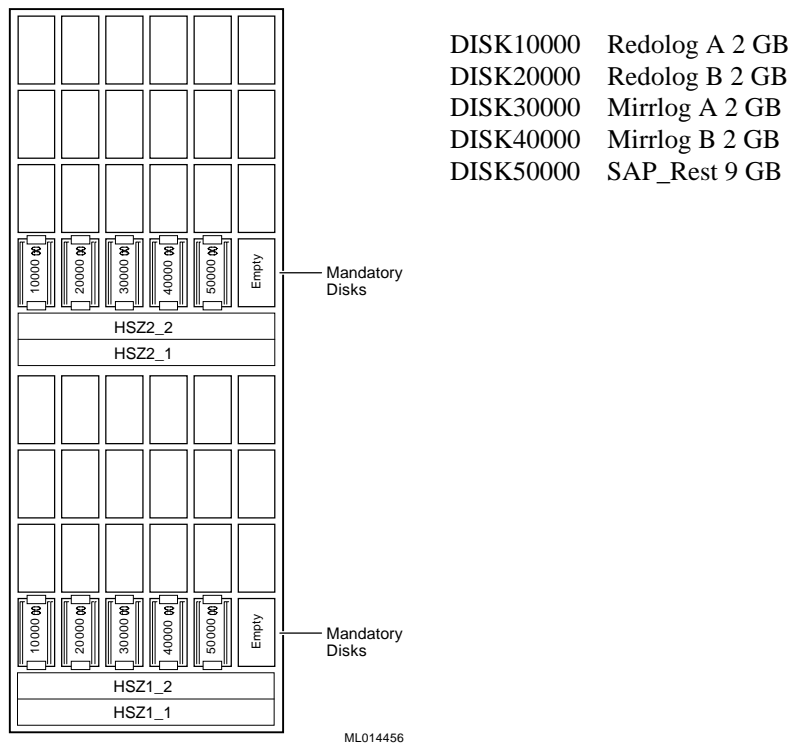
Note

The disk layout depends on how many disks for the database you have ordered, 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 database (sapdata1 - sapdata6).

If you ordered two ESA 10000, please read the section Setup Database Disks on HSZ3 and HSZ4 and the section Performance Hints. Depending on the amount of the disks you have ordered, you have to setup the disks.

The following disks are mandatory regardless of how many database disks you have chosen for the database. These disks are only installed on the first ESA 10000 (Oracle Logfiles and SAP-Rest are only used once). The second ESA 10000 contains only data disks (sapdata).

Figure 4-1: ESA 10000 Frontview - Mandatory Disks



```

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 d401 DISK50000 writeback_cache
maximum_cached_transfer=1024

```

System Installation and Setup

The disk DISK60000 is not used in the first two HSZs (HSZ1 and HSZ2). This slot is reserved for the future, if the customer needs more space for the archive logfiles in the productive environment. The disk can easily be mirrored by LSM and added to the disk SAP-Rest. The following shows the setup for the database disks. Depending on the ordered disks choose the correct part, which is valid.

Database Disks

Depending on the amount of disks and the database size the customer needs for his SAP R/3 System you have to setup the database disks. The database is located on stripesets. The upper three shelves in the HSZ1 and HSZ2 are used for them. Use DISK60100 (right most disk in the second shelf from the bottom), DISK60200 and DISK60300 as the startpoint. Add more disks beside these disks to define the stripesets you need. The minimum amount of disks should be two regardless which disktype you are using.

4 GB Disk Stripeseets:

3 Stripeseets with 2 disks each:	24 GB
3 disks each:	36 GB
4 disks each:	48 GB
5 disks each:	60 GB
6 disks each:	72 GB

9 GB Disk Stripeseets:

3 Stripeseets with 2 disks each:	54 GB
3 disks each:	81 GB
4 disks each:	108 GB
5 disks each:	135 GB
6 disks each:	162 GB

Note

If you need more then 162 GB you have to order and install a second ESA 10000 with two HSZ70. Each HSZ70 contains two controllers.


Figure 4-2 shows nine database disks and the mandatory disks. Start with DISK60100 to create three stripeset with three disks each. If your database needs more space, create stripesets with more disks. The size of the database depends on the amount of disks in a stripeset. Create always stripesets with the same amount of disks in each to be sure to get a well distributed performance regardless which stripeset is used by SAP R/3. The setup of the second ESA 10000 is done later in this manual.

Example:

If the database size should be 36 GB, you need 9 times 4 GB disks to create this size. Start with the rightmost disks. You will get 3 stripesets with the size of 12 GB each. To increase a database just add one additional disk to each stripeset . (The minimum are 2 disks and the maximum are 6 disks in one stripeset).

Do not create a 36 GB database with 9 GB disks, because you need only 4 disks and you are not able to create 3 stripesets with the same amount of disks in each.

Figure 4-2: ESA 10000 - Database with Nine Times 4 GB Disks



ML014457

```

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


```

Stripeset s1 is the lower set and s3 the upper set.

The operating system will see eight disks on HSZ1 and on HSZ2. Five mandatory disks and three database disks.

Example with Eighteen Disks for the Database

Figure 4-3: ESA 10000 Frontview - Database with Eighteen Disks with 4 GB Disks



```

HSZ> add stripe s1 DISK60100 DISK50100 DISK40100
DISK30100 DISK20100 DISK10100
HSZ> add stripe s2 DISK60200 DISK50200 DISK40200
DISK30200 DISK20200 DISK10200
HSZ> add stripe s3 DISK60300 DISK50300 DISK40300
DISK30300 DISK20300 DISK10300
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
    
```

Setup the second HSZ (HSZ2_1 and HSZ2_2) in the ESA 10000 as you did it with the first one. Both HSZs have to be setup identically, because the disks in these HSZs are shared on the buses and mirrored via LSM.

Setup the Second ESA 10000

Set the prompt, date, failover, id's and preferred id's on HSZ3 and HSZ4. These two HSZs are only to extend the database from 162 GB to 370 GB. The differences between the first two HSZs and the second HSZs are:

- No mandatory disks for redologs and SAP_Rest
 - Prompts HSZ3_1, HSZ3_2 and HSZ4_1, HSZ4_2
 - Four stripesets instead of three stripesets and mandatory disks (see section Performance Hints).
1. Set the prompt, date, failover, id's and preferred id's on HSZ3 :


```
HSZ1> set this prompt="HSZ3_1> "
HSZ1_1> set failover copy=this
HSZ1_1> set other prompt="HSZ3_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=09-apr-1998:15:30:00
```
 2. Do this setup on the HSZ4 too, using the previous steps except the prompt name.
 3. Check the settings on both HSZs, as recommended during the setup of the first HSZ.

Setup the Database Disks on HSZ3 and HSZ4 (Optional)

As described during setup of HSZ1 and HSZ2, the database disks are stripesets. But contrary to HSZ1 and HSZ2 you have no mandatory Disks on HSZ3 and HSZ4. You have space to create four stripesets. Start from the right most slots in the four shelves to build the stripesets. The minimum amount of disks in a single stripeset are two Disks.

Performance Hints

If all HSZ's are not filled up with database disks but you need more than 162 GB of space you can easily spread the stripesets equally over the ESA 10000s.

For example if your database needs a space of approximately 210 GB, create 3 stripesets with 5 times 9 GB Disks on HSZ3 and HSZ4 which gives you a size of 135 GB, and create 3 stripesets with 3 times 9 GB disks on HSZ1 and HSZ2 which gives you another 81 GB diskspace. Try to create more stripesets on the second two HSZs instead on the first two HSZs to get a equal load to the database, because the Oracle logfiles and their mirror volumes are on HSZ1 and HSZ2 (remember that the Oracle logfiles and their mirrorpart are written permanently).

Use either 3 or 6 logical disks on the operating system level, because the disklayout of SAP needs 6 filesets (sapdata1 to sapdata6). Use a fourth stripeset on HSZ3 and HSZ4 only if your database needs a diskspace of more than 324 GB. This is the upper limit with three stripesets made of 6 times 9 GB disks in each of the four HSZs.

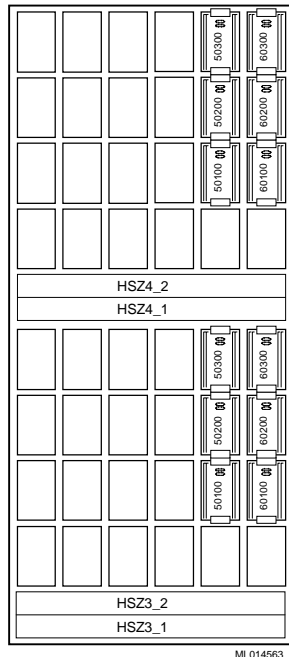
If you need this fourth stripeset, you have to change the SAP disklayout later on during the SAP installation.

You also have to create a seventh directory (sapdata7) and spread the SAP data across these 7 filesets.

System Installation and Setup

Setup with Three Stripesets and Two Disks Each on the HSZ3 and HSZ4

Figure 4-4: ESA 10000 Frontview HSZ3 and HSZ4




```
HSZ> add stripe s1 DISK60100 DISK50100
HSZ> add stripe s2 DISK60200 DISK50200
HSZ> add stripe s3 DISK60300 DISK50300
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
```

Do the same setup on HSZ4.

Setup with Three Stripsets and Six Disks Each on the HSZ3 and HSZ4

Figure 4-5: ESA 10000 Frontview HSZ3 and HSZ4 with 3 Stripsets (6 Disks)



ML014564

```

HSZ> add stripe s1 DISK60100 DISK50100 DISK40100
DISK30100 DISK20100 DISK10100
HSZ> add stripe s2 DISK60200 DISK50200 DISK40200
DISK30200 DISK20200 DISK10200
HSZ> add stripe s3 DISK60300 DISK50300 DISK40300
DISK30300 DISK20300 DISK10300
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


```

Do the same setup on HSZ4.

System Installation and Setup

Setup with Four Stripsets and Six Disks Each on the HSZ3 and HSZ4

Figure 4-6: ESA 10000 Frontview HSZ3 and HSZ4 with 4 Stripsets (6 Disks)



```
HSZ> add stripe s1 DISK60000 DISK50000 DISK40000
DISK30000 DISK20000 DISK10000
HSZ> add stripe s2 DISK60100 DISK50100 DISK40100
DISK30100 DISK20100 DISK10100
HSZ> add stripe s3 DISK60200 DISK50200 DISK40200
DISK30200 DISK20200 DISK10200
HSZ> add stripe s4 DISK60300 DISK50300 DISK40300
DISK30300 DISK20300 DISK10300
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
```

Do the same setup on HSZ4.

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 `rz0` as root disk and `rz6` as mirror of `rz0`.

1. Change to the `/dev` directory, as follows:
cd /dev
2. Use the `MAKEDEV` command to create disk device special files for `rz0`:
./MAKEDEV rz0
MAKEDEV: special file(s) for rz0:
`rz0a rz0b rrz0b rz0c rrz0c rz0d rrz0d rz0e rrz0e rz0f rrz0f`
`rz0g rrz0g rz0h rrz0h`
3. Check and write the disk label information on `rz0`:
disklabel -wr rz0 XYZ
You can use `XYZ` as the `disktype` option, because `disklabel` will check the diskheader anyway.
4. Start the disk label editor and label as follows:
disklabel -e rz0
`/dev/rrz0a:`
`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`

System Installation and Setup

```
8 partitions:
#   size      offset  fstype  [fsize bsize cpg]
a:  524288      0  unused
      # (Cyl.  0 - 156*)
b: 13052836  524288  unused
      # (Cyl.156*- 4040*)
c: 17773524      0  unused    0    0  # (Cyl. 0 - 5289)
d: 4194304 13577124  unused
      # (Cyl.4040*- 5289*)
e:      0      0  unused    0    0  # (Cyl.0- -1*)
f:      0      0  unused    0    0  # (Cyl.0- -1)
g:   2048 17769428  unused    0    0  # (Cyl.5288*- 5289*)
h:   2048 17771476  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 n) Custom Installation.
2. Follow the display and enter data when prompted. The hostname for the first HiTest system was tcr003.
3. Select a disk for DIGITAL UNIX installation. The root file system will be placed on the "a" partition of this disk.

	Disk Type	Device Name	Device Number	Controller Name	Controller Number
1)	RZ1DB-CA	rz0	0	SCSI	0
2)	RZ1DB-CA	rz6	6	SCSI	0

Enter your choice: 1

The rz0 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	13575124	4194304	17769427	c
e	0	0	-1	
f	0	0	-1	
g	17769428	2048	17771475	
h	17771476	2048	17773523	

4. 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**

5. 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**

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

Disk	Device	Device	Controller	Controller
Type	Name	Number	Name	Number
1) RZ1DB-CA	rz0	0	SCSI	0
2) RZ1DB-CA	rz6	6	SCSI	0

Enter your choice: **1**

7. Select the rz0 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**

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

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

System Installation and Setup

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

Disk Type	Device Name	Device Number	Controller Name	Controller Number
1) RZ1DB-CA	rz0	0	SCSI	0
2) RZ1DB-CA	rz6	6	SCSI	0

Enter your choice: **1**

Swap1 will be on the b partition of the rz0 because all of its other partitions are already in use.

10. You may choose to have a second swapping area (swap2).
Do you want a second swapping area (y/n) ?**n**
11. 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**
12. Check the file system:
You have requested this file system layout:
* root file system on rz0a, type AdvFS
* /usr file system on rz0d, type AdvFS
* /var will be on /usr
* first swapping area (swap1) will be on rz0b
* no second swapping area (swap2)
Is this the correct file system layout (y/n)? **y**
13. Install all mandatory and optional subsets.
14. Set the console boot variables, as follows, then boot:
>>> **set boot_osflags A**
>>> **set bootdef_dev dka0**
>>> **boot**
15. When prompted, select all kernel options. At this point, the kernel is built and the system is rebooted.
16. Do not edit the configuration file.
17. When the system is up, start to modify the disklabel of disk rz6 to prepare it for mirroring the boot disk.
18. Write a disklabel to rz6:
#disklabel -wr rz6 XYZ

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

19. Edit the disklabel of disk rz6 and set the following partition values:

```
#disklabel -e rz6
/dev/rrz6a:
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*- 4041*)
c: 17773524      0  unused    0    0  # (Cyl.  0 - 5289)
d:  4194304 13577172  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:      2048 17771476  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 swap space. The space of the g: partition is used to extend the swap space 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.

20. Encapsulate the boot disk:

```
#volencap rz0
Setting up encapsulation for rz0.
- Disk rz0 is the system boot disk and LSM is not initialized.
Creating simple disk rz0e to initialize LSM and rootdg.
- Partition rz0a is the root partition which requires 2 passes to
encapsulate and the temporary use of a free partition.
Using partition rz0f for temporary root encapsulation.
- Creating nopriv disk for primary swap device rz0b.
- Creating nopriv disk for rz0d.
```

The following disks are queued up for encapsulation or use by LSM. You must reboot the system to perform the actual encapsulations. rz0e rz0a rz0f rz0b rz0d

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.

System Installation and Setup

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 rz6
Mirroring system disk rz0 to disk rz6.

This operation will destroy all contents on disk rz6.
The partition map from rz0 will be copied to rz6 and all volumes associated with rz0 will be mirrored.

Do you want to continue with this operation? (y or n) y
Initializing rz6.

Mirroring rootvol to rz6a.
Mirroring swapvol to rz6b.
Mirroring vol-rz0d to rz6d.

2. Now the disklabel of both disks rz0 and rz6 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 rz0 is mirrored by disk rz6.
Check that with the following command:

```
#volprint
```

The output should look like this:

TYPE	NAME	ASSOC	KSTATE	LENGTH	COMMENT
dg	rootdg	rootdg	-	-	
dm	rz6a	rz6a	-	524288	
dm	rz6b	rz6b	-	13051860	
dm	rz6d	rz6d	-	4194304	
dm	rz6e	rz6e	-	0	
dm	rz0a	rz0a	-	524288	
dm	rz0b	rz0b	-	13051860	
dm	rz0d	rz0d	-	4194304	
dm	rz0e	rz0e	-	0	
sd	rz6a-01p	rootvol-02	-	16	
sd	rz6a-01	rootvol-02	-	524272	
sd	rz6b-01	swapvol-02	-	13051860	
sd	rz6d-01	vol-rz0d-02	-	4194304	
sd	rz0a-01p	rootvol-01	-	16	
sd	rz0a-01	rootvol-01	-	524272	
sd	rz0b-01	swapvol-01	-	13051860	
sd	rz0d-01	vol-rz0d-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-rz0d-01	vol-rz0d	ENABLED	4194304	
plex	vol-rz0d-02	vol-rz0d	ENABLED	4194304	
vol	rootvol	root	ENABLED	524288	
vol	swapvol	swap	ENABLED	13051860	
vol	vol-rz0d	fsngen	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 -rz0d:

```
# voldisksetup -i rz0g publen=512 privlen=512
# voldisksetup -i rz0h publen=512 privlen=512
# voldisksetup -i rz6g publen=512 privlen=512
# voldisksetup -i rz6h publen=512 privlen=512
# voldg -g rootdg adddisk rz0g=rz0g
# voldg -g rootdg adddisk rz0h=rz0h
# voldg -g rootdg adddisk rz6g=rz6g
# voldg -g rootdg adddisk rz6h=rz6h
# volmake -g rootdg sd rz0g-01 dm_name=rz0g dm_offset=0 len=1
# volmake -g rootdg sd rz0h-01 dm_name=rz0h dm_offset=0 len=1
# volmake -g rootdg sd rz6g-01 dm_name=rz6g dm_offset=0 len=1
# volmake -g rootdg sd rz6h-01 dm_name=rz6h dm_offset=0 len=1
# volsd -g rootdg aslog vol-rz0d-01 rz0g-01
# volsd -g rootdg aslog vol-rz0d-02 rz6g-01
```

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

System Installation and Setup

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 -rz0d.
`#volprint -Ath`
or
`#dxlsm`

The devices rz0h-01 and rz6h-01 are reserved if you want to enable BCL on the /device rootvol. However, if enabled, there is performance disadvantage.

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 lists the settings used for this HiTest System as installed in the test laboratory. According to the SAP rules, the hostname should not be longer than eight characters. See also *Check list - Installation Requirements: Oracle*, which comes with the SAP software kit. The virtual hostname should be at least eight characters. These hostnames are set in the section Postinstallation for TruCluster V1.5 in a SAP R/3 Environment. 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 tcr003.
IP address for interface fta0 is 1.0.0.3.
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 tcr003e.
The IP address for interface tu0 is 155.56.201.103.
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.

Copy and install the patches:

1. Get `duv40das00001-19980312.tar` and all available readme files.
2. Get `patches.osf-tcr1.5-17mar1998.tar.Z` and all available readme files.
Install the TCR 1.5 patch after you have installed TCR 1.5.
3. If you are using Network Save and Restore, install the following Patch:
`GURU://usr/guest/misc/osf/marshall/TCR150/libbaseapi.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
```

Note

If you have chosen a different SAP System ID, use this instead of TCR.

2. Start nfssetup as follows:

```
tcr003> 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: tcr004
Enter the directory pathname: /usr/sap/trans
Netgroup/Machine name: tcr004
```

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: tcr004
/usr/sap/trans exported to: tcr004
```

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.

System Installation and Setup

Hosts and Rhosts

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

```
tcr003 root
tcr003e root
tcr004 root
tcr004e 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	
1.0.0.3	tcr003	← FDDI host name DB Server
1.0.0.30	vtcr003	← Virtual host name for TCR150 DB Server
155.56.201.103	tcr003e	← Ethernet name DB Server
155.56.201.111	vtcr003e	← Virtual name for TCR150 PC Connection DB Server
1.0.0.4	tcr004	← FDDI host name Appl. Server
1.0.0.40	vtcr004	← Virtual host name for TCR150 Appl. Server
155.56.201.104	tcr004e	← Ethernet host name DB Server
155.56.201.112	vtcr004e	← Virtual name for TCR150 PC Connection Appl Server

The user's PC connects via the virtual name `vtcr003e` and `vtcr004e` to the SAP R/3 System. The communication between DB Server and Application Server is done via `vtcr003` and `vtcr004`. 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 `vtcr003` and `vtcr003e` will fail over to the Application Server and the PC's can connect to `vtcr003e` which is physically on `tcr004`. This is also valid for `vtcr004` and `vtcr004e` on `tcr003`.

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/rz1c 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/rz2c swap3 ufs sw 0 2
- If you are using 8 GB to 10 GB of memory, add another 9 GB disk next to swap3 and modify the /etc/fstab file:
/dev/rz3c swap4 ufs sw 0 2

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

```
# disklabel -wr rz1 XYZ
# disklabel -wr rz2 XYZ
# disklabel -wr rz3 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

For our HiTest environment we do not need 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/<cdrom> /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_Uutilities/doc
```

The documentation is available in Text, HTML and Postscript format.

System Installation and Setup

3. Change to the directory and install the software:
`cd /mnt/AdvFS_Advanced_Uutilities/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 [tcr003]:**tcr003**
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.[TCR003]: **TCR003**
Do you want to edit the configuration file ? (y/n) [n]: **no**
Select the controllers that define the shared ASE I/O Buses.
Enter your choice (comma or space separated): **1 2**
(If you are using two ESA1000 with four HSZs, add 3 and 4)
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)	scsi1	scsi1	pza0	0	pci0	6
2)	scsi2	scsi2	pza1	0	pci1	5
3)	scsi3	scsi3	pza1	0	pci1	9
4)	scsi4	scsi4	pza2	0	pci2	0

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**
6. The kernel is rebuilt after exiting from the menu. Check the kernel files, then move them and reboot as follows:
tcr003> `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:


```
tcr003> mv /vmunix /vmunix.before.ase
tcr003> mv /sys/TCR003/vmunix /vmunix
tcr003> reboot
```
8. Repeat the previous steps on the Application Server (tcr004 in this HiTest Suite), except labeling the shared disk. This has to be done only once.
9. Install available TCR patches.

Notes

If the configuration in /sys/conf/TCR003 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 disks you have ordered (4.3 GB disks or 9.1 GB disks and the amount of disks). The Mandatory Part is shown first and is valid for all installations. The step on depending the database disks are shown step by step later on. This setup is done only on tcr003.

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 named HSZ1 contains disk rz9, rz10, rz11, rz12, and rzd12.
 The second HSZ70 named HSZ2 contains their equivalent mirroredisks rz17, rz18, rz19, rz20 and rzd20.


```
# voldisksetup -i rz9c nlog=1 nconfig=1
# voldisksetup -i rz10c nlog=1 nconfig=1
# voldisksetup -i rz11c nlog=1 nconfig=1
# voldisksetup -i rz12c nlog=1 nconfig=1
# voldisksetup -i rzd12c nlog=1 nconfig=1
# 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
```
2. Create the LSM disk group sapdg:


```
# voldg init sapdg rz9=rz9c
```

System Installation and Setup

3. Add all other mandatory disks to LSM group sapdg:

```
# voldg -g sapdg adddisk rz10=rz10c
# voldg -g sapdg adddisk rz11=rz11c
# voldg -g sapdg adddisk rz12=rz12c
# voldg -g sapdg adddisk rzd12=rzd12c
# voldg -g sapdg adddisk rz17=rz17c
# voldg -g sapdg adddisk rz18=rz18c
# voldg -g sapdg adddisk rz19=rz19c
# voldg -g sapdg adddisk rz20=rz20c
# voldg -g sapdg adddisk rzd20=rzd20c
```
4. 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 rz9 rz9c simple 1024 8376488 /dev/rrz9c
dm rz10 rz10c simple 1024 8376488 /dev/rrz10c
dm rz11 rz11c simple 1024 8376488 /dev/rrz11c
dm rz12 rz12c simple 1024 8376488 /dev/rrz12c
dm rzd12 rzd12c simple 1024 8376488 /dev/rrzd12c
```
 - b) Next look for the disks rz9, rz10, rz11, rz12 and rzd12. The disks rz17 to rzd20 are created automatically during the mirroring.
 - c) Evaluate the size of the subdisk. Use the values you got from the volprint command for the rz9. This example shows a 4 GB disk. In this configuration, the rz9 to the rz12 and the rz17 to rz20 are 2 GB disks. The rzd12 and rzd20 are 9 GB disks. Your display can differ from this. Subtract 2 from this value. You get 8376486 for the example of the rz9.

```
#valassist -g sapdg -U fsgen make redoA 8376486
alloc=0 align=0 rz9
```
 - 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 rz9
#volassist -g sapdg -U fsgen make redoB $SIZE
alloc=0 align=0 rz10
#volassist -g sapdg -U fsgen make mirrA $SIZE
alloc=0 align=0 rz11
#volassist -g sapdg -U fsgen make mirrB $SIZE
alloc=0 align=0 rz12
#volassist -g sapdg -U fsgen make SAPRest $SIZE
alloc=0 align=0 rzd12
```
 - e) Mirror the LSM disks in the background:

```
#volassist -g sapdg mirror redoA rz17 &
#volassist -g sapdg mirror redoB rz18 &
#volassist -g sapdg mirror mirrA rz19 &
#volassist -g sapdg mirror mirrB rz20 &
#volassist -g sapdg mirror SAPRest rzd20 &
```

- 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 rz9-01
#volsd -g sapdg aslog redoA-02 rz17-01
#volsd -g sapdg aslog redoB-01 rz10-01
#volsd -g sapdg aslog redoB-02 rz18-01
#volsd -g sapdg aslog mirrA-01 rz11-01
#volsd -g sapdg aslog mirrA-02 rz19-01
#volsd -g sapdg aslog mirrB-01 rz12-01
#volsd -g sapdg aslog mirrB-02 rz20-01
#volsd -g sapdg aslog SAPRest-01 rzd12-01
#volsd -g sapdg aslog SAPRest-02 rzd20-01
```

### First ESA 10000 Installation

The HSZ1 contains the disks rzb9, rzb10, and rzb11. The equivalent mirroredisks on the HSZ2 are rzb17,rzb18, and rzb19.

- Initialize all database disks and their equivalent mirror disks, which should be used by LSM:
 

```
voldisksetup -i rzb9c nlog=1 nconfig=1
voldisksetup -i rzb10c nlog=1 nconfig=1
voldisksetup -i rzb11c nlog=1 nconfig=1
voldisksetup -i rzb17c nlog=1 nconfig=1
voldisksetup -i rzb18c nlog=1 nconfig=1
voldisksetup -i rzb19c nlog=1 nconfig=1
```
- Add all other disks to diskgroup sapdg:
 

```
voldg -g sapdg adddisk rzb9=rzb9c
voldg -g sapdg adddisk rzb10=rzb10c
voldg -g sapdg adddisk rzb11=rzb11c
voldg -g sapdg adddisk rzb17=rzb17c
voldg -g sapdg adddisk rzb18=rzb18c
voldg -g sapdg adddisk rzb19=rzb19c
```
- 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 rzb9-01 dm_name=rzb9
 dm_offset=0 len=1k
volmake -g sapdg sd rzb10-01 dm_name=rzb10
 dm_offset=0 len=1k
volmake -g sapdg sd rzb11-01 dm_name=rzb11
 dm_offset=0 len=1k
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
```
- Create the LSM data volumes. Get the size of the LSM subdisk with volprint as described before and create the disks:
 

```
volassist -g sapdg -U fsgen make SAP1_S1 $SIZE
 alloc=0 align=0 rzb9
volassist -g sapdg -U fsgen make SAP1_S2 $SIZE
 alloc=0 align=0 rzb10
volassist -g sapdg -U fsgen make SAP1_S3 $SIZE
 alloc=0 align=0 rzb11
```
- Mirror the LSM Disks in the background:
 

```
volassist -g sapdg mirror SAP1_S1 rzb17 &
volassist -g sapdg mirror SAP1_S2 rzb18 &
volassist -g sapdg mirror SAP1_S3 rzb19 &
```

## System Installation and Setup

- 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 rzb9-01
volsd -g sapdg aslog SAP1_S1-02 rzb17-01
volsd -g sapdg aslog SAP1_S2-01 rzb10-01
volsd -g sapdg aslog SAP1_S2-02 rzb18-01
volsd -g sapdg aslog SAP1_S3-01 rzb11-01
volsd -g sapdg aslog SAP1_S3-02 rzb19-01
```

### Second ESA 10000 Installation (Optional if the Database is Bigger than 164 GB)

The HSZ3 contains the disks rzb25, rzb26, and rzb27. Their equivalent mirror disks on HSZ4 are rzb33, rzb34, and rzb35.

- Initialize all database disks and their equivalent mirror disks, which should be used by LSM:

```
voldisksetup -i rzb25c nlog=1 nconfig=1
voldisksetup -i rzb26c nlog=1 nconfig=1
voldisksetup -i rzb27c nlog=1 nconfig=1
voldisksetup -i rzb33c nlog=1 nconfig=1
voldisksetup -i rzc34c nlog=1 nconfig=1
voldisksetup -i rzc35c nlog=1 nconfig=1
```
- Add all other diskblocks to diskgroup sapdg:

```
voldg -g sapdg adddisk rzb25=rzb25c
voldg -g sapdg adddisk rzb26=rzb26c
voldg -g sapdg adddisk rzb27=rzb27c
voldg -g sapdg adddisk rzb33=rzb33c
voldg -g sapdg adddisk rzc34=rzc34c
voldg -g sapdg adddisk rzc35=rzc35c
```
- 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 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 rzb33-01 dm_name=rzb33
dm_offset=0 len=1k
volmake -g sapdg sd rzc34-01 dm_name=rzc34
dm_offset=0 len=1k
volmake -g sapdg sd rzc35-01 dm_name=rzc35
dm_offset=0 len=1k
```
- Create the LSM data volumes. Get the size of the LSM subdisk with volprint as described before and create the disks:

```
volassist -g sapdg -U fsgen make SAP2_S1 $SIZE
alloc=0 align=0 rzb25
volassist -g sapdg -U fsgen make SAP2_S2 $SIZE
alloc=0 align=0 rzb26
volassist -g sapdg -U fsgen make SAP2_S3 $SIZE
alloc=0 align=0 rzb27
```

5. Mirror the LSM Disks in the background:
- ```
# volassist -g sapdg mirror SAP2_S1 rzb33 &
# volassist -g sapdg mirror SAP2_S2 rzc34 &
# volassist -g sapdg mirror SAP2_S3 rzc35 &
```

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 SAP2_S1-01 rzb25-01
# volsd -g sapdg aslog SAP2_S1-02 rzb33-01
# volsd -g sapdg aslog SAP2_S2-01 rzc26-01
# volsd -g sapdg aslog SAP2_S2-02 rzc34-01
# volsd -g sapdg aslog SAP2_S3-01 rzc27-01
# volsd -g sapdg aslog SAP2_S3-02 rzc35-01
```

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 4-7: Example for Volume SAP01



ML014463

SAP R/3 Installation

This section describes how to install SAP R/3. The current release at installation time was 4.0B.

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).

We used the name TCR for our SAP System ID. You can use another name for SID, but check the *SAP R/3 Installation Guide* which names are not supported.

- Installation directory: /sapmnt/TCR/install

SAP and Oracle Directory Structure

SAP enforces a rigid naming scheme concerning the entry points for the R/3 structures. See also the manual *SAP R/3 Installation on UNIX - Oracle Database*. The directories /sapmnt/TCR and /usr/sap/trans should already exist. They are created during the nfssetup. If not, create them now. All other mountpoints, which are on the shared buses are only mounted when the procedure for postinstallation of TruCluster V1.5 in a SAP environment is performed.

Map the SAP structures on the available physical structure.

Mandatory Directories for all Database Options

The following directories are mandatory for all Database Options:

```
mkdir /oracle
mkdir /oracle/TCR
mkdir /oracle/stage
mkdir /oracle/stage/stage_803
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

```

Database with Disks with One ESA 10000 (includes HSZ1 and HSZ2):

This section describes the the setup with one ESA 100000:

```

mkfdmn /dev/vol/sapdg/SAP1_S1 SAP1_S1
mkfset SAP1_S1 sapdata1
mkfset SAP1_S1 sapdata2
mkfdmn /dev/vol/sapdg/SAP1_S2 SAP1_S2
mkfset SAP1_S2 sapdata3
mkfset SAP1_S2 sapdata4
mkfdmn /dev/vol/sapdg/SAP1_S3 SAP1_S3
mkfset SAP1_S3 sapdata5
mkfset SAP1_S3 sapdata6
mount SAP1_S1#sapdata1 /oracle/TCR/sapdata1
mount SAP1_S1#sapdata2 /oracle/TCR/sapdata2
mount SAP1_S2#sapdata3 /oracle/TCR/sapdata3
mount SAP1_S2#sapdata4 /oracle/TCR/sapdata4
mount SAP1_S3#sapdata5 /oracle/TCR/sapdata5
mount SAP1_S3#sapdata6 /oracle/TCR/sapdata6

```

System Installation and Setup

Database Disks with Two ESA 10000, Six Stripesets

This section describes the setup for the database disks with two ESA 10000 (includes HSZ1, HSZ2, HSZ3 and HSZ4 and six stripesets):

```
mkfdmn      /dev/vol/sapdg/SAP1_S1  SAP1_S1
mkfset      SAP1_S1      sapdata1
mkfdmn      /dev/vol/sapdg/SAP1_S2  SAP1_S2
mkfset      SAP1_S2      sapdata2
mkfdmn      /dev/vol/sapdg/SAP1_S3  SAP1_S3
mkfset      SAP1_S3      sapdata3
mkfdmn      /dev/vol/sapdg/SAP2_S1  SAP2_S1
mkfset      SAP2_S1      sapdata4
mkfdmn      /dev/vol/sapdg/SAP2_S2  SAP2_S2
mkfset      SAP2_S1      sapdata5
mkfdmn      /dev/vol/sapdg/SAP2_S3  SAP2_S3
mkfset      SAP2_S3      sapdata6
mount       SAP1_S1#sapdata1      /oracle/TCR/sapdata1
mount       SAP1_S2#sapdata2      /oracle/TCR/sapdata2
mount       SAP1_S3#sapdata3      /oracle/TCR/sapdata3
mount       SAP2_S1#sapdata4      /oracle/TCR/sapdata4
mount       SAP2_S2#sapdata5      /oracle/TCR/sapdata5
mount       SAP2_S3#sapdata6      /oracle/TCR/sapdata6
```

Database Disks with Two ESA 10000, Seven Stripesets

This section describes the setup for the database disks with two ESA 10000 (includes HSZ1, HSZ2, HSZ3 and HSZ4) and seven stripesets. This setup is needed if the database is bigger than 324 GB:

```
mkfdmn      /dev/vol/sapdg/SAP1_S1  SAP1_S1
mkfset      SAP1_S1      sapdata1
mkfdmn      /dev/vol/sapdg/SAP1_S2  SAP1_S2
mkfset      SAP1_S2      sapdata2
mkfdmn      /dev/vol/sapdg/SAP1_S1  SAP1_S3
mkfset      SAP1_S3      sapdata3
mkfdmn      /dev/vol/sapdg/SAP2_S1  SAP2_S1
mkfset      SAP2_S1      sapdata4
mkfdmn      /dev/vol/sapdg/SAP2_S2  SAP2_S2
mkfset      SAP2_S2      sapdata5
mkfdmn      /dev/vol/sapdg/SAP2_S3  SAP2_S3
mkfset      SAP2_S3      sapdata6
mkfdmn      /dev/vol/sapdg/SAP2_S4  SAP2_S4
mkfset      SAP2_S4      sapdata7

mount       SAP1_S1#sapdata1      /oracle/TCR/sapdata1
mount       SAP1_S2#sapdata2      /oracle/TCR/sapdata2
mount       SAP1_S3#sapdata3      /oracle/TCR/sapdata3
mount       SAP2_S1#sapdata4      /oracle/TCR/sapdata4
mount       SAP2_S2#sapdata5      /oracle/TCR/sapdata5
mount       SAP2_S3#sapdata6      /oracle/TCR/sapdata6
mount       SAP2_S4#sapdata7      /oracle/TCR/sapdata7
```


Starting the SAP R/3 Installation

For the SAP R/3 installation a certified SAP installer is a prerequisite.

The SAP Notes will be available at SAP and must be read before the installation. Get them with the Fax template which will be delivered with each installation kit or via a connection to SAP (for example OSS).

Always get the newest Notes from SAP because they will be updated frequently. All the information you need is delivered by SAP in the R/3 Installation Kit.

Install as follows:

1. Install Central System on tcr003
Use 00 as the Instance Number.
2. Order the SAP License with the Fax template which is delivered with the R/3 installation kit. The temporary license will work for four weeks.
3. Install Dialog System on tcr004
Use /sapmnt/TCR/install as the installation directory.
Use 01 as the instance number.
4. After the installation of the dialog instance you have to order a second SAP License. The reason is, that after failover of the central instance the message server is started on the second system which needs the new license.
Use the Fax template which is delivered with the R/3 Installation kit.
The temporary license will work for four weeks.

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 through the 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 through 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

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.

System Installation and Setup

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.

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.

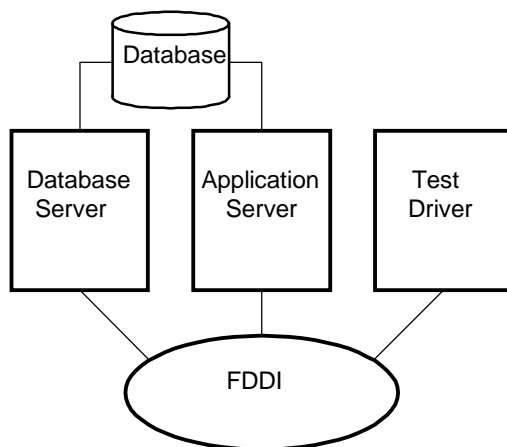
Overview of Results

Interoperability testing was performed successfully on the SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 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 AS test environment.

Figure 5-1: Test Environment



ML014464

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 vtrcr004 to vtrcr003 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 vtrcr004 back to the Application Server System and rerun 100 SD User (50/50).
6. Relocate the Service DB Server from vtrcr003 to vtrcr004 to check if both instances can run on the Application Server in 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 vtrcr003 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 shadow copy.
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.

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 Operating System

The following problems were identified:

Error in UERF

Problem	Disk errors in the uerf (UNIX Error Formatter).
Impact	/usr partition is getting full.
Solution	<p>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 disks in the /usr/var/opt/advfsd/disks.ignore file as described in the advfs_support notesfile. The only solution is to stop the advfsd with the following command:</p> <pre>/sbin/rc3.d/S53advfsd stop</pre> <p>If you want to use the dtadvfs again you can restart it:</p> <pre>/sbin/rc3.d/S53advfsd start</pre> <p>To prevent the system starting the daemon during system startup, remove the link to the original file or move it to a different name.</p>

AppSet Software

The following problems were identified:

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	<p>If you answer with default, you have to create the directory bin manually.</p> <pre># mkdir /usr/local/bin</pre>

SM51 Failed

Problem	<p>Errors in Syslog and SM51 failed with Dialog Instance</p> <pre>----- 14:14:38 tcr003_TCR_00 RD SAP* Q0I K Betriebssystem Aufruf getservbyname gescheitert (errno 0) 14:14:38 tcr003_TCR_00 RD SAP* Q0A K Service sapgw01 nicht bekannt 14:14:38 tcr003_TCR_00 RD SAP* S0T S NiConnect mißlungen, rc: -0003 14:14:38 tcr003_TCR_00 RD SAP* S0R S > Host tcr004e 14:14:38 tcr003_TCR_00 RD SAP* S0S S > Service: sapgw01 14:14:38 tcr003_TCR_00 RD SAP* Q0I K Betriebssystem Aufruf getservbyname gescheitert (errno 0) 14:14:38 tcr003_TCR_00 RD SAP* Q0A K Service sapdp01 nicht bekannt 14:14:38 tcr003_TCR_00 RD SAP* S0C K Verbindung kann nicht aufgebaut werden, RC=-01 14:14:38 tcr003_TCR_00 D0 000 SAP* SM51 R49 K Kommunikations-Fehler, CPIC-RC=020, SAP-RC=639 14:14:38 tcr003_TCR_00 D0 000 SAP* SM51 R64 K > CPIC- Funktion: CMINIT(SAP) </pre>
Impact	SAP R/3 does not work.
Solution	<p>Check in both systems that all sockets are set in /etc/services. For example: sapdp00, sapdp01, ... sapgw00, sapgw01, ...</p>

Entries in TCPDB

Problem	<p>If you check the SAP R/3 installation and you look at the syslog with transaction SM21, you will see the following error: In table TCPDB are 0 entries.</p>
Impact	SAP R/3 installation has an error message.
Solution	Apply SAP note 15023.

Import of Online Documentation Failed

Problem The import of the online documentation failed. You get the following messages:
tcr003e_tcradm> ./STARTUP

```
-----
Enter "SERVER" or "SCROLL" [SERVER] := scroll
      R3up message
-----
*** END OF UPGRADE PHASE  INITPUT ***
*** START OF UPGRADE PHASE DBCHK ***
running /usr/sap/TCR/SYS/exe/run/tp
pf=/usr/sap/put/bin/TPPARAM getdbinfo TCR
Warning: The return code of tp is 16
      R3up error message
R3up> SEVERE ERROR: unable to read from the database.
R3up> Call "tp getdbinfo TCR" and analyze the output.
      >>> PLEASE READ THE HELP TEXT <<<
Enter <CR> to continue, "help", or "cancel"...
```

```
tcr003e_tcradm> more /usr/sap/put/log/SLOGVAR
PANIC PANIC PANIC 19980428162246
ERROR: SAP-System not installed properly. Systems name
(TCR) is different from name in TADIR (SAP)
ERROR: EXIT(16) -> process ID is: 17595
PANIC PANIC PANIC 19980428162246
ERROR: SAP-System not installed properly. Systems name
(TCR) is different from name in TADIR (SAP)
ERROR: EXIT(16) -> process ID is: 17598
ERROR: tp_connect:  Cannot read 'tnsnames.ora' in
/usr/sap/trans
ERROR: tp_connect:  Cannot read 'tnsnames.ora' in
/usr/sap/trans
ERROR: tp_connect:  Cannot read 'tnsnames.ora' in
/usr/sap/trans
```

Impact Import of the online documentation is not possible.

Solution In the documentation, the import of the online documentation is done before the setup of the transport and correction system. But the import only works with the setup transport. Please setup the transport and correction system before trying to import the documentation. How to do that see the documentation on CD-ROM R/3 SYSTEM RELEASE 4.0A ONLINE DOCUMENTATION.

Errors in dev_w files in the Instance Work Directory

Problem After R/3 Installation you will find the following errors in the dev_w files in the Instance work directory:

```
D *** ERROR -> invalid parameter (diagoutn 2925)
```

Impact No Impact, just error message. But the Hot Package must be installed.

Solution Import the Hot Packages 4.0B/1 which will be delivered with the SAP R/3 Kit. See also SAP notes 102452, 104127, and 104664.

APPL-SERVER not found

Problem	During the installation of the Dialog Instance on the Application Server the R3SETUP 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.

Error in dev_w

Problem	The following errors are in the dev_w files in the instance work directory: E Info: Enqueue using fast table implementation E Info: Factor 4 larger enqueue/table_size required than in Rel 3.xx
Impact	No impact
Solution	Ignore error.

Error During Making Oracle Intelligent Agent

Problem	During Installation of Oracle with the tool orainst the procedure fails with the following error message. Error during action Making Oracle Intelligent Agent (target=install Command: make -f ins_oemagent.mk install
Impact	Oracle installation corrupted
Solution	Get SAP note 104316 and apply the necessary actions.

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 = 'TCR' "
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 = TCR

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.

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. <pre># perl -v</pre> This is perl, version 5.003 with EMBED build under dec_osf at Sep 20 1996 13:47:02 + suidperl security patch Copyright...

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 <i>not</i> use the hosttype PR in the apl file of the local simulation directory. This is not supported.

Error After Relocating with TruCluster AS

Problem	After relocating an instance with TruCluster AS, you will see the following error in the SAP syslog. <pre> E07 K Fehler 00000 : 3 No such process in Modul rslgsmcc (071) E07 K Fehler 00004 : 8 Address already in use in Modul rslgsend (002)</pre>
Impact	No impact.
Solution	Ignore error.

Error in Deassign TCP/IP alias

Problem	If the database and the central instance as well as the dialog instance are running on one system and the central system will be relocated to the 2 nd system you will get the following error in the logfile of the Application Server in /usr/local/ase on the system where you started the relocate. You will get a mail as user root: <pre>ASE-Action: De-assign TCP/IP alias vtcr004 ifconfig: ioctl (SIOCDIFADDR): Can't assign requested address ASE-Error: De-assign TCP/IP alias vtcr004 failed</pre>
Impact	No impact
Solution	Ignore the message and the mail

Problems and Solutions

Error for Each Simulated User

Problem	The benchmark environment will generate errors for each simulated user on the Application Server.
Impact	No impact
Solution	Ignore the errors. The benchmark is working properly. A fix will be available soon.

Problems with Sapgui

Problem	If you start the database service on the application server you will get a popup if you have started a sapgui on one of the systems. English: Your system is not correctly installed! Deutsch: Ihr System ist nicht korrekt installiert!
Impact	No impact, just a warning.
Solution	See the documentation how to modify R/3 in an TruCluster AS environment in Appendix A.

Detailed Hardware Configuration

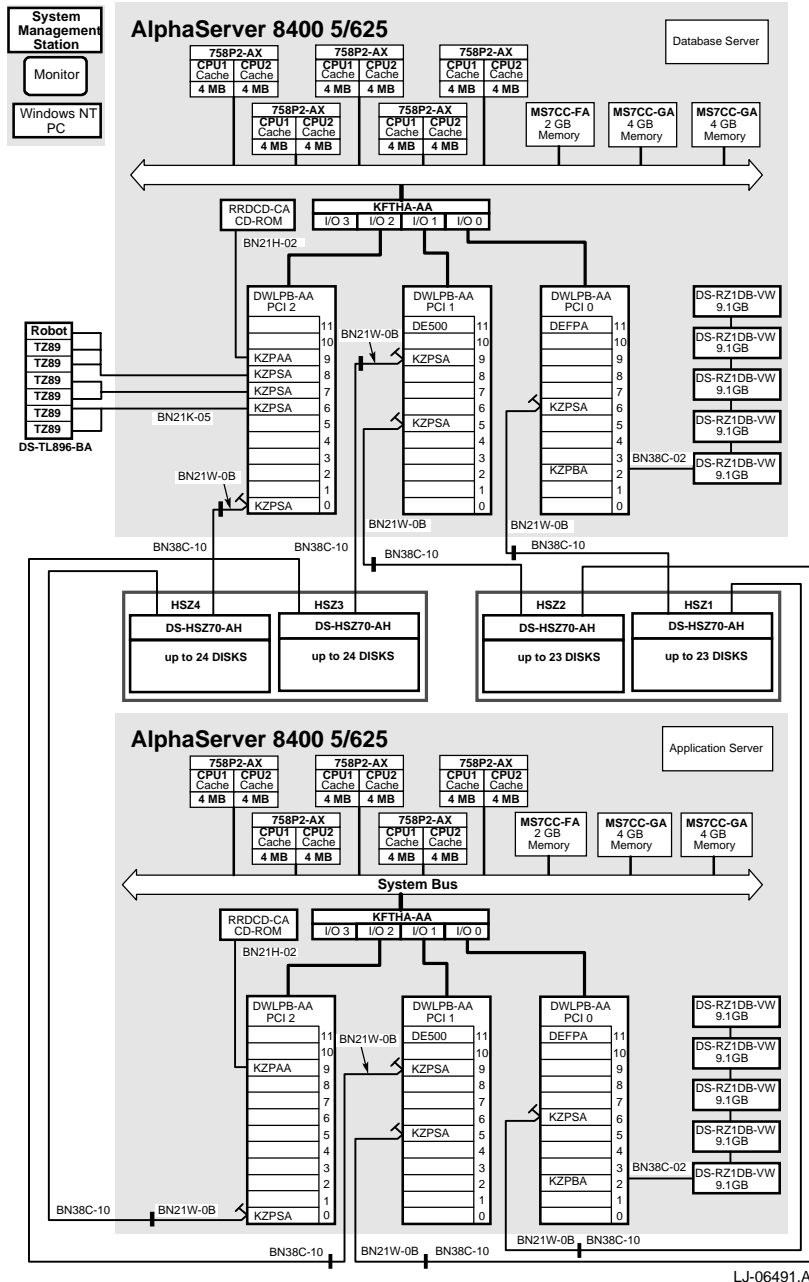
This chapter describes the minimum and maximum hardware configuration for the SAP Oracle TruCluster AS DIGITAL UNIX AlphaServer 8400 HiTest Suite by providing the following:

- 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 rz0 and rz6 during the UNIX Installation.

Table 7-1: Configuration Cabling

Part Number	Qty	Description	From	To
BN21K-15	3	SCSI bus	KZPSA-BB	TL896-BA
BN21W-0B	8	SCSI bus	KZPSA-BB	BN38C-10
BN38C-10	8	SCSI bus	BN21W-0B	ESA 10000
BN34B-20	2	SCSI bus	DEFPA	DEFPA
BN21H-02	2	SCSI bus	KZPAA-AA	CD-ROM
BN38C-02	2	SCSI bus	KZPBA-CA	BA670

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

Front		Rear
	Centerplane	4
3		5
2		6
1		7
0		8

ML014124

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

Slot	Minimum Option	Maximum Option	Description
0	CPU	CPU	Dual CPU processor 5/625
1		CPU	Dual CPU processor 5/625
2		CPU	Dual CPU processor 5/625
3		CPU	Dual CPU processor 5/625
4		CPU	Dual CPU processor 5/625
5		MS7CC-GA	4 GB Tlaser memory option
6		MS7CC-GA	4 GB Tlaser memory option
7	MS7CC-FA	MS7CC-FA	2 GB Tlaser memory option
8	KFTHA-AA	KFTHA-AA	I/O module, 4 channel

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

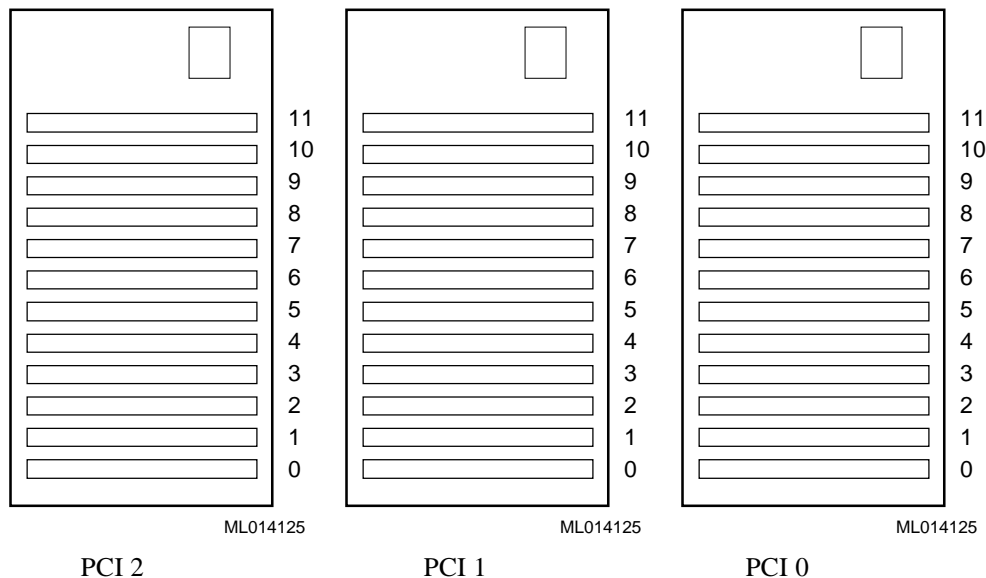


Table 7-3: I/O Slot Usage (Minimum and Maximum Configurations)

Slot	Minimum Configuration Options	Maximum Configuration Options	Description
PCI 2:			
9	KZPAA-AA	KZPAA-AA	SE SCSI controller
8	KZPSA-BB [*]	KZPSA-BB [*]	FWD SCSI controller
7	KZPSA-BB [*]	KZPSA-BB [*]	FWD SCSI controller
6	KZPSA-BB [*]	KZPSA-BB [*]	FWD SCSI controller
0	KZPSA-BB	KZPSA-BB	FWD SCSI controller
PCI 1:			
11	DE500-AA	DE500-AA	10/100 Ethernet controller
9	KZPSA-BB	KZPSA-BB	FWD SCSI controller
5	KZPSA-BB	KZPSA-BB	FWD SCSI controller
PCI 0:			
11	DEFPA-DA	DEFPA-DA	FDDI controller
6	KZPSA-BB	KZPSA-BB	FWD controller
3	KZPBA-CA	KZPBA-CA	UltraSCSI controller
[*] This option is present in Database Server only.			

A

Instructions to Modify the TCR Available Server SAP R/3 Installation

Stop R/3

Stop R/3 on the Application Server and stop R/3 and Oracle on the Database Server.

Setting up the XNTPD Time Daemon

Install the time daemon XNTPD on both systems. You can do the ntpsetup or modify the /etc/ntp.conf file manually for implementing the time daemon. Run ntpsetup to do a basic configuration of ntp and to add the startup of the ntp daemons to the init.d files. All commands should be executed as user root.

An example of setting up the XNTPD time daemon:

1. At the server side (the machine which should provide the system time):

```
vi /etc/ntp.conf peer 127.127.1.1
rcmgr set XNTPD_CONF YES
/sbin/init.d/xntpd start
```

2. At the client side:

```
vi /etc/ntp.conf server <server-name>
                        (in our example tcr003)
rcmgr set XNTPD_CONF YES
rcmgr set XNTP_SERV1 <servername>
rcmgr set XNTP_SERV2 ""
rcmgr set XNTP_SERV3 ""
/sbin/init.d/xntpd start
```

After the TruCluster AS Software installation is complete, install the applications, you want highly available, including mount-points, application-specific modifications, and creating the action scripts. Setup the services, using the asemgr utility. This part is described later in this installation documentation. For more information refer to the *Guide to the DECsafe Available Server*.

Modifications to Provide a Failover without Database Dependencies

After you have installed the R/3 system, you have to make both systems identical with the same mount-points and directory structure for the system. In our case, all operating system dependencies which belong to the R/3 system and to the database system.

General Modifications

Perform the following general modifications to the systems:

1. It is recommended for the next steps to save all modified files to files with the extension for example .prease. This makes it easier to restore the original state of the installation.
2. Each of these services must have a TCP/IP address in the network in the same subnet like the physical addresses. Ask your system administrator to provide the new addresses. You have to add the new hostnames and addresses to the /etc/hosts file of both machines.
For example:

```
/etc/hosts
1.1.1.101 dbserver
1.1.1.102 appserver
1.1.1.110 virtual
           # virtual name for the database (vtcr003)
1.1.1.120 virtual-app
           # virtual name for the appl. server (vtcr004)
```
3. Modify the /etc/services file for the SAP and DATABASE communication ports on the failover server, if not already set from the Application Server installation.
4. All directories for all SAP instances in the TruCluster AS environment have to be under /usr/sap/ on the shared device. As user root, copy the directories of the Application Server to the shared directory on the Database Server.
For example:

```
# cd /usr/sap/<SID>
# tar cvf D01.tar D01 SYS
# rcp D01.tar tcr003:/usr/sap/<SID>/
```

Extract the file on the Database Server:

```
# cd /usr/sap/<SID>
# tar xpvf D01.tar
```
5. Create the /usr/sap/<SID> directory for the Application Server on both machines locally, including all files. Thus the dialog instance is able to run on the Database Server without shared disks and central instance. The creation of the local /usr/sap/<SID>/D01 directory on the Database Server can only be done with unmounted disks or after the relocation of the central instance to the Application Server.
Now the directory /usr/sap/<SID> is installed three times, one time on the shared disks and two times on both machines locally.
6. Edit the /etc/passwd file on the Application Server and complete it with the ora<sid> entry of the Database Server. You do not need to add the <sid>adm entry because it already exists. Do the same with the /etc/group file. Edit this file and add group ID oper and dba. To activate the modifications use the mkpasswd /etc/passwd command after the modification or use vipw for modifying.
7. Add user ora<sid> to group oper on the Application Server.

8. Modify the environment files for the <sid>adm user and for the database user to run the database on the virtual name. First modify the Database Server files. There are several files where the environment of the database is set and depends on hostname entries. First copy the following listed files as described:

Replace the string virtual with the virtual name of the Database Server and the string virtual-app with the virtual name of the Application Server:

```
# cp ~<sid>adm/.dbenv_<hostname>.csh ~<sid>adm/.dbenv_virtual.csh
# cp ~<sid>adm/.sapenv_<hostname>.csh ~<sid>adm/.sapenv_virtual.csh
# cp ~ora<sid>/.dbenv_<hostname>.csh ~ora<sid>/.dbenv_virtual.csh
# cp ~ora<sid>/.sapenv_<hostname>.csh ~ora<sid>/.sapenv_virtual.csh
```

For example:

```
# cp ~tcradm/.dbenv_tcr003.csh ~tcradm/.dbenv_vtcr003.csh
```

Do the same with the environment files for the <sid>adm on the Application Server:

```
# cp ~<sid>adm/.dbenv_<hostname>.csh ~<sid>adm/.dbenv_virtual-app.csh
# cp ~<sid>adm/.sapenv_<hostname>.csh ~<sid>adm/.sapenv_virtual-app.csh
```

For example:

```
# cp ~tcradm/.dbenv_tcr004.csh ~tcradm/.dbenv_vtcr004.csh
```

9. Edit the file ~<sid>adm/.cshrc, ~ora<sid>/cshrc on the Database Server and modify the if clause under the following line, replacing the hostname entry:

```
# SAP environment
if ( -e $HOME/.sapenv_virtual.csh ) then
    source $HOME/.sapenv_virtual.csh
else if ( -e $HOME/.sapenv.csh ) then
    source $HOME/.sapenv.csh
endif
# RDBMS environment
if ( -e $HOME/.dbenv_virtual.csh )then
    source $HOME/.dbenv_virtual.csh
else if ( -e $HOME/.dbenv.csh ) then
    source $HOME/.dbenv.csh
endif
```

Be sure that the filesystem /oracle/<SID> is mounted on the Database Server.

10. Edit the file ~<sid>adm/.cshrc on the Application Server and modify the if clause under the following line, replacing the hostname entry:

```
# SAP environment
if ( -e $HOME/.sapenv_virtual-app.csh ) then
    source $HOME/.sapenv_virtual-app.csh
else if ( -e $HOME/.sapenv.csh ) then
    source $HOME/.sapenv.csh
endif
# RDBMS environment
if ( -e $HOME/.dbenv_virtual-app.csh )then
    source $HOME/.dbenv_virtual-app.csh
else if ( -e $HOME/.dbenv.csh ) then
    source $HOME/.dbenv.csh
endif
```

Do this as the particular user to keep the right owner and group.

Instructions to Modify the TCR Available Server SAP R/3 Installation

11. Change the SAP profiles in the /sapmnt/<SID>/profile directory, and copy the physically named instance profiles files to virtual named profiles.

For example:

```
# cp TCR_DVEBMGS00_ 'hostname' TCR_DVEBMGS00_virtual
# cp START_DVEBMGS00_ 'hostname' START_DVEBMGS00_virtual
# cp TCR_D01_ 'hostname' TCR_D01_virtual-app
# cp START_D01_ 'hostname' START_D01_virtual-app
```

12. Modify the START..... profiles replacing all physical hostname entries with the virtual hostname entry:

```
# vi START_DVEBMGS00_virtual
# vi START_D01_virtual-app
```

It is recommended to save the original files for switching to the original installation state.

13. Add the SAPLOCALHOSTFULL and SAPLOCALHOST parameter. Set the virtual name of the Database Server in the profile of the central instance and set the virtual name of the Application Server in the profile of the application-server file <SID>_<INSTANCE>, also change physical hostname entries in DEFAULT.PFL.

- a) DEFAULT.PFL:

```
SAPSYSTEMNAME = <SID>
SAPDBHOST      = virtual
rdisp/mshost   = virtual
rdisp/sna_gateway = virtual
rdisp/sna_gw_service = sapgw00
rdisp/vbname   = virtual_<SID>_<INSTANCE_No>
rdisp/enqname  = virtual_<SID>_<INSTANCE_No>
rdisp/btcname  = virtual_<SID>_<INSTANCE_No>
```

- b) <SID>_DVEBMGS_virtual:

```
SAPLOCALHOSTFULL = virtual.frs.dec.com # with bind domain
SAPLOCALHOST     = virtual
```

- c) <SID>_D01_virtual-app:

```
SAPLOCALHOSTFULL = virtual-app.frs.dec.com
SAPLOCALHOST     = virtual-app
```

The DB-reconnect is not supported yet, therefore see notes 24806 and 41678 in SAP info system OSS.

14. Create new start and stop scripts in the home directory of the <sid>adm user. In the following example, the physical and virtual names are used as described above:

- a) On the Database Server side:

```
# cp startsap_ "hostname" _00 startsap_virtual_00
# cp stopsap_ "hostname" _00 stopsap_virtual_00
```

Insert the virtual name of the Database Server in the scripts to start and stop the database and the central instance and right after that insert the virtual name of the Application Server in the scripts to start and stop the Application Server.

```
# vi startsap_virtual_00
and
```

```
# vi stopsap_virtual_00
START_PROFILE="START_DVEBMGS00_virtual"
```

Copy these files to the Application Server, and check the correct ownership and permissions.

- b) On the Application Server side:

```
# cp startsap_appname_01 startsap_virtual-app_01
```

Instructions to Modify the TCR Available Server SAP R/3 Installation

```
# cp stopsap_appname_01 stopsap_virtual-app_01
```

Do the same modifications like described before, and copy the files to the Database Server.

Do not forget to check the ownership and permissions.

The new start and stop scripts should now be present on both machines in the directory of the <sid>adm user and called:

startsap_virtual_00 - file to start the database and central instance

stopsap_virtual_00 - file to stop the database and central instance

startsap_virtual-app_01 - file to start the application instance

stopsap_virtual-app_01 - file to stop the application instance

- c) Change the startsap and stopsap aliases in the .sapenv_virtual.csh file on the Database Server and .sapenv_virtual-app.csh on the Application Server in the <sid>adm home directory. Therefore, change the hostname with virtual on the Database Server and hostname with virtual-app on the Application Server. Please keep a copy of these files.

```
# vi .sapenv_virtual.csh (on the Database Server)
```

```
alias startsap „$HOME/startsap_virtual_00“
```

```
alias stopsap „$HOME/stopsap_virtual_00“
```

```
# vi .sapenv_virtual-app.csh (on the Application Server)
```

```
alias startsap „$HOME/startsap_virtual-app_01“
```

```
alias stopsap „$HOME/stopsap_virtual-app_01“
```

15. If you have a multiple Application Server environment outside the DECsafe cluster you have to edit the .rhosts file for user <sid>adm, to allow rsh commands for restarting the external instances. This has to be done on both systems:

```
# vi ~<sid>adm/.rhosts
```

```
dbserver root <sid>adm # physical name
```

```
appserver root <sid>adm # physical name
```

16. Do not use the aliases stopsap or startsap, if one system is down and the service is on the other system.

Modify the /etc/exports-files

Modify the /etc/exports-files on both systems as follows:

1. First remove all entries in /etc/exports on both systems which belong to the TruCluster AS environment like /sapmnt/<SID> and /usr/sap/trans.

Because we are exporting two directories the following file /etc/exports.ase.virtual is used and has to be created on both machines. Please do that as user root.

2. The /var/ase/sbin/nfs_mountd uses this file to export the two directories.

Insert the following lines to the file /etc/exports.ase.virtual.

```
#vi /etc/exports.ase.vtcr003 (vtcr003 is our virtual name of the Database Server)
```

```
#ASE exports file for service db_alias_host
```

```
#sap#trans exports (after this line) Do not delete this line.
```

```
/usr/sap/trans -root=dbserver-hostname:appserver-hostname
```

```
# sap#sapmnt exports (after this line) - Do not delete this line.
```

```
/sapmnt/<SID> -root=dbserver-hostname:appserver-hostname
```

3. Because this file can take effect, you have to include the following line to the /etc/export.ase file, if it does not exist, create it on both systems:

```
.INCLUDE /etc/exports.ase.vtcr003.
```

4. Be sure that NFS (network file system) is installed on both systems for mounting and exporting nfs directories.

Instructions to Modify the TCR Available Server SAP R/3 Installation

5. Copy the actual software package with the scripts and the documentation from the DIGITAL internal WEB page <http://www.fra.dec.com/sap-eng/available/cluster.html> to the `/usr/local/ase` directory on both machines. The name of the package is `R3_DECSAFE_V*.tar`, where `*` stands for the version number.

6. Create `/usr/local/ase` with protection 755 on both systems.

7. Extract the `R3_DECSAFE_V*.tar` :

```
# cd /usr/local/ase/  
# tar xvf R3_DECSAFE_V*.tar
```

Contents of the `R3_DECSAFE_V*.tar` is:

<code>install.doc</code>	Installation documentation in word and
<code>install.ps</code>	postscript
<code>db_serv.ksh</code>	Script to start the disk-service
<code>as_serv.ksh</code>	Script to start the usr-defined-service
<code>app_restar.ksh</code>	Script for restarting external SAP applications
<code>rc_serv.conf</code>	File for the environment parameters that you have to modify for your environment
<code>/bin</code>	Binaries used within the scripts

8. Copy the binary file `/sapmnt/<SID>/exe/cleanipc` on both machines into `/usr/local/ase/bin`. This file is used to clean the shared memory on the specified instance. If the directory `/sapmnt/<SID>/exe` is not mounted via nfs on the Application Server you can copy the file with rcp.

In the `rc_serv.conf` file, all the environment variables and parameters for the two service scripts have to be set. In section Parameter in the `rc_serv.conf` file is an extract of the parameters to be modified in this script, after you installed it on the two systems.

9. Change TPPARAM

The `/usr/sap/trans/bin/TPPARAM` file contains settings for the transport control program `tp`. This file has the parameter `<SID>/dbhost` which has to be changed to the virtual name of the dbhost. Be sure that all hard-coded hostname entries are changed to the virtual name. It is recommended to save the original files before modifying them.

Modifications for the Oracle Database

The failover system has to have the same mount-points and directories which are necessary for the R/3 system. Be sure that the mount-points have the correct protections and the correct owner and group definitions.

1. Create the `/oracle/stage/` directory if it doesn't exist:

- a) Copy `/etc/oratab` to `/etc` on the application.

- b) Change the hostname in the `tnsnames.ora` in directory `/oracle/<SID>/network/admin/` on the DB Server and on the Application Server, and remove the hostname in `listener.ora` in the directory `/oracle/<SID>/network/admin/` on the DB Server. Be sure to keep a backup of the original files.

Instructions to Modify the TCR Available Server SAP R/3 Installation

- c) Save the following directories on the Application Server, and copy them to the Database Server, to be able to run the dialog instance on the Database Server and the central instance plus database on the Application Server. Be sure that all devices that belong to the TruCluster AS environment are dismounted, because the needed Oracle network files will need to be copied to `/oracle/<SID>` which is mounted under normal circumstances. These files are hidden if the database is started on the Database Server.
- d) Login on the Application Server as user root:

```
# cd /oracle/<SID>
# tar cvf appl.tar lib network ocommon rdbms
# rcp appl.tar <db-server-hostname>:/oracle/<SID>
```
- e) Login on the Database Server as user root:

```
# cd /oracle/<SID>
# ls (no files should be there except appl.tar)
# tar xpvf appl.tar
```
2. Delete the entries for the mount-points that you are using for the disk-service in the `/etc/fstab` on the DB Server.
In our example, you have to remove all disks that belong to the shared buses.
3. Dismount all directories which belong to the disk-service you want to add, otherwise the add of the disk service will fail.
4. Delete the entries for the mount-points that you are using for the user-specified-service in the `/etc/fstab` on the Application Server. In our example you have to remove all filesets which are mounted via nfs from the Database Server.

Create a Disk Service

Add all AdvFS domains which belong to the SAP shared disks as follows:

1. Test both start and stop scripts manually.
2. Test the relocation of the services manually.
3. Add the scripts into the ASE services with a time-out higher than 600 sec.
If this is not sufficient for your application choose a higher time-out. (Refer to the section How to Set Up a Disk Service)

Create a disk-service using the `asemgr-tool`. Create the disk-service on the Database Server with the following mount-points:

Instructions to Modify the TCR Available Server SAP R/3 Installation

LSM-Volumegroup	Mount-points
AdvFS saprest#oraclestage	NONE
AdvFS saprest#oracle	NONE
AdvFS saprest#usrsap	NONE
AdvFS saprest#saparch	NONE
AdvFS saprest#sapbackup	NONE
AdvFS saprest#sapreorg	NONE
AdvFS redologA#logA	NONE
AdvFS redologB#logB	NONE
AdvFS mirrlogA#mlogA	NONE
AdvFS mirrlogB#mlogB	NONE
AdvFS SAP1_S1#sapdata1	NONE
AdvFS SAP1_S2#sapdata2	NONE
AdvFS SAP1_S3#sapdata3	NONE
AdvFS SAP2_S1#sapdata4	NONE
AdvFS SAP2_S2#sapdata5	NONE
AdvFS SAP2_S3#sapdata6	NONE
AdvFS saprest#sapmnt	NONE
AdvFS saprest#saptrans	NONE

SAP2_S1 to SAP2_S3 are valid if you are using HSZ3 and HZS4 (2nd ESA10000).
With only one ESA10000 you have SAPS1_S1 to SAP1_S6.

Create a User-Defined Service

Create a user-defined service for the Application Server:

1. Test both start and stop scripts manually.
2. Test the relocation of the services manually.
3. Add the scripts into the TruCluster AS services with a time-out higher than 600 sec. If this is not sufficient for your application, choose a higher time-out.

Task to Do after Relocation of the Database

Since SAP R/3 version 3.0 the saplicense check is based on the message server and multiple licenses will be accepted for one system. In this case, you just have to insert the two licenses once, and the message server will automatically select the appropriate license.

1. Ensure that a temporary license is installed, because the system only comes up with the temporary license.
2. Order a second license for the failover machine, because the SAP licensing procedure is based on the MAC (Media Access Control) address of the system, and in our case the MAC address is changed after the failover.
3. Normally the license check is based on the CPU ID of the machine, but in case of DIGITAL the CPU does not have an ID.
saplicense -get
4. Fill out the saplicense formula to order the second license with the new key.
5. Install the second license with saplicense-install.

Procedure How to Put (SAP R/3 upgrade) a R/3 System in a Failover Environment

In the case of using virtual names and IP-addresses for the R/3 system, SAP recommends that the system should run on the physical hostname instead of the virtual hostname.

Steps for your upgrade procedure:

1. Shutdown the R/3 system and the database.
2. Save all the files which you have modified to run SAP in the TruCluster AS environment.
3. Restore all original files which you have saved before the modification of SAP in the TruCluster AS environment.
4. Startup R/3 and check all functions (SM51, SM21, ...)
5. Start the Upgrade (PUT) of SAP.
6. Shutdown the R/3 and the database.
7. Restore all the files which you have modified to run SAP in the TruCluster AS environment.
8. Startup R/3 and check all functions (SM51, SM21, ...)

Up to now it is not supported to start the put with virtual names.

TruCluster AS Manager Utility

The TruCluster AS Manager utility provides an easy-to-use interactive interface that allows you to set up and manage the TruCluster AS from any system in the DECsafe environment.

Start the TruCluster AS manager utility, to set up services in ASE, to make systems, disks, and applications highly available:

```
# asemgr
```

Before you set up these services modify all members where the application should run. Each member must have the same conditions to start the service.

When you create a service in an ASE, the asemgr utility prompts you for the following information:

- A unique name
- Either an application or storage configuration to make highly available
- Information about which member system you want to run the service and how you want the service to behave when a failure occurs

DECsafe provides support for three types of service:

- NFS-service - provide highly available access to exported disk
- Disk-service - provide highly available access to disks and disk based applications
- User-defined-service - provide highly available access to none disk based applications

In this configuration, a disk service was used for the database, central instance, and a user-defined service for the Application Server.

How to Set Up a Disk Service

These steps show how to implement the disk service using the asemgr utility:

1. Select managing ASE:

```
TruCluster Available Server (ASE)

ASE Main Menu

a) Managing the ASE    -->
m) Managing ASE Services -->
s) Obtaining ASE Status -->

x) Exit                ?) Help

Enter your choice:
```

The input is (m) for managing the ASE

2. Select service configuration:

```
Managing ASE Services

c) Service Configuration -->
r) Relocate a service
on) Set a service on line
off) Set a service off line
res) Restart a service
s) Display the status of a service
a) Advanced Utilities    -->

x) Exit to the Main Menu    ?) Help

Enter your choice [x]:
```

The input is (c) for service configuration

Instructions to Modify the TCR Available Server SAP R/3 Installation

3. Add a new service:

```
Service Configuration

a) Add a new service
m) Modify a service
d) Delete a service
s) Display the status of a service

x) Exit to Managing ASE Services  ?) Help

Enter your choice [x]:
```

The input is (a) for add a new service.

4. Select the service type:

```
Adding a service

Select the type of service:

1) NFS service
2) Disk service
3) User-defined service

x) Exit to Service Configuration  ?) Help

Enter your choice [2]:
```

The input is (2) for Disk service.

Instructions to Modify the TCR Available Server SAP R/3 Installation

5. Specify name of the disk service.

You are now adding a new disk service to ASE.

A disk service consists of a disk-based application and disk configuration that are failed over together. The disk configuration can include UFS filesystems, AdvFS filesets, LSM volumes, or raw disk information.

Disk Service Name

The name of a disk service must be a unique service name. Optionally, an IP address may be assigned to a disk service. In this case, the name must be a unique IP host name set up for this service and present in the local hosts database on all ASE members.

Enter the disk service name ('q' to quit):

The next state is the name of the disk-service. The virtual name of the Database Server was used as the name for the service (vtcr003).

Choose **No** if you are asked if you want to associate this name with an IP-address. This will be done in the start and stop scripts for the services.

6. Now you have to add all the disks which belong to the disk-service. Add all AdvFS filesets which belong to the shared buses.

Specifying Disk Information

Enter one or more device special files, AdvFS filesets, or LSM volumes to define the disk storage for this service.

For example: Device special file: /dev/rz3c
 AdvFS fileset: domain1#set1
 LSM volume: /dev/vol/dg1/vol01

To end the list, press the Return key at the prompt.

Enter a device special file, an AdvFS fileset, or an LSM volume as storage for this service (press Return to end):

You are then asked for the mount-points of these disks.

Do not insert any mount-points, this will also be done by the service scripts.

Answer the question with: **NONE**

If all filesets are added, answer the question with the Return key.

Instructions to Modify the TCR Available Server SAP R/3 Installation

7. Modify the user-defined scripts for the start and stop of this disk service.

Modifying user-defined scripts for `vtcr003`:

- 1) Start action
- 2) Stop action
- 3) Add action
- 4) Delete action

- x) Exit - done with changes

Enter your choice [x]: 1

The input is (1) for Start action.

8. In the menu Modifying user-defined scripts you can insert the start and stop action script.

Modifying the start action script for `vtcr003`:

- a) Add a start action script
-) Edit the start action script
-) Modify the start action script arguments []
-) Modify the start action script timeout [300]
-) Remove the start action script
- x) Exit - done with changes

Enter your choice [x]:

Choosing (a) for adding a start script.

Now you will be asked to enter the full pathname of your start action script or "default" for the default script (x to exit):

Choose **default** and press Return, enter the argument list for the start action script with **NONE** and set the time-out in seconds to **300**.

9. Now you can edit the start script and place the call of the service script:

```
a) RUNNING)          # This is a true stop of ${svcName}.
    ;;
    esac
    /usr/local/ase/db_serv.ksh      start
```

b) Choose (x) for exiting start action.

c) Choose (2) for entering stop action.

d) Choose (a) to add a stop script.

e) Now you will be asked to enter the full pathname of your start action script or "default".

Choose **default** and press Return, enter the argument list for the start action script with **NONE** and set the time-out in seconds to **300**.

f) Choose (e) to edit the stop action script.

Instructions to Modify the TCR Available Server SAP R/3 Installation

- g) Now you can edit the start script and place the call of the service script:

```
case "${MEMBER_STATE}" in
BOOTING)          # Stopping ${svcName} as ASE member boots.
    ;;
RUNNING)          # This is a true stop of ${svcName}.
    /usr/local/ase/db_serv.ksh stop
    ;;
esac
```

It is best to do a manual test first without the implementation of the start and stop scripts.

- h) Choose `Exit` to leave stop action.
- i) Choose `Exit` Modifying user-defined scripts.
- j) Decide if the service should be balanced, favored or restricted to a member system which is shown in the following menu.
- k) Choose `f` for Favor Members.
- l) Set the values `1, 2` in the next question.
- m) Choose `No` if ASE should relocate to a more highly favored member.
- n) Choose `Yes` to add the disk service.
- o) Check, if all devices will be mounted and Oracle and SAP R/3 will be started. If not, check the `log.*` files in `/usr/local/ase`.

How to Set Up a User-Defined Service

1. In the menu Adding a service create an user-defined service.

Adding a service

Select the type of service:

- 1) NFS service
- 2) Disk service
- 3) User-defined service

x) Exit to Service Configuration ?) Help

Enter your choice [1]:

Choose (3) for user-defined service.

2. State the name of the user-def-service. The virtual name of the Application Server was used as the name of the service (vtcr004).
3. Choose Start action.
4. Choose Edit the start action script.

Instructions to Modify the TCR Available Server SAP R/3 Installation

5. Place the start script as follows:

```
RUNNING)          # This is a true stop of ${svcName}.  
  ;;  
esac  
/usr/local/ase/as_serv.ksh start
```

6. Choose Exit.

7. Choose Stop action.

8. Choose edit the stop action script

```
case "${MEMBER_STATE}" in  
  BOOTING)          # Stopping ${svcName} as ASE member boots.  
    ;;  
  RUNNING)          # This is a true stop of ${svcName}.  
    /usr/local/ase/as_serv.ksh stop  
    ;;  
esac
```

9. Choose favored and set the following values: 2,1 in the next question.

10. Choose No, if ASE should relocate to a more highly favored member.

11. Under the menu service configuration you can insert the start and stop action scripts by choosing the following points:

- Modify service
- Service name
- Miscellaneous modifications
- User-defined action script

You have the possibility to choose the default option for the start and stop scripts, there.

12. You can insert your service script call.

Modifying user-defined scripts for `vtcr004`:

- 1) Start action
- 2) Stop action
- 3) Add action
- 4) Delete action

- x) Exit - done with changes

Enter your choice [x]:

Enter menu option (1) and (2) to implement the start and stop scripts.

How to Relocate the Services

The following describes the inputs to relocate the services:

1. Select to Relocate a Service:

```
Managing ASE Services

c) Service Configuration -->
r) Relocate a service
on) Set a service on line
off) Set a service off line
res) Restart a service
s) Display the status of a service
a) Advanced Utilities -->

x) Exit to the Main Menu      ?) Help

Enter your choice [x]:
```

The input is (r) for relocate.

2. Enter the service you want to relocate to the other machine. The service will stop first on the machine where it ran using the stop scripts. After this, if the service is a disk-service ASE dismounts the shared disks and starts to mount them on the target machine. Then the ASE software will activate the start scripts and the relocation will be successfully completed.

Transactions and Tests to Check the R/3 System

The transactions and tests to check the R/3 system are the following:

- SM59
Insert the virtual address into the server list.
With this transaction you can test the internal communication of the R/3 system.
- SM12
With this transaction you can check the enqueue and the update processes of your system:
For the enqueue server choose the menu point ZUSÄTZE -> DIAGNOSIS
and for the update choose ZUSÄTZE -> DIAGNOSIS IN UPDATE
Both transaction should run without any errors.
- SM51 shows the applications servers which are running in the system.
- SM21 shows the sys-log of the central and local system.
- SRZL Monitor => SAPoptions => startup/shutdown => control
- DPMON monitor for sap instance
- /usr/sap/<SID>/DVEBMGS/work => logfiles for the central instance
- /usr/sap/<SID>/D01/work => logfiles for the dialog instance

Modifications in R/3 to Prevent the Database Running on the Application Server to Pop-Up a Sapgui Error

If you test the relocation of the database and the central instance to the Application Server, you will see a popup after login on to a sapgui. The message says:

English: Your system is not correctly installed.
Deutsch: Ihr System ist nicht korrekt installiert

This is a new feature in the R/3 installation procedure. During installation, the executable R3SETUP creates a table in the database. This table is called INSTVERS and contains several information about the hostname, System ID, successful installation and so on. The logon to an instance will popup the message previously mentioned because the table INSTVERS does not contain the hostname of the Application Server. This table can be updated with the transaction SE16. How to do that is explained here:

1. Logon to R/3 as user "sap*"
2. Run transaction SE16

Table name: INSTVERS (Do not hit Return)

3. Choose menu button Table – Create entries

4. Fill in the following values:

ID	1 (or any other free ID)
R3RELEASE	40B
SAPSID	TCR (or the SID you chose)
SYSID	01 (or the System ID you chose)
OPSYS	OSF1
HOSTNAME	tcr004 (hostname of the Appl. Server)
DBHOSTNAME	tcr004 (hostname of the Appl. Server)
DBTYPE	ORACLE
DATUM	26.05.1998
TIME	14:35:11
ACTION	N
STATUS	0
TOOL	R3SETUP
INFO	"write any information you like"

If you want to see which ID is free, choose menu button Table- Table Contents

5. Save the settings and check them with menu button Table - Table Contents.
6. Check that the failover is working properly without any popup.

Testing R/3 Basic Functionality

The following are SAP R/3 Services:

- Message Service
- Gateway Services
- Enqueue Service
- Update Services
- Dialog Services
- Batch Services
- Spool Services

Testing Message and Gateway

A simple test for Message Server and Gateway can be done using transaction SM51. The first list that shows up is the Application Server list, sent by the Message Server. This list might look like this:

Name	Host	Type
vtrcr003_TCR_00	vtrcr003	Dialog Update Lock Batch Spool
vtrcr004_TCR_01	vtrcr004	Dialog

If you double-click on a specific Application Server you'll get the list of work processes running on this Application Server. This list is received via the local gateway of your Application Server (the Application Server you're working on).

The Message-Server is also involved in the communication between Dialog WP and Enqueue Server (when locking objects) and Dialog WP and Update WP (when updating asynchronously). This can be tested by processing a business transaction (see Application Tests).

The Gateway Services are also used for CPIC and RFC Connections. The RFC table for logical destinations can be configured with transaction SM59. There is one destination for each system, defined as a R/3 Connection and one internal connection for each Application Server. Currently you might see a connection time-out when using the relocatable hostname for the R/3 Connection to your system. This is caused by your IP address and subnetwork settings. The gateway will not pick up the relocatable address as a local gateway address. A workaround is setting the parameter gw/netstat to space and gw_alternative_hostnames to your relocatable hostname. You will see the error in the gateway trace file (/usr/sap/SID/DVEBMGS00/work/dev_rd). To test a RFC from the operating system level use startRFC.

If the RFC package has been installed, it is located in /sapmnt/SID/exe/rfcsdk/bin. To test a RFC, issue the following command:

```
startRFC -d <SID> -u <USER> -p <PASS> -t -3 -h <host> -g  
<gwhost>  
-x sapgw00 -F RFC_CUSTOMER_GET -E KUNNR=\* -T  
CUSTOMER_T,197,w=result
```

where : <SID> is your SAP SYSTEM ID
<USER> is a valid user in client 000
<PASS> is the users password
<host> is the Application Server host (can be relocatable)
<gwhost> is your Gateway host (can be relocatable, see comment above)

On success the command will not display anything. The file 'result' will contain the values returned by the call. On error, the file dev_rfc will contain trace information.

Testing Enqueue and Update Services

The Enqueue WP locks objects on the SAP level using a enqueue table in main memory. A basic function test for locking can be done in transaction SM12.

1. On the first screen select Details-Diagnosis to test locking.
2. Select Details-Diagnosis to test locking with unlocking RSENQ* in the Program field and press Return.
3. Press F8 on the following screen and double click on the line 'Display programs regardless of class'. You'll get a list of ABAP's for lock testing.
4. RSENQTS1 creates a number of enqueue entries. Call the ABAP, enter a number of lock entries.

Instructions to Modify the TCR Available Server SAP R/3 Installation

5. Press F8 to execute. You'll get a list with all the enqueue requests. Do not exit this screen.
6. Open another session (System>Create Session) and call SM12 (press Return). You'll see the locks which have been created by RSENQTS1.RSENQTS1 tests the abortion of an asynchronous update.
7. Open 3 sessions.
8. Call SM12 (and press Return) in session #1.
9. Call SM50 in session #2 and SE38 in session #3.
Session #1 will show no lock entries.
10. Call the ABAP RSENQTS1 in session #3.
Session #1 will show 1 lock entry for your Dialog WP and 1 lock entry, as soon as the Update WP is running. For 20 seconds you will see a program running in an Update WP. After 20 seconds the Update process creates an error. The 2 lock entries should disappear and you'll get a popup information telling that you received an express document (parameter rdisp/vbmail must be set to 1).
11. Go to your inbox (on the main menu select Office-Inbox). Transaction SM13 now shows a terminated update request (call SM13 and hit return). Delete the request by selecting Administration>Delete-All records.

Further lock tests can be done using online transaction (see Application Tests).

Batch Services

This section is a description how to handle the Batch Service.

1. Schedule a batch job with transaction SM36:
 - a) Enter a jobname and a job class (A, B or C). Do not enter a destination.
 - b) Press button Steps and enter the program to run (for example RSPFFPAR shows parameter values).
 - c) Save the step and go back.
 - d) Press Start time to select a start time and date.
 - e) Select Immediate and Periodic job.
 - f) Enter the execution interval for the job (Period values). Choose a small interval (for example 5 minutes).
 - g) Save the Start time settings and return to the entry screen of transaction SM36.
When saving the job now, it will be scheduled to run immediately. The Application Server that processes the batch job will be the one configured with the parameter rdisp/btcname (usually in DEFAULT.PFL).
After the job finished it is scheduled to run again after the defined interval. It then will be scheduled on an Application Server that has a free Batch-WP (for example time based scheduled batch jobs are distributed over all Batch Services).
2. To get an overview of all scheduled batch jobs, use transaction SM37.
 - a) Enter the jobname and press F8. The list shows jobs in their current state.
 - b) To see additional information on a job, place the cursor on the line and select Display.
 - c) On the following screen select Detailed information. This will also show the host and the WP number the process ran on.

Instructions to Modify the TCR Available Server SAP R/3 Installation

3. Trigger a batch job from outside by raising an event.
 - a) First, define an appropriate event using transaction SM62.
 - b) Schedule a job with transaction SM36 to be run when the event occurs (this is defined as 'Start time').
 - c) Raise the event by calling sapevt (located in /sapmnt/SID/exe).
sapevt USER_EVENT -t pf=/sapmnt/SID/profile/DEFAULT.PFL

Testing Spool Service

To test the Spool Service you have to define a printer with transaction SPAD, do the following:

1. Select Printer names and press Return.
2. Select Create to define a printer.
3. Enter the name of your printer and select the appropriate device type. For now you have to use the stationary hostname for the printer. Host printer is the name of the printer as configured on OS level. Access method can be L or N for locally defined printers (local to the Spool WP).
4. Save your printer and return to the list of printers.
5. To test printing, select Output device-Print list. Enter the name of your printer and select Print immed.
6. To view the status of your print request, call transaction SP01 (Spool overview).
 - a) Press Enter to get the list of your spool requests.
 - b) Select Output requests to see the status of the corresponding print request.
 - c) Choose Display log to view a log for your print request (if available).

Application Tests

To test Dialog, Message, Enqueue and Update Service in a business transaction, do the following (all tests should be done in a client copied from client 000):

1. Material Master Change
 - a) Call Transaction MM02.
 - b) Enter Material `DE-HAWA1'.
 - c) Select views (for example Engineering and Purchasing) and press Return.
 - d) Enter Plant `0001' and press Return.
 - e) You should see a lock entry with transaction SM12.
 - f) Change material (for example gross weight on Engineering view).
 - g) Press Return on all screens in MM02.
 - h) On the last screen press Yes to save your changes.
The material is updated asynchronously.
The lock entries should disappear after the material has been changed.
2. Customer Master Change.
3. Call transaction FD02.
4. Enter Customer AT002.

Instructions to Modify the TCR Available Server SAP R/3 Installation

5. Enter Company Code AT01.
6. Select views (address) and press Return.
7. Change Address and save.
8. Create Purchase Order.
9. Call Transaction ME21.
10. Enter Vendor DE-001.
11. Enter Purchasing Organization 0001.
12. Enter Purchasing group 001 and press Return.
13. Enter Material DE-HAWA1.
14. Enter quantity 10.
15. Enter net price 100.
16. Enter a delivery date.
17. Enter Plant 0001.
18. If you get warnings (a W-..... message): press Return again.
19. Save at the end.
20. Check purchase order with transaction ME23.

Additional Tests

The following sections are descriptions of additional tests for the Transport System and the DECsafe Configuration.

Test the Transport System

Create a transport request using transaction SE01:

1. Switch on radio button `Transport request` and press button `Create`. Enter a Short text like `Test transport system(general system check)`.
2. Select T as transport type. If you do not have a dummy system set up, use transaction SE06 (as user DDIC) to set one up.
3. In SE06, select `Change`. Confirm the popup with `Yes`. On the following screen select `Continue` and confirm again with `Yes`.
4. Enter you system parameters (name=DUM) and save.
5. Back in SE01, use system DUM as target system. Save your transport request and press `Editor`.
6. To document your tests use the test evaluation list in the attachment.

Testing the Installed DECsafe Configuration

1. Use the `asemgr` to check out service and member status.
2. Make sure that all members see all shared disks (UERF and `show device` at the console).
3. Make sure each member can ping each other member over the primary network.
4. Make sure that the ase logging is turned on.
5. Read through the `/var/adm/syslog.dated/"DATE"/daemon.log` file to find problems concerning ASE.
6. If the normal logging level does not show any problems, try to switch the mode to informational. This will log much more information, but too much for normal control.
7. Resetting ASE will restart all the ASE daemons. This is sometimes helpful.
8. The command `assetup -f set ASE database back` to the installation state, then you have to reconfigure your services.

Debugging and Diagnosing Tool `sys_check`

This tool is very useful for debugging and diagnosing system problems, it is an ksh shell script that generates a .HTML file as output. This output file shows information and system configuration data.

The scripts and patches resides on the WWW page:
http://www-unix.zk3.dec.com/tuning/tools/sys_check or
`UNIX:://public/tuning/tools/sys_check`

Usage:

- # `sys_check -help` -prints the basic help message.
- # `sys_check -all` -performs all output functions

`sys_check` is run as root, redirecting the output to a file.

To run the command:

- # `./sys_check > file.html`

Tasks for Setting the ASE Off-line

To set the ASE off line, you have to remove the ASE start daemons from the `/sbin/rc3.d` directory on each system in the cluster.

- # `cd /sbin/rc3.d`

and rename the startup scripts for the ASE daemons:

- # `mv S00aseam .S00aseam`
- # `mv S14asemember .S14asemember`

From now on the ASE daemons will not be started.

Restrictions and Limitations

The following are limitations of the ASE environment:

- The power down of the Application Server and automatic relocation proceed with problems.
First it seems to be OK, because all processes of both instances are running, but if you try to work with the Application Server instance you do not have the connection to the message server. In the `/usr/sap/<SID>/DV*/work/disp_ms` file following error occurs:
`client not unique`
This means that the message server did not recognize that the Application Server is down, and the Application Server wants to connect the message server a second time. This is not a failure from DECSafe but from SAP. The same situation will occur in a none DECSafe environment when the Application Server crashes and will be restarted. To solve this problem all Application Server and the central instance has to be restarted. (SAP knows this problem and it is fixed in version 30E and higher).
- The failure of the enqueue service (for example due to failure of its machine) currently forces a restart of all application hosts. If the enqueue service fails, all locks for transactions which have not yet been committed are lost, for example, the user context in such a transaction is lost. To avoid the problem of multiple users then concurrently updating the same data after the enqueue service has been restarted, all running transactions would need to be reset. The only solution here is to stop the host supporting the R/3 instance, aborting running transactions. SAP is currently working on a solution to avoid the restart of all R/3 instances.
Using the DB-Reconnect worked fine when the enqueue server did not fail. If the enqueue server fails, all Application Servers have to be restarted. Failing of the enqueue server may result an inconsistency of the database, if the Application Servers resides in the db-reconnect status. This is a problem that has to be fixed from SAP.
- Released and running spool jobs has to be restarted after a failover, you cannot continue the aborted jobs, the same problem occurs with running and released batch jobs.
- The DB-reconnect is not released for a consistent lock management of the database, because if the enqueue server fails, there are no lock entries on the tables anymore and this means, that the exclusive write to one table of a transaction is disabled. In this case, the workprocesses of all instances has to be restarted again.
There is the situation where CI/DB run together on one host machine and DB-reconnect may be used by external AS. This works only for R/3 release 30E, and must be ensured that the external AS re-attaches the CI/DB instances within 300 sec after CI restart. Consequently, DB reconnect can be used if DB and CI are split on different hosts and only the DB host fails. In this case, the CI and any external AS can use DB-reconnect. Because of insecure restart behaviors, the DB-reconnect should not be used right now. This is a recommendation from SAP.

Parameter in the rc_serv.conf file

The following example shows the /usr/local/ase/rc_serv.conf file which is delivered with the service package.

Only the parts that have to be changed were printed.

```
#####
#
# 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 parameters 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=vtcr003                 # name of the disk-database service in the asemgr
DB1_ALIAS_HOST=vtcr003           # virtual name of the db and central instance for net-
                                # adapter 1
DB1_INSTANCE=DVEBMGS00           # instance name of the central instance
DB1_ALIAS2_HOST=vtcr003e         # 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($0,n-
1)}'`
DB1_STARTSAP="startsap_vtcr003_00" # start script for the central instance
DB1_STOPSAP="stopsap_vtcr003_00"   # stop script for the central instance
TIMEOUT=200                        # timeframe in sec. to stop the R/3 system +
                                # database
DB1_START_LOG="${ASEDIR}/log.start_db1_service # logfile for the start action
DB1_STOP_LOG="${ASEDIR}/log.stop_db1_service" # logfile for the stop action
#####
#
# Global parameters 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.
```


Instructions to Modify the TCR Available Server SAP R/3 Installation

```
#####
AS1_SERV=vtcr004           # name of the user-defined-application service in the asemgr
AS1_ALIAS_HOST=vtcr004    # virtual name of the appl. server instance for net-adapter 1
AS1_ALIAS2_HOST=vtcr004e  # virtual name of the appl. server instance for net-adapter 2
AS1_INSTANCE=D01         # instance name 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($0,n-1)}'`
AS1_STARTSAP="startsap_vtcr004_01" # start script for the appl. server instance
AS1_STOPSAP="stopsap_vtcr004_01"   # stop script for the appl. server instance
AS1_START_LOG="${ASEDIR}/log.start_as1_service" # logfile for the start action
AS1_STOP_LOG="${ASEDIR}/log.stop_as1_service"   # logfile for the stop action
#####
```

more text ...

```
#####
DB_TYPE=ORACLE
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 \
SAP1_S1#sapdata1         /oracle/TCR/sapdata1 \
SAP1_S2#sapdata2         /oracle/TCR/sapdata2 \
SAP1_S3#sapdata3         /oracle/TCR/sapdata3 \
SAP2_S1#sapdata4         /oracle/TCR/sapdata4 \
SAP2_S2#sapdata5         /oracle/TCR/sapdata5 \
SAP2_S3#sapdata6         /oracle/TCR/sapdata6 \
“
```

Note

SAP2_S1 to SAP2_S3 are only used for the second ESA10000.
With one ESA10000 you will have SAP1_S1 to SAP1_S6.
