

HP Integrity Data Warehouse White Paper

Integrated Solutions Deployment Lab

Revision 1.0



Executive Summary	3
Terms and Definitions.....	3
Data Warehouse Solution Overview	4
Marketing Introduction: Why use hp Integrity for a Data Warehouse implementation?	4
Software Components.....	4
Hardware Components.....	4
Application Environment.....	4
Test Procedures	5
“Market Basket Analysis” Simulation Setup	5
Simulation Execution.....	5
Conclusions.....	6
Test Results	7
CPU Utilizations in the Different Scenarios.....	7
Response Times	10
Known Problems and Tips	11
General Tips	11
Known Problems	11
Appendix A – Hardware Environment	12
External Storage	12
Internal Devices and IO Cards.....	13
Appendix B – Software Environment	14
Appendix C – Patches	15

Figure 1: CPU Utilization for 16 CPUs, 16 Degrees and 16 CPUs, 32 Degrees.....	7
Figure 2: Disk Access Rates for 16 CPUs, 16 Degrees and 16 CPUs, 32 Degrees.....	8
Figure 3: Load Averages for 16 CPUs, 16 Degrees and 16 CPUs, 32 Degrees.....	9
Figure 4: Server Cabling to External Storage	12
Figure 5: rx8620 Internal Peripheral and HBA Layout.....	13
Table 1: Test Response Times (hh:mm:ss).....	10
Table 2: Software Products Installed.....	14

Executive Summary

The purpose of this white paper is to help sales representatives and solution architects understand how to implement a data warehousing solution on a specific HP Integrity platform. By leveraging the work that HP performed in this implementation, customers can reduce their deployment time, which improves the Total Customer Experience for their HP-UX 11i HP Integrity implementation. Customer Focused Testing reduces deployment time by validating and documenting all major hardware and software elements in a given configuration.

This document demonstrates to customers and HP field personnel that this configuration information has been successfully validated in the test environment. Please note that while these configurations have been effectively validated, they have not been optimized for specific customer workloads and may therefore need further refinement.

Terms and Definitions

- RDBMS: Relational Database Management System – a type of database management system (DBMS) that stores data in the form of related tables.
- Data Warehouse: A collection of data designed to support management decision-making. Development of a Data Warehouse includes development of systems to extract data from operating systems plus installation of a warehouse database system that provides managers flexible access to the data.

Data Warehouse Solution Overview

Marketing Introduction: Why use hp Integrity for a Data Warehouse implementation?

HP Integrity servers have demonstrated outstanding performance in many areas, including Data Warehousing. In the latest TPC-H results at www.tpc.org, a HP Superdome running HP-UX 11i and Oracle 10g with a 10-terabyte (TB) database had a performance result of 49,104 Composite Query per Hour (QphH) and a price/performance ratio of \$118/QphH (from www.tpc.org). This was one of the only two 10TB results and the one with the best price performance. This level of outstanding performance would be observed in all HP Integrity servers, from 2-way to 128-way.

Outstanding performance, excellent pricing, ease of implementation and remarkable stability are among the many reasons to use HP Integrity servers. This document exemplifies why HP Integrity servers are the best choice for our customers.

Software Components

The software stack consists of the Oracle database, the HP-UX 11i Enterprise Operating Environment, and a series of test scripts that simulate Data Warehouse applications and users.

The software components are:

- HP-UX 11i v2 Enterprise Operating Environment
- HP C/ANSI C Developer 's Bundle
- HP aC++ Compiler
- HP Itanium Fortran 90 Compiler
- Java2 1.4 SDK for HP-UX
- Java2 1.3 RTE for HP-UX
- HP GlancePlus/UX Pak
- HP StorageWorks Command View
- Oracle 9i Database

Hardware Components

Oracle 9.2.0.2 was installed on a single-partition, 16-CPU, rx8620 Integrity server. A set of seven HP StorageWorks va7410 disk arrays was used as storage.

Application Environment

The market basket simulation test was based upon the business requirements and data models used by a number of our large supermarket chain customers. The test package employs two fact tables and five dimension tables, and it exercises a variety of time-consuming queries from simulated users (5-125). All of the data was managed by an Oracle 9.2.0.2 64-bit RDBMS, which functions as the Data Warehouse server.

Test Procedures

“Market Basket Analysis” Simulation Set-up

It took approximately 28 hours to complete the test setup. The following tasks had to be performed manually:

- Create directories for the executable scripts, queries and suites.
- Create directories for the database and output.
- Test executables had to have variables initialized to point to the proper working, output and result directories.
- Data files had to be “gunzip’ed”.
- Create logical volumes for the test database via a setup script (lvm_db.sh). The volumes were created as:
 - System file (816MB)
 - Log files (4300MB x 2)
 - Data files (10896MB x 52)
 - Index files (4080MB x 52)
 - Temp space files (14304MB x 30)
 - Default space file (2MB)
- The logical volumes were linked to database file names via a second script (lvm_link.sh).
- A final script (create.sh) performed the actual database initialization:
 - Create the test database
 - Create public rollback segments
 - Create tablespaces for default space, temporary space, fact tables, dimension tables and all the indexes. (6 hours 42 minutes)
 - Build data dictionary (follow the setup of TPC-D)
 - Create the test user and set the default table space and temporary table space for the user
 - Create dimension tables and fact tables (partitioned)
 - Disable constraints
 - Load data to all the tables (2 hours 37 minutes)
 - Gather global statistics, follow the setup of TPC-D
 - Ensure that any bitmap indexes subsequently created will be as small (compressed) as possible (2 hours 38 minutes)
 - Create indexes for each table (11 hours 44 minutes)
 - Add primary keys for fact tables
 - Enable the constraints back (4 hours 9 minutes)
 - Set degree of Parallelism for tables and cache the dimension tables

Simulation Execution

The top level test script (multi.tst) invokes the test driver (run_multi.sh) on the test suites with multiple users (scaling from 5 to 125). It takes about 24 hours to finish all of the workload testing.

The test suites were classified into the following five categories:

- Easy_xxx: The test suite contains multiple (such as 75) easy queries, each of which will finish in ≤ 5 minute. Easy_75 represents that 75 users access the database simultaneously with each user invoking one corresponding SQL script in the test suite
- Quick_xxx: The test suit contains multiple (such as 125) quick queries, each of which will finish in a very short time (i.e. ≤ 1 minute). Quick_125 represents that 125 users access the database simultaneously with each user invoking one corresponding SQL script in the test suite.
- Medium_xxx: The test suite contains multiple (such as 50) medium queries, each of which will finish in ≤ 20 minutes. Medium_50 represents that 50 users access the database simultaneously with each user invoking one corresponding SQL script in the test suite.
- Intense_xxx: The test suite contains multiple (such as 25) intense queries, each of which will finish in ≤ 1 hour. Intense_25 represents that 25 users access the database simultaneously with each user invoking one corresponding SQL script in the test suite.
- Mixed test suites: Each mixed test suite is named as QxxExxMxxlxx. An example may be "q25e15m10i5". This indicates that the test suite contains 25 quick queries, 15 easy queries, 10 medium queries, 5 intense queries and 55 users access the database simultaneously with each user invoking one corresponding SQL script in the test suite.

Conclusions

The HPUX Enterprise Operating Environment Application Infrastructure proved to be both stable and reliable when subjected to various load levels over an extended period of time.

- 100% application availability for 24 CHO with variable user load
- Transaction rates and response times well within acceptable tolerances
- No hardware failures
- No software failures

Test Results

CPU Utilizations in the Different Scenarios

The following two parameters were used in the simulation:

- cpu_count
- parallel_degree

The performance of the workload testing might vary depending on the settings of the above two parameters. The default for "Degrees of Parallelism" (DOP) is "2 * cpu_count".

Figure 1: CPU Utilization for 16 CPUs, 16 Degrees and 16 CPUs, 32 Degrees

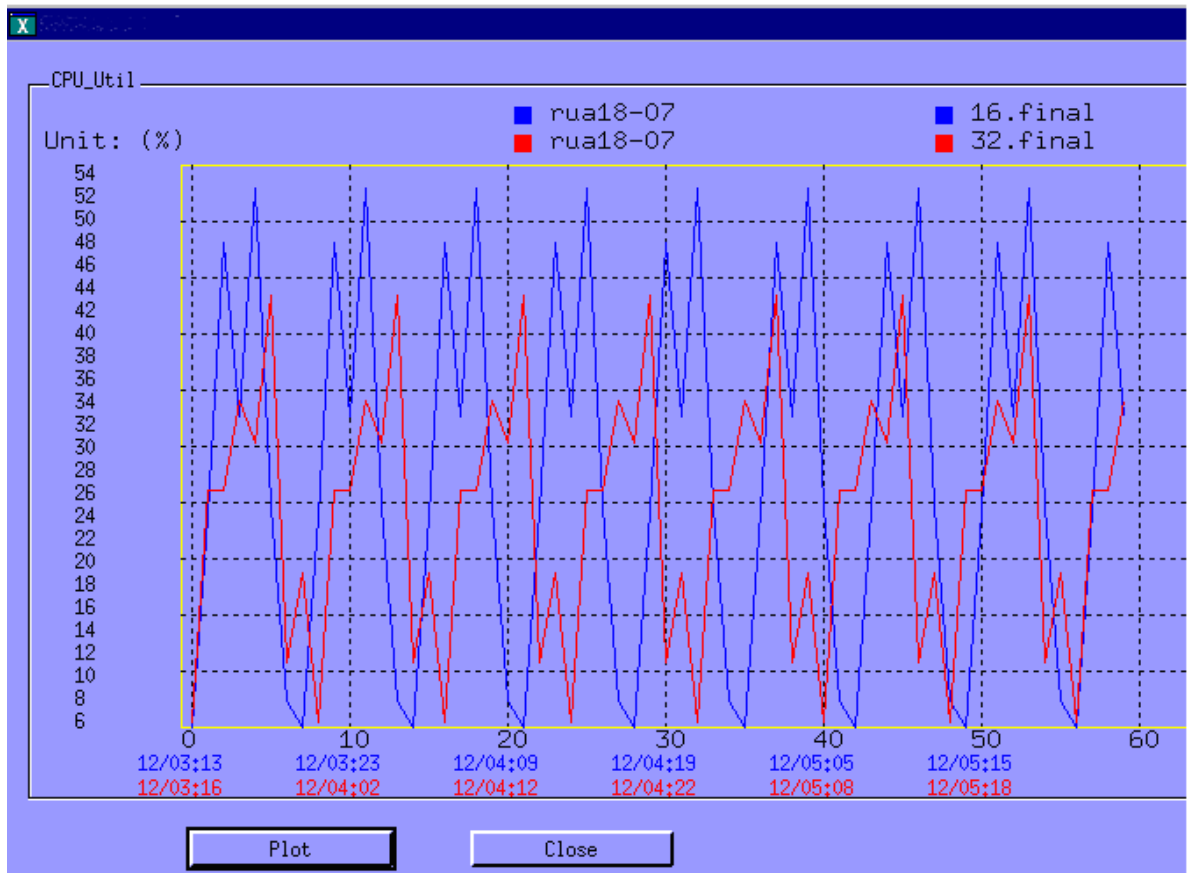


Figure 2: Disk Access Rates for 16 CPUs, 16 Degrees and 16 CPUs, 32 Degrees


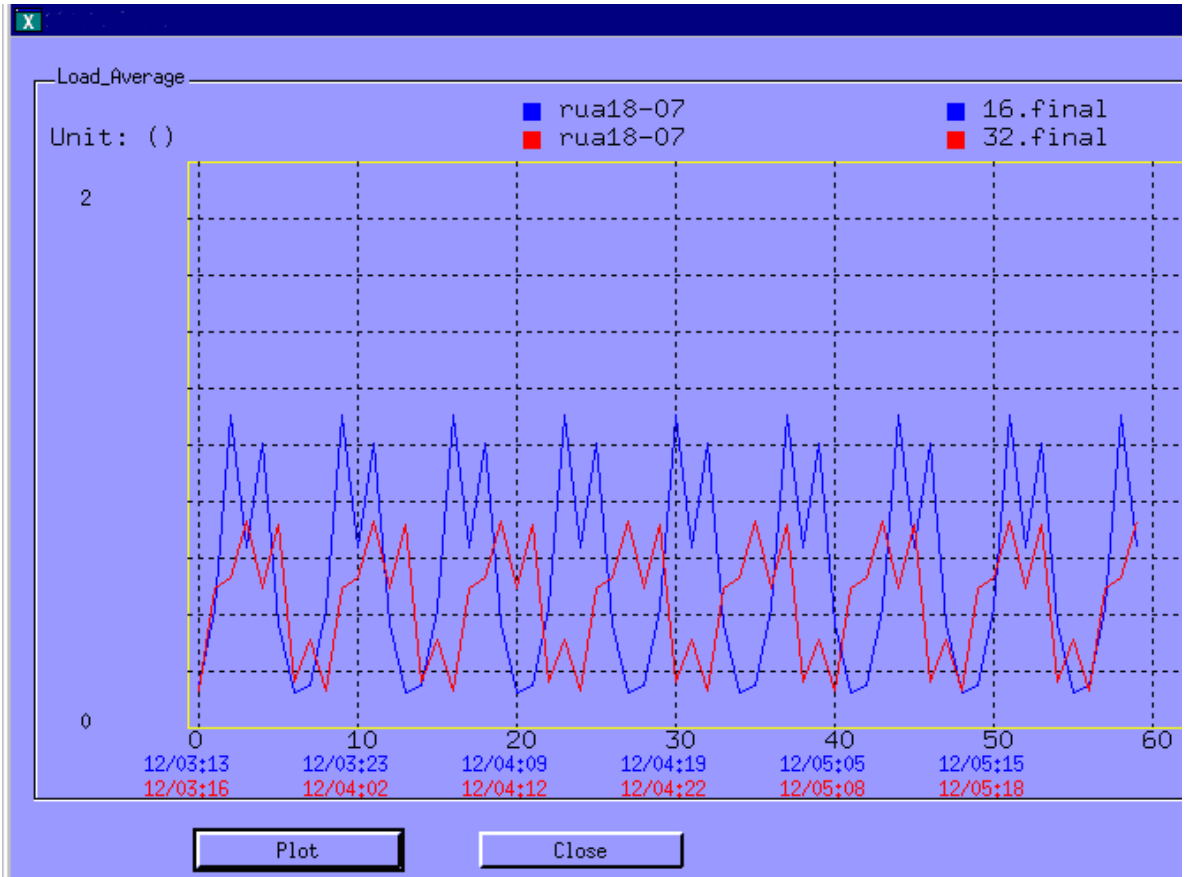


Figure 3: Load Averages for 16 CPUs, 16 Degrees and 16 CPUs, 32 Degrees



Response Times

Table 1: Test Response Times (hh:mm:ss)

Test Suite	Active CPU count/ Oracle CPU count/ Parallel Degree 16/08/16	Active CPU count/ Oracle CPU count/ Parallel Degree 16/16/16	Active CPU count/ Oracle CPU count/ Parallel Degree 16/16/32	Active CPU count/ Oracle CPU count/ Parallel Degree 16/16/48
Easy_15	00:01:36	00:01:35	00:02:10	00:02:17
Easy_25	00:02:00	00:01:56	00:02:11	00:02:00
Easy_35	00:02:08	00:02:03	00:02:11	00:02:12
Easy_45	00:02:16	00:02:10	00:02:17	00:02:03
Easy_55	00:02:12	00:02:13	00:02:20	00:02:19
Easy_65	00:02:29	00:02:38	00:02:24	00:02:37
Easy_75	00:02:35	00:02:39	00:02:52	00:02:40
Quick_25	00:01:13	00:01:24	00:01:19	00:01:14
Quick_50	00:00:29	00:00:29	00:00:33	00:00:31
Quick_75	00:01:23	00:01:15	00:01:15	00:01:14
Quick_100	00:01:57	00:01:49	00:01:53	00:01:50
Quick_125	00:01:34	00:01:45	00:01:34	00:01:52
Medium_10	00:06:38	00:06:26	00:10:21	00:09:11
Medium_15	00:07:13	00:07:22	00:07:46	00:07:29
Medium_20	00:08:05	00:08:08	00:10:40	00:29:53
Medium_25	00:10:09	00:17:27	00:09:59	00:09:15
Medium_30	00:11:15	00:10:22	00:30:12	00:30:28
Medium_35	00:17:05	00:18:34	00:11:37	00:11:51
Medium_40	00:12:56	00:12:48	00:33:12	00:19:33
Medium_45	00:20:18	00:18:58	00:16:35	00:15:37
Medium_50	00:19:21	00:18:58	00:23:53	00:43:05
Intense_5	00:11:02	00:11:07	00:12:29	00:11:42
Intense_10	00:18:19	00:18:23	00:20:27	00:18:43
Intense_15	00:19:16	00:18:48	00:21:29	00:19:18
Intense_20	00:23:37	00:23:50	00:24:07	00:23:29
Intense_25	00:24:31	00:24:07	00:24:42	00:23:36
q25e15m10i5	00:12:13	00:12:15	00:13:44	00:12:17
q50e25	00:02:12	00:02:22	00:02:28	00:02:21
q50m25	00:29:47	00:19:15	00:32:39	00:30:04

q50e15m10i5	00:21:46	00:21:57	00:21:31	00:21:51
q50e15m10	00:07:18	00:07:17	00:07:29	00:07:37

Known Problems and Tips

General Tips

- The Oracle oerr ora <error code> command can be used to find the probable Cause and recommended Action for Oracle error messages.
- Many “Out of resource” problems can be fixed by retuning kernel parameters.
- The web site <http://metalink.oracle.com> provides access to the Oracle bug database. Known problems always contain fix information or workarounds.

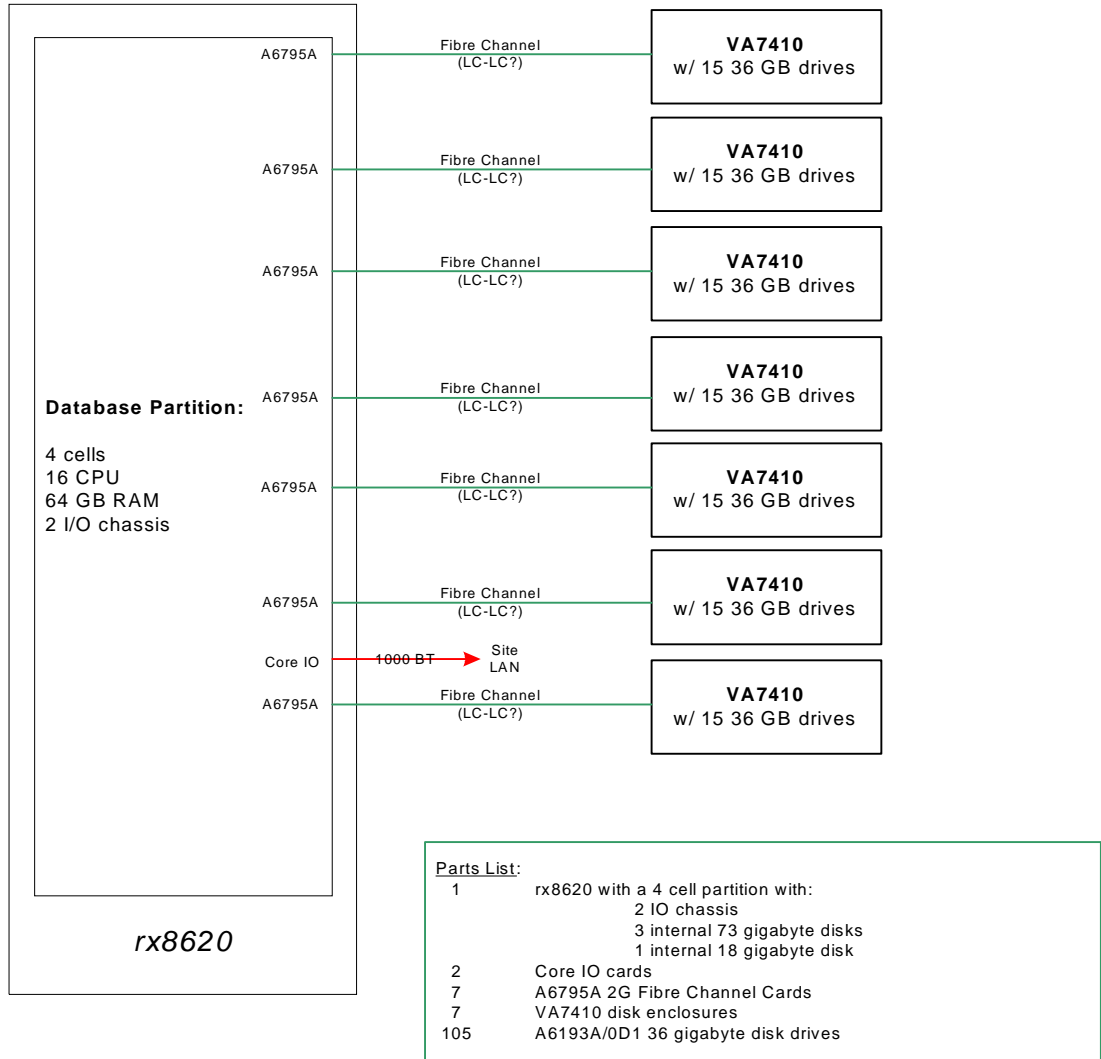
Known Problems

- The command vgcreate will hang if the new volume group’s metadata will not fit into a single physical extent. This will occur whenever the physical volume(s) are very large and the physical extent size is defaulted. In this case, the 134 GB LUN was enough to generate the problem. To work around the problem, we modified the command from “vgcreate –p 20” to “vgcreate –p 14 –s 32”. A patch, PHKL_30686, is available to address this.
- There is a new synchronization problem such that multiple simultaneous calls to pvcreate will result in all of the disks being assigned the same physical identifier (PVID). We worked around this by modifying our setup script to loop through all of the physical volumes sequentially.
- ORA-00020: Maximum number of processes (600) exceeded. This error may occur during workload testing. We had to increase the parameter “processes” in “init.tmp” at setup time and in “run_init.pops” at run time. (The original script set the value of 200 for version 8.1.5. It is recommended to set the value to 300 at setup time and 400 at runtime for Oracle version 8.1.7 and Oracle 9.2.0.2).
- The command swapinfo –tm may be used to check whether the system has enough swap space. This simulation required at least six gigabytes of swap.
- The kernel parameters nproc and swchunk must be increased to handle the configuration given to Oracle. For this simulation, the values were 20840 and 10240 respectively.
- ORA-27142: Could not create new processes. This error will occur whenever the kernel parameter nproc is not big enough.
- ORA-01652: Unable to extend temp segment by XXX bytes in tablespace SYSTEM. Instead of increasing the size of SYSTEM tablespace, we performed the DDL command alter user sys temporary tablespace ts_temp.

Appendix A – Hardware Environment

External Storage

Figure 4: Server Cabling to External Storage



Internal Devices and IO Cards

Figure 5: rx8620 Internal Peripheral and HBA Layout

Hostname	Model	# of Processors	CPU Speed
rua18-07: Partition 0 rua18-06: Mgmt Processor	ia64 hp server rx8620	16 CPU	1.5 GHz
Backplane Layout			
Bay 0 Chassis 0		Bay 0 Chassis 1	
Slot 1: 0/0/8/0/0 -> va7410	Slot 2: 0/0/10 -> va7410	Slot 3: 0/0/12/0/0 -> va7410	Slot 4: 0/0/14 -> va7410
Slot 5: 0/0/6/0/	Slot 6: 0/0/4	Slot 7: 0/0/2/0/0	Slot 8: 0/0/1/0/0
Slot 9: 0/0/14 A6795A HP Tachyon XL2 FCMS -> va7410	Slot 10: 0/0/10 A6795A HP Tachyon XL2 FCMS -> va7410	Slot 11: 0/0/8/0/0 A6795A HP Tachyon XL2 FCMS -> va7410	Slot 12: 0/0/4 A6795A HP Tachyon XL2 FCMS -> va7410
Core I/O 0	0/0/0/1/0 A7109-60001 PCI 1000Base-T Core 0/0/0/3/1 SCSI C1010 Ultra160 Wide LVD	Core I/O 1	1/0/0/1/0 A7109-60001 PCI 1000Base-T Core 1/0/0/3/1 SCSI C1010 Ultra160 Wide LVD

Front Panel Layout			
Core I/O 0	0/0/0/2/1 SCSI C1010 Ultra Wide Single-ended 0/0/0/2/1.2 DVD ROM	0/0/0/2/0 SCSI C1010 Ultra Wide Single-ended 0/0/0/2/0.6.0 73G HP 73.4GST373453LC	0/0/0/3/0 SCSI C1010 Ultra Wide Single-ended 0/0/0/3/0.6.0 18G HP 18.2GATLAS10K3_18_SCA
Core I/O 1	1/0/0/2/1 SCSI C1010 Ultra Wide Single-ended	1/0/0/2/0 SCSI C1010 Ultra Wide Single-ended 1/0/0/2/0.6.0 73G HP 73.4GST373453LC	1/0/0/3/0 SCSI C1010 Ultra Wide Single-ended 1/0/0/3/0.6.0 73G HP 73.4GST373453LC

Appendix B – Software Environment

Table 2: Software Products Installed

Products	Version	Description
B3701AA	C.03.71.23	HP GlancePlus/UX Pak for s800 11.23
B3901BA	C.05.50	HP C/ANSI C Developer's Bundle (S800)
B6834AA	B.01.03.01	HP-UX Security Patch Check Tool
B3909DB	B.11.23	HP Itanium Fortran 90 Compiler and associated products (S800)
B3913DB	C.05.50	HP aC++ Compiler (S800)
B6848BA	1.4.gm.46.3	Ximian GNOME 1.4 GTK+ Libraries for HP-UX 11iV1.6
B6849AA	B.02.01.02	Bastille Security Hardening Tool
B8339BA	B.03.00.03	servicecontrol manager Server and Agent Bundle
B8465BA	A.01.05.05	HP WBEM Services for HP-UX
B9073BA	B.06.00	HP-UX iCOD (Instant Capacity on Demand)
B9788AA	1.3.1.09.07	Java2 1.3 SDK for HP-UX
B9789AA	1.3.1.09.07	Java2 1.3 RTE for HP-UX
B9901AA	A.03.05.06	HP IPFilter 3.5alpha5
Base-VXVM	B.03.50.IA.00 3	Base VERITAS Volume Manager Bundle 3.5 for HP-UX
CDE-English	B.11.23	English CDE Environment
CMDVIEWSDM	A.1.07.00	hp StorageWorks Command View SDM
FDDI-00	B.11.23.00	PCI FDDI;Supptd HW=A3739B;SW=J3626AA
FibrChanl-00	B.11.23.01	PCI FibreChannel;Supptd HW=A6795A,A5158A
GigEther-00	B.11.23.01	PCI GigEther;Supptd HW= A4926A/A4929A/A6096A;SW=J1642AA
GigEther-01	B.11.23.01	PCI GigEther;Supptd HW=A6825A/A6794A/A6847A/ A8685A/A9782A/A9784A/A7109A
HPUX11i-OE- Ent	B.11.23	HP-UX Enterprise Operating Environment Component
HPUXBaseAux	B.11.23	HP-UX Base OS Auxiliary
HPUXBaseOS	B.11.23	HP-UX Base OS
IEther-00	B.11.23.01	PCI IEther;Supptd HW=A6974A
Judy	B.11.11.04.1 5	Judy Library - development and runtime libra ries for handling dynamic arrays
MOZILLA	1.2.1.03.00	Mozilla 1.2 for HP-UX

MySQL	3.23.54a.00	MySQL open-source database
NPar	B.01.00	nPartition Provider - HP-UX
OnlineDiag	B.11.23.01.2 1	HPUX 11.23 Support Tools Bundle, Jun 2003
ParMgr	B.11.23.01.0 0	Partition Manager - HP-UX
Sec00Tools	B.01.00.03	Install-Time security infrastructure.
T1456AA	1.4.1.03.01	Java2 1.4 SDK for HP-UX
T1457AA	1.4.1.03.01	Java2 1.4 RTE for HP-UX
T1471AA	A.03.10.007	HP-UX Secure Shell
USB-00	B.11.23	Object Oriented USB Driver
hpuxwsApache	A.1.0.06.01	HP-UX XML Web Server Tools
hpuxwsTomcat	A.1.0.06.01	HP-UX Tomcat-based Servlet Engine
hpuxwsWebmin	A.1.0.06.01	HP-UX Webmin-based Admin
hpuxwsXml	A.1.0.06.01	HP-UX XML Web Server Tools
perl	D.5.8.0.A	Perl Programming Language
scsiU320-00	B.11.23	PCI SCSI U320; Supptd HW=CoreIO
HPOVLIC	4.31.000	HPOVLIC
HPOVSAMDA	03.10.00.01 13	hp OpenView storage area manager HostAgent DIAL
HPOVSAMHA	03.10.00.01 13	hp OpenView storage area manager HostAgent Framework
HPOVSAMJR	1.4.1.3	hp OpenView storage area manager embedded JRE
HPOVSAMSG	03.10.00.01 13	hp OpenView storage area manager HostAgent SCSI Gateway
SMI-S_VA	1.02.02	hp StorageWorks SMI-S VA
SSLIC	1.0.0.0	HP Storage Software License Manager
Oracle RDBMS	9.2.0.2	Oracle Relational Database Management System

Appendix C – Patches

During testing, there was an lvm problem. At that time, the R&D lab provided a workaround. Now, however, the patch PHKL_30686 has been released to fix this problem.

© 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

