



Oracle Rdb7
OpenVMS AlphaServer 4x00

DIGITAL HiTest Notes

Part Number: EK-HORVA-HN. A01

January 1998

Revision/Update Information: This is a new document.

**Digital Equipment Corporation
Maynard, Massachusetts**

January 1998

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

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

© Digital Equipment Corporation 1998. All rights reserved.

The following are trademarks of Digital Equipment Corporation: AlphaServer, DEC C, DECevent, DIGITAL, eXcursion, HSJ, OpenVMS, PATHWORKS, POLYCENTER, ServerWORKS, StorageWorks, VAX, VMScluster, and the DIGITAL logo.

The following are third-party trademarks:

Intel is a registered trademark of Intel Corporation.

Oracle is a registered trademark and Oracle Rdb7 is a trademark of Oracle Corporation.

Windows 95 and Windows NT are trademarks of Microsoft Corporation.

All other trademarks are the property of their respective owners.

Table of Contents

1 Introduction

DIGITAL HiTest Suite and Its Advantages	1-1
Overview of This DIGITAL HiTest Suite	1-2

2 Configuration Data

Hardware and Software Components	2-1
Special Configuration Rules.....	2-5

3 System Installation and Setup

Hardware Installation.....	3-1
AlphaServer Systems	3-1
Operating System Installation	3-1
Layered Products	3-1
Application.....	3-2
CPG Database.....	3-2

4 Interoperability Tests and Results

Overview of Results.....	4-1
Test Environment	4-2
SYSGEN Parameters	4-3
SETPARAMS.DAT Values for the Minimum Configuration	4-4
SETPARAMS.DAT Values for the Maximum Configuration.....	4-6
Storage Subsystem.....	4-8
Storage Management	4-9
Disk Setup	4-9
Test Database Structure	4-10
Test Data Distribution.....	4-10
Global and Row Cache Buffers.....	4-12
Test Database Creation	4-13
Test Scenarios and Scripts	4-15
VLM-Enabled Test Scenarios	4-15
Test Scenario 1: Single-User Tests (Cold Cache)	4-15
Test Scenario 2: Single-User Tests (Warm Cache).....	4-15
Test Scenario 3: Multi-User Tests	4-16
Non-VLM Enabled Test Scenarios.....	4-16
Test Scenario 4: Distributed 10-User Load Tests.....	4-16
Test Scenario 5: Distributed 20-User Load Tests.....	4-16

Contents

RMU Backup and Restore and Fault Insertion/Error Recovery Test Scenarios	4-16
Test Scenario 6: RMU Backup and Restore Tests	4-16
Test Scenario 7: Fault Insertion/Error Recovery Tests	4-17
Purpose of Query Scripts	4-17
Test Tools	4-17
Test Configuration	4-17
Test Process	4-18
Query 1	4-18
Query 2	4-20
Query 3	4-21
Query 4	4-23
Query 5	4-24
Test Results	4-26
Test 1: Single-User Test (Cold Cache)	4-26
Test 2: Single-User Test (Warm Cache)	4-27
Test 3: Multi-User Load Tests	4-27
Non-VLM Enabled Database Query Tests	4-28
Test 4: Distributed 10-User Load Test	4-28
Test 5: Distributed 20-User Load Test	4-28
RMU Backup and Restore Tests	4-29
Test 6: Error Recovery	4-29
Disk Failure	4-29
HSJ Failure	4-30
System Power Failure	4-31

5 System Limits and Characterization Data

Comparison of Run Times	5-1
Database Buffering Schemes	5-2
Workload Characterizations	5-3
RDB Account Quotas	5-5
SYSTEM Account Quotas	5-5
Accounting Statistics from a One-User Load Test	5-6

6 Problems and Resolutions

Hardware	6-1
Operating System	6-1
Application	6-1

A Detailed Hardware Configuration

System Diagram	A-2
AlphaServer 4100 Configurations	A-3
AlphaServer 4100 PCI Slot Usage	A-4
Database Storage Maps	A-5

B Operational Command and Supporting Files

Figures

Figure 4-1: Test Environment	4-2
Figure 4-2: Oracle Rdb7 OpenVMS Storage Layout	4-8
Figure A-1: System Diagram	A-2
Figure A-2: AlphaServer 4100 Motherboard	A-3
Figure A-3: AlphaServer 4100 PCI Slot Usage	A-4

Tables

Table 2-1: Oracle Rdb7 OpenVMS AlphaServer 4x00 DIGITAL HiTest Template.....	2-2
Table 2-2: System Management Station Template	2-4
Table 2-3: Component Revision Levels	2-5
Table 4-1: CPG Database.....	4-10
Table 4-2: Data Distribution Across Channels	4-11
Table 4-3: Allocation of Row Cache Buffers	4-12
Table 4-4: Test Conditions.....	4-26
Table 4-5: Single-User, Cold Cache Test Results.....	4-26
Table 4-6: Single-User, Warm Cache Test Results.....	4-27
Table 4-7: 5- and 10-User Load Test Results for the Minimum Configuration	4-27
Table 4-8: 5- and 10-User Load Test Results for the Maximum Configuration.....	4-28
Table 4-9: Non-VLM Enabled 10-User Load Test Results	4-28
Table 4-10: RMU Backup and Restore Test Results for the Maximum Configuration	4-29
Table 4-11: Actual Elapsed Query Times.....	4-30
Table 4-12: Actual Elapsed Query Times.....	4-30
Table 4-13: Actual Elapsed Query Times.....	4-31
Table 5-1: Sample of RDB Account Quotas.....	5-5
Table 5-2: Sample of SYSTEM Account Quotas.....	5-5
Table 5-3: Accounting Statistics from a Cold Cache One-User Load Test.....	5-6
Table 5-4: Accounting Statistics from a Warm Cache One-User Load Test	5-6
Table A-1: AlphaServer 4100 Usage (Minimum and Maximum Configurations)	A-3
Table A-2: PCI Slot Usage (Minimum and Maximum Configurations)	A-4
Table A-3: Database Storage Map for HSJ001	A-5
Table A-4: Database Storage Map for HSJ002	A-6
Table A-5: Database Storage Map for HSJ003	A-7

Preface

This document provides an overview of DIGITAL HiTest Suites and detailed technical information about interoperability test results for the Oracle Rdb7 OpenVMS AlphaServer 4x00 HiTest Suite.

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.

Road Map

This document contains the following chapters:

1. Introduction – Provides a brief summary of the benefits of DIGITAL HiTest Suites and an overview of the Suite covered in this document.

2. Configuration Data – Includes tables of configuration data about the hardware and software components that define the Template, and special configuration rules if any.

3. System Installation and Setup – Presents useful information for installing and setting up this DIGITAL HiTest Suite.

4. Interoperability Tests and Results – Describes how the tests were set up (including database organization), what data and programs were placed on what disks, and how the tests were run.

5. System Limits and Characterization Data – Summarizes any system limitations or characterization data that were identified during testing.

6. Problems and Resolutions – Discusses any problems and resolutions that were discovered during testing.

Appendix A: Detailed Hardware Configuration – Contains more detailed information about the hardware and software components listed in the Configuration Data chapter.

Appendix B: Operational Command and Supporting Files – Describes the command procedures and SQL scripts that were used to set up, load, or run tests on the CPG database.

Feedback and Ordering Information

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 document title and part number (EK-HORVA-HN. A01) in your correspondence about this document.

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

DIGITAL HiTest Suite and Its Advantages

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 for customers with an existing VMScluster investment.

DIGITAL Product Management and Engineering select the components and design the configurations in each HiTest Suite to ensure high system reliability, application performance, and upgradability. A Suite's hardware and software components have been successfully tested for interoperability.

A HiTest Suite specifies allowed ranges of hardware and software components, as well as each component's part number, description, and revision information. These specifications are listed in the *DIGITAL HiTest Template*.

The components in a HiTest Suite are organized into two groups, the *DIGITAL HiTest Foundation* and the *DIGITAL HiTest AppSet*. The HiTest Foundation includes the hardware, operating system, middleware, and database software. The HiTest Foundation can be used as a base on which any customer-desired applications can be installed. The HiTest AppSet includes the software specific to one class of customer solutions.

Configuring a DIGITAL HiTest Suite is straightforward. 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.

The HiTest Suite is documented in the *DIGITAL HiTest Notes*. The HiTest Notes list the HiTest Foundation and HiTest AppSet components. HiTest Notes also describe the testing of the Suite and include configuration details, installation instructions, tuning parameters, problems encountered and their solutions, and system diagrams.

Some components listed in the HiTest Foundation or AppSet may be optional. If the minimum quantity is zero (0), then the component is optional. If the minimum quantity is one or more, then you must order at least the minimum quantity.

The maximum quantities represent the largest group of components that were tested for interoperability with all the other components in the Suite. Although it may be possible to place more than the specified maximum quantity of a component on a DIGITAL system, extensive interoperability testing was not done at that level and such a system would not be considered a DIGITAL HiTest System.

Introduction

You can select any combination of components with quantities ranging from the minimum to the maximum specified. Occasionally, special configuration rules give further guidance or restrict configurations. These rules appear in the Configuration Data chapter of the HiTest Notes.

A customer can include the Suite-specified hardware and software they need and then layer on additional software. Other types of hardware, called *add-on hardware*, can also be added to a DIGITAL HiTest System. Further information is available through the following URLs:

<http://cosmo.tay.dec.com> (Intranet)

<http://www.partner.digital.com:9003/cgi-bin/comet> (Internet)

Even though the customer may install application software that is not specified in the Suite, the customer and DIGITAL still experience the advantages of knowing that all of the Suite-based hardware and software interoperates correctly. Of course, the full benefit of configuring a system from a HiTest Suite is obtained when the system includes only specified HiTest Foundation and AppSet components.

Overview of This DIGITAL HiTest Suite

The Oracle Rdb7 OpenVMS AlphaServer 4x00 HiTest Suite consists of the following software components:

- Oracle Rdb7, Version 7.0.1.0
- OpenVMS, Version 7.1

This Suite will meet the add-on Data Warehouse needs of customers with an existing VMSccluster investment.

Configuration Data

This chapter describes the tested DIGITAL HiTest Configuration Suite including the hardware, software, and firmware components, and their revision levels. Special configuration rules are explained if required.

Hardware and Software Components

Table 2-1 identifies the range of hardware and software components that can be configured using the Oracle Rdb7 OpenVMS AlphaServer 4x00 HiTest Suite. This range includes both the AlphaServer 4000 and the AlphaServer 4100 systems with two to four central processing unit (CPU) options and 2 GB to 4 GB of memory.

Table 2-2 lists the System Management Station hardware and software.

Table 2-3 lists the revision levels of the components.

The HiTest Template (Table 2-1) consists of three categories:

- **AppSet Software** – Includes software specific to one class of customer solutions, in this case add-on Data Warehousing for customers with an existing VAX VMScluster investment.
- **Foundation Hardware** – Includes the base system, storage, and other hardware options.
- **Foundation Software** – Includes the operating system, middleware, and database software.

When ordering an item from a HiTest Template, select a quantity that is within the minimum/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, then order at least the minimum quantity, but not more than 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 Appendix A.

Table 2-1: Oracle Rdb7 OpenVMS AlphaServer 4x00 DIGITAL HiTest Template

Oracle Rdb7 HiTest AppSet				
OpenVMS AlphaServer 4x00 HiTest Foundation				
For documentation and updates: http://cosmo.tay.dec.com and http://www.partner.digital.com:9003/cgi-bin/comet For hardcopy of this Suite's HiTest Notes, order EK-HORVA-HN.				
Line Item	Description	Part Number	HiTest Range	
			Min	Max
AppSet Software				
1	Oracle Rdb7 To order, call 1-800-ORACLE1 or send E-mail to: infodec@us.oracle.com	Oracle	1	1
Foundation Hardware				
2	<i>Select just one base system:</i> AlphaServer 4000 5/466 Drawer, OpenVMS, 2 GB AlphaServer 4100 5/466 Drawer, OpenVMS, 2 GB <i>Hardware includes:</i> <ul style="list-style-type: none"> • 5/466 MHz CPU with 4 MB cache • 2-GB of memory • PB2GA-JB TRIO64 1 MB Graphics • DE500-AA 10/100 Mbit Fast Ethernet • KZPSA-AA FW SCSI and cable • SCSI CD-ROM drive • RX23L-AB 1.44 MB Floppy drive • LK47W-A2 PS/2 style keyboard • Three-button PS/2 compatible mouse <i>Software includes:</i> <ul style="list-style-type: none"> • OpenVMS operating system license and media • DIGITAL EIP, consisting of the following services: <ul style="list-style-type: none"> – Management – Network Transport and Connectivity – File and Print – Internet – Database – Mail – Application Development – Presentation 	DY-53JEB-GA DY-51JAB-GB	1	1
3	Pedestal with StorageWorks shelf <i>Hardware includes:</i> <ul style="list-style-type: none"> • 4.3 GB Wide Disk • Pedestal with one StorageWorks shelf and mounting kit 	BA30P-AB	1	1
4	466 MHz CPU OpenVMS SMP UPG	KN304-DC	1	3
5	2 GB Memory Option	MS330-GA	0	1
6	PCI to CI Adapter	CIPCA-AA	1	2
7	CI Bus cable	BNCIA-20	1	2
8	4.3 GB 7200 RPM UltraSCSI Disks	DS-RZ1CB-VW	1	3
9	PCI 10/100-Mbit Fast Ethernet controller	DE500-AA	0	1
10	<i>Select one high-resolution color monitor (the 21" is recommended):</i> 15-inch monitor 17-inch monitor 21 inch monitor	SN-VRCX5-WA SN-VRCX7-WA SN-VRCX1-WA	1	1
11	System Management Station	See Table 2-2	0	1

Oracle Rdb7 HiTest AppSet				
OpenVMS AlphaServer 4x00 HiTest Foundation				
For documentation and updates: http://cosmo.tay.dec.com and http://www.partner.digital.com:9003/cgi-bin/comet For hardcopy of this Suite's HiTest Notes, order EK-HORVA-HN.				
Line Item	Description	Part Number	HiTest Range	
			Min	Max
Foundation Software				
12	OpenVMS, V7.1	Included with item 2	–	–
13	DIGITAL Enterprise Integration Package	Included with item 2	–	–
14	OpenVMS Media and documentation on CD-ROM Note: Media and documentation are required for the first system at the site.	QA-MTIAA-H8	0	1
15	OpenVMS Cluster Software License	QL-MUZAC-AA	1	1
16	OpenVMS Alpha Software Product Library Package Note: The OpenVMS Alpha Software Product Library Package is required for the first system at the site. The OpenVMS Alpha Software Product Library Package includes items 19-22 on this table.	QA-03XAA-W8	0	1
17	Oracle SQL/Services To order, call 1-800-ORACLE1 or send E-mail to: infodec@us.oracle.com	Oracle	1	1
18	Oracle Parallel Query & Distributed Option	Included with item 1	1	1
19	License for DEC C for OpenVMS Alpha V5.5 Note: The asterisk (*) in the part number for this line item denotes variant fields. For additional information about available licenses, services, and media, refer to the appropriate price book.	QL-015A*-**	0	1
20	License for DEC C++ for OpenVMS Alpha V5.5 Note: The asterisk (*) in the part number for this line item denotes variant fields. For additional information about available licenses, services, and media, refer to the appropriate price book.	QL-MJ1A*-**	0	1
21	License for UCX V4.1 (for PCM use only) Note: The asterisk (*) in the part number for this line item denotes variant fields. For additional information about available licenses, services, and media, refer to the appropriate price book. Note: UCX, also known as DIGITAL TCP/IP Services for OpenVMS, is included with item 13.	QL-0LXA*-**	0	1
22	DECevent V2.3	Included with item 2	0	1

Table 2-2: System Management Station Template

Oracle Rdb7 HiTest AppSet System Management Station				
Line Item	Description	Part Number	HiTest Range	
			Min	Max
System Management Station Hardware				
Note: This HiTest Suite supports the use of a system management station. When the management 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 PWS 200i Personal Workstation for Windows NT <i>Hardware includes:</i> <ul style="list-style-type: none"> • 200 MHz CPU with 256 KB cache • 64 MB memory • Matrox Millenium 3D graphics • 10BaseT/10Base2 Ethernet • 2.0 GB UW disk • EIDE CD-ROM • 1.44 MB floppy • PS/2 style keyboard • Two-button PS/2 compatible mouse <i>Software includes:</i> <ul style="list-style-type: none"> • Windows NT 4.0 Note: A functionally equivalent X86 system may be substituted without invalidating this HiTest Template.	SN-B3KAP-EL	1	1
2	<i>Select one high-resolution color monitor (the 21" is recommended):</i> 15-inch monitor 17-inch monitor 21-inch monitor	SN-VRCX5-WA SN-VRCX7-WA SN-VRCX1-WA	1	1
System Management Station Software				
3	Windows NT Client, Version 4.0 media, documentation and license	Included with item 1	1	1
4	Windows NT Service Pack 3	Microsoft	1	1
5	DIGITAL PATHWORKS 32, Version 7.0A Media and Documentation	Included with item 13 in Table 2-1	1	1
6	DIGITAL PATHWORKS 32, Version 7.0A system license <i>Includes:</i> <ul style="list-style-type: none"> • PATHWORKS 32 networking components • eXcursion for Windows 95 and Windows NT 4.0 • PowerTerm 525 terminal emulator 	Included with item 13 in Table 2-1	1	1
7	OpenVMS Management Tools (OMT) V2.0	Included with base system	1	1

The following table lists the revision levels of the components.

Table 2-3: Component Revision Levels

Hardware Component	Hardware	Firmware	Software
SRM Console	–	V4.9-20	–
Fast/Wide SCSI Controller (KZPSA-AA)	–	2	–
CI Controller (HSJ50)	–	V3.1J	–
4.3 GB 7200 RPM UltraSCSI disk (DS-RZ1CB-VW)	A01	–	–
StorageWorks Shelf Power Supply (DS-BA35X-HH)	B01 or greater	–	–
466 MHz CPU OpenVMS SMP UPG (KN304-DC)	A02	–	–
2 GB Memory Option (MS330-GA)	E	–	–
PCI-CI Host Bus Adapter (CIPCA-AA)	1	A315	–
Software Component		Version/Revision	Patch Level
Oracle Rdb7		7.0.1.0	–
OpenVMS		7.1	Alpcpuc02_071
DEC C for OpenVMS Alpha		5.5	–
DEC C++ for OpenVMS Alpha		5.5	–
DECevent		2.3	–
UCX		4.1	–

Special Configuration Rules

There are no special configuration rules for the Oracle Rdb7 OpenVMS AlphaServer 4x00 HiTest Suite.

System Installation and Setup

This chapter presents information that is useful when installing and setting up a DIGITAL HiTest System configured from this DIGITAL HiTest Suite. System preparation includes installation of the hardware, operating system, and applications.

Hardware Installation

The hardware was installed and interconnected as shown in Appendix A.

AlphaServer Systems

DIGITAL recommends the installation of the AlphaServer 4x00 system be performed by authorized DIGITAL Service personnel. Installation procedures are provided with the systems when shipped. Installation manuals are also provided with the optional components ordered, when appropriate.

Appendix A contains the detailed hardware configuration for the systems, beginning with a detailed graphical view of the test cluster configuration.

Operating System Installation

Install standard OpenVMS VAX Version 7.1 and OpenVMS Alpha 7.1 operating systems using all the default values. Refer to Chapter 4 for suggestions on system-specific tuning of account quotas and SYSGEN parameters.

Layered Products

Install the following versions of layered products in the order shown, using all the standard defaults:

1. DECEvent 2.3
2. DEC C 5.5
3. DEC C++ 5.5
4. UCX 4.1
5. Oracle Rdb7 Release 7.0.1.0, Patch Level 10
6. Oracle SQL/Services
7. Oracle Parallel Query & Distributed Option

System Installation and Setup

Application

Install Oracle Rdb7 as described in the installation guide that is supplied with the Oracle Rdb7 CD-ROM.

CPG Database

For the testing of this HiTest Suite, DIGITAL HiTest Engineering used the Consumer Packaged Goods (CPG) database from Oracle Corporation. During the testing of this HiTest Suite, parameters were modified as described in Chapter 4 of this document.

Interoperability Tests and Results

This chapter describes how the tests were set up (including database organization), what data and programs were placed on what disks, and how the tests were run.

This chapter describes:

- Test Environment
- Test Tools
- Test Configuration
- Test Process and Results

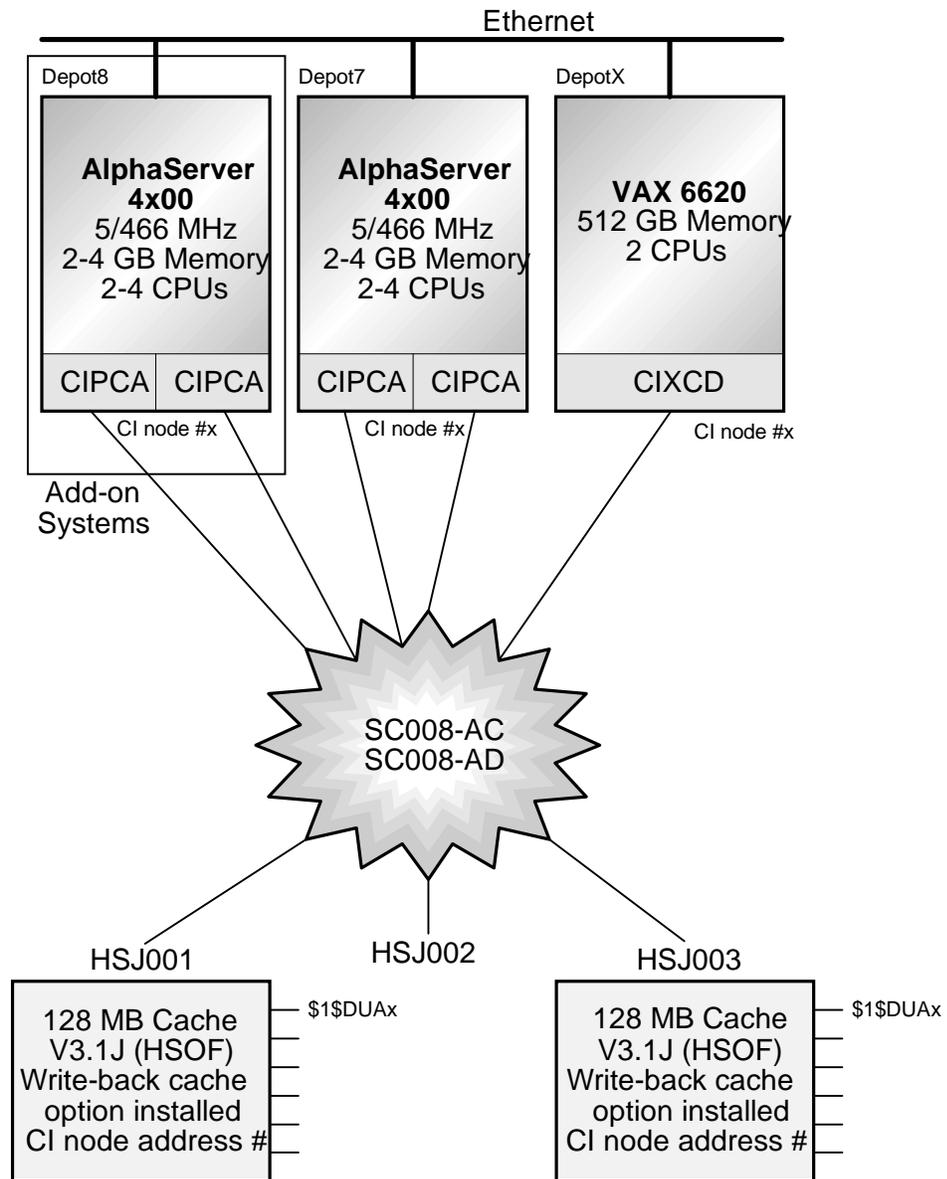
Overview of Results

Interoperability testing was performed successfully on the Oracle Rdb7 OpenVMS AlphaServer 4x00 HiTest Suite. The tests verified that the business processes functioned correctly.

Test Environment

Figure 4-1 shows the Oracle Rdb7 test environment. The test cluster consisted of either one or two AlphaServer 4100s connected to an existing CI-based VMScluster with a VAX 6620 systems.

Figure 4-1: Test Environment



ML014189

SYSGEN Parameters

The following MODPARAMS.DAT file was used by AUTOGEN to set up system parameters for the initial testing of this HiTest Suite.

Note

The initial testing of this suite included running a 10-user load every four hours for a 24-hour period. After the initial testing of this suite, AUTOGEN was run with feedback enabled on a 4100 system in the minimum configuration and a 4100 in the maximum configuration. This produced the output listed in the “SETPARAMS.DAT Values for the Minimum Configuration” and “SETPARAMS.DAT Values for the Maximum Configuration” sections.

MODPARAMS.DAT File for AlphaServer 4100

```

SCSSYSTEMID=64519                                min_PQL_DBIOLM=16384
SCSNODE="DEPOT7  "                                min_PQL_MBIOLM=16384
VAXCLUSTER=2                                      min_PQL_DDIOLM=16384
EXPECTED_VOTES=1                                  min_PQL_MDIOLM=16384
VOTES=1                                           min_PQL_DFILLM=16384
RECNXINTERVAL=20                                  min_PQL_MFILLM=16384
DISK_QUORUM=" "                                    min_PQL_DPRCLM=16384
QDSKVOTES=1                                       min_PQL_MPRCLM=16384
QDSKINTERVAL=10                                  min_PQL_DTQELM=16384
ALLOCLASS=1                                       min_PQL_MTQELM=16384
LOCKDIRWT=3                                       min_PQL_DENQLM=16384
NISCS_CONV_BOOT=0                                 min_PQL_MENQLM=16384
NISCS_LOAD_PEA0=1                                 min_PQL_DPGFLQUOTA=500000
NISCS_PORT_SERV=0                                 min_PQL_MPGFLQUOTA=500000
MSCP_LOAD=1                                       min_s2_size=500
MSCP_SERVE_ALL=1
min_NPAGEDYN=50000000 !increased on 17-Feb-1997
min_ERLBUFFERPAGES=24 !increased on 05-Mar-1997
PAGEFILE=0 !set to hold file size with reduced memory 30-Apr
SWAPFILE=0 !set to hold file size with reduced memory 30-Apr
DUMPFIL=0 !set to hold file size with reduced memory 30-Apr
CHANNELCNT=16384
VCC_FLAGS=1
VCC_MAXSIZE=204800
min_GBLPAGES=512000
min_GBLPAGFIL=512000
min_PQL_DASTLM=16384
min_PQL_MASTLM=16384

```

Interoperability Tests and Results

SETPARAMS.DAT Values for the Minimum Configuration

The following SETPARAMS.DAT file was generated as output from AUTOGEN during the initial testing of the minimum configuration of this HiTest Suite.

```
use current                               set PQL_MBYTLM 100000
write sys$system:alphavmssys.old         set PQL_DDIOLM 16384
use default                               set PQL_MDIOLM 16384
set PFCDEFAULT 64                        set PQL_DFILLM 16384
set GBLSECTIONS 600                      set PQL_MFILLM 16384
set GBLPAGES 8999571                     set PQL_DPGFLQUOTA 500000
set GBLPAGFIL 512000                     set PQL_MPGFLQUOTA 500000
set MAXPROCESSCNT 674                    set PQL_DPRCLM 16384
set MULTITHREAD 2                         set PQL_MPRCLM 16384
set PROCSECTCNT 64                       set PQL_DTQELM 16384
set MINWSCNT 20                           set PQL_MTQELM 16384
set PAGFILCNT 4                           set PQL_DWSDEFAULT 5680
set SWPFILCNT 2                           set PQL_MWSDEFAULT 5680
set SYSMWCNT 10250                        set PQL_DWSQUOTA 11360
set BALSETCNT 672                         set PQL_MWSQUOTA 11360
set WSMAX 786432                          set PQL_DWSEXTENT 786432
set NPAGEDYN 50000000                     set PQL_MWSEXTENT 786432
set NPAGEVIR 200000000                   set PQL_DENQLM 16384
set PAGEDYN 6356992                       set PQL_MENQLM 16384
set QUANTUM 20                             set ACP_MAPCACHE 420
set MPW_WRTCLUSTER 64                     set ACP_HDRCACHE 1682
set MPW_HILIMIT 65536                     set ACP_DIRCACHE 1682
set MPW_LOLIMIT 2016                      set ACP_DINDXCACHE 420
set MPW_IOLIMIT 4                          set ACP_QUOCACHE 674
set MPW_THRESH 4032                       set ACP_SYSACC 18
set MPW_WAITLIMIT 65792                   set ACP_SWAPFLGS 14
set MPW_LOWAITLIMIT 65280                 set VAXCLUSTER 2
set S2_SIZE 500                            set EXPECTED_VOTES 1
set PFRATL 0                               set VOTES 1
set PFRATH 8                              set RECNXINTERVAL 20
set WSINC 2400                             set DISK_QUORUM " "
set WSDEC 4000                             set ALLOCLASS 1
set LONGWAIT 30                           set LOCKDIRWT 3
set ERRORLOGBUFFERS 6                     set NISCS_LOAD_PEA0 1
```

Interoperability Tests and Results

```
set DUMPSTYLE 9
set MVTIMEOUT 3600
set FREELIM 692
set FREEGOAL 2688
set GROWLIM 692
set BORROWLIM 692
set CLISYMTBL 512
set LOCKIDTBL 32512
set RESHASHTBL 32768
set SCSBUFFCNT 512
set SCSCONNCNT 18
set SCSRESPCNT 300
set SCSSYSTEMID 64519
set SCSNODE "DEPOT7"
set PRCPOLINTERVAL 30
set PASTDGBUF 16
set LNMSHASHTBL 512
set ACP_MULTIPLE 0
set RMS_DFMBC 16
set RMS_DFMBFSKD 0
set RMS_DFMBFSMT 0
set RMS_DFMBFSUR 0
set RMS_DFMBFREL 0
set RMS_DFMBFIDX 0
set RMS_DFMBFHSH 0
set RMS_DFNBC 8
set PQL_DASTLM 16384
set PQL_MASTLM 16384
set PQL_DBIOLM 16384
set PQL_MBIOLM 16384

set NISCS_PORT_SERV 0
set MSCP_LOAD 1
set TMSCP_LOAD 0
set MSCP_SERVE_ALL 1
set MSCP_BUFFER 128
set MMG_CTLFLAGS 3
set ERLBUFFERPAGES 24
set SHADOWING 2
set SHADOW_MAX_COPY 18
set ZERO_LIST_HI 4096
set GH_EXEC_CODE 1024
set GH_EXEC_DATA 96
set GH_RES_CODE 1024
set GH_RES_DATA 0
set NISCS_MAX_PKTSZ 1498
set NPAG_BAP_MIN 565248
set NPAG_BAP_MAX 8519680
set NPAG_BAP_MAX_PA 1073741824
set PIXSCAN 10
set VIRTUALPAGECNT 2147483647
set CHANNELCNT 16384
set PIOPAGES 336
set CTLPAGES 100
set PHYSICAL_MEMORY %XFFFFFFFF
set LOAD_SYS_IMAGES 7
set VCC_FLAGS 1
set VCC_MAXSIZE 204800
write current
write sys$system:autogen.par
exit
```

Interoperability Tests and Results

SETPARAMS.DAT Values for the Maximum Configuration

The following SETPARAMS.DAT file was generated as output from AUTOGEN during the initial testing of the maximum configuration of this HiTest Suite.

```
use current                set PQL_MBYTLM 100000
write sys$system:alphavmssys.old  set PQL_DDIOLM 16384
use default                set PQL_MDIOLM 16384
set PFCDEFAULT 64         set PQL_DFILLM 16384
set GBLSECTIONS 600       set PQL_MFILLM 16384
set GBLPAGES 8999571     set PQL_DPGFLQUOTA 500000
set GBLPAGFIL 512000     set PQL_MPGFLQUOTA 500000
set MAXPROCESSCNT 674    set PQL_DPRCLM 16384
set MULTITHREAD 4        set PQL_MPRCLM 16384
set PROCSECTCNT 64       set PQL_DTQELM 16384
set MINWSCNT 20          set PQL_MTQELM 16384
set PAGFILCNT 4          set PQL_DWSDEFAULT 10064
set SWPFILCNT 2          set PQL_MWSDEFAULT 10064
set SYSMWCNT 10250       set PQL_DWSQUOTA 20128
set BALSETCNT 672        set PQL_MWSQUOTA 20128
set WSMAX 1091584        set PQL_DWSEXTENT 1091584
set NPAGEDYN 50000000    set PQL_MWSEXTENT 1091584
set NPAGEVIR 200000000   set PQL_DENQLM 16384
set PAGEDYN 6356992     set PQL_MENQLM 16384
set QUANTUM 20           set ACP_MAPCACHE 420
set MPW_WRTCLUSTER 64    set ACP_HDRCACHE 1682
set MPW_HILIMIT 131072   set ACP_DIRCACHE 1682
set MPW_LOLIMIT 2016    set ACP_DINDXCACHE 420
set MPW_IOLIMIT 4        set ACP_QUOCACHE 674
set MPW_THRESH 4032     set ACP_SYSACC 17
set MPW_WAITLIMIT 131328 set ACP_SWAPFLGS 14
set MPW_LOWAITLIMIT 130816 set VAXCLUSTER 2
set S2_SIZE 500          set EXPECTED_VOTES 1
set PFRATL 0             set VOTES 1
set PFRATH 8             set RECWXINTERVAL 20
set WSINC 2400           set DISK_QUORUM "
set WSDEC 4000           set ALLOCLASS 1
set LONGWAIT 30         set LOCKDIRWT 3
set ERRORLOGBUFFERS 10  set NISCS_LOAD_PEA0 1
```

Interoperability Tests and Results

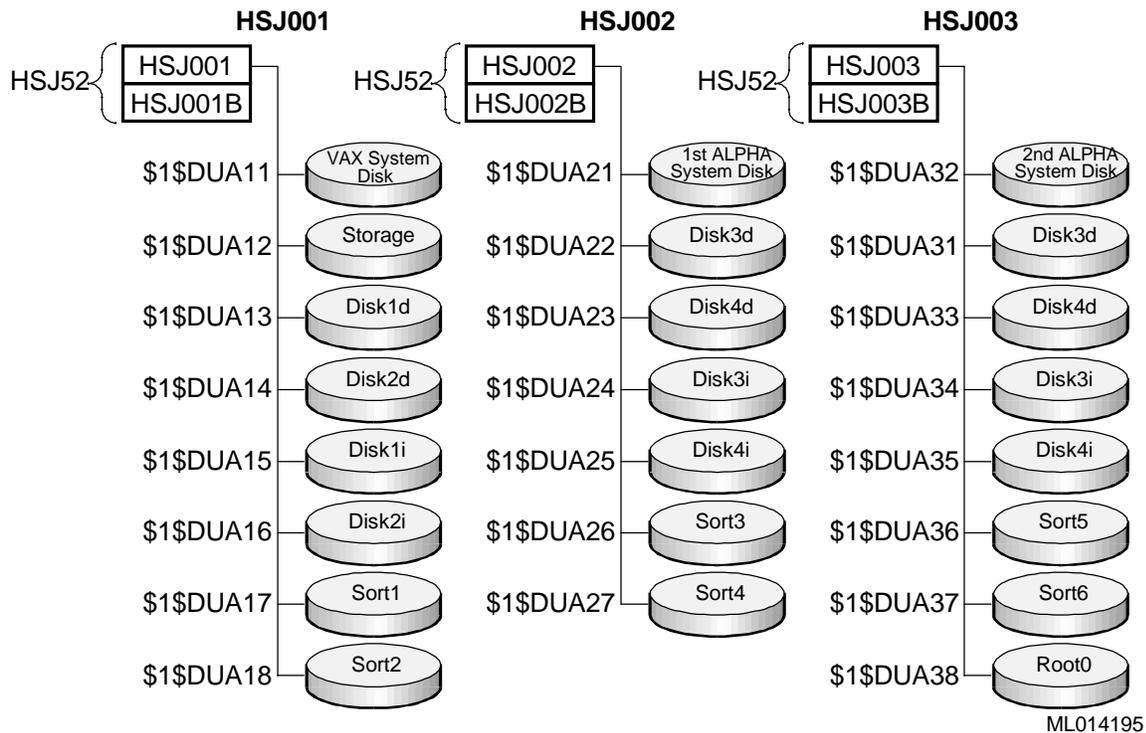
```
set DUMPSTYLE 9
set MVTIMEOUT 3600
set FREELIM 692
set FREEGOAL 2688
set GROWLIM 692
set BORROWLIM 692
set CLISYMTBL 512
set LOCKIDTBL 32512
set RESHASHTBL 32768
set SCSBUFFCNT 512
set SCSCONNCNT 18
set SCSRESPCNT 300
set SCSSYSTEMID 64519
set SCSNODE "DEPOT7"
set PRCPOLINTERVAL 30
set PASTDGBUF 16
set LNMSHASHTBL 512
set ACP_MULTIPLE 0
set RMS_DFMBC 16
set RMS_DFMBFSKD 0
set RMS_DFMBFSMT 0
set RMS_DFMBFSUR 0
set RMS_DFMBFREL 0
set RMS_DFMBFIDX 0
set RMS_DFMBFHSH 0
set RMS_DFNBC 8
set PQL_DASTLM 16384
set PQL_MASTLM 16384
set PQL_DBIOLM 16384
set PQL_MBIOLM 16384

set NISCS_PORT_SERV 0
set MSCP_LOAD 1
set TMSCP_LOAD 0
set MSCP_SERVE_ALL 1
set MSCP_BUFFER 128
set MMG_CTLFLAGS 3
set ERLBUFFERPAGES 24
set SHADOWING 2
set SHADOW_MAX_COPY 18
set ZERO_LIST_HI 8192
set GH_EXEC_CODE 1024
set GH_EXEC_DATA 96
set GH_RES_CODE 1024
set GH_RES_DATA 0
set NISCS_MAX_PKTSZ 1498
set NPAG_BAP_MIN 565248
set NPAG_BAP_MAX 8519680
set NPAG_BAP_MAX_PA 1073741824
set PIXSCAN 20
set VIRTUALPAGECNT 2147483647
set CHANNELCNT 16384
set PIOPAGES 336
set CTLPAGES 100
set PHYSICAL_MEMORY %FFFFFFFF
set LOAD_SYS_IMAGES 7
set VCC_FLAGS 1
set VCC_MAXSIZE 204800
write current
write sys$system:autogen.par
exit
```

Storage Subsystem

Figure 4-2 provides a high-level view of the storage disk layout for the test cluster. The storage disk layout consisted of three Dual-Redundant HSJ50 Controllers and 108 4.3 GB disks in an SW800 cabinet and 18 shelves.

Figure 4-2: Oracle Rdb7 OpenVMS Storage Layout



For more details, see the database storage maps in Appendix A.

The storage subsystem consisted of the following components:

- Star Coupler, eight-node with cabinet
- Star Coupler, eight-node add on
- Three basic controller shelves for HSJ50 Controllers
- Six StorageWorks Array Controllers, 6 channel CI, 36–42 devices, 32 MB read cache
- 108 4.3 GB fixed wide disks
- 18-shelf expansion units
- One SW800 Data Center Cabinets
- 18 8-bit I/O modules for BA356 shelf
- 18 SCSI-2 hi-density “a” cables, 1.5 m
- Six CI bus cables, 20 m

Storage Management

It is understood that a customer's database and disk layout may already exist. The following information describes the layout of the test environment disk storage.

POLYCENTER Console Manager (PCM) was used in the test environment to perform the following functions:

- Shutting down systems
- Rebooting systems
- Running standalone diagnostics
- Installing layered products
- Configuring HSJ controllers
- Monitoring system and HSJ controller status

PCM was also used to manage the test HSJ configuration and capture failure messages, alerting the need to repair or replace faulty disk drives. Alarm conditions were reported to a single workstation using PCM. HSJ50 consoles were connected to terminal servers, which were reported as nodes to the PCM, simplifying the management of a large configuration of over 100 disk drives.

Disk Setup

PCM was used to set up disks as follows:

- The test database data was stored on six five-member RAID 5 arrays.

Note

Write-back cache was enabled for all drives attached to the HSJ50 Controllers.

The Index drives consisted of six three-member RAID 5 arrays. The Storage and Root drives each consisted of a single four-member RAID 5 array and a single three-member RAID 5 array respectively.

- The OpenVMS VAX and OpenVMS Alpha operating systems were located on separate six-member RAID 5 RAIDsets.
- There were six sort drives made up of six four-member RAID 5 RAIDsets.
- The remaining spindles were used as sparesets for the RAID 5 arrays.

Refer to Appendix A for detailed storage maps. For more information, refer to *StorageWorks Array Controllers: HS Family of Array Controllers User Guide* (EK-HSFAM-UG).

Test Database Structure

The database was built using the Consumer Packaged Goods (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 includes one year of sales data in a star schema data warehouse optimized for decision support.

The CPG database, constructed with Oracle Rdb7, consisted of six tables:

- SALES_FACTS table, which consists of the bulk of the database
- Five dimension tables: CHANNEL, DAILY_PERIOD, MONTHLY_PERIOD, MARKET, and PRODUCT.

Table 4-1 shows the five table spaces in the CPG database:

Table 4-1: CPG Database

Table Space	Contents	Size
FACTS	SALES_FACTS Table 41 data files (varying sizes)	147.2 GB 2,799,497,440 rows
FACTSINDEX	SF_KEY Index on SALES_FACTs Table 522 data files	137 GB
DIMENSION	1 Data File Dimension Tables: CHANNEL DAILY_PERIOD MONTHLY_PERIOD MARKET PRODUCT	316 KB 41 rows 2189 rows 72 rows 1002 rows 522 rows
DIMINDEX	1 Data File Index for Dimension Tables	138 KB
TEMPFILE	24 data files (5.5 MB each) required for SF_KEY index build	137.6 GB

Test Data Distribution

Originally, storage areas were created assuming a uniform distribution of data by channel. The initial database design was to partition by channel for data storage, but to partition indexes by product. This benefited the time, channel, and queries by product.

Initial database loading, however, showed that data distribution was not uniform across the 41 channels. There were eight instances each of five channel groups (Discount, Drug, Warehouse, Supermarket, Convenience) and one “Total” channel. The following table summarizes this distribution and the sizes of per-channel datasets as record count, storage page count, and MB size.

Table 4-2: Data Distribution Across Channels

Group	Records/Mo.	Pages/Mo.	MB/Mo.	Pages/Yr.	MB/Yr.
Warehouse	433,690	1,456	23	17,470	273
Discount	1,794,528	6,024	94	72,289	1,130
Drug	1,794,528	6,024	94	72,289	1,130
Convenience	2,661,908	8,936	140	107,230	1,675
Supermarket	7,951,376	26,692	417	320,307	5,005
“Total”	1,794,528	6,024	94	72,289	1,130
Total by group (41) assume 31 days	118,882,768	399,082	6,236	4,788,979	74,828
Data by year assume 31 days	1,426,593,216	4,788,979	74,828		
Index by year assume 31 days	1,426,593,216	4,788,979	74,828		
2-year small DBMS	2,853,186,432	9,577,957	299,311		

Global and Row Cache Buffers

The query scripts were run in different buffer configurations:

- Global buffers set to 20,000, 100 users, 200 buffers/user, and VLM row caches enabled on the AlphaServer systems
- Global buffers only, 2500 total, maximum of 7 users, 100 buffers/user, on the VAX system

The row cache buffers were designed to hold all small tables and all referenced data in the larger tables and indexes as shown in Table 4-3.

Table 4-3: Allocation of Row Cache Buffers

Name	Record Length	Cache	Rows	Used	Bytes Allocated	Bytes Used
SYSTEM_CACHE		1,000	500	484	500,000	484,000
CHANNEL*	226	226	41	40	9,266	9,040
C_CHANNEL	960	960	10	2	9,600	1,920
C_CHANNEL_ID	960	960	10	0	9,600	0
C_CHNLGRP	215	216	10	3	2,160	648
DAILY_PERIOD*	70	70	2,189	2,189	153,230	153,230
D_DAY	960	960	100	0	96,000	0
D_MONTH	215	216	100	17	21,600	3,672
D_QUARTER	215	216	100	0	21,600	0
MARKET*	161	164	1,002	140	164,328	22,960
MONTHLY_PERIOD*	52	52	72	0	3,744	0
M_DISTRICT	215	216	100	18	21,600	3,888
M_MARKET	960	960	100	2	96,000	1,920
M_MARKET_ID	960	960	100	0	96,000	0
M_REGION	215	216	100	2	21,600	432
PRODUCT	332	332	522	21	173,304	6,972
P_BRAND	215	216	100	0	21,600	0
P_MFR	215	216	100	0	21,600	0
P_PRODUCT	960	960	100	3	96,000	2,880
P_PRODUCT_ID	960	960	100	0	96,000	0
SALES_FACTS*	55	56	10,000,000 [†]	2,394,477	560,000,000	134,090,712
SF_KEY	960	960	1,000,000 [†]	38,512	960,000,000	36,971,520
				Totals	1,521,634,832	171,753,794
* These are tables.						
[†] These numbers represent the size of row caches and should be compared to the “Bytes Allocated” column to gauge the actual memory utilization.						

Test Database Creation

The Oracle Rdb7 OpenVMS HiTest System was designed so that a series of custom .COM, .SQL, and other supporting files could be used to set up and configure the test system. These files were created by Oracle Corporation. Descriptions of these files are provided in Appendix B.

The custom files performed the following sequential functions:

1. Allocating database storage
2. Creating the database storage files
3. Compiling and linking flat file builder program
4. Generating 12, 24 or 72 input streams
5. Loading the database
6. Creating indexes and journals (needed for VLM)
7. Executing queries

The rest of this section explains these functions in the order performed. Unless otherwise specified, the command files, SQL scripts, and supporting files shown on the following pages are located in the top-level [RDB] account directory. See also Appendix B.

Step 1: Allocating database storage

The database was not automatically brought on line at boot time. When bringing the system up for the first time, the following command files were executed from the RDB account:

```
$ @INIT DISKS.DAT
$ @LOGICALS DISKS.DAT
```

For use in initial creation of the database only, the following was executed:

```
$ @INIT DISKS.TMP
$ @LOGICALS DISKS.TMP
```

These commands created and initialized the TEMP space needed to hold the flat file data.

If the system needed to be rebooted after it was set up, the MOUNT.COM file was used in place of the INIT.COM previously shown.

To establish the search list logicals, the following was executed:

```
$ @CPG.COM
```

Step 2: Creating the database storage files

For each AlphaServer system, the following command was used to create the files for an empty database:

```
$ SQL @CPG.SQL (executed on maximum configured AlphaServer 4100)
```

Note

The preceding commands also created the small (empty) dimension tables in addition to the sales facts table. The schema design for the database was influenced by the queries because they performed query comparisons by year of an item (for example, a channel) to a group, and unioned the two result sets together. Data was partitioned by channel and the index was partitioned by product.

Step 3: Compiling and linking flat file builder

These files were located in the [.LOAD] subdirectory.

Two variants were used, the original and the binary version. The RMU load utility could read a binary or ASCII file. If ASCII was used, RMU would have to convert it to binary. The more efficient method is to use the binary loader. Using a binary input stream also resulted in a 25% space savings (55 bytes vs. 72 bytes input record length) and reduced use of CPU resources. These savings were needed to fit three months of data per each temp device. In all, eight temp devices were used to represent two years worth of data (24 months).

The ASCII and binary flat file builder programs were:

```
FACT_FILE_BLD.C          (ASCII)
FACT_FILE_BIN.C          (binary)
```

The following commands compiled and linked either program:

```
$ CC file_name
$ LINK file_name
```

Step 4: Generating the input streams

These files were located in the [.LOAD] subdirectory.

Just as with the flat file builders, there were two versions of the load programs. Either loader program was executed and used the FACT_FILE_Bxx.EXE generated in step 3. The loader programs were:

```
$ @LOADER          (ASCII loader)
$ @BINARY          (Binary loader)
```

BINARY.COM is recommended since it is an improved (faster) version of LOADER.COM and uses less storage space.

These command files called the appropriate FACT_FILE_Bxx file. If BINARY.COM was executed with no parameters, 24 load files (months) would be created by default.

The input to the previous command files was the DAILYnn.DAT files that were provided with the CPG database. The .COM files read the DAILYxx.DAT files (where xx = the month number) and created a DAILYxx.UNL file. The .DAT files were nearly identical to the .UNL files except that the date was in a different format than the original (Oracle7) format. The date format for the RMU Loader enforced a 4-digit year as in 'YYYYMMDD' format vs. 'DD-MMM-YY' format used by the Oracle7 SQL loader input files.

Step 5: Loading the database

These files were located in the [.LOAD] subdirectory.

The small dimension tables were loaded by executing:

```
$ @LOAD-DIM
```

There were two variants of command files that were used to load the sales facts table depending on whether binary or ASCII load files were created. These were:

```
$ @RMUBLD.COM (ASCII)
$ @RMUBIN.COM (binary)
```

Again, execute only one of these, not both. The recommended file is RMUBIN.COM. There are multiple variants of the RMUBIN.COM. These variants end in 4, 8, 17, and 41, and represent the number of parallel loader processes to use to load the flat files. The RMUBIN17.COM file was found to be the most efficient.

Step 6: Creating the indexes and journals

The indexes were created as follows:

```
$ SQL @INDEXES
```

The journals were created as follows:

```
$ @JOURNALS
```

Step 7: Executing the queries

These files were located in the [.QUERIES] subdirectory.

Queries could be executed on either the 1994 or 1995 portion of the database. They could be executed as an individual query (1-5) or as a single-user load (parallel execution of queries 1-5), using the following:

```
$ SQL @QUERY-9x-n.SQL
```

(where x = year and n = query number)

```
$ @QUERIES.COM 9x batch_queue
```

```
$ @QUERY_N_USER n 9x batch_queue
```

(where batch_queue is the queue of the system to run on, for example, DEPOT8_BATCH)

Test Scenarios and Scripts

This section describes the test scenarios that were used to generate test data. The test scenarios consisted of VLM-enabled tests, non-VLM enabled tests, RMU Backup and Restore tests, and Fault Insertion/Error Recovery tests.

The VLM-enabled test scenarios and the Non-VLM Enabled test scenarios used five query scripts during testing. The five query scripts are described in the Test Process section of this chapter.

VLM-Enabled Test Scenarios

Test Scenarios 1 through 3 were the VLM-enabled test scenarios. These test scenarios used the large memory system sharing row cache. This scheme is limited to a system's 64-bit address space or amount of physical memory available on the system.

Test Scenarios 1 through 3 were run on the HiTest System in both the minimum and maximum configurations.

Test Scenario 1: Single-User Tests (Cold Cache)

This test scenario involved the simultaneous, single-user execution of query scripts 1 through 5 in parallel with cold cache.

The purpose of this test was to determine the system performance of a basic load.

This test was also run on the VAX node DEPOTX.

When this test was run on the VAX node DEPOTX, it was run without VLM capabilities. This is because VLM capabilities do not exist on VAX systems.

Test Scenario 2: Single-User Tests (Warm Cache)

This test scenario involved the simultaneous, single-user execution of query scripts 1 through 5 in parallel with warm cache.

The purpose of this test was to determine the system performance of a basic load.

Interoperability Tests and Results

Test Scenario 3: Multi-User Tests

This test scenario involved five and ten users simultaneously executing query scripts 1 through 5.

The purpose of this test was to determine the system performance under an increased load.

Note

Tests in excess of a 10-user load fail due to resource problems on a single VLM-enabled AlphaServer.

Non-VLM Enabled Test Scenarios

Test Scenarios 4 and 5 were the Non-VLM Enabled Test Scenarios. These test scenarios involved turning off VLM to allow the same database to be opened and simultaneously queried from multiple nodes.

Test Scenarios 4 and 5 were run on the HiTest System in both the minimum and maximum configurations.

Test Scenario 4: Distributed 10-User Load Tests

This test scenario involved running two 5-user test loads from two separate AlphaServers.

The purpose of this test was to determine the difference in system performance that results when a 10-user load is split between two AlphaServers accessing a non-VLM enabled database rather than run on a single AlphaServer with a VLM-enabled database.

Test Scenario 5: Distributed 20-User Load Tests

This test scenario involved running two 10-user test loads from two separate AlphaServers.

The purpose of this test was to determine whether, by adding a second AlphaServer system, the total number of users accessing the database can be doubled.

RMU Backup and Restore and Fault Insertion/Error Recovery Test Scenarios

This section describes the RMU Backup and Restore Test Scenario and the Fault Insertion/Error Recovery Test Scenario.

Note

The RMU Backup and Restore Test and the Fault Insertion/Error Recovery Test were run in RMU-enabled mode on the AlphaServer 4100.

Test Scenario 6: RMU Backup and Restore Tests

This test involves performing an RMU backup and restore of the one-year database.

The purpose of this test was to do the following:

- Determine the approximate backup and restore throughput times for a multifile/multidisk database.
- Verify the procedure for backing up and restoring a multifile/multidisk database.
- Identify potential problems in backing up and restoring a multifile/multidisk database.

Test Scenario 7: Fault Insertion/Error Recovery Tests

This test simulates potential hardware problems during an active query. The simulations of hardware problems include the following:

- Loss of a disk in a RAID 5 array
- Loss of an HSJ50
- Loss of system power

The purpose of this test was to determine the impact of each simulated problem on the database.

Purpose of Query Scripts

Functional verification of the ability to perform query operations was demonstrated using five SQL join scripts. These queries exercised functionality of the RDB7 server and SQL Utility.

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

The typical question asked was:

“What percentage of product share in dollars and units did a particular product have in a particular area as compared to competitive products in the same area?”

The queries were designed to search the database in varying ways to exercise the database.

Test Tools

The tools used for interoperability testing include the test scenarios described in the Test Process section of this chapter. The purpose of the test scenarios was to test a demonstration database under various conditions. Query scripts were written to search the database for various types of historical data.

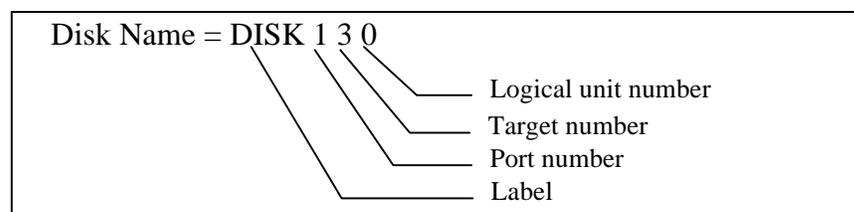
Each query script was submitted to a batch queue and each batch entry then generated a log file. The log file contained process statistics such as elapsed time which is the run time shown in the Test Results section. The run time also included any batch process overhead time.

Test Configuration

The test load was generated by using the five query scripts. The same disk storage configuration of 108 4.3 GB RZ29B-VW disks was used for both the minimum and maximum configurations.

Data was distributed across three dual-redundant HSJ controllers located in a SW800 cabinet. The cabinet housed three HSJ52 controllers and 18 BA356 shelves such that each controller supported six shelves. The 108 RZ29B-VW disks were evenly distributed across the six ports of each HSJ50 Controller.

Each disk was given a disk name consisting of a label, port number, target number, and logical unit number (LUN). Note that, in the following example, the disk name is “DISK.”



Interoperability Tests and Results

Disks were grouped into RAIDsets that spanned several ports of each controller. Each set was assigned a container name such R2 (RAIDset) that corresponded to an OpenVMS device name such as \$1\$DUA12 and an HSJ LUN such as D12.

For more information about data storage, see Appendix A.

OpenVMS Device Name	LUN	Container Name	Disks
			<input type="checkbox"/> DISK 110
			<input type="checkbox"/> DISK 210
\$1\$DUA12	⇒ D12	⇒ R2	⇒ <input type="checkbox"/> DISK 310
			<input type="checkbox"/> DISK 410

Test Process

The test process consisted of five queries. Each query was supported by two scripts, one for a 1994 portion of the database (for example, QUERY-94-1.SQL) and one for a 1995 portion of the database (for example, QUERY-95-1.SQL). For this round of testing, only the 1995 database existed.

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

Script for Query 1 (QUERY-9x-1.SQL)

```
set default date format 'SQL92';
set quoting rules 'SQL92';
set date format DATE 001;
set dialect 'oracle level11';
set line length 255;
set transaction read only;

/* 1. Star -- Product Share of Brand */

SELECT 'All Wheat Flakes' Product, AL2.MONTH,
       sum(AL5.UNIT_SALES) UNITS, sum(AL5.DOLLAR_SALES) DOLLARS, count(*),
       DISTRICT, CHANNEL_GROUP CHNL
```

```

FROM PRODUCT AL4,
     SALES_FACT AL5,
     CHANNEL AL1,
     DAILY_PERIOD AL2,
     MARKET AL3
WHERE (AL5.PRODUCT_ID=AL4.PRODUCT_ID
      AND AL5.MARKET_ID=AL3.MARKET_ID
      AND AL5.CHANNEL_ID=AL1.CHANNEL_ID
      AND AL5.DAY=AL2.DAY)
      AND (district='Connecticut'
          AND CHANNEL_GROUP in('Supermarket')
          AND BRAND in ('Quellogs Wheat Flakes')
          AND AL2.DAY BETWEEN '01-JAN-199x' AND '31-DEC-199x')
                                     (where x = 4 or 5)
GROUP
  BY DISTRICT, CHANNEL_GROUP, AL2.MONTH

UNION

SELECT '20 Oz Wheat Flakes' Product, AL2.MONTH,
       sum(AL5.UNIT_SALES) UNITS, sum(AL5.DOLLAR_SALES) DOLLARS, count(*),
       DISTRICT, CHANNEL_GROUP CHNL
FROM PRODUCT AL4,
     SALES_FACT AL5,
     CHANNEL AL1,
     DAILY_PERIOD AL2,
     MARKET AL3
WHERE (AL5.PRODUCT_ID=AL4.PRODUCT_ID
      AND AL5.MARKET_ID=AL3.MARKET_ID
      AND AL5.CHANNEL_ID=AL1.CHANNEL_ID
      AND AL5.DAY=AL2.DAY)
      AND (district = 'Connecticut'
          AND CHANNEL_GROUP in('Supermarket')
          AND PRODUCT='QLGS WHT FLK 20 OZ'
          AND AL2.DAY BETWEEN '01-JAN-199x' AND '31-DEC-199x')
                                     (where x = 4 or 5)
GROUP
  BY DISTRICT, CHANNEL_GROUP, AL2.MONTH
;

```

Interoperability Tests and Results

Query 2

Query 2 compared 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 was 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?*”

Script for Query 2 (QUERY-9x-2.SQL)

```
set default date format 'SQL92';
set quoting rules 'SQL92';
set date format DATE 001;
set dialect 'oracle levell1';
set line length 255;
set transaction read only;

/* 2. Star -- Channel share of all channels */

SELECT  'All Channels' CHNL,  AL2.MONTH,
        sum(AL5.UNIT_SALES) Units, sum(AL5.DOLLAR_SALES) Dollars, count(*),
        'NY + PA' DISTRICT, PRODUCT
FROM    PRODUCT AL4,
        SALES_FACT AL5,
CHANNEL AL1,
        DAILY_PERIOD AL2,
        MARKET AL3
WHERE   (AL5.PRODUCT_ID=AL4.PRODUCT_ID
        AND AL5.MARKET_ID=AL3.MARKET_ID
        AND AL5.CHANNEL_ID=AL1.CHANNEL_ID
        AND AL5.DAY=AL2.DAY
        AND DISTRICT in ('New York', 'Pennsylvania')
        AND CHANNEL_GROUP in ('Supermarket', 'Convenience',
                              'Warehouse', 'Drug', 'Discount')
        AND PRODUCT= 'QLGS WHT FLK 15 OZ'
        AND AL2.DAY BETWEEN '01-JAN-199x' AND '31-DEC-199x')
                                (where x = 4 or 5)

GROUP
        BY PRODUCT, AL2.MONTH

UNION
```

```

SELECT CHANNEL CHNL, AL2.MONTH,
       sum(AL5.UNIT_SALES) Units, sum(AL5.DOLLAR_SALES) Dollars, count(*),
       'NY + PA' DISTRICT , PRODUCT
FROM   PRODUCT AL4,
       SALES_FACT AL5,
       CHANNEL AL1,
       DAILY_PERIOD AL2,
       MARKET AL3
WHERE  (AL5.PRODUCT_ID=AL4.PRODUCT_ID
       AND AL5.MARKET_ID=AL3.MARKET_ID
       AND AL5.CHANNEL_ID=AL1.CHANNEL_ID
       AND AL5.DAY=AL2.DAY
       AND DISTRICT in ('New York' , 'Pennsylvania')
       AND CHANNEL='Safeway'
       AND PRODUCT='QLGS WHT FLK 15 OZ'
       AND AL2.DAY BETWEEN '01-JAN-199x' AND '31-DEC-199x')
                                           (where x = 4 or 5)
GROUP
      BY CHANNEL, PRODUCT, AL2.MONTH
;

```

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?*”

Script for Query 3 (QUERY-9x-3.SQL)

```

set default date format 'SQL92';
set quoting rules 'SQL92';
set date format DATE 001;
set dialect 'oracle level11';
set transaction read only;

/* 3. Star Market share of Region */

SELECT 'Northeast Total' MARKET, AL2.MONTH,
       sum(AL5.UNIT_SALES) Units, sum(AL5.DOLLAR_SALES) Dollars, count(*),
       CHANNEL_GROUP, PRODUCT

```

Interoperability Tests and Results

```
FROM PRODUCT AL4,
    SALES_FACT AL5,
    CHANNEL AL1,
    DAILY_PERIOD AL2,
    MARKET AL3
WHERE (AL5.PRODUCT_ID=AL4.PRODUCT_ID
    AND AL5.MARKET_ID=AL3.MARKET_ID
    AND AL5.CHANNEL_ID=AL1.CHANNEL_ID
    AND AL5.DAY=AL2.DAY
    AND REGION='Northeast'
    AND CHANNEL_GROUP in ('Convenience')
    AND PRODUCT= 'QLGS WHT FLK 10 OZ'
    AND AL2.DAY BETWEEN '01-JAN-199x' AND '31-DEC-199x')
                                                    (where x = 4 or 5)
GROUP
    BY CHANNEL_GROUP, PRODUCT , AL2.MONTH

UNION

SELECT MARKET, AL2.MONTH,
    sum(AL5.UNIT_SALES) Units, sum(AL5.DOLLAR_SALES) Dollars, count(*),
    CHANNEL_GROUP, PRODUCT
FROM PRODUCT AL4, SALES_FACT AL5,
    CHANNEL AL1, DAILY_PERIOD AL2, MARKET AL3
WHERE (AL5.PRODUCT_ID=AL4.PRODUCT_ID
    AND AL5.MARKET_ID=AL3.MARKET_ID
    AND AL5.CHANNEL_ID=AL1.CHANNEL_ID
    AND AL5.DAY=AL2.DAY
    AND MARKET='Bridgeport'
    AND CHANNEL_GROUP in ('Convenience')
    AND PRODUCT='QLGS WHT FLK 10 OZ'
    AND AL2.DAY BETWEEN '01-JAN-199x' AND '31-DEC-199x')
                                                    (where x = 4 or 5)
GROUP
    BY MARKET, CHANNEL_GROUP, PRODUCT, AL2.MONTH
;
```

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?*”

Script for Query 4 (QUERY-9x-4.SQL)

```

set default date format 'SQL92';
set quoting rules 'SQL92';
set date format DATE 001;
set dialect 'oracle levell1';
set transaction read only;

/* 4. Star -- Product share of SubCategory -all competitive prods */

SELECT 'All Wheat Products' Product, AL2.MONTH,
       sum(AL5.UNIT_SALES) Units, sum(AL5.DOLLAR_SALES) Dollars, count(*),
       DISTRICT, CHANNEL_GROUP CHNL
FROM   PRODUCT AL4,
       SALES_FACT AL5,
       CHANNEL AL1,
       DAILY_PERIOD AL2,
       MARKET AL3
WHERE  (AL5.PRODUCT_ID=AL4.PRODUCT_ID
       AND AL5.MARKET_ID=AL3.MARKET_ID
       AND AL5.CHANNEL_ID=AL1.CHANNEL_ID
       AND AL5.DAY=AL2.DAY
       AND district='Connecticut'
       AND CHANNEL_GROUP in('Supermarket')
       AND BRAND IN ('Quellogs Wheat Flakes', 'Boast Weeties', 'Boast Oatey
Rounds',
                    'Quellogs Wheaten Rye')
       AND AL2.DAY BETWEEN '01-JAN-199x' AND '31-DEC-199x')
                                           (where x = 4 or 5)
GROUP
      BY DISTRICT, CHANNEL_GROUP, AL2.MONTH

UNION

```

Interoperability Tests and Results

```
SELECT  '20 Oz Wheat Flakes' Product, AL2.MONTH,
        sum(AL5.UNIT_SALES) Units, sum(AL5.DOLLAR_SALES) Dollars, count(*),
        DISTRICT, CHANNEL_GROUP CHNL
FROM    PRODUCT AL4,
        SALES_FACT AL5,
        CHANNEL AL1,
        DAILY_PERIOD AL2,
        MARKET AL3
WHERE   (AL5.PRODUCT_ID=AL4.PRODUCT_ID
        AND AL5.MARKET_ID=AL3.MARKET_ID
        AND AL5.CHANNEL_ID=AL1.CHANNEL_ID
        AND AL5.DAY=AL2.DAY
        AND district = 'Connecticut'
        AND CHANNEL_GROUP in('Supermarket')
        AND PRODUCT='QLGS WHT FLK 20 OZ'
        AND AL2.DAY BETWEEN '01-JAN-199x' AND '31-DEC-199x')
                                (where x = 4 or 5)
GROUP
        BY DISTRICT, CHANNEL_GROUP, AL2.MONTH
;
```

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?*”

Script for Query 5 (QUERY-9x-5.SQL)

```
set default date format 'SQL92';
set quoting rules 'SQL92';
set date format DATE 001;
set dialect 'oracle level1';
set transaction read only;

/* 5. Star -- Product share of brand in 10 test markets aggregated */
/* include the star hint - 1/29/97 JMM */

SELECT  'All Wheat Flakes' Product, AL2.MONTH,
        sum(AL5.UNIT_SALES) Units, sum(AL5.DOLLAR_SALES) Dollars, count(*),
        CHANNEL_GROUP CHNL, '10-States'
```

```

FROM PRODUCT AL4,
    SALES_FACT AL5,
    CHANNEL AL1,
    DAILY_PERIOD AL2,
    MARKET AL3
WHERE (AL5.PRODUCT_ID=AL4.PRODUCT_ID
    AND AL5.MARKET_ID=AL3.MARKET_ID
    AND AL5.CHANNEL_ID=AL1.CHANNEL_ID
    AND AL5.DAY=AL2.DAY
    AND district in ('Connecticut', 'Delaware','Maine','Pennsylvania','New
York',
    'Oregon', 'Alaska', 'CA North', 'CA South','Washington')
    AND CHANNEL_GROUP in ('Supermarket')
    AND BRAND IN ('Quellogs Wheat Flakes')
    AND AL2.DAY BETWEEN '01-JAN-199x' AND '31-DEC-199x')
                                                    (where x = 4 or 5)
GROUP
    BY AL2.MONTH, CHANNEL_GROUP

UNION

SELECT '20 Oz Wheat Flakes' Product, AL2.MONTH,
    sum(AL5.UNIT_SALES) Units, sum(AL5.DOLLAR_SALES) Dollars, count(*),
    CHANNEL_GROUP CHNL, '10-States'
FROM PRODUCT AL4, SALES_FACT AL5,
    CHANNEL AL1, DAILY_PERIOD AL2, MARKET AL3
WHERE (AL5.PRODUCT_ID=AL4.PRODUCT_ID
    AND AL5.MARKET_ID=AL3.MARKET_ID
    AND AL5.CHANNEL_ID=AL1.CHANNEL_ID
    AND AL5.DAY=AL2.DAY
    AND district in ('Connecticut', 'Delaware','Maine','Pennsylvania','New
York',
    'Oregon', 'Alaska', 'CA North', 'CA South','Washington')
    AND CHANNEL_GROUP in ('Supermarket')
    AND PRODUCT='QLGS WHT FLK 20 OZ'
    AND AL2.DAY BETWEEN '01-JAN-199x' AND '31-DEC-199x')
                                                    (where x = 4 or 5)
GROUP
    BY AL2.MONTH, CHANNEL_GROUP
;

```

Test Results

Initial tests on the maximum VLM platform verified the ability to process SQL queries, in which all query data resides on the local system in a single database. The set of five query scripts, previously described, was executed on each system and run times were derived from query log files.

The following table shows test conditions that were used for Tests 1, 2, and 3.

Table 4-4: Test Conditions

Condition	AlphaServer 4100		VAX 6620
	Minimum	Maximum	–
Configuration:	Minimum	Maximum	–
Database opened:	1 year	1 year	1 year
Global buffers:	20,000	20,000	2,500
Buffers per user:	200	200	100
User slots:	100	100	7
CPUs:	2	4	2
Memory:	2 GB	4 GB	512 MB

Test 1: Single-User Test (Cold Cache)

This test performed the five SQL queries in parallel using cold cache conditions on the AlphaServer 4100 in the minimum configuration, the maximum configuration, and on the VAX 6620. The following table summarizes the run-time results on the test systems.

Table 4-5: Single-User, Cold Cache Test Results

Query	AlphaServer 4100 (Min. Config.) Run Times (hr:min:sec)	AlphaServer 4100 (Max. Config.) Run Times (hr:min:sec)	VAX 6620 Run Times (hr:min:sec)
Query 1	00:19:52.17	00:20:02.30	01:47:31.40
Query 2	00:33:38.49	00:34:17.85	01:58:48.28
Query 3	00:04:40.84	00:04:38.64	00:40:12.18
Query 4	02:07:40.67	02:07:48.43	04:17:47.06
Query 5	00:38:57.99	00:38:56.60	04:26:27.35

The VAX run times show that the VAX system did not perform as well as the AlphaServer system. Among other things, this was due to the AlphaServer VLM capabilities.

Test 2: Single-User Test (Warm Cache)

This test performed the five SQL queries in parallel using warm cache conditions on the AlphaServer 4100 in the minimum configuration and the maximum configuration. The following tables summarize the run-time results on the test systems.

Note

For the VAX 6620 platform, there is no substantial difference between the single-user, cold-cache test results and the single-user, warm-cache test results. For this reason, there are no test results listed for the single-user warm cache test on the VAX 6620.

Table 4-6: Single-User, Warm Cache Test Results

Query	AlphaServer 4100 (Min. Config.) Run Times (hr:min:sec)	AlphaServer 4100 (Max. Config.) Run Times (hr:min:sec)	VAX 6620 Run Times (hr:min:sec)
Query 1	00:01:56.72	00:01:48.60	–
Query 2	00:02:14.06	00:02:03.29	–
Query 3	00:02:51.50	00:02:25.71	–
Query 4	00:05:09.02	00:04:36.14	–
Query 5	00:22:02.22	00:21:20.38	–

Test 3: Multi-User Load Tests

This test consisted of 5- and 10-user load tests. The 5-user test was comprised of a total of 25 queries issued simultaneously with 5 users each issuing 5 queries. The 10-user test was comprised of a total of 50 queries issued simultaneously with 10 users each issuing 5 queries.

Table 4-7 and Table 4-8 show the range of run times for each query in the 5- and 10-user load tests. Table 4-7 shows the range of multiuser load test run times for the minimum configuration of this HiTest Suite. Table 4-8 shows the range of multiuser load test run times for the maximum configuration of this HiTest Suite.

In the following tables, the columns labeled “Shortest Run Time” and “Longest Run Time” list the shortest and longest run times for each individual query in a multi-load user test. For example, Table 4-7 lists 4 minutes, 52.46 seconds as the shortest run time for Query 1 in the 5-user load test run under the minimum configuration; Table 4-7 lists 7 minutes, 3.46 seconds as the *longest* run time for Query 1 in the 5-user load test run under the same configuration.

Table 4-7: 5- and 10-User Load Test Results for the Minimum Configuration

Query	Run Time for 5 Users (hr:min:sec)		Run Time for 10 Users (hr:min:sec)	
	Shortest Run Time	Longest Run Time	Shortest Run Time	Longest Run Time
Query 1	00:04:52.46	00:07:03.46	00:19:35.99	00:20:24.33
Query 2	00:07:19.63	00:08:53.85	00:19:46.03	00:40:33.94
Query 3	00:10:04.60	00:10:32.29	00:20:18.88	00:34:20.24
Query 4	00:19:37.34	00:20:01.86	00:25:56.30	01:27:13.48
Query 5	01:05:33.66	01:07:09.05	01:53:17.60	02:25:04.87

Table 4-8: 5- and 10-User Load Test Results for the Maximum Configuration

Query	Run Time for 5 Users (hr:min:sec)		Run Time for 10 Users (hr:min:sec)	
	Shortest Run Time	Longest Run Time	Shortest Run Time	Longest Run Time
Query 1	00:04:58.56	00:06:49.30	00:06:54.03	00:13:08.04
Query 2	00:06:53.21	00:07:25.81	00:12:19.40	00:14:16.65
Query 3	00:08:08.20	00:08:17.29	00:15:04.10	00:16:15.43
Query 4	00:14:23.77	00:14:52.45	00:25:55.28	00:28:34.18
Query 5	00:48:31.19	00:48:53.83	01:22:05.75	01:33:38.32

Non-VLM Enabled Database Query Tests

The following section provides the results of the non-VLM enabled database query tests.

Test 4: Distributed 10-User Load Test

This test involved splitting the 10-user load test across two AlphaServers; a 5-user load was run on each AlphaServer. The purpose of this test was to determine the difference in performance, if any, that results from running a 10-user load test across two AlphaServers in non-VLM enabled mode rather than on one VLM-enabled AlphaServer.

Table 4-9 shows the range of run times for each query in the non-VLM enabled 10-user load test. In Table 4-9, the columns labeled “Shortest Run Time” and “Longest Run Time” list the shortest and longest run times for each individual query in the non-VLM enabled 10-user load test. For example, Table 4-9 lists 1 hour, 22 minutes, 6.03 seconds as the shortest run time for Query 1 in the non-VLM enabled 10-user load test run under the minimum configuration; Table 4-9 lists 2 hours, 21 minutes, 4.29 seconds as the *longest* run time for Query 1 in the non-VLM enabled 10-user load test run under the minimum configuration.

Table 4-9: Non-VLM Enabled 10-User Load Test Results

Query	Run Time on AlphaServer 4100 (Minimum Configuration) (hr:min:sec)		Run Time on AlphaServer 4100 (Maximum Configuration) (hr:min:sec)	
	Shortest Run Time	Longest Run Time	Shortest Run Time	Longest Run Time
Query 1	01:22:06.03	02:21:04.29	01:20:18.63	01:58:45.60
Query 2	01:53:41.67	02:56:23.24	01:51:48.19	02:27:51.59
Query 3	00:22:42.00	01:29:03.81	00:15:56.94	01:07:34.13
Query 4	05:16:19.77	05:18:18.08	04:37:05.42	04:38:31.62
Query 5	04:10:51.19	04:42:22.43	02:59:24.59	03:55:28.01

Test 5: Distributed 20-User Load Test

This test consisted of running two 10-user load tests across two AlphaServers. The purpose of this test was to verify that the number of users can be doubled by adding a second AlphaServer system.

The test completed successfully.

RMU Backup and Restore Tests

The RMU Backup and Restore tests consisted of an RMU backup and restore of the one-year database. These tests determined the approximate throughput times for RMU backup and restore operations on a multifile, multidisk database.

The one-year database is described in Table 4-1 and Table 4-2.

Table 4-10 lists the results of the RMU Backup and Restore tests for this HiTest Suite.

Table 4-10: RMU Backup and Restore Test Results for the Maximum Configuration

	Backup Statistics (hr:min:sec)	Restore Statistics (hr:min:sec)
Buffered I/O count	679	1598
Direct I/O count	10678179	10682380
Page faults	1818	4011
Charged CPU time	0 02:23:07.25	0 01:49:09.93
Peak working set size	17408	34032
Peak virtual size	203936	228688
Mounted volumes	3	3
Elapsed time	0 13:58:28.90	0 16:43:31.33

Test 6: Error Recovery

This test simulated various potential hardware problems during an active query. The following fault scenarios were tested with the AlphaServer 4100 system in the maximum configuration:

- Loss of a disk in a RAID 5 array
- Loss of an HSJ50
- Loss of system power

The following sections discuss each simulated failure.

Disk Failure

Test Description: DISK150 on HSJ002, which was part of the RAIDset known as sort4, was removed after a set of queries were started.

Note

A disk was removed from the sort4 raidset because the sort4 raidset had a relatively high amount of I/O in a single-user warm cache test.

Expected results: Queries would run to successful completion with little to no performance impact.

Actual results: Results were as expected. Queries ran to successful completion with no performance impact.

Interoperability Tests and Results

Table 4-11 shows the actual elapsed query times for the disk failure tests and the single-user, warm cache tests without disk failure on the AlphaServer 4100 in the maximum configuration.

Table 4-11: Actual Elapsed Query Times

Query	Disk Failure Tests Run Times (hr:min:sec)	Tests without Disk Failure Run Times (hr:min:sec)
Query 1	00:02:03.60	00:01:48.60
Query 2	00:02:20.03	00:02:03.29
Query 3	00:02:41.13	00:02:25.71
Query 4	00:04:55.88	00:04:36.14
Query 5	00:22:12.31	00:21:20.38

HSJ Failure

Test description: The flashram card on HSJ002B was removed after a set of queries were started.

Expected results: Queries would run to successful completion with little to no performance impact.

Actual results: As expected, queries ran to successful completion with no overall significant performance impact.

Table 4-12 shows the actual elapsed query times for the HSJ failure tests and the single-user, warm cache tests without HSJ failure on the AlphaServer 4100 in the maximum configuration.

Table 4-12: Actual Elapsed Query Times

Query	HSJ Failure Tests Run Times (hr:min:sec)	Tests without HSJ Failure Run Times (hr:min:sec)
Query 1	00:01:48.78	00:01:48.60
Query 2	00:02:04.49	00:02:03.29
Query 3	00:02:29.33	00:02:25.71
Query 4	00:04:38.34	00:04:36.14
Query 5	00:21:18.35	00:21:20.38

Note

The HSJ failure test was run under warm cache conditions.

System Power Failure

Test description: The system was power failed (via key) after a set of queries had been started.

Expected results: Queries would not be completed, however, no database corruption should be experienced as a result. Reissuing the queries after the system booted back up and the database was remounted should result in successful completion of all queries.

Actual results: The system was rebooted and the queries ran successfully but more slowly than they had run before the power failure occurred.

Note

Refer to Chapter 6 for information on preventing the degradation of query performance after a power failure.

Once the queries ran, the database was shut down and reopened and the queries ran as expected.

Table 4-13 shows the actual elapsed query times for the system power failure tests and the single-user, cold cache tests without system power failure on the AlphaServer 4100 in the maximum configuration.

Table 4-13: Actual Elapsed Query Times

Query	System Power Failure Tests Run Times (hr:min:sec)	Tests without System Power Failure Run Times (hr:min:sec)
Query 1	00:19:48.45	00:20:02.30
Query 2	00:34:22.02	00:34:17.85
Query 3	00:04:42.47	00:04:38.64
Query 4	02:08:04.57	02:07:48.43
Query 5	00:39:01.05	00:38:56.60

System Limits and Characterization Data

This chapter describes any system limits that may have been determined as a result of the testing, along with information about the system characterization during testing. Areas covered include:

- Comparison of run times
- Database buffering schemes
- Workload characterizations
- RDB account quotas
- System RDB account quotas
- Accounting information from a one-user load test

Comparison of Run Times

The tests demonstrated the interoperability of the AlphaServer systems in a VMScluster. The following comparisons were made of the AlphaServer and VAX systems:

- In the tests of SQL queries in parallel, the cold cache test times for the AlphaServer 4100 system ranged from 4 minutes, 38.64 seconds to 2 hours, 7 minutes, 48.43 seconds. For the same tests, the VAX systems ranged from 40 minutes, 12.18 seconds to 4 hours, 26 minutes, 27.35 seconds.
- In the VLM-enabled 5- and 10-user load tests, the test times for the AlphaServer 4100 system ranged from 4 minutes, 52.46 seconds to 2 hours, 25 minutes, 4.87 seconds. The VLM-enabled 5- and 10-user load tests were not run on the VAX systems because comparably longer run times were expected.
- In the non-VLM enabled 10-user load tests, the test times for the AlphaServer 4100 system ranged from about 15 minutes, 56.94 seconds to 5 hours, 18 minutes, 18.08 seconds. The non-VLM enabled 10-user load tests were not run on the VAX systems because comparably longer run times were expected.
- In the HSJ Failure Test, the elapsed query times for the AlphaServer 4100 system ranged from 1 minute, 48.78 seconds to 21 minutes, 18.35 seconds.
- In the Disk Failure Test, the elapsed query times for the AlphaServer 4100 system ranged from 2 minutes, 3.60 seconds to 22 minutes, 12.31 seconds.
- In the System Power Failure Test, the elapsed query times for the AlphaServer 4100 system ranged from 4 minutes, 42.47 seconds to 2 hours, 8 minutes, 4.57 seconds.

System Limits and Characterization Data

- When comparing VAX and Alpha architectures, a query on the VAX node demonstrated that the comparison was so overwhelmingly in favor of the AlphaServer nodes that any further testing of the VAX nodes was deemed unnecessary. (See Table 4-5 for actual data.)
- The test configuration was not fully optimized for performance. An extensive tuning effort would have obtained better run-time results. The 64-bit and VLM capabilities of the AlphaServer systems were clearly evident with the run times that were observed.

Database Buffering Schemes

In conducting the tests on the VAX and AlphaServer systems, certain architectural differences were uncovered, the primary one being the difference in memory addressing between a 32-bit system (VAX system) and a 64-bit system (AlphaServer system).

In designing the database buffering schemes for both platforms, these differences had to be considered. The Oracle Rdb7 database provided several different buffer management styles to support 32- and 64-bit systems.

These buffer styles included the following:

- Conventional process sharing local buffers (available on VAX and AlphaServer systems). This scheme is limited to a process' 32-bit address space and was not used in the testing of this HiTest Suite.
- Conventional process sharing global buffers (available on VAX and AlphaServer systems). This scheme is limited to a process' 32-bit address space (1 GB) and was used in the non-VLM enabled testing of this HiTest Suite.
- System sharing global buffers (available on AlphaServer systems only). This scheme is limited to a system's 32-bit address space (1.8 GB) and was not used in the testing of this HiTest Suite.
- System sharing row cache (available on AlphaServer systems only). This scheme is limited to a system's 32-bit address space (1.8 GB) and was not used in the testing of this HiTest Suite.
- Large memory system sharing row cache (for example, Very Large Memory or VLM, available on AlphaServer systems only). This scheme is limited to a system's 64-bit address space or amount of physical memory available on the system and was used in the VLM-enabled testing of this HiTest Suite.

The first three buffering styles are additive and must all reside within the limits of conventional 32-bit addressing. The latter three styles are only available on 64-bit AlphaServer systems and account for a vast performance difference between the platforms.

Workload Characterizations

The following excerpt of accounting information from the query log files shows the consumption of resources on a single-user cold cache (worst case) test. In the following example, the platform was the AlphaServer 4100 in the maximum configuration.

The following example includes batch processing resource usage.

The following is the accounting information from a cold cache 1-user load test on an AlphaServer 4100 in the maximum configuration with VLM enabled:

Query 1:

Accounting information:

Buffered I/O count:	854	Peak working set size:	48496
Direct I/O count:	44754	Peak virtual size:	255184
Page faults:	4465	Mounted volumes:	0
Charged CPU time:	0 00:01:00.89	Elapsed time:	0 00:20:02.30

Query 2:

Accounting information:

Buffered I/O count:	921	Peak working set size:	55952
Direct I/O count:	81183	Peak virtual size:	262160
Page faults:	4878	Mounted volumes:	0
Charged CPU time:	0 00:01:40.31	Elapsed time:	0 00:34:17.85

Query 3:

Accounting information:

Buffered I/O count:	950	Peak working set size:	67872
Direct I/O count:	9333	Peak virtual size:	273584
Page faults:	5158	Mounted volumes:	0
Charged CPU time:	0 00:01:46.10	Elapsed time:	0 00:04:38.64

Query 4:

Accounting information:

Buffered I/O count:	830	Peak working set size:	113632
Direct I/O count:	351977	Peak virtual size:	319440
Page faults:	8575	Mounted volumes:	0
Charged CPU time:	0 00:05:16.52	Elapsed time:	0 02:07:48.43

Query 5:

Accounting information:

Buffered I/O count:	852	Peak working set size:	385344
Direct I/O count:	55563	Peak virtual size:	590576
Page faults:	25666	Mounted volumes:	0
Charged CPU time:	0 00:19:21.42	Elapsed time:	0 00:38:56.60

System Limits and Characterization Data

The following is the accounting information from a warm cache 1-user load test on an AlphaServer 4100 in the maximum configuration with VLM enabled:

Query 1:

Accounting information:

Buffered I/O count:	853	Peak working set size:	48608
Direct I/O count:	4039	Peak virtual size:	255120
Page faults:	4506	Mounted volumes:	0
Charged CPU time:	0 00:00:34.21	Elapsed time:	0 00:01:48.60

Query 2:

Accounting information:

Buffered I/O count:	936	Peak working set size:	56240
Direct I/O count:	4609	Peak virtual size:	262160
Page faults:	4873	Mounted volumes:	0
Charged CPU time:	0 00:00:42.26	Elapsed time:	0 00:02:03.29

Query 3:

Accounting information:

Buffered I/O count:	931	Peak working set size:	67056
Direct I/O count:	5684	Peak virtual size:	273584
Page faults:	5098	Mounted volumes:	0
Charged CPU time:	0 00:00:56.10	Elapsed time:	0 00:02:25.71

Query 4:

Accounting information:

Buffered I/O count:	830	Peak working set size:	113040
Direct I/O count:	9712	Peak virtual size:	319440
Page faults:	8530	Mounted volumes:	0
Charged CPU time:	0 00:02:14.92	Elapsed time:	0 00:04:36.14

Query 5:

Accounting information:

Buffered I/O count:	847	Peak working set size:	384608
Direct I/O count:	30618	Peak virtual size:	590576
Page faults:	25538	Mounted volumes:	0
Charged CPU time:	0 00:14:47.06	Elapsed time:	0 00:21:20.38

RDB Account Quotas

Table 5-1 contains a list of account quotas from the User Authorization file of this HiTest System. The account quotas in Table 5-1 were used to run this HiTest Suite from the RDB account.

Use the values in Table 5-1 as the initial account quotas when running this HiTest Suite from the RDB account.

Table 5-1: Sample of RDB Account Quotas

Account Quota	Value	Account Quota	Value	Account Quota	Value
Maxjobs:	0	Fillm:	8000	Bytlm:	512000
Maxacctjobs:	0	Shrfillm:	0	Pbytlm:	0
Maxdetach:	0	BIOlM:	8000	JTquota:	8192
Prclm:	20	DIOlM:	8000	WSdef:	4096
Prio:	4	ASTlm:	8000	WSquo:	8192
Queprio:	4	TQElm:	8000	WSextent:	16384
CPU:	(none)	Enqlm:	32767	Pgflquo:	1024000

SYSTEM Account Quotas

Table 5-2 contains a list of account quotas from the User Authorization file of this HiTest System. The account quotas in Table 5-2 were used to run this HiTest Suite from the SYSTEM account.

Use the values in Table 5-2 as the initial account quotas when running this HiTest Suite from the SYSTEM account.

Table 5-2: Sample of SYSTEM Account Quotas

Account Quota	Value	Account Quota	Value	Account Quota	Value
Maxjobs:	0	Fillm:	1000	Bytlm:	65536
Maxacctjobs:	0	Shrfillm:	0	Pbytlm:	0
Maxdetach:	0	BIOlM:	200	JTquota:	8192
Prclm:	10	DIOlM:	200	WSdef:	1024
Prio:	4	ASTlm:	300	WSquo:	2048
Queprio:	0	TQElm:	200	WSextent:	16384
CPU:	(none)	Enqlm:	2000	Pgflquo:	300000

Accounting Statistics from a One-User Load Test

The following tables contain accounting statistics from one-user load tests on AlphaServer in the maximum configuration with VLM enabled. Table 5-3 contains the accounting statistics from the cold cache one-user load test. Table 5-4 contains the accounting statistics from the warm cache one-user load test.

Table 5-3: Accounting Statistics from a Cold Cache One-User Load Test

	Query 1	Query 2	Query 3	Query 4	Query 5
Buffered I/O count	854	921	950	830	852
Direct I/O count	44754	81183	9333	351977	55563
Page faults	4465	4878	5158	8575	25666
Charged CPU time	0 00:01:00.89	0 00:01:40.31	0 00:01:46.10	0 00:05:16.52	0 00:19:21.42
Peak working set size	48496	55952	67872	113632	385344
Peak virtual size	255184	262160	273584	319440	590576
Mounted volumes	0	0	0	0	0
Elapsed time	0 00:20:02.30	0 00:34:17.85	0 00:04:38.64	0 02:07:48.43	0 00:38:56.60

Table 5-4: Accounting Statistics from a Warm Cache One-User Load Test

	Query 1	Query 2	Query 3	Query 4	Query 5
Buffered I/O count	853	936	931	830	847
Direct I/O count	4039	4609	5684	9712	30618
Page faults	4506	4873	5098	8530	25538
Charged CPU time	0 00:00:34.21	0 00:00:42.26	0 00:00:56.10	0 00:02:14.92	0 00:14:47.06
Peak working set size	48608	56240	67056	113040	384608
Peak virtual size	255120	262160	273584	319440	590576
Mounted volumes	0	0	0	0	0
Elapsed time	0 00:01:48.60	0 00:02:03.29	0 00:02:25.71	0 00:04:36.14	0 00:21:20.38

6

Problems and Resolutions

This chapter describes any problems that may have been encountered during the testing that have not been fixed and eliminated. A resolution for each problem is given. The resolution provides the system manager or user with a fix or workaround for the problem.

The following problems were identified during testing:

Hardware

There were no hardware problems discovered during testing.

Operating System

There were no operating system problems discovered during testing.

Application

Problem When the database was remounted after a power failure, query performance decreased significantly. The decrease in query performance was due to a problem in the database recovery process: the deletion of row cache backing store files before the database recovery process completes.

Resolution To prevent this problem from occurring, upgrade to Version 7.0.1.1 of Oracle Rdb. To obtain Version 7.0.1.1 of Oracle Rdb, contact Oracle by calling 1-800-ORACLE1 or sending E-mail to infodec@us.oracle.com.

Short of upgrading to Version 7.0.1.1 of Oracle Rdb, the workaround for this problem is to open, close, and reopen the database upon power failure. This clears the row cache backing store files.

A

Detailed Hardware Configuration

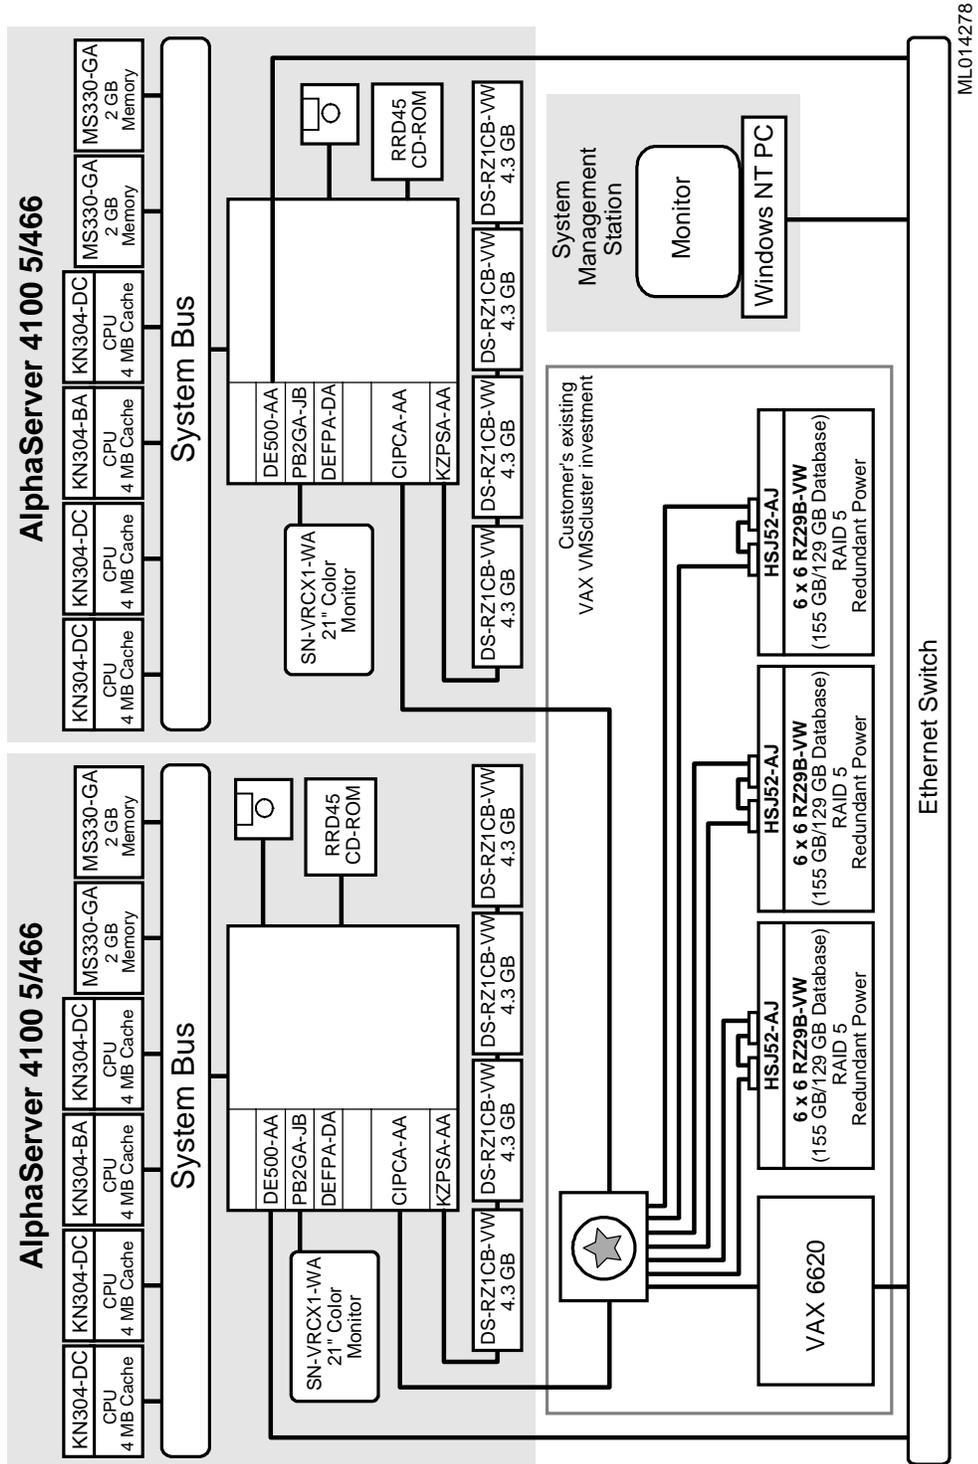
This appendix describes the minimum and maximum hardware configuration for the following:

- System Diagram
- AlphaServer 4100 configurations, including:
 - System motherboard
 - PCI backplanes (slot usage)
- Database Storage Maps

System Diagram

Figure A-1 shows a diagram of the entire HiTest Suite.

Figure A-1: System Diagram



AlphaServer 4100 Configurations

Figure A-2 and Table A-1 show the AlphaServer 4100 system motherboard and describe the minimum and maximum hardware configurations used in this HiTest Template.

Figure A-2: AlphaServer 4100 Motherboard

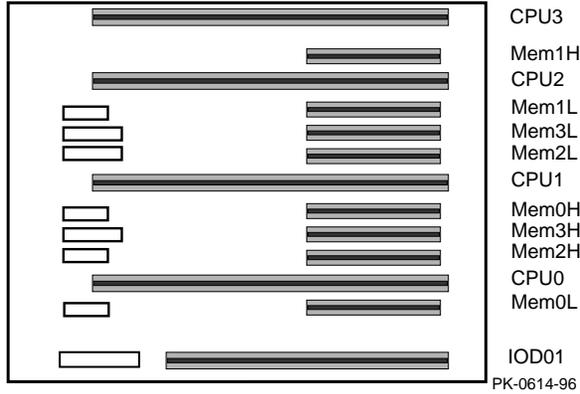


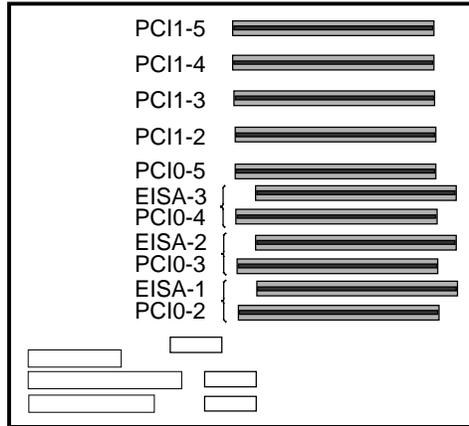
Table A-1: AlphaServer 4100 Usage (Minimum and Maximum Configurations)

Slot	Minimum Configuration Options	Maximum Configuration Options	Description
CPU3	open	KN304-DC	466 MHz CPU
Mem1H	MS330-FA	MS330-FA	1 GB memory
CPU2	open	KN304-DC	466 MHz CPU
Mem1L	MS330-FA	MS330-FA	1 GB memory
Mem3L	open	MS330-FA	1 GB memory
Mem2L	open	MS330-FA	1 GB memory
CPU1	KN304-DC	KN304-DC	466 MHz CPU
Mem0H	MS330-FA	MS330-FA	1 GB memory
Mem3H	open	MS330-FA	1 GB memory
Mem2H	open	MS330-FA	1 GB memory
CPU0	KN304-DC	KN304-DC	466 MHz CPU
Mem0L	MS330-FA	MS330-FA	1 GB memory
IOD01	open	open	

AlphaServer 4100 PCI Slot Usage

Figure A-3 and Table A-2 show the PCI slot usage for the minimum and maximum configurations of this HiTest Template.

Figure A-3: AlphaServer 4100 PCI Slot Usage



ML013980

Table A-2: PCI Slot Usage (Minimum and Maximum Configurations)

Slot	Minimum Configuration Options	Maximum Configuration Options	Description
PCI1-5	DE500-AA	DE500-AA	PCI 10/100-Mbit Fast Ethernet controller
PCI1-4	open	open	
PCI1-3	KZPSA-AA	KZPSA-AA	Fast Wide Single-ended SCSI controller
PCI1-2	CIPCA-AA	CIPCA-AA	Port - PCI to CI host bus adapter
PCI0-5	open	CIPCA-AA	Port
EISA-3/ PCI0-4	CIPCA-AA	CIPCA-AA	Link
EISA-2/ PCI0-3	open	CIPCA-AA	Link
EISA-1/ PCI0-2	open	open	

Database Storage Maps

The following tables map each OpenVMS device name to the HSJ physical devices in each RAIDset (RAID 5).

In Table A-3, the HSJ Container Names that use the “R*n*” naming scheme are RAIDset names. For example, “R1” is RAIDset 1, “R2” is RAIDset 2, and so on.

Table A-3: Database Storage Map for HSJ001

OpenVMS Device Name	Volume Label	Volume Description	HSJ LUN	HSJ Container Name	HSJ Physical Devices	Chunk Size (blocks)
\$1\$DUA11	VAXVMS071	VAX System Disk	D11	R1	DISK100 200 300	256
\$1\$DUA12	STORAGE	Storage for Kits and RBD Account	D12	R2	DISK110 210 310 410	256
\$1\$DUA13	disk1d	Data Disk	D13	R3	DISK120 220 320 420 520	256
\$1\$DUA14	disk2d	Data Disk	D14	R4	DISK130 230 330 430 530	256
\$1\$DUA15	disk1i	Index Disk	D15	R5	DISK400 500 600	256
\$1\$DUA16	disk2i	Index Disk	D16	R6	DISK150 510 610	256
\$1\$DUA17	sort1	Sort/Temp Disk	D17	R7	DISK140 240 340 630	256
\$1\$DUA18	sort2	Sort/Temp Disk	D18	R8	DISK250 440 540 640	256
–	–	–	–	SPARESET	DISK350 450 550 650	256

Detailed Hardware Configuration

In Table A-4, the HSJ Container Names that use the “Rn” naming scheme are RAIDset names. For example, “R1” is RAIDset 1, “R2” is RAIDset 2, and so on.

Table A-4: Database Storage Map for HSJ002

OpenVMS Device Name	Volume Label	Volume Description	HSJ LUN	HSJ Container Name	HSJ Physical Devices	Chunk Size (blocks)
\$1\$DUA21	AlphaSYS071	First Alpha System Disk	D21	R1	DISK100 200 300 400 500 600	256
\$1\$DUA22	disk3d	Data Disk	D22	R2	DISK120 220 320 420 520	256
\$1\$DUA23	disk4d	Data Disk	D23	R3	DISK130 230 330 430 530	256
\$1\$DUA24	disk3i	Index Disk	D24	R4	DISK110 210 310	256
\$1\$DUA25	disk4i	Index Disk	D25	R5	DISK410 510 610	256
\$1\$DUA26	sort3	Sort/Temp Disk	D26	R6	DISK140 240 340 640	256
\$1\$DUA27	sort4	Sort/Temp Disk	D27	R7	DISK150 250 440 540	256
–	–	–	–	SPARESET	DISK350 450 550 650	256

In Table A-5, the HSJ Container Names that use the “Rn” naming scheme are RAIDset names. For example, “R1” is RAIDset 1, “R2” is RAIDset 2, and so on.

Table A-5: Database Storage Map for HSJ003

OpenVMS Device Name	Volume Label	Volume Description	HSJ LUN	HSJ Container Name	HSJ Physical Devices	Chunk Size (blocks)
\$1\$DUA32	AlphaSYS071	Second Alpha System Disk	D32	R1	DISK120 220 320 420 620	256
\$1\$DUA31	disk3d	Data Disk	D31	R2	DISK150 250 340 440 540	256
\$1\$DUA33	disk4d	Data Disk	D33	R3	DISK130 230 330 430 530	256
\$1\$DUA34	disk3i	Index Disk	D34	R4	DISK110 210 630	256
\$1\$DUA35	disk4i	Index Disk	D35	R5	DISK410 510 610	256
\$1\$DUA36	sort5	Sort/Temp Disk	D36	R6	DISK140 240 520 640	256
\$1\$DUA37	sort6	Sort/Temp Disk	D37	R7	DISK310 400 500 600	256
\$1\$DUA38	root0	Oracle RDB Root Disk	D38	R8	DISK100 200 300	256
–	–	–	–	SPARESET	DISK350 450 550 650	256

B

Operational Command and Supporting Files

The following command procedures and SQL scripts were used to set up, load, or run tests on the CPG database. Details about file contents are available upon request.

BINARY.COM

A DCL command procedure to create binary input load files. BINARY.COM created 12 files, each 5.5 GB per month, that constituted a year's worth of sales data.

CACHES.SQL

An SQL script to create VLM row caches. This script is included in the CPGxx.SQL script (where "xx" is the year number).

CPG.COM

A DCL command procedure to create RMS search list logicals for all keys directories (for example, CPG all database storage directories, TMP – all temporary files, SRT – all sort devices, AIJ ñ all after-image journal files).

CPG.SQL

An SQL script that builds the blank database.

DISKS.DAT

A device list file detailing the devices comprising the database.

DISMOUNT.COM

A DCL command procedure to dismount files.

INDEXES-94.SQL

Creates the indexes for the 1994 database on the AlphaServer 4100 system.

INDEXES-95.SQL

Creates the indexes for the 1995 database on the AlphaServer 4100 system.

Operational Command and Supporting Files

INIT.COM

A command file that initializes the database after a system boot.

JOURNALS-94.COM

Creates the journals for the 1994 database on the AlphaServer 4100 system.

JOURNALS-95.COM

Creates the journals for the 1995 database on the AlphaServer 4100 system.

LOAD-DIM.COM

A command file that loads the small dimension tables.

LOADER.COM

A DCL command procedure to create ASCII input load files. LOADER.COM created 12 files, each 6.8 GB per month, that constituted a years worth of sales data.

LOGICALS.COM

A DCL command procedure to define device logical names.

LOGIN.COM

A login command procedure to establish the VMS\$MENU logical names.

MOUNT.COM

A DCL command procedure to mount disk devices.

RMU-COPY95.COM

A DCL command procedure to perform an RMU copy of the 1995 database.

RMUBIN.COM

A DCL command file that serially loads the sales facts table from binary load files.

RMUBIN17.COM

A DCL command file that loads the sales facts table in parallel using 17 threads.

RMUBLD.COM

A DCL command file that loads the sales facts table from ASCII load files.

SF_KEY-94.SQL

An SQL script used for contingencies to build the index on the sales facts table for 1994.

SF_KEY-95.SQL

An SQL script used for contingencies to build the index on the sales facts table for 1995.

