digital

SAS Oracle
DIGITAL UNIX AlphaServer 4100

DIGITAL HiTest Notes

Part Number: EK-HSSUA-HN. C01

August 1998

Revision/Update Information: This is a revised document.

Digital Equipment Corporation
Maynard, Massachusetts

August 1998

Table of Contents

5 Tests and Results

6 Problems and Solutions

7 Detailed Hardware Configuration

Figures

Tables

Preface

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

Audience

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

Organization

This document is organized as follows:

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

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-HSSUA-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`, then select Technical Support (Internet)

Related Documents

This document references the following manuals:

- *AlphaServer 4100 Configuration and Installation Guide*
- *AlphaServer 4100 System Drawer User's Guide*
- *Getting Started: HSZ70 Solutions Software V7.0B for DIGITAL UNIX Installation Guide*
- *CLI Reference Manual*
- *DIGITAL UNIX Installation Guide for Version 4.0D*
- *DIGITAL UNIX System Configuration and Tuning Guide*
- *DIGITAL UNIX Version 4.0D Release Notes*
- *Networker Save and Restore for DIGITAL UNIX Installation Guide*
- *Networker Save and Restore for DIGITAL UNIX Administrator's Guid*
- *Oracle8 Server for DIGITAL UNIX Installation and Configuration Guide*
- *Oracle for UNIX Performance Tuning Tips*
- *Installation Instructions for the SAS System for DIGITAL UNIX Release 6.12*
- *ServerWORKS Manager Overview and Installation* Guide

1
Advantages of DIGITAL HiTest Suites

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

What Is a DIGITAL HiTest Suite?

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

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 SAS Oracle DIGITAL UNIX AlphaServer 4100 HiTest Suite contains three groups of components: the *DIGITAL HiTest Foundation*, the *DIGITAL HiTest AppSet* and the *DIGITAL 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. The System Management Station is an optional standalone personal computer system containing software used to manage the HiTest system.

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

Additional Hardware and Software

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

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

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

2
About This DIGITAL HiTest Suite

The SAS Oracle DIGITAL UNIX AlphaServer 4100 HiTest Suite provides a combination of industry-leading database technology, a robust operating system, and a fast, reliable hardware platform. This suite uses Oracle8 for high-speed query-intensive data warehousing and the DIGITAL AlphaServer 4000/4100, a scalable, dependable, open system that delivers high productivity at an affordable price.

The testing of this HiTest suite focused primarily on the interoperability of Oracle, DIGITAL UNIX with SAS, and the AlphaServer 4100.

The SAS Oracle DIGITAL UNIX AlphaServer 4100 HiTest Suite includes the following components:

- SAS
- Oracle8 Server
- DIGITAL UNIX
- AlphaServer 4100
- StorageWorks ESA 10000

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

- Availability
- Installability
- Interoperability
- Manageability
- Scalability
- Services
- Year 2000 Compliance

Availability

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

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

Features of SAS Oracle DIGITAL UNIX AlphaServer 4100 HiTest Suite

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: SAS Oracle DIGITAL UNIX AlphaServer 4100 Availability Features

Availability Feature	Enabling Technology	Always Included	Customer Optional
Data Protection	Backup and restore	Yes	
Data Availability	Redundant disk storage (RAID)		Yes*
	Redundant power for disk storage		Yes
	Redundant controllers for fault tolerance	Yes	
Platform Availability	Redundant power for system	Yes	
*RAID is customer optional if you require redundancy.			

Recommendations for SAS Oracle DIGITAL UNIX AlphaServer 4100 HiTest Suite

This HiTest Suite provides high availability while maintaining high performance because hardware with a high mean-time-between-failure (MTBF) is used. Use of fault-tolerant and redundant hardware avoids single points of failure, and 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. Three small tabletop tape systems were chosen providing up to 192 GB of compressed storage.
- Data Availability
  - Data availability is implemented with RAID 5 striping. Allocate spareset units to ensure recovery of stripeset members.
  - The HSZ70s include a redundant controller that ensure fault tolerance (read-only mode).
- Platform Availability
  - Redundant power supplies are used to avoid a single point of failure.

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 is protected over time. 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.

---

[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.

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

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

RAID Levels 2 and 4 are not in general use.

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

Figure 2-1: RAID Level Summary

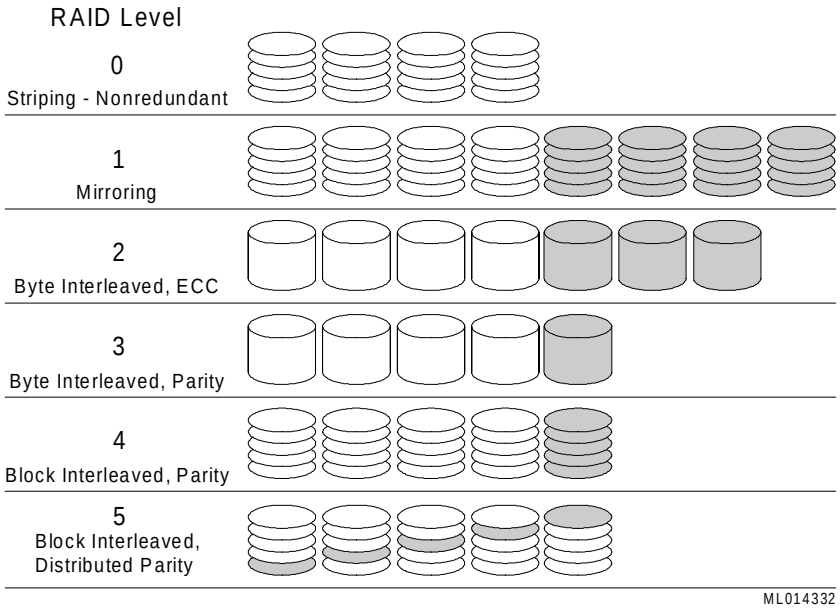

DIGITAL and other companies also use the terms RAID 0+1, RAID 1+5 and Adaptive 3/5 to refer to the combinations of these, and other, storage technologies. Table 2-2 describes the RAID types to consider when choosing a RAID configuration.

Table 2-2: RAID Levels and Descriptions

RAID Level	Description	Advantages/Disadvantages
0	• Striping • Data segmented and distributed across several disks	+ increase in performance due to parallelism in read and write - no fault tolerance (*not* a high availability solution)
1	• Hardware Mirroring • Data written twice to different disk spindles within the disk array	+ good performance in read-intensive applications (data can be read in parallel from several disks) - slower in writes (multiple writes required) - spindle costs doubled
0+1	• Striped Mirroring • Combined level 0 and 1 • Data mirrored onto and striped across several disks	+ good performance in reads (RAID 1) + write performance improved versus RAID 1 due to parallelism + adequate response maintained in event of disk

RAID Level	Description	Advantages/Disadvantages
	• Best for performance-critical, fault-tolerant environments	failure - spindle costs doubled - recovery is I/O intensive
2	• Parallel access array • Striped • ECC on separate drives	+ high data transfer rate + ECC detects and corrects errors - low I/O request rate - not appropriate with modern drives
3	• Parallel access array • Small minimum chunk size • Check bit calculated from data • Parity bits on dedicated disk, data striped across remaining disks	+ good performance in reads due to parallelism (like RAID 0) + costs only slightly increased compared to disks without high availability solutions + good performance with long records (high data transfer rate) - write performance penalty due to check bit calculation - cannot overlap I/O (low I/O request rate)
4	• Independent access array • Parity disk	+ processes multiple requests simultaneously - parity disk is a bottleneck on writes
5	• Independent access array • Parity Bit • Check bit and data distributed (striped) across multiple disks • Best in environments that are mostly read and are not performance sensitive	+ good performance in reads due to parallelism (like RAID 0) + costs only slightly increased compared to disks without high availability solutions + overlapped I/O - write performance penalty due to check bit calculation
1+5	• RAID 5 combined with mirroring • Mirroring provided by LSM or Volume Shadowing • Most reliable and highest performance solution	+ good performance in reads due to parallelism (like RAID 0) + double redundancy makes disk failure barely noticeable - spindle costs more than double - write performance penalty due to check bit calculation
Adaptive 3/5	• The best features of 3 and 5 • Adapts between Level 3 and Level 5 in response to changes in the application's workload	+ good performance in reads due to parallelism (like RAID 0) + costs only slightly increased compared to disks without high availability solutions + performs well with a wide variety of I/O loads even when load characteristics change minute by minute. - write performance penalty due to check bit calculation

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. When these procedures and the product installation guides are followed, no problems should be encountered. Chapter 6 describes some configuration problems that may be encountered if the installation procedures are not followed.

Within the HiTest environment, after removing the system from the shipping skid, it required approximately 8 hours to install and configure the hardware for the maximum configuration. Installation and configuration of the software took approximately 8 hours. Additional time may be required for any reconfiguration. Expect installation times to vary significantly in other environments depending on factors such as the expertise of the installer and the environment in which the installation occurs.

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.

There were no major interoperability issues discovered during the testing of this suite.

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 2-2, 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. Provides the data from which this graph is derived.

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 2-2: SAS Oracle DIGITAL UNIX AlphaServer 4100 HiTest Suite Scalability

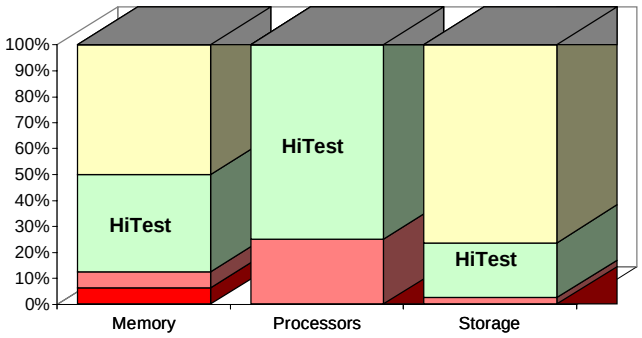

	Systems configured within this range involve tradeoffs between these, and sometimes other components.
	Systems configured within this range generally exceed the requirements of the market.
HiTest	Systems configured within this range meet the requirements of most customers.
	Systems configured within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

Other HiTest Suites support the use of SAS and Oracle on alternative configurations, including DIGITAL's AlphaServer 8400, clustered configurations and OpenVMS.

Services

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

Proof of Commitment: The DIGITAL Uptime Guarantee

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

Portfolio of Business Critical Services

The three vital elements of DIGITAL Business Critical Services are:

- Availability Review

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

- Business Critical Gold Support

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

  – Notification of software patches as soon as they become available

  – Notification of known problems and fixes

  – Monthly service activity review

  – Operating system upgrade impact planning

  – Bi-annual System Healthcheck assessments. These are conducted using advanced system-based tools to assess the performance and security of systems. The collected data is analyzed against accepted practices, and the findings, together with recommendations for corrective action, are documented in a summary report.

- Availability Partnership

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

  – Configuration and topology documentation

  – Availability status reporting

  – Change impact analysis

  – Proactive problem avoidance based on proactive patch/FCO/firmware management

  – Periodic detailed data collection and analysis

  – Availability model update

  – Contingency planning

  – Service planning and advising

Complementary Support Services

The three key Business Critical Services are augmented by:

- On-Site Parts Service

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

- Installation and Startup

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

Meeting Client Needs Locally or Globally

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

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

For More Information

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

Year 2000 Compliance

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

While HiTest does not explicitly test for Year 2000 compliance in the components of this Suite, HiTest does check the published status of components where Year 2000 compliance would be a concern. The Year 2000 information presented here is accurate as of June 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: SAS Oracle DIGITAL UNIX AlphaServer 4100 Year 2000 Compliance

Component	Year 2000 Status
SAS System 6.12	Blue
Oracle8 for DIGITAL UNIX	Blue
TruCluster Available Server for DIGITAL UNIX	Blue
DIGITAL UNIX	Blue
AlphaServer 4100	Blue
StorageWorks ESA 10000	Blue

3
Configuration Data

This chapter describes the SAS Oracle DIGITAL UNIX AlphaServer 4100 DIGITAL 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 SAS Oracle DIGITAL UNIX AlphaServer 4100 HiTest Suite. These two tables form the DIGITAL HiTest Template.

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

Table 3-4 lists the revision levels of the components. To reduce the chance of interoperability issues, order the exact software version listed in the HiTest Templates. Hardware components can be ordered as listed, and the system installer should ensure that the hardware components version level is the same version or later than those tested.

The DIGITAL HiTest Template consists of three categories:

- AppSet Software – Includes software specific to one class of customer solutions, in this case SAS
- Foundation Hardware – Includes the base system, storage, and other hardware options
- Foundation Software – Includes the operating system 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 – SAS Oracle UNIX AlphaServer 4100

SAS HiTest AppSet Oracle UNIX AlphaServer 4100 HiTest Foundation Hardware				
For documentation and updates: http://cosmo.tay.dec.com and http://www.businesslink.digital.com				
Line Item	**Description**	**Part Number**	**HiTest Range**	
			Min	**Max**
AppSet Software				
1	**The SAS System, Release 6.12** Contact SAS at: `http://www.sas.com` **Note:** This AppSet is not required when the foundation hardware and software is ordered for use with a non-HiTest application.	**SAS Institute**	1	1
Foundation Hardware				
2	*Select one base system*: **AlphaServer 4100 5/533 DIGITAL UNIX 1 GB** **AlphaServer 4000 5/533 DIGITAL UNIX 1 GB** *Hardware includes*: • 5/533 MHz CPU with 4 MB cache • 1 GB Memory • PB2GA-JC TRIO64 1 MB Graphics • DE500-AA 10/100 Mbit Fast Ethernet • KZPBA-CA UWSE SCSI and BN38C-02 cable • SCSI CD-ROM drive • RX23L-AB 1.44 MB Floppy drive • LK47W-A2 PS/2 style keyboard • Three-button PS/2 compatible mouse *Software includes*: • DIGITAL UNIX Operating System base license • Unlimited User license • Server Extension license • Internet AlphaServer Admin Software • ServerWORKS • DECevent • BMC Patrol Agent • OnliNet Power Management	**DA-51KAC-FB** **DA-53KEC-FA**	1	1
3	*Select a pedestal or cabinet enclosure:* **Pedestal Enclosure** **Cabinet Enclosure**	**BA30P-AA** ❶ **H9A10-EF**	1	1
4	*Select a pedestal or cabinet UltraSCSI shelf:* **Single-bus UltraSCSI shelf for Pedestal** **Single-bus UltraSCSI shelf for Cabinet**	**BA36R-SD** **BA36R-RC**	1	1
5	**4.3 GB 7200 RPM UltraSCSI disk**	**DS-RZ1CB-VW**	1	1
6	*Order one kit for each system drawer in a cabinet enclosure:* **System Drawer Mounting Kit**	**CK-BA30A-BA** ❶	1	1
7	**533 MHz CPU DIGITAL UNIX SMP UPG** **Note:** The 4000 holds a maximum of two CPUs.	**KN305-DB**	1	3
8	**450 Watt Power Supply**	**H7291-AA**	0	2
9	**1 GB Memory Option** **2 GB Memory Option** **Note:** This HiTest Template supports a memory range from 1 GB to 4 GB per system. When selecting memory options, stay within the Template's 4 GB maximum. The 4000 holds a maximum of two memory options.	**MS330-FA** **MS330-GA**	1	See Note
❶ Indicates that geography-specific part number variants are available. Check the appropriate price book for details.				

SAS HiTest AppSet
Oracle UNIX AlphaServer 4100 HiTest Foundation Hardware

For documentation and updates: http://cosmo.tay.dec.com and http://www.businesslink.digital.com

Line Item	Description	Part Number	HiTest Range	
			Min	Max
10	**PCI one-port FWD SCSI Controller, HD64**	**KZPBA-CB**	1	3
Note: Line items 11 and 12 show the storage arrays. Select line item 11 for the minimum configuration and line item 12 for maximum configuration.				
11	**Enterprise Storage Array 10000** *Hardware includes*: • Dual DS-HSZ70-AH (1) • BN37A-10 (1), BN38E-0B (1) • DS-BA370-AA (1)	**DS-SWXES-DA** ❶	1	0
12	**Enterprise Storage Array 10000** *Hardware includes*: • Dual DS-HSZ70-AH (1) • BN37A-10 (2), BN38E-0B (2) • DS-BA370-AA (2)	**DS-SWXES-AA** ❶	0	3
13	**4.3 GB UltraWide SCSI disks**	**DS-RZ1CB-VW**	15	144
14	**PCI to FDDI adapter SAS, MMF**	**DEFPA-AB**	1	1
15	**20-m SC to SC dual fiber optic cable**	**BN34B-20**	1	1
16	**10BaseT Twisted-Pair Ethernet cable**	**BN25G-07**	1	1
17	*Select one high-resolution color monitor:* **15-in Flat-square with 0.28 dot pitch** **17-in Trinitron aperture grille, 0.26mm** **21-in Diamondtron aperture grille, 0.30mm**	**SN-VRCX5-WA** ❶ **SN-VRCX7-WA** ❶ **SN-VRCX1-WA** ❶	1	1
18	**32/64 GB DAT autoloader, tabletop** Includes BN23G-01 HD50 to LD50 SCSI cable	**TLZ9L-DB**	3	3
19	*Order one adapter for the TLZ9L-DB* **PCI SCSI Controller FNSE**	**KZPAA-AA**	1	1
20	*Order cables to interconnect the tape loaders.* **SCSI Cable, 3 meter**	**BC19J-03**	2	2
21	**System Management Station**	**See Table 3-3**	0	1
❶ Indicates that geography-specific part number variants are available. Check the appropriate price book for details.				

Table 3-2: DIGITAL HiTest Template – Foundation Software

SAS Oracle UNIX AlphaServer 4100 HiTest Foundation Software						
For documentation and updates: http://cosmo.tay.dec.com and http://www.businesslink.digital.com						
Line Item	**Description**	**Part Number**	**HiTest Range**		**Required By**	
			Min	**Max**	**Fnd†**	**App†**
Foundation Software						
1	**DIGITAL UNIX Operating System, Version 4.0D**	**Included with item 2 of Table 3-1**	1	1	Yes	Yes
2	**DIGITAL UNIX Operating System, V4.0D Software Media Kit**	**QA-MT4AA-H8**	1	1	Yes	Yes
3	**DIGITAL UNIX Server Extensions**	**Included with item 2 of Table 3-1**	1	1	Yes	Yes
4	**ServerWORKS**	**Included with item 2 of Table 3-1**	1	1	Yes	Yes
5	**DIGITAL CD-ROM Software Library for DIGITAL UNIX Systems layered products**	**QA-054AA-H8**	1	1	Yes	Yes
6	**Oracle8 Server for DIGITAL UNIX V8.0.3, including:** Oracle Server Manager 3.0.3 Parallel Query Option 8.0.3 PL/SQL 8.0.3 Oracle TCP/IP Protocol Adapter 8.0.3 SQL*Plus 8.0.3		1	1	Yes	Yes
7	**HSZ70 Solutions Software for DIGITAL UNIX includes:** • PC Card containing the storage controller software • StorageWorks Command Console • Licenses and documentation	**QB-5SBAB-SA**	1	1	Yes	Yes
8	**HSZ70 Solutions Software for DIGITAL UNIX** including everything listed in item 7, except documentation.	**QB-5SBAB-SB**	1	5	Yes	Yes
9	**StorageWorks PLUS, which includes:** • Networker Save and Restore for DIGITAL UNIX V4.3 • DIGITAL UNIX Logical Storage Manager (LSM) • AdvFS Utilities	**QB-5RYAG-AA**	1	1	Yes	Yes
10	**NetWorker Save and Restore Database Module for Oracle, Version 2.0 license** Media and doc kit part number is QA-3P4AA-H8	**QL-3P4AG-AA**	1	1		
11	*Order one license when ordering a TLZ9L-DB* **NSR Jukebox Tier 1 License (1–16 slots)**	**QL-04UAL-3B**	1	3	Opt'l	Opt'l
12	**SCSI CAM Layered Components for DIGITAL UNIX** Right to use is included with item 8	**Included with item 8**	1	1	Yes	Yes
13	**Hard copy of this Suite's HiTest Notes**	**EK-HSSUA-HN**	1	1	Yes	Yes
†Fnd = Foundation, App = AppSet						

Table 3-3: System Management Station Template

SAS Oracle HiTest AppSet System Management Station				
For documentation and updates: http://cosmo.tay.dec.com and http://www.partner.digital.com:9003/cgi-bin/comet				
Line Item	**Description**	**Part Number**	**HiTest Range**	
			Min	**Max**
Management Station Hardware				
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.				
1	**DIGITAL PC 5100 ST System** *Hardware includes*: • 233 MHz Pentium CPU with MMX • 512 KB secondary cache • 32 MB memory • Integrated Fast Ethernet (10/100) • 24X CD-ROM • 3D Video • 3.2 GB disk drive • 1.44 MB floppy • Integrated Audio • Country Kit, North American *Software includes*: • Windows NT Workstation 4.0 (factory installed) **Note:** A functionally equivalent 80X86 system may be substituted without invalidating this HiTest Template.	**FR-DAB04-AJ**	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	*Select one high-resolution monitor:* **21” (19.6” view) 1600 x 1200 @ 75 Hz** **19” (18” view) 1600 x 1200 @ 75 Hz** **17” (16” view) 1280 x 1024 @ 75 Hz**	 **FR-PCXAV-WZ** **FR-PCXAV-TZ** **FR-PCXAV-YZ**	1	1
Software Installed on Management Station				
5	**Windows NT Workstation 4.0**	**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 (800) 360-7561	**Microsoft**	1	1
7	**DIGITAL ServerWORKS Manager, Version 3.0**	**Included with the base system**	1	1
8	**StorageWorks Command Console, Version 2.0**	**Included with StorageWorks kit**	1	1
Software Installed on Managed System(s)				
9	**Base UNIX systems management tools**	**Included with DIGITAL UNIX**	1	1
10	**DIGITAL UNIX Management Agent**	**Included with item 7**	1	1
11	**StorageWorks Command Console Agent**	**Included with item 8**	1	1

Table 3-4: Component Revision Levels

Hardware Component	Hardware	Firmware
FWD SCSI controller (KZPBA-CB)	A01	5.53
RAID Array controller (HSZ70-AH)	H01	HSOF V70Z-0
4.3 GB disks (DS-RZ1CB-VW)	A01	LYJ0
533 MHz CPU (KN305-DB)	A03	–
Memory (MS330-FA)	A03	–
Fast Ethernet adapter (DE500-AA)	2.0	–
PCI to FDDI Adapter (DEFPA-AB)	B02	V3.1
Software Component	**Version**	**Patch Level**
DIGITAL UNIX	4.0D	–
DIGITAL UNIX Server Extensions	–	–
Oracle8 Server	8.0.3	–
SQL*Plus	8.0.3	–
PL/SQL	8.0.3	–
Networker Save and Restore for DIGITAL UNIX	4.3	–
Logical Storage Manager (LSM)	4.0D	–
AdvFS Utility	4.0D	–
SCSI CAM Layered Components for DIGITAL UNIX	3.1C	–
System Management Station Software Component	**Version**	**Patch Level**
Windows NT Workstation	4.0	–
DIGITAL ServerWORKS Manager	3.0	–
StorageWorks Command Console	2.0	–

Special Configuration Rules

There are no special configuration rules for this HiTest suite. For specific information regarding configuring components, refer to the appropriate document listed in the Preface.

4
System Installation and Setup

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

Hardware Installation

The *AlphaServer 4100 Configuration and Installation Guide* provides instructions for installing your AlphaServer 4100 system. Chapter 7 of this HiTest Note provides additional information specific to the minimum and maximum configurations of this HiTest Suite.

RAID Controller Installation

Disk storage is configured using the Command Line Interpreter (CLI) to set up the RAID Array Controller and for configuring the disk storage. Refer to *the HSZ70 Solutions Software V7.0B for DIGITAL UNIX Installation Guide* for detailed information.

HSZ70 Array Controller Setup

The HSZ70 controllers can be configured using either a maintenance terminal with the Command Line Interpreter (CLI) or a graphical user interface available through the StorageWorks Command Console (SWCC). For instructions on installing the StorageWorks Command Console on a system management station, see the section System Management Station Installation later in this chapter. While the CLI provides the most detailed level of subsystem control, the StorageWorks Command Console replicates most of the functions available within the CLI in graphic form and provides a user-friendly method of executing CLI commands.

Before you can use the StorageWorks Command Console, you must complete the following installation procedures:

- Install the DIGITAL UNIX operating system on the AlphaServer
- Install StorageWorks Command Console on the optional System Management Station specified in Table 3-3
- Install the StorageWorks and HSZ70 Agents on the AlphaServer
- Set up network connections between the AlphaServer and the System Management Station

If you choose not to include a management station, you can access the CLI by connecting a maintenance terminal to the maintenance port on the front of the controller. After you have initially configured the controller, making it visible to the host, you can perform all other configuration tasks through a remote connection.

The following procedure describes how to set up the HSZ70 Array Controllers using CLI:

1. Set the number of SCSI target IDs (0 through 3):

   ```
   CLI > SET THIS_CONTROLLER ID = (0, 1, 2, 3)
   ```

2. Set failover to place the controllers in a dual-redundant configuration.

   ```
   CLI > SET FAILOVER COPY=THIS_CONTROLLER
   ```

3. Restart both controllers by pressing the OCP reset button or entering the following commands:

   ```
   CLI > RESTART OTHER_CONTROLLER
   CLI > RESTART THIS_CONTROLLER
   ```

4. Enter the following commands to determine whether the preceding parameters are set correctly:

   ```
   CLI > SHOW THIS_CONTROLLER FULL
   CLI > SHOW OTHER_CONTROLLER FULL
   ```

5. Set preferred paths to balance the load and improve the performance of the HSZ array controller pair:

   ```
   CLI > SET THIS_CONTROLLER PREFERRED_ID = (0, 1)
   CLI > SET OTHER_CONTROLLER PREFERRED_ID = (2, 3)
   ```

6. Run the CONFIG program to locate and add devices to the array controller's configuration:

   ```
   CLI > RUN CONFIG
   ```

7. Set up the stripesets using the following commands:

   ```
   CLI > ADD RAIDset name POLICY=BEST_PERFORMANCE
   RECONSTRUCTION=NORMAL disk disk disk disk
   ```

8. Initialize the stripesets:

   ```
   CLI > INITIALIZE name CHUNKSIZE = 256
   ```

   Refer to the Disk Storage Configuration section for more information.

9. If a host-accessible logical unit needs to be created from the stripeset, do the following:

   ```
   CLI > ADD UNIT "unit name" "RAIDset name" WRITEBACK_CACHE
   ```

Disk Storage Configuration

Configure the storage in RAID 5 stripesets using between 3 and 14 disks per RAIDset. RAID 0 can be used if redundancy is not required. For this test configuration, each RAIDset consisted of six disks, at 4.3 GB per device, for a total of 25.8 GB of disk space with approximately 21.9 GB exported usable space per RAIDset.

Where *n* is the number of drives:

(*n* - 1) * (device size) = disk space per RAIDset

For configurations other than the minimum and maximum tested by HiTest, see the *Getting Started: HSZ70 Solutions Software V7.0B for DIGITAL UNIX Installation Guide.*

Calculating Maximum Possible Chunk Size

The maximum chunk size possible for a RAID 5 set is determined using the number members in the RAIDset. Beginning with the maximum chunk size setting for HSZ controllers (2048 disk blocks of 512 bytes each), and based on a db_block_size of 32 KB, the maximum possible chunk size for our configuration was calculated as outlined in Table 4-1.

Calculate maximum possible chunk size using Table 4-1 as an example.

Table 4-1: Calculation for Maximum possible Chunk Size

Variable	Calculation	Comment
db_block_size	= 32 Kbytes	= 32,768 bytes.
max chunk size	= 2048 / (*n*-1) = 2048 / (6-1) = 409.6	*n* = number of RAIDset disks. Maximum HSZ chunk size, divided by RAID 5 RAIDsets of 6 disks each, less 1.
max chunk bytes	= disk bytes per block * max chunk size = 512 * 409 = 209,408 bytes	Disk block size (512 bytes), times our maximum chunk size.
Whole db blocks/chunk	= max chunk bytes / db_block_size = 209,408 / 32,768 = 6.39 = 6	Maximum chunk size possible in whole db blocks per chunk (rounded down).

Setting Chunk Size on HSZ Console

For best performance, compute optimal chunk size based on db_block_size, and db_blocks per chunk using the following formula:

$$\frac{(\text{db\_block\_size}) * (\#\text{ of db\_blocks per chunk})}{512}$$

The optimal chunk size will be within the maximum possible chunk size calculated in Table 4-1.

A chunk size of 256 KB was chosen for this test configuration based on the maximum possible I/O size of Oracle during a sequential scan. The chunk size of 128 KB, or four db blocks/chunk, is within the maximum possible calculated in Table 4-1.

Enabling Write-Back Cache

Enabling HSZ write-back cache is recommended. The write-back cache allocates cache memory to both read and write operations, permitting processing to continue without waiting for I/O completions.

Operating System Installation

This section describes the installation and configuration of the DIGITAL UNIX Version 4.0D operating system on the AlphaServer 4100 systems. It includes:

- Configuring DIGITAL UNIX
- Installing Licenses

Configuring DIGITAL UNIX

The configuration in Table 3-1 comes with the DIGITAL UNIX Version 4.0D operating system already installed. Configure your system as shown in the following sections. For detailed installation information, refer to the *DIGITAL UNIX Installation Guide.*

Install with all kernel options and set the kernel parameters as shown in the section UNIX Kernel Parameters.

Swap Space

There are two methods for allocating swap space on DIGITAL UNIX systems:

- Immediate Mode

  This swap mode is the default mode. Swap space is reserved whenever virtual address space is created. Immediate mode requires more swap space overhead, but ensures that swap space is available when needed.

- Deferred Mode (lazy mode)

Swap space is not reserved until the virtual memory subsystem needs to write modified virtual pages to swap space. Deferred mode does not allocate swap space in advance, which may result in swap space not being available when needed. When this occurs, the process may be terminated asynchronously.

To implement deferred mode swapping which is recommended when using very large memory systems, move the symbolic link for sbin/swapdefault as follows:

```
mv /sbin/swapdefault /sbin/swapdefault.old
```

UNIX Kernel Parameters (/etc/sysconfigtab)

The parameter settings used for the UNIX kernel were based on a configuration with 1 GB (Minimum Configuration) and 4 GB (Maximum Configuration) of memory.

Maximize the Unified Buffer Cache (parameters ubc-minpercent and ubc-maxpercent) to allow a larger Oracle buffer cache.

Depending on specific customer requirements for balancing a system for using both Oracle and an application such as the SAS System, ubc-minpercent, ubc-maxpercent, and gh-chunks may be set differently, to make more physical memory available to the application.

Installing Licenses

Register the DIGITAL licenses using the License Management Facility (LMF).

The licenses required for the SAS Oracle DIGITAL UNIX AlphaServer 4100 HiTest Suite are specified in Table 3-2.

Application Installation

This section describes installing the Foundation Software and Appset Software including middleware, database software, and data warehouse application software.

DIGITAL Layered Products

The following DIGITAL Layered Products are required for the SAS Oracle DIGITAL UNIX AlphaServer 4100 HiTest Suite:

- Networker Save and Restore (NSR) for DIGITAL UNIX

The media for these products is provided on the *DIGITAL Software Library for DIGITAL UNIX Systems Layered Products* CD-ROMs. Except where noted, install the products using standard defaults.

This section discusses, in detail, the installation and configuration of Networker Save and Restore (NSR).

Networker Save and Restore

This section describes the installation and setup of Networker Save and Restore.

Installing Networker Save and Restore for DIGITAL UNIX

Install NSR with the help of the *Networker Save and Restore for DIGITAL UNIX Installation Guide.* Except as noted, take the standard defaults.

Setting Up the Tape Library

Set up the automated tape library (jukebox) for backup, using the *Networker Save and Restore for DIGITAL UNIX Administrator's Guide*:

1. Execute the command `MAKEDEV.MC mcxx` to produce the `/dev/mcxx` files.

   Where *xx* = (SCSI Bus ID * 8) + y
   y = value of the LUN

   For this HiTest Suite, the SCSI Bus ID is 6 and the LUN value is 0. Therefore the value of *xx* is equal to 48 and the command is:

   ```
   ./MAKEDEV.MC mc42
   /MAKEDEV.MC mc42
   /MAKEDEV.MC mc44
   ```

2. Add the following entry to the Medium Changer Capabilities Database /etc/micap:

   ```
   Mc48|juke_1:\
   :mc=/dev/mc48:\
   :nd#4:dn=nrmt0h,nrmt1h,nrmt2h=tlz9
   ```

3. To use the jukebox with NSR, configure it by executing the NetWorker `jb_config` command. This utility creates and stores the necessary resource data that NetWorker requires to manipulate the jukebox.

   The jb_config command prompts you for jukebox values including the jukebox name and the device path, tape drive path, and tape drive type for each of the tape drives.

4. Label the tapes using the following command. All 24 tapes are labeled.

   ```
   nsrjb -L -j tz40 -v -Y -f /dev/nrmt0h -b default -S "1-8"
   ```

   Once this command has been issued, the tape library will appear when the NSR Juke Box Labeling dialog box is displayed.

Set Up Networker Save and Restore (NSR)

Use the NetWorker Administrator's utility, `nwadmin`, to set up NSR according to the *NetWorker Save and Restore for DIGITAL UNIX Administrators Guide*.

- Setting Up NetWorker Directives

  Set the NetWorker Directives for Oracle using the Directives window of the Customize pull-down menu in the NetWorker Administrator window.

  To set Oracle directives, create the following directive named Oracle:

  ```
  <</>>
  forget
  ignore
  +obkasm:*
  ```

  The Oracle directive will be associated with the client performing the backup/restore. The `obkasm` is a DMO-specific program that specifies how a set of files is to be backed up and recovered. This directive instructs the client to run `obkasm` on all files specified in the backup/restore.

  The Oracle directive's instructions tell the process to:

<</>>	Start at root
forget	Forget previous directives
ignore	Ignore directives in the file system (.nsr files)
+obkasm:*	Run obkasm on all files under root (the + indicates all subdirectories)

- Setting Up a NetWorker Group

  Set up a single NetWorker Group for Oracle from the Groups window of the NetWorker Administrator Customize pull-down menu, according to the *NetWorker Save and Restore Database Module for Oracle Administrator's Guide.*

- Setting Up a NetWorker Client

  Set up NetWorker Client for Oracle8 database backup from the Clients window of the Client Setup pull-down menu.

  In the HiTest configuration, NetWorker Save and Restore and the database resided on a single system, therefore it was defined using NetWorker Administrator. `Depot 3` was specified as the client system name (this is where the database resides); and `default` was specified as the directive; and as the group.

- Setting Up Backup Schedules

  From the Schedule window of the Customize pull-down menu, set up backup schedules for NetWorker Save and Restore.

- Setting Up NetWorker Volume Pools

  From the Pools window of the Customize pull-down menu, set up a volume pool to include all tape volumes.

- Setting Up NetWorker Save and Restore Server

  To set up the NetWorker Save and Restore Server, follow the instructions in the *NetWorker Save and Restore for DIGITAL UNIX Administrator's Guide*.

Oracle Installation

Install the Oracle8 Server, Version 8.0.3; including the following products from the Oracle8 Server CD: Parallel Query Option, PL/SQL, SQL*Plus, and SQL*Net, using Oracle Installer from the Oracle8 Server CD.

Refer to the Oracle8 *Server Installation Guide* for detailed installation information.

Enabling Process Limits for Oracle DBA User Account

Using the UNIX C Shell, ensure that the following settings are in the `.login` file of the Oracle DBA user account:

```
limit datasize unlimited
limit stacksize unlimited
limit memoryuse unlimited
limit addresspace unlimited’
```

The following table contains UNIX Kernel parameter settings used during the testing of this configuration. These settings were for testing purposes only and are not tuning recommendations. As with any system, the user requirements as well as response time metrics need to be evaluated prior to setting UNIX, Oracle and SAS tuning parameters.

Table 4-2: UNIX Kernel Parameters

Parameter	Minimum Configuration Value	Maximum Configuration Value
vm parameters:		
vm-maxvas	1073741824	4292967296
ubc-minpercent	1	1
ubc-maxpercent	2	2
gh_chunks	0	0
rt parameters:		
aio-max-num	1024	1024
aio-max-percent	2	2
proc:		
max-proc-per-user	1024	267
max-threads-per-user	1024	532
max-per-proc-data-size	1073741824	4292967296
max-per-proc-address-space	1073741824	4292967296
ipc parameters:		
shm-max	1073741824	2139095040
shm-seg	32	32
msg-max	8192	8192
msg-mnb	16384	16384
msg-mni	1024	1024
msg-tql	4096	4096
sem-aem	16384	16384
sem-mni	200	256
num-of-sems	200	200
sem-opm	200	200

Table 4-3: Oracle 64-bit Option Parameters

Parameter	Minimum Configuration Value	Maximum Configuration Value	Comment
db_block_size	32768	32768	DSS and data warehouse applications benefit from a Big Oracle Block (BOB).
db_block_buffers	10000	100000	This number should be maximized to provide the highest cache hit ratio without adversely affecting the memory requirements of other Oracle and system processes.
disk_asynch_io	true	true	Can be used to control whether I/O to data files, control files and log files are asynchronous. If a platform supports asynchronous I/O to disk, it is recommended that this parameter is left to its default.
sort_area_size		104857600	Memory area allocated to each process/thread performing sorts. This parameter should be maximized during intensive sort operations, such as index creation on large tables. This is highly sensitive to the parallel degree of the tables being sorted, as each thread will allocate this space.
sort_direct_writes	auto	auto	Can improve sort performance if memory and temporary space are abundant on your system. This parameter controls whether sort data will bypass the buffer cache to write intermediate sort results to disk. When set to the default of AUTO, and the value of the sort area size is greater than ten times the block size, memory is allocated from the sort area to do this. When SORT_DIRECT_WRITES is TRUE, additional buffers are allocated from memory during each sort.
sort_write_buffer_size	131072	32768	When sort_direct_writes is set to true, this parameter must be set to a value between 32768 and 131072, that is a multiple of 32768.
cache_size_threshold	180000	180000	Specifies the maximum size of a cached partition of a table split among the caches of multiple instances. If the partition is larger than the value of this parameter, the table is not split among the instances' caches. The default value of this parameter is 1/10 the number of database blocks in the buffer cache. This parameter can also specify the maximum cached partition size for a single instance. As of Release 8.0.3, the CACHE_SIZE_THRESHOLD parameter is being denigrated.
db_file_multiblock_ read_count	8	8	This parameter is used for multi-block I/O and specifies the maximum number of blocks read in one I/O operation during a sequential scan. The default is 8. DSS (data warehouse) database environments tend to get the most benefit from maximizing the value for this parameter.

Parameter	Minimum Configuration Value	Maximum Configuration Value	Comment
parallel_max_servers	72	72	Setting this parameter greater than 1 allows the Oracle RDBMS to take advantage of the parallel query option, for tables where the parallel degree has been set. This parameter should minimally be set to the largest degree specified on any table. Index creation on a table with a parallel degree set will also use this option; however 2 threads will be created for each degree. Therefore, this parameter should be set to twice the degree for full benefit on index creation.
shared_pool_size	18000000	18000000	This parameter affects the performance of the library cache (shared SQL and PL/SQL areas) and the dictionary cache. This parameter may be reduced if the cache hit ratio is not adversely affected. If cache misses are prevalent, the open_cursors parameter may need to be increased.
Log_checkpoint_interval	99999999	99999999	Set sufficiently large number to disable time-based checkpoints forcing checkpoints only to occur on log switches.

Database Creation and Storage Considerations

Creating a database that ensures proper and efficient use of Oracle8 requires special considerations. This section describes these considerations and summarizes how they were used.

Plan your database tables and indexes by estimating how much space they will require to support existing and future needs based upon expected growth.

When creating a database using either a file system or raw devices, configure the file system and database storage based on the characteristics of the database to be installed. Allocate disk space to optimize performance and allocate storage to minimize Input/Output (I/O) contention and maximize I/O throughput.

For better performance distribute tablespaces across as many devices as possible, and separate tablespaces onto separate disks.

The file system and database storage map for the maximum and minimum configurations of this HiTest Suite are shown and discussed in Chapter 5. The processes used to create the maximum and minimum configuration databases for this HiTest Suite are also discussed in Chapter 5.

Raw Disk Handling

Oracle supports the use of file system and raw devices. Implement the database to take advantage of DIGITAL UNIX asynchronous I/O capability. The amount of overall disk space required depends on the volume of data in the database, the database processes during peak load, as well as the expected growth and expansion over time. To optimize throughput and minimize I/O contention, distribute Oracle database files over as many disks as possible. This can be accomplished by the use of RAID set volumes.

Consider the following when setting up the raw disk:

- Use RAID controllers. This Suite had the disks set up as RAID 5 sets at the hardware level using the RAID array controllers. RAID 5 provides disk striping with distributed parity, but provides only the capacity of *n*-1 devices, where *n* = Number of RAIDset members. For example, five 4.3 GB drives have a 21.9 GB RAID 5 capacity.

- When possible, spread system disk, database tables, indexes, and Oracle redo logs across separate disks.
- Plan to place data to raw disks on different SCSI buses to reduce I/O contention.
- Use symbolic links to assign meaningful names to all the raw disks to improve manageability.

An example of how the raw disk was setup and partitioned is shown in Chapter 5.

Temporary Tables

Consider the following when setting up temporary tables:

- Create temporary tablespaces on their own disk.
- For large sort operations such as initial index creation of FACTS table, increase the database initialization parameter SORT_AREA_SIZE accordingly, based on the amount of memory available.
- Reuse temporary table space when possible. For example, to create FACTS TABLE INDEX required approximately 50 GB of temporary space to support sorting during index creation. Once the index was created, the temporary tablespace was dropped and recreated with a much smaller size. This frees additional storage which can be used for other temporary storage requirements, such as maintaining Oracle export dumps on disk prior to them being backed up to tape.

SAS Installation

Install the SAS System, Release 6.12 using the SAS System Installation Manager from the SAS System distribution media, according to *Installation Instructions for the SAS System for DIGITAL UNIX, Release 6.12, TS040*, (Cary, NC:SAS Institute Inc., 1997).

Using the Link Method

Use the link method to make the SAS System available to users, as described in the *Installation Instructions for the SAS System for DIGITAL UNIX*.

Define a symbolic link between the SAS command and the SAS root directory which must be defined in the search path:

```
ln -s /sas /sas612/sas
```

Note

When issuing the link command, you must be root user.

Editing the SAS Config File

Edit the SAS configuration file, `config.sas612,` to provide sufficient memory for the large SAS dataset size as indicated in Table 4-4.

Table 4-4: SAS Configuration File Settings

Setting	Comment
-memsize 256m	Limits amount of memory allocated by the SAS System. Default is 32 MB.
-sortsize 256m	Limits amount of memory allocated during sorting operations. Default is 16 MB.
-work /sas2	Specifies where the SAS work library is created. This is a temporary work library: SAS data sets created there are deleted when the system terminates. Default is `/usr/tmp`.

System Management Station Installation

The inclusion of a system management station in this HiTest Suite is optional. If you choose not to include a management station, the configuration of storage can be performed using the Command Line Interface (CLI). You can access the CLI by connecting a maintenance terminal to the maintenance port on the front of the controller. After you have initially configured the controller, making it visible to the host, you can perform all other configuration tasks through a remote connection.

Installation of software components can be performed using a VT220 terminal or similar X-windows capable terminal.

If you have selected a system management station, use the instructions in this section to install and configure the hardware and software on the management station and the software on the systems being managed.

Hardware Installation

Install the system management station hardware using the instructions in the documentation provided with your DIGITAL PC 5100 ST system or comparable 80x86 system.

Software Installation

This section describes how to install the system management station software, including the operating system, the management applications installed on the management station, and the clients and agents installed on the systems being managed.

Operating System Installation

Install and configure the Windows NT Workstation operating system using the documentation provided with your PC.

Install and configure the SNMP agent on the system management station as follows:

1. Select the Network item from the Windows NT Control panel.
2. Click the Services tab of the Network property page to install SNMP and configure it with the IP addresses or name of the AlphaServer system.

Install Windows NT Service Pack 3 (SP3).

ServerWORKS

To install and configure ServerWORKS, you must:

- Install ServerWORKS Manager on the system management station.
- Install ServerWORKS Agents on the managed systems. The AlphaServer systems specified in this HiTest configuration come with the ServerWORKS Agents factory installed.
- Define the managed systems to the system management station.

Installing ServerWORKS Manager

Install ServerWORKS Manager, which is provided with your AlphaServer 4100 system, using the instructions provided in the *ServerWORKS Manager Overview and Installation* Guide.

StorageWorks Command Console

To install and configure StorageWorks Command Console, you must:

- Install the StorageWorks Command Console Client on the system management station
- Install and configure the HSZ70 Agent on the AlphaServer systems

The StorageWorks Command Console has recently been upgraded to provide the capability to both configure and manage HSZ70 storage. If the software is not provided with your HSZ70 Solutions kit, you can access it by downloading it from the Web.

StorageWorks Command Console Client

Install and configure the StorageWorks Command Console Client on the system management station using the instructions provided in the *DIGITAL StorageWorks Command Console Getting Started* Guide.

After the Client is installed, you can access online help for information about configuring and using the Command Console.

Install the HSZ70 Agent on the AlphaServer system.

StorageWorks HSZ70 Agents

The StorageWorks HSZ70 Agent works with the StorageWorks Command Console Client to configure, operate, and monitor the AlphaServer 4100 storage subsystems.

Install the HSZ70 Agent on the AlphaServer 4100 system using the instructions provided in the *DIGITAL StorageWorks Command Console Getting Started* Guide.

All Agent installations must be performed locally, do not attempt to install the Agent over the network.

5
Tests and Results

The DIGITAL HiTest program tests for several types of conditions 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 SAS Oracle DIGITAL UNIX AlphaServer 4100 HiTest Suite.

Oracle database tests were executed on the AlphaServer 4100 system to verify the ability to load data, build index structure, partition database components, and process SQL based queries. All users were local terminal users.

Subsequently, SAS system database tests were conducted with the SAS system running scripts to test forecasting and statistical modeling functions.

Test Environment

The following diagram illustrates a conceptual data warehouse implementation. The circled areas indicate the focus of the testing of this HiTest Suite.

Figure 5-1: Conceptual Overview of the Data Warehouse

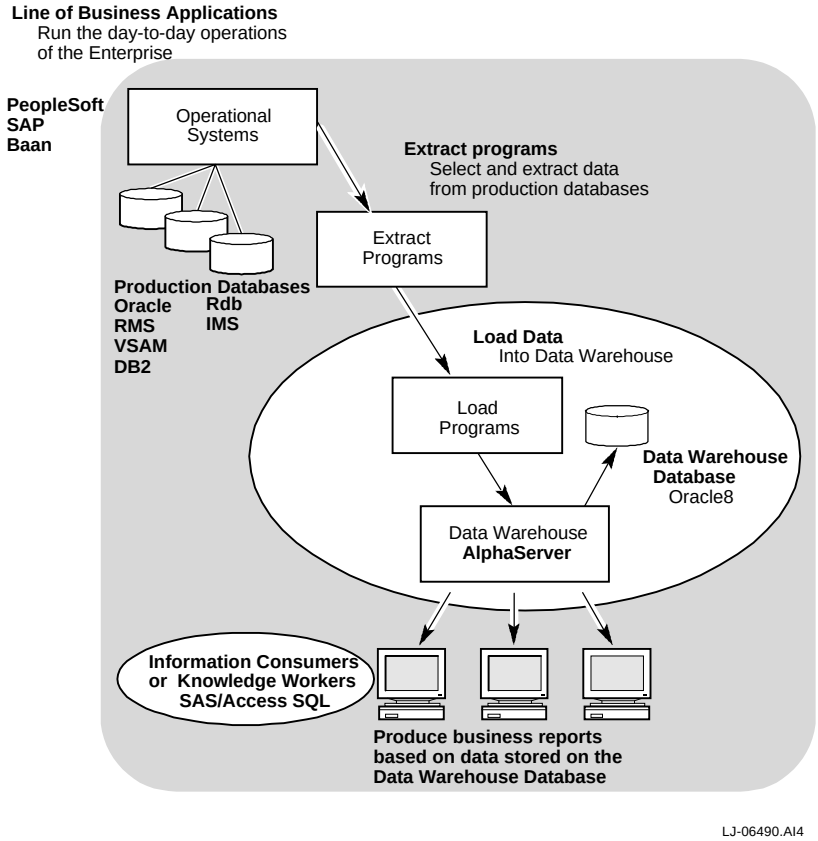

Detailed drawings showing the test configurations of the DIGITAL UNIX AlphaServer 4100 foundation hardware are shown in Chapter 7.

Test Tools

This section describes the test tools used for interoperability testing.

Consumer Packaged Goods Database

The Consumer Packaged Goods (CPG) database was used as the test database. It contained data that represents sales information for a variety of products across geographical areas of the US, including sales volume, channels, and other characteristics. The minimum configuration database used consists of 79,249,896 rows of sales data. The maximum configuration database consists of 876,551,880 rows of sales data.

The databases were built using the CPG Database Demo scripts provided by Oracle Corporation. The CPG Database represents typical marketing and sales data for a consumer products manufacturing firm. The data used was organized for decision support.

Test Scripts

Test scripts were used to demonstrate and test the interoperability of the SAS Oracle DIGITAL UNIX AlphaServer 4100 HiTest Suite, in a Data Warehousing and Decision Support environment. The scripts were designed to emulate typical decision support questions about the historical activity of a product sales environment.

Test Configuration

During testing of minimum and maximum configurations the Oracle databases used were placed on 6 disk RAID 5 stripesets, and were implemented using raw devices.

Maximum Configuration

The maximum configuration database was made up of 12 data files comprising the large SALES_FACT table. The index for the SALES_FACT table consisted of eight data files. Four volumes of TEMP space were required to build the SALES_FACT_INDEX. The database used to test the maximum configuration contained 876,551,880 rows of sales data.

Minimum Configuration

The minimum configuration database was made up of two RAID 5 data files comprising the large SALES_FACT table. The index for the SALES_FACT table consisted of two data files. Two volumes of TEMP space were required to build the SALES_FACT_INDEX. The database used to test the minimum configuration contained 79,249,896 rows of sales data.

File System and Database Storage Map (Maximum Configuration)

The file system and database storage map for the maximum configuration tested is provided in Table 5-1. Performance and storage considerations guided the placement of the data files. As shown in the file system and database storage map, the database consists of approximately 140 GB of storage.

Table 5-1: File System/Database Storage Map, Maximum Configuration

SCSI BUS #1	
HSZ70 Controller	
/dev/rrzb8c:	FACTS1
/dev/rrz8c:	FACTS4
/dev/rrzb9c:	/sas1
/dev/rrz9c:	/sas2
HSZ70 Controller	
/dev/rrzb10c:	TEMP1
/dev/rrz10c:	/sas2
/dev/rrzb11c:	FACTS INDEX3
/dev/rrz11c:	/sas1

SCSI BUS #2	
HSZ70 Controller	
/dev/rrz16a:	system
/dev/rrz16b:	swap1
/dev/rrz16g:	usr
/dev/rrzb16c:	Oracle
/dev/rrzc16c:	FACTS4
/dev/rrz17c:	var
/dev/rrzb17a:	DIMENSIONS
/dev/rrzb17b:	loga_1
/dev/rrzb17c:	logb_1
/dev/rrzb17d:	logc_1
/dev/rrzb17e:	DIMENINDEX
/dev/rrzb17f:	SYSTEM
HSZ70 Controller	
/dev/rrz18c:	swap2
/dev/rrzb18c:	/data1
/dev/rrz19c:	FACTSINDEX2
/dev/rrzb19c:	TEMP4
SCSI BUS #3	
HSZ70 Controller	
/dev/rrzb24c:	FACTS3
/dev/rrz24c:	TEMP2
/dev/rrzb25c:	/sas1
/dev/rrz25c:	/sas2
HSZ70 Controller	
/dev/rrzb26c:	FACTSINDEX1
/dev/rrz26c:	TEMP3
/dev/rrzb27c:	FACTS INDEX4
/dev/rrz27c:	/data2

An attempt was made to optimize I/O access paths, while minimizing I/O contention for database load and data retrieval.

One 21.9 GB RAID 5 stripeset was subdivided to seven smaller equal partitions of 1,785,490 KB each to store the SYSTEM (190 MB), DIMENSION (10 MB), and DIMINDEX (10 MB) tablespaces, and the redo logs (1000 MB/log).

File System and Database Storage Map (Minimum Configuration)

The file system and database storage map for the minimum configuration tested is provided in Table 5-2.

Table 5-2: File System Database Storage Map, Minimum Configuration

SCSI BUS #1	
HSZ70 Controller	
/dev/rrz8a	/SYSTEM
/dev/rrz8b	/LOGA_1
/dev/rrz8c	/LOGB_1
/dev/rrz8d	/DIMENSIONS
/dev/rrz8e	/FACTSINDEX1
/dev/rrz8f	/TEMP1
/dev/rrz8g	/TEMP2
/dev/rrz8h	/DIMENINDEX
SCSI BUS #2	
HSZ70 Controller	
/dev/rrz10a	/FACTS1
/dev/rrz10b	/FACTS2
/dev/rrz10c	/FACTS3
/dev/rrz10d	/FACTSINDEX3
/dev/rrz10e	/FACTS4
/dev/rrz10f	/TEMP4
/dev/rrz10g	/FACTSINDEX2
/dev/rrz10h	/TEMP3

During minimum configuration testing, storage considerations were the primary factor in determining the placement of data files. The available storage configured as part of the minimum configuration limited the amount of data stored within the database, as well as the ability to store online data which was used for loading.

System Limits and Characterization Data

As with any configuration, the physical limits of the system can effect the performance of the system. In planning a data warehouse or decision support system factors such as storage requirements, potential growth, query and reporting strategy must all be considered. During this testing, the minimum configuration was restricted in the following areas:

- Memory

  With 1 GB of physical memory, the system will be limited in the number of users it can support, as well as the amount of data which can be cached within the database buffer cache.

- CPU

  The minimum configuration single CPU environment will not be able to take advantage of SMP, Oracle parallel processing, as well as large volumes of concurrent users.

Test Process and Results

Interoperability of the SAS Oracle8 and Digital Unix AlphaServer 4100 was tested using the following processes:

- Creating a database
- Loading a database using SQL loader
- Creating, dropping and re-creating indexes required to access the stored data
- Executing test scripts (which represented typical queries against the database) using SQL and SAS/System 12

To test the Oracle8 functionality, which allows for the partitioning of data, an additional FACT table was created with four partitions based upon yearly date range. Data was loaded using an SQL loader based upon ranged partition. Indexes were then created.

To demonstrate the ability to archive data using partitions the following process was performed:

- Yearly partitions were exported using the Oracle8 export utility.
- A single years partition was dropped, effectively removing the earliest years worth of data.
- A single partition was restored using the Oracle8 import utility.
- The indexes associated with that table were dropped and re-created.

Oracle Database Backup and Restore

The following process was used to perform a backup and restore:

- The partitioned FACTS table was exported to a compressed UNIX TAR file, using the Oracle8 export utility.
- Additional export dmp files were created to compressed TAR files for each of the four Sales_fact partitions.
- All UNIX TAR files containing the Oracle files were stored off-line using the NSR backup utility to backup to tape.
- The compressed TAR files were restored to the system, using the NSR Restore utility.
- After dropping the partitions, the partitions were restored to the SALES_FACT partitioned table using the Oracle8 import utility.
- Oracle indexes for the partitioned FACTs table were rebuilt.

Recoverability was tested using the NSR Network Save and Restore software. Application and database file systems were backed up to tape. File systems were restored to the system to test the NSR restore functionality.

Database Creation

The process used to create the database consisted of creating the disk space, creating tablespaces, altering the tablespaces to expand the database, creating the tables, loading the data, and creating indexes.

The CPG data model consists of six tables including: the SALES_FACTS table, which contained the bulk of the database, and five dimension tables: PRODUCT, MARKET, CHANNEL, DAILY_PERIOD, and MONTHLY_PERIOD as shown in Figure 5-2.

Figure 5-2: Logical Database Design

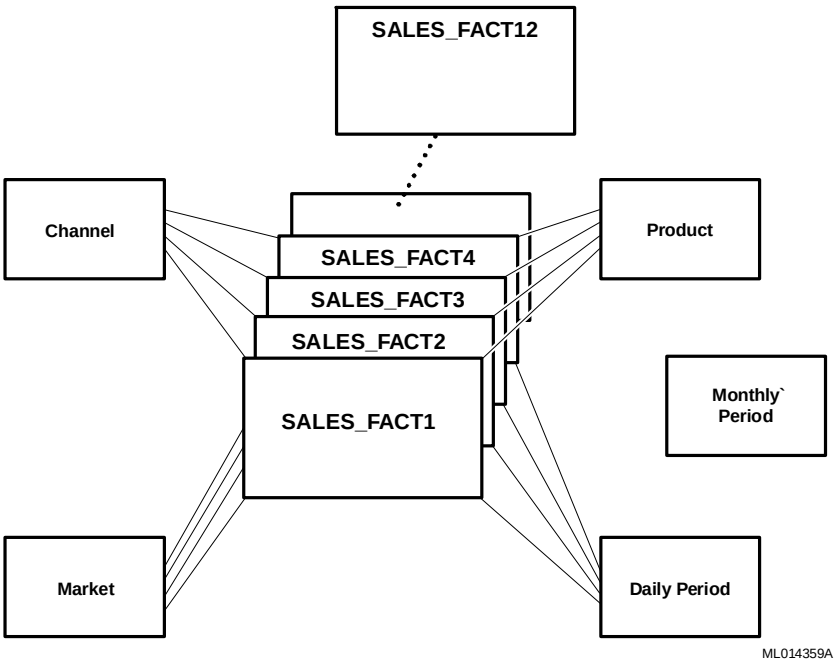

Tablespace Configurations

Both the FACTS and FACTSINDEX tablespaces used the storage parameters shown in Table 5-3. The FACTS tablespace contained the SALES_FACT table, which was populated using `sqlldr`. The initial extent of a tablespace is not used by `sqlldr`. Therefore, the initial extent was set to 64K to limit this unusable space. Table 5-3 illustrates that the only difference between the tablespace configurations for the maximum and minimum configurations was the datafile size.

Table 5-3: Database Tablespace Storage Parameters

Storage Parameters	FACTS	FACTSINDEX	TEMP
datafile size (Max. Configuration)	5112 MB	5112 MB	5112 MB
(Min. Configuration)	4,670 MB	4,670 MB	4,670 MB
initial_extent	64 KB	1,000 MB	200 MB
next_extent	1,000 MB	1,000 MB	200 MB
max_extent	2,041 MB	2,041 MB	2,041 MB
pctincrease	0	0	0
pctfree	0	0	N/A

Tablespace Creation

Each tablespace was created sequentially using one Oracle data file. Once created, the tablespaces were altered using the `alter tablespace` command, which added additional data files to each tablespace.

Table 5-4 shows six Oracle tablespaces which comprised the CPG database. The free space associated with each tablespace is shown to indicate the extent to which the various tablespaces were used.

Table 5-4: Oracle8 CPG Database, Maximum Configuration

Tablespace	Contains	Size	Free Space %
FACTS	SALES_FACTS Table 12 Data Files 5112 MB	49,076 MB 61344 (876,551,880)	28%
FACTSINDEX	SF_KEY Index on SALES_FACTs Table 10 Data Files 5112 MB	51,120 MB	18%
DIMENSION	1 Data File, Dimension Tables: CHANNEL DAILY_PERIOD MARKET PRODUCT	1024 MB	92%
DIMINDEX	1 Data File Index for Dimension Tables	1024 MB	88%
TEMPFILE	10 Data Files (5112 MB each) required for SF_KEY index build	51,120 MB	N/A
SYSTEM	1 Data File	190 MB	N/A

Database Load, SQL Loader (Maximum Configuration)

Data was loaded to the database using the SQL Loader. Table 5-5 shows different load scenarios that were used to load the facts data. In each method, one months data in a flat file was used as input to the sqlldr.

Table 5-5: Facts Database Load Rates, Maximum Configuration

Data Files	Rows	Time MM:ss	Avg Rows/Minute	Load Scenario
FACTS1	37223436	35:05	2122001	Serial Creation, 1 Stream
FACTS2	37223436	32:38		
FACTS3	37223436	37:08	4009705	Serial Creation, 4 Stream
FACTS4	37223436	35:11		
FACTS5	37223436	35:59		
FACTS6	37223436	34:40		
FACTS7	37223436	51:09	4345148	Serial Creation, 6 Stream
FACTS8	37223436	51:21		
FACTS9	37223436	49:49		
FACT10	37223436	51:08		
FACTS11	36022680	51:24		
FACTS12	37223436	51:09		
FACTS13	37223436	35:48	2079521	Serial Creation, 2 Stream
FACTS14	33621168	31:58		
FACTS15	37223436	36:48	4036884	Serial Creation, 4 Stream
FACTS16	37223436	34:49		
FACTS17	37223436	36:50		
FACTS18	37223436	36:53		
FACTS19	37223436	54:54	35993652	Serial Creation, 6 Stream
FACTS20	37223436	53:54		
FACTS21	37223436	52:30		
FACTS22	37223436	54:48		
FACTS23	37223436	53:00		
FACTS24	37223436	62:03		

Database Load, SQL Loader (Minimum Configuration)

For the minimum configuration database, data files were loaded to the SALES_FACT table. The number of rows loaded and the load times are specified in Table 5-6.

Table 5-6: Facts Database Load Rates, Minimum Configuration

Data Files	Rows	Time	Rows/Sec	Load Scenario
FACTS1	123208316	13:11	155762	Sequential load
FACTS2	123208316	27:13	150898	2 way parallel
FACTS3	123208316	27:15		
FACTS4	123208316	40:34	151609	3 way parallel
FACTS5	123208316	40:38		
FACTS6	123208316	40:44		

Creating the Index

Creating the index was accomplished by allocating temporary storage to allow the index to be built and then using the Oracle Parallel Query option to build the index. During index

creation, it is important to allocate sufficient space for temporary storage for the index to build successfully. Due to the time required to build an index on a data warehouse, it is important to calculate these parameters correctly.

Calculating Required Index Space

Temporary space requirements should be estimated at a minimum of 10% above the required index space.

Calculate initial index and temporary space required, to allow creation of a concatenated index SF_KEY on the large SALES_FACT table.

Calculate the space requirements using the following procedure:

1. Obtain, or if necessary estimate, the total number of rows in the SALES_FACT table:

   ```
   sales_fact = 876551880 rows (Maximum Configuration)
   ```

2. Calculate the index block header size, using the formula:

block header	= **fixed header + variable header** = 113 bytes + initrans*24 bytes = 113 bytes + 2*24bytes = 161 bytes

3. Calculate available data space per data block as follows:

available data space	= **DB block size - block header** = 32768 bytes - 161 bytes = 32607 bytes

4. Calculate the entry column size by summing the length of the columns in the index. The four columns used in this index are product_id, chan_id, market_id, and day, which are declared as 3 char (6) and a date field (7).

entry column size	= **sum (length of 4 columns)** = 6 + 6 + 6 + 7 = 25 bytes

5. Calculate the bytes per entry as follows:

bytes per entry	= **entry column size + entry header + rowid + F + V** Where, entry column size = 25 entry header = 2 bytes rowid = 6 bytes F = fixed length bytes/entry (# of columns 127 bytes) = 5 (4 specified columns plus 1 for rowid) V = variable length byte (0 for all calculations)
Note: bytes per entry = 25 + 2 + 6 + 5 + 0 = 38 bytes	

6. Calculate total index blocks required as follows: (Maximum configuration example)

index blocks	1.05 * ((#rows * **bytes per entry**) = ---------- ((**available data space/bytes per entry**)***byte per entry**)) = 1.05 * ((876551880 * 38) / ((32607/38) * 38)) = 1021528.2 index blocks **Note:** The multiplier 1.05 is used to account for the index branch nodes and the use of parallel index create will most likely result in files that are not 100% utilized.
Index bytes	= **index blocks** * **db block size** = 1021528.2 * 32768 = 33473437487.2 bytes
index MB	= **index bytes** / (1024 * 1024) = 31922.7MB

7. The parallelism and extent sizes used to enhance performance will adversely affect efficiency of the storage utilization for the index. To account for this behavior a 10% safety margin is built into the calculation of final index space as follows:

Final index space (MB)	= **index MB** * 1.1 = 31922.7 MB * 1.1 = 35115 MB

8. Determine (whole number) the maximum number of extents in a single datafile. For this database, the maximum number of extents is calculated as follows:

Max. number of extents per datafile	= **datafile size** / **extent size** = 5112MB/1048576000 = 5.1

9. Determine the total number of files required to hold the SF_KEY index as follows:

index files	**final index space** (bytes) = ---------- (**extent size** * **number of extents per file**) = **35115 MB**/(**1048576000*5**) = 6.3 (always rounded up) = 7

Index Creation

Parallel index creation, provided by the Oracle Parallel Query option, is an important database management function. Parallel index creation was used during creation of the index on the large SALES_FACT table.

The concatenated SF_KEY index was created with a parallel degree of two for the minimum configuration tested and a parallel degree of six for the maximum configuration tested, which required that the parallel_max_servers parameter be set to a minimum of twice the parallel degree value. This was necessary because the degree of parallelism on index creation was implemented by two cooperating query servers. The parallel degree for index creation was enabled by altering the default degree on the SALES_FACT table.

Use the `svrmgr` commands to perform this:

```
svrmgr> CONNECT cpg/cpg
```

For Minimum Configuration Tested:

```
svrmgr> alter TABLE SALES_FACT parallel (degree 2);
```

For Maximum Configuration Tested:

```
svrmgr> alter TABLE SALES_FACT parallel (degree 6);
```

During *index create*, the system global area (SGA) size was significantly reduced as the SORT_DIRECT_WRITES parameter allows the SGA to be bypassed when performing index creation. By decreasing the SGA size, additional memory was provided for sorting by each parallel query server, reducing the elapsed time required to create the index.

The changes specified in Table 5-7 were made to the standard database parameters to facilitate index creation.

Table 5-7: Index Parameters

Parameter	Minimum Configuration Value	Maximum Configuration Value
db_block_buffers	2000	2000
sort_area_size	83886080	536870912
sort_write_buffer_size	131072	131072

Index creation was performed for 164061480 rows of SALES_FACT data to support the minimum tested configuration and 928643603 rows of SALES_FACT data to support the maximum configuration tested.

The actual *create index* statement used was:

```
svrmgr> create unique index CPG.SF_KEY on
CPG.SALES_FACT(PRODUCT_ID, CHANNEL_ID, MARKET_ID, DAY)
tablespace FACTSINDEX unrecoverable;
```

Oracle Database Queries

Functional verification of the ability to perform query operations was demonstrated using five SQL join scripts and a full table scan. These queries exercised functionality of the RDBMS server, SQL*Plus, and Parallel Query.

The scripts were designed to emulate typical decision support questions about the historical activity of a product sales environment. In most cases, the result of these queries would be used to generate sales trends.

The following five queries were designed to search the database in various ways to exercise the database. All queries returned results grouped by month.

Query 1

Query 1 asked “What was the product share of a specific brand of cereal as compared to other cereals in the same product category, in a particular state, in a particular type of store.” The information was grouped by month to show market trends.

The business question asked is:

“How did 20 oz. Wheat Flakes do in 1995 as compared to all types of wheat flakes in supermarkets in the state of Connecticut?”

Query 2

Query 2 compares the sales of a specific product, in a particular outlet in a region, against the sales of the same product through all channel outlets. The information is grouped by month to show market trends.

The business question asked was:

"What percentage of sales of 15 oz. Wheat Flakes were made in the Safeway stores in NY and PA as compared to all outlets in the NY and PA areas?"

Query 3

Query 3 compared the market share of a product in a particular type of store, in a particular market location, to sales of all types of outlets in the region. The information was grouped by month to show market trends.

The business question asked was:

"How are 10 oz. Wheat Flakes doing in convenience stores in Bridgeport Connecticut as compared to the entire northeast region?"

Query 4

Query 4 compared the market share of a particular product, in a particular type of store, in a particular market location to all sales of competitive products in the same market location. The information was grouped by month to show market trends.

The business question asked was:

"What was the market share of 20 oz. Wheat Flakes in Connecticut supermarkets in 1995?"

Query 5

Query 5 compared the product share of a given product, combining several areas, to total sales across the same areas.

The business question asked was:

"What was the market share of 20 oz. Wheat Flakes across 10 test market areas?"

Oracle Database Tests

Three Oracle Database tests were performed to test interoperability. The tests and test results are discussed in the following paragraphs.

Test One – Process Executing Queries 1–5 Sequentially

The first test performed five sequential queries (SQL joins). The performance of a query is dependent upon the optimizer choosing an efficient query plan. In general, analyzing the tables and indexes enable the optimizer to produce effective plans. Table 5-8 lists the results of test one (maximum configuration). Table 5-9 lists the results of test one (minimum configuration).

Table 5-8: Performance Results for Test One, Maximum Configuration

Maximum Configuration (928643603 Rows)			
	1 USER	2 USERS	4 USERS
	Cold Cache	Warm Cache	
Query 1	134	22	33
Query 2	125	50	70
Query 3	198	51	67
Query 4	373	59	89
Query 5	405	283	289
Total	1235	465	548

Table 5-9: Performance Results for Test One, Minimum Configuration

Minimum Configuration (164061480 Rows)		
	1 USER	2 USERS
Query 1	57.08	116
Query 2	39.05	131
Query 3	37.9	133
Query 4	115.31	204
Query 5	189.62	310
Total	438.96	894

Test Two – Executing Queries 1–5 in Parallel

The second test executed all five queries in parallel. Parallel queries more stressfully exercise the database versus sequential queries performed in test one.

Tables 5-10 and 5-11 list performance results for the parallel query tests for the minimum and maximum configurations.

Table 5-10: Performance Results for Test Two, Maximum Configuration

Maximum Configuration (928643603 Rows)			
	1 USER (5 Processes)	2 USERS (10 Processes)	4 USERS (20 Processes)
	Cold Cache	Warm Cache	
Query 1	100	78.5	138.5
Query 2	124	133	257
Query 3	193	124	247
Query 4	276	146.5	270
Query 5	269	278	494
Total	962	760	1406.5

Table 5-11: Performance Results for Test Two, Minimum Configuration

Minimum Configuration (164061480 Rows)		
	1 USER (5 processes)	**2 USERS (10 Processes)**
Query 1	60	115
Query 2	74	131
Query 3	82	133
Query 4	125	204
Query 5	217	310
Total	558	893

Test Three – Testing Parallel Queries Using Degrees of Parallelism

The degree of parallelism used in parallel query operations effects level of performance during parallel query operations. The impact to overall performance was tested using several different degrees of parallelism in both the minimum and maximum test configurations.

Tests were performed using various degrees of parallelism.

An example of how to set degree of parallelism is as follows:

```
svrmgr> CONNECT cpg/cpg
svrmgr> alter TABLE SALES_FACT parallel (degree #);
```

The Parallel Query Option was tested using the following query to obtain the number of rows in the SALES_FACT table:

```
select count(*) from sales_fact;
```

SAS System Database Tests

Functional testing of SAS included interactive script execution, where the interactive mode used the X Window display manager. Scripts retrieved data from the existing Oracle Consumer Packaged Goods Data Warehouse via SAS/Access. Forecasting and statistical modeling test scripts were executed.

Forecasting Tests

Scripts were executed to create datasets for forecasting results on a daily, weekly, and monthly level.

Test One

Daily forecast for a single market, single channel, single product, and eight months of daily data to produce a 30-day ahead forecast on a daily level.

Test Two

Weekly forecast for a single market, single channel, single product and eight months of data aggregated to week-ending levels to produce a 12-week ahead forecast. The results contained data for the last 12 weeks and the forecast 12 weeks.

Statistical Modeling Tests

Scripts were executed to create datasets for statistical modeling to detect differences between stores, channels groups, products, and so on.

Test One

Build a statistical model of how channel and monthly factors effect daily sales using daily level data for a single market, single product and eight months of data. Examine the interaction between channel and month.

Test Two

Build two statistical models looking at channel effects and channel and month effects with no interactions using eight months of data aggregated to the monthly channel for a single market and a single product.

Test Three

Build a repeated measure analysis to examine channel group differences over time using eight months data aggregated to the monthly channel level for a single product and a single market.

Redundant Component Testing

This section describes redundant component and storage device failure testing.

Redundant Power Supply Failure Test

In order to verify the available redundant power supply units which were part of the minimum configuration, one of three H7291-AA 450 Watt power supplies were disconnected while executing database loads. During this time there was no impact to the system, application or database processing.

Storage Device Failure Test

During testing, a member of a RAID 5 stripeset contained the SALES_FACTS table was removed to simulate a disk failure. Status of the failed disk changed to ‘failed set’ and the spareset was incorporated into the RAIDset. During this event, there was no impact on the system and database processes running at the time.

6
Problems and Solutions

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

Foundation Hardware

No problems were encountered.

Foundation Software

The following problem was identified:

Problem When installing SAS 6.12, the environmental variable SASORA must be set to establish the version of Oracle to which it is being connected. For Oracle 7 the variable was set to V7. Attempts to set that variable to V8 while testing with Oracle8 were unsuccessful.

Impact There was no impact on the configuration being tested.

Solution Set the variable to V7. Setting the variable to V7 does not effect the execution of the SAS queries. Consult Oracle for additional information.

7
Detailed Hardware Configuration

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

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

System Diagrams

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

Figure 7-1: System Diagram, Maximum Configuration

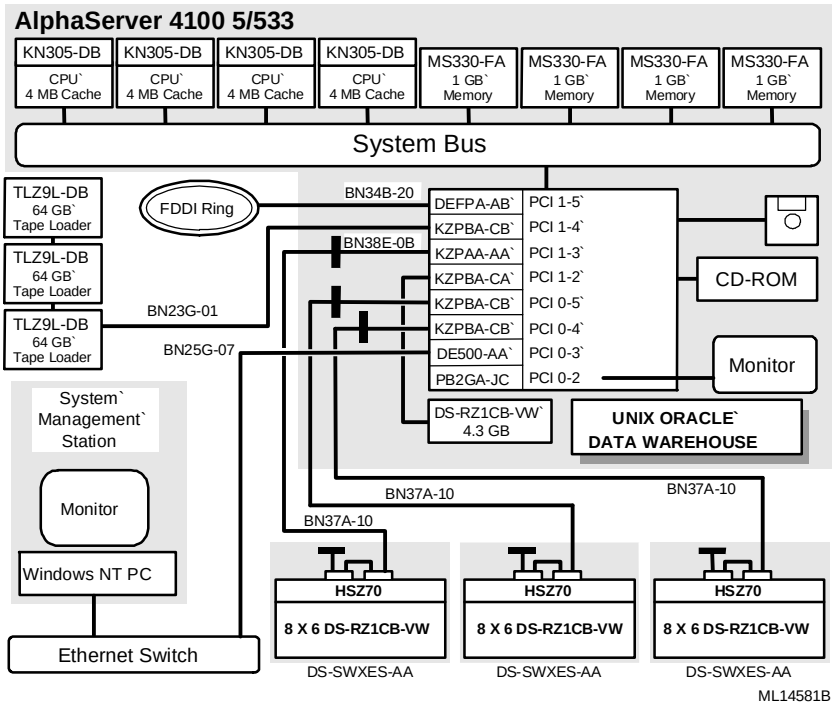

Table 7-1: Configuration Cabling

Part Number	Qty	Description	From	To
BN38E-0B	3	SCSI cable	KZPBA-CB	BN37A-10
BN34B-20	1	Dual fiber optic cable	DEFPA-AB	FDDI Ring
BN25G-07	1	Ethernet cable	DE500-AA	Ethernet Switch
BN37A-10	3	SCSI cable	BN38E-0B	HSZ70
BN19J-03	2	SCSI cable	TLZ9L-DB	TLZ9L-DB

HiTest System Slot Configurations

Figure 7-2 shows the HiTest System Slot Usage and Table 7-1 describes the minimum and maximum hardware configurations used in this HiTest Template.

Figure 7-2: HiTest System Slot Usage

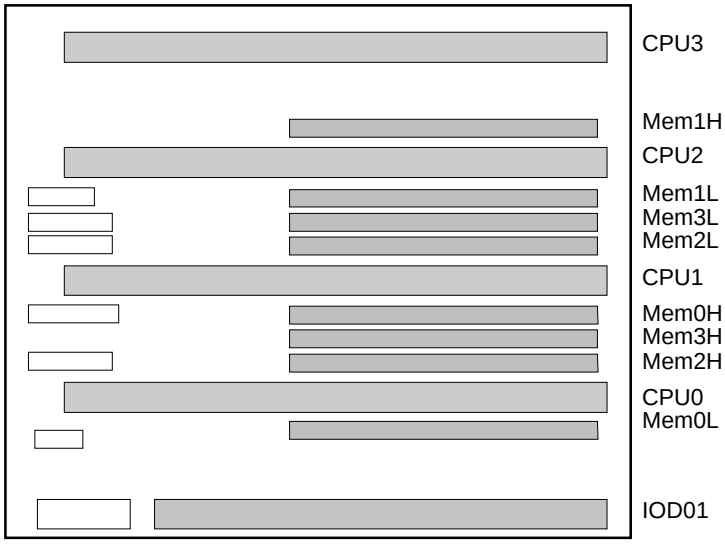

LJ-06489.AI4

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

Slot	Minimum Configuration	Maximum Configuration	Description
CPU3	–	KN305-DB	533 MHz CPU 4 MB cache
Mem1H	–	MS330-FA	Memory pair 1 (2 of 2)
CPU2	–	KN305-DB	533 MHz CPU 4 MB cache
Mem1L	–	MS330-FA	Memory pair 1 (1 of 2)
Mem3L	–	MS330-FA	Memory pair 3 (1 of 2)
Mem2L	–	MS330-FA	Memory pair 2 (1 of 2)
CPU1	–	KN305-DB	533 MHz CPU 4 MB cache
Mem0H	MS330-FA	MS330-FA	Memory pair 0 (2 of 2)
Mem3H	–	MS330-FA	Memory pair 3 (2 of 2)
Mem2H	–	MS330-FA	Memory pair 2 (2 of 2)
CPU0	KN305-DB	KN305-DB	533 MHz CPU 4 MB cache
Mem0L	MS330-FA	MS330-FA	Memory pair 0 (1 of 2)
IOD01	Bridge	Bridge	System Bus to PCI bus bridge module

Input/Output Slot Usage

Figure 7-3 and Table 7-2 show the input/output slot usage for the minimum and maximum configurations of this HiTest Template.

Figure 7-3: I/O Slot Usage

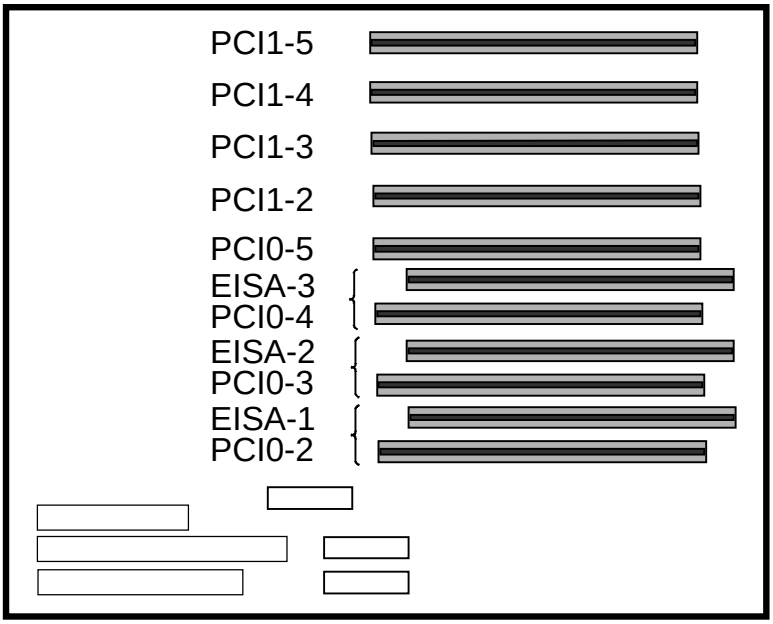

ML013980

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

Slots	Minimum Configuration	Maximum Configuration	Description
PCI1-5	DEFPA-AB	DEFPA-AB	PCI-FDDI adapter
PCI1-4	–	KZPAA-AA	FWD SCSI controller
PCI1-3	–	KZPBA-CB	FWD SCSI controller
PCI1-2	–	KZPBA-CA	FWSE SCSI adapter
PCI0-5	KZPBA-CB	KZPBA-CB	FWD SCSI controller
EISA-3/ PCI0-4	KZPBA-CB	KZPBA-CB	FWD SCSI controller
EISA-2/ PCI0-3	DE500-AA	DE500-AA	Ethernet controller
EISA-1/ PCI0-2	PB2GA-JC	PB2GA-JC	S3 TRIO64 1 MB Graphics Adapter