Best practices for replication of Oracle storage area networks white paper

Executive summary

Data recovery and protection are the top most priorities of business today and essential elements of any business continuity strategy—the bigger the enterprise, the more challenging the problem—requiring robust and reliable solutions. Final choices and IT decisions depend on how a business perceives its tolerance for data inaccessibility or loss. As such, protection can be as simple as backup of the information locally to a secondary storage device, or it can be complex, involving clustered application and database servers from multiple sites replicating data continuously to other secure remote sites over a long distance.

In response to the dilemma of choice, this document presents an easy-to-mange failover and failback recovery option available for Oracle® 10g data stored in storage area networks (SANs) and looks at three strategies for replicating Oracle databases over long distances between primary and remote HP StorageWorks XP arrays. The strategies focus on two different technologies and approaches—one using a real-time array-based replication of data across long distances, the other using a host and log-based replication system that updates data at intervals. Each strategy has it own advantages and disadvantages that address enterprise storage and replication requirements. This paper also looks at a third approach that integrates the two technologies into a single replication solution and considers the benefits of this combined strategy.

To validate these solutions, the HP StorageWorks Customer Focused Testing team (CFT) has tested Oracle 10g databases with Real Application Clusters (RAC) for replication, failover, and recovery across primary and remote SAN sites running Red Hat Enterprise Linux Advanced Server. The first phase of the testing used the array-based replication capabilities of HP StorageWorks Continuous Access XP along with other HP component software and firmware. The second phase of the tests used Oracle Data Guard, a host and network, and the third phase used a combination of both. This paper presents the results of these tests, along with best practices and recommendations for "best fit" scenarios.

Note that the configuration used in the test environment is only one of many possibilities that could have been selected. Oracle was selected as being the most popular and widely deployed enterprise database software, and HP Storage Works XP arrays as the "best-of-breed" storage arrays available in the market today. Both are well tested and well known for their reliability and dependability.

For your convenience, a list of all hardware and software equipment appears in Appendix A. All recommendations and best practices are derived from the test procedures and results presented in this document.

Introduction

Each of the following replication technologies offers specific advantages as part of a complete SAN replication solution. The following section reviews the HP Continuous Access XP and Oracle Data Guard products and how they replicate Oracle 10g data. It also describes the role of other software products used in the test configurations and test procedures.

HP Continuous Access XP solution

HP StorageWorks Continuous Access XP software is a high-performance, real-time remote data mirroring solution for replication between XP disk arrays. It is platform, operating system, and application agnostic. Integrated with Oracle 10g RAC, Continuous Access provides fast bi-directional failover and failback recovery that results in reliable data protection and sustained business continuity.

The Continuous Access XP solution for replication of Oracle SANs, as presented in this document, is hardware/array-based and utilizes block-level replication over Fibre Channel—as this is fast and provides high-level data consistency. Using Continuous Access XP connectivity, the HP StorageWorks XP12000 Disk Array is capable of running nearly full speed even while a remote mirroring operation is simultaneously in progress on the array. Array-based replication has little impact on server performance because the data is replicated directly between the arrays without requiring CPU compute cycles.

When replicating data over long distances, the choice of asynchronous and synchronous copy modes enables Continuous Access XP to meet the most demanding requirements for data currency and system performance. In asynchronous mode, Continuous Access XP completes the write I/Os at the local site without waiting for it to be written to the cache of the remote XP disk array. This optimizes performance at the local site, especially when replicating data over long distances. In contrast, synchronous mode requires that the write I/O be written to the remote cache before I/O completion occurs at the local host. While this method improves data integrity, in long-distance SANs the latency period between I/O completions can interfere with application performance. The advantage of Continuous Access XP is the flexibility it gives to administrators whose job it is to decide which mode is best for their environment. In addition, when Continuous Access XP is combined with Continuous Access XP Extension, it becomes a strategic tool for cost-effective continental-distance data replication and disaster recovery.

The Continuous Access XP test configurations developed for this paper include the following HP StorageWorks XP products.

Continuous Access XP Extension

Continuous Access XP utilizes Continuous Access XP Extension when functioning in asynchronous mode. Continuous Access XP Extension creates the asynchronous copy when the primary/local site server writes to cache and data volumes. In this process, the I/O is immediately acknowledged by the XP array to the server, and the cache data is written asynchronously to the remote XP array. The remote site XP array then acknowledges the write to the primary array, and the cache is cleared.

A definite advantage in using Continuous Access XP Extension is the reduced response time for primary site operations. There is also remote sequence stamping to ensure continuous remote-mirror consistency and efficient link-optimized remote copy operations.

Business Copy XP

HP StorageWorks Business Copy XP replicates data within HP XP disk arrays. When used with Continuous Access XP, it enables a wide range of critical enterprise solutions, including disaster recovery mirroring and zero downtime split-mirroring backups.

Business Copy XP was utilized in these tests (on the remote site) to maintain a mirror image of the data replicated from the primary site. The data could then be used for reporting, data mining, or data backup on the remote site after momentarily stopping the replication and then mirroring to obtain a known point-in-time data boundary. This also provided for the better utilization of the server resources at the remote site.

Command View XP

HP StorageWorks Command View XP is a web-based application interface for setting up and managing one or more XP disk arrays. It centralizes array administration, as well as providing advanced Fibre Channel and I/O management, and enables a single common management station for managing HP StorageWorks XP arrays across the organization. It has a modular architecture that also allows for future growth. Both Command View XP and Business Copy XP user interfaces are part of the Command View XP management framework.

RAID Manager XP for Replication Control

HP StorageWorks RAID Manager XP provides a command line interface (CLI) for configuring and managing Continuous Access XP and Business Copy XP volumes. It uses a specially designed XP command device to communicate commands to the XP disk array. In the tests described in this paper, the team used RAID Manager to manage failover, failback, pairing-splitting, and re-synchronization operations. The software allowed the flexibility of controlling both Continuous Access XP and Business Copy XP using a single CLI that could then be integrated into scripts and cluster software.

Refer to Appendix D for the RAID Manager configuration files used in the tests.

Performance Advisor XP

HP StorageWorks Performance Advisor XP is an invaluable tool used in tests to monitor the front-end (CHIP), back-end (ACP), Cache Memory (CM), Shared Memory (SM), and I/O performance and throughput for the XP array. The tool monitors cache efficiency and gives administrators the ability to identify and fix storage bottlenecks quickly. For more information, see the Performance monitor tool section.

Oracle Data Guard (ODG) solution

Oracle Data Guard (ODG) utilizes a primary Oracle database and one or more Oracle standby databases, with software to monitor, manage, and automate data replication, such that if the production database is taken offline or becomes damaged, the standby database can take over. In this way, ODG protects critical data from being interrupted or lost.

Within each Oracle configuration, standby databases are maintained as transactionally consistent copies of the production database. This consistency is maintained using Oracle logs. When data changes occur in the primary database, they are written to online redo logs. These changes are then transported to the standby databases and applied. Other Oracle management services facilitate switch-over and failover operations as needed.

Physical and logical standby databases

Oracle standby databases can be either a *physical standby* database or a *logical standby* database. A physical standby database has on-disk database structures that are identical block by block to the primary database; the standby database is "read-only." A logical standby database is an independent database that contains the same data as the primary database. Data is updated using SQL statements. The advantage of a logical standby database is that it can be used by other applications at the same time that it is being used for data recovery. ODG also supports both asynchronous and synchronous transaction transfers.

Like Continuous Access XP, ODG offers a centralized easy-to-manage solution for high availability and disaster recovery. In addition, its failover and switch-over capabilities allow easy role reversal between the primary and standby databases. As required, the standby database can become the primary database, and vice versa. This potentially minimizes downtime at the primary site as a result of planned (upgrades) or unplanned events.

ODG can include up to nine standby databases, consisting of both logical and physical standby databases that are configured as separate and independent. Primary and standby databases can run on a single node or in an Oracle RAC environment. For testing, ODG primary and standby databases were configured to run on an Oracle 10g RAC.

Protection modes and redo services

ODG provides three levels of data protection (maximize protection, performance, and availability) that determine whether redo log shipping is asynchronous or synchronous. When a business selects to "maximize protection" or "availability," all log transports are synchronous. If they select to maximize "performance," it can be both synchronous and asynchronous depending on the configuration.

On the standby site, ODG uses the Remote File Server Process (RFS) to receive redo logs from the primary and uses the Managed Recovery Process (MRP) to apply the redo information to the physical standby database.

Integrated Continuous Access XP and ODG solution

Both the Continuous Access XP and ODG solutions integrate seamlessly with the Oracle 10g RAC and protect clustered data by ensuring high availability and information integrity from network failures, site outages, and disastrous events. When configured together in the same environment, the Continuous Access XP and ODG products work together to replicate **all** data (including non-Oracle data, such as flat files, and other application data, like MS Exchange databases) stored in XP arrays. In such a configuration, Continuous Access XP provides data protection for Oracle and non-Oracle data, while ODG provides an additional layer of protection and replication for Oracle data only.

Test scenarios

The Oracle SAN tests were divided into three phases that correspond to the three objectives described in the Executive summary. They are:

- Phase 1—Continuous Access XP **only**
- Phase 2—ODG **only**
- Phase 3—Continuous Access XP and ODG together

The test configurations used for each phase are described in the Test configurations section. Utilizing these configurations, the HP CFT team planned and performed the following tests:

- Validation of bi-directional replication for Phase 1 and Phase 2 configurations.
- Oracle database recovery response from a single-node failure in a Phase 3 configuration only.
- Operating system and Oracle 10g RAC recovery from data path failure (inter- and intra-site) in a Phase 3 configuration.
- Application recovery time from network failure (cluster interconnect) for Phase 1, Phase 2, and Phase 3 configurations.
- Application recovery time from inter-site connection failure for Phase 1, Phase 2, and Phase 3 configurations.
- Recovery period after site outage (loss of power to production site) for Phase 1, Phase 2, and Phase 3.
- Capture of data accessibility, information/statistics at remote site for Continuous Access XP and ODG (all test phases).

The validation and performance data from these tests provide valuable comparative information for determining which replication strategy is best suited for a specific Oracle SAN application.

The two technologies—Continuous Access XP and ODG—have different functionality and benefits. As such, the scenarios test the effective recovery response times for several common failure cases and provide information that helps to determine if one or both technologies have significant performance impact within a range of similar hardware and software configurations for Oracle 10g clusters.

Note that bi-directional replication was tested for two phases: Phase 1, which tested a "Continuous Access XP only" configuration, and Phase 2, which tested an "ODG only" configuration.

Table 1 describes each test scenario in more detail. For step-by-step descriptions of the procedures for each test, refer to the Test procedures section. Results of these tests are provided in the Test results section.

Table 1. Test Scenarios for Replicated Oracle SANs

Test scenario	Configuration type	Description
Performance of bi-directional replication	Phase 1: Continuous Access XP only Phase 2: ODG only	Tests the replication of an Oracle database cluster over long distance between two identically configured SAN sites—primary and remote.
Recovery response time of Oracle 10g databases—array failure	Phase 3: Continuous Access XP and ODG together	Tests the response of an Oracle 10g database running on a single node to the failover of the XP array from primary to remote site containing the replicated storage.
Recovery response time of Oracle 10g database failure—node failure	Phase 3: Continuous Access XP and ODG together	Tests the response of an Oracle 10g database cluster on a two node OCFS cluster to a single node failure.
Recovery response time of OS and Oracle 10g RAC—data path failure	Phase 3: Continuous Access XP and ODG together	Tests the recovery response of the Red Hat Linux operating system and the Oracle 10g RAC to a data path failure (inter-site and intra-site).
Application recovery time—cluster network failure	Phase 1: Continuous Access XP only Phase 2: ODG only Phase 3: Continuous Access XP and ODG together	Tests the recovery response of the application to network failure related to the cluster interconnect.
Application recovery time—site-to-site network failure	Phase 1: Continuous Access XP only Phase 2: ODG only Phase 3: Continuous Access XP and ODG together	Tests the recovery response of the application response to ISL failure related to the site-to-site interconnect.
Recovery response time—power loss	Phase 1: Continuous Access XP only Phase 2: ODG only Phase 3: Continuous Access XP and ODG together	Tests recovery performance after a site outage that results in loss of power to production site.

Test configurations

There are three test configurations—one for each of the three phases of testing previously described. All test configurations consisted of two separate SAN sites that simulated one primary production site separated physically over a simulated distance (50km) from a second, remote or standby, site.

Each SAN site required identical setups for Red Hat Enterprise Linux Advanced Server on all servers, and connected to a dual-node Oracle 10g RAC on the site. Redundant Fibre Channel switches were used to enhance the availability of the SAN.

Because testing objectives focused on replication validation and comparative performance, the three test configurations differed in only four ways:

1. Replication software installed to replicate data between the two sites
2. Simulated length of distance between sites
3. Database size
4. Type of data replicated

Note that ODG replicates only Oracle databases across SAN sites (see Table 2), whereas Continuous Access XP replicates any data stored within a SAN.

Table 2: Test Phases and Data Replicated

Test	Uses…	To replicate…
Phase 1	Continuous Access XP only	Any data across the primary and remote SAN sites
Phase 2	ODG only	Only Oracle data across the primary and remote SAN sites
Phase 3	Both Continuous Access-XP and ODG	Any data across the primary and remote SAN sites

Refer to Appendix A for a complete list of software and hardware components used to configure the test environments for the test phases described in this next section.

Phase 1: Continuous Access XP test environment

Phase 1 testing utilized Continuous Access XP to replicate data stored in an Oracle 10g RAC across a simulated SAN. It also used Continuous Access XP Extension, Command View XP, and Performance Advisor XP to set up and monitor replication activity. Figure 1 illustrates the full physical test configuration for Phase 1 testing.

In this configuration, the primary site used one HP StorageWorks XP12000 Disk Array to simulate an OLTP production system that contained the critical business data. The remote site was similarly configured with one HP StorageWorks XP1042 Disk Array and did not provide access to the replicated data until a failover occurred. The distance between the two sites was simulated using a 50-km Fibre Channel spool.

To enhance the availability of the SAN, the SAN fabric was separated from the replication traffic by using physical switches for the HOSTS-SAN connectivity and SAN-SAN replication traffic.

Figure 1. Test configuration for Continuous Access XP with Oracle 10g

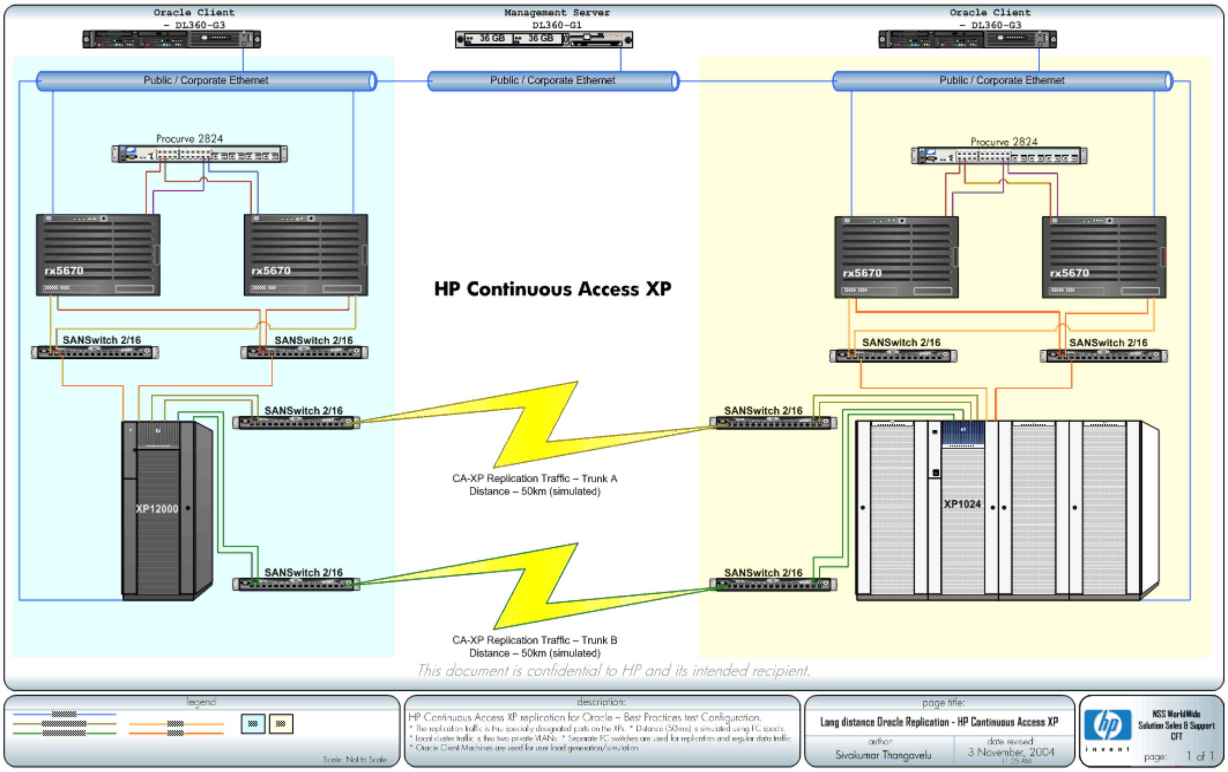

HP Continuous Access XP replication for Oracle – Best Practices test Configuration.
* The replication traffic is thru specially designated ports on the XPs. * Distance (50ms) is simulated using FC spools.
* Local cluster traffic is thru two private VLANs. * Separate FC switches are used for replication and regular data traffic
* Oracle Client Machines are used for user load generation/simulation

Phase 2: ODG test environment

The Phase 2 test configuration used ODG to replicate Oracle 10g databases across a SAN. The configuration supported data replication testing between a primary site and a remote site over a 10-km/50-km distance using Fibre Channel spools connected to the HP ProCurve 2824 network switches, which are GigE-capable and support long-distance SFP ports and single-mode Fibre connections. Figure 2 illustrates the full test configuration for Phase 2 testing.

In the Phase 2 configuration, the primary site was configured with one HP StorageWorks XP12000 Disk Array that operated a simulated OLTP production database continuously updated by ODG. The remote site was configured with one HP StorageWorks XP1042 Disk Array.

The nodes on the remote site were configured to act as a standby database updated by ODG.

Figure 2. Test configuration for ODG with Oracle 10g

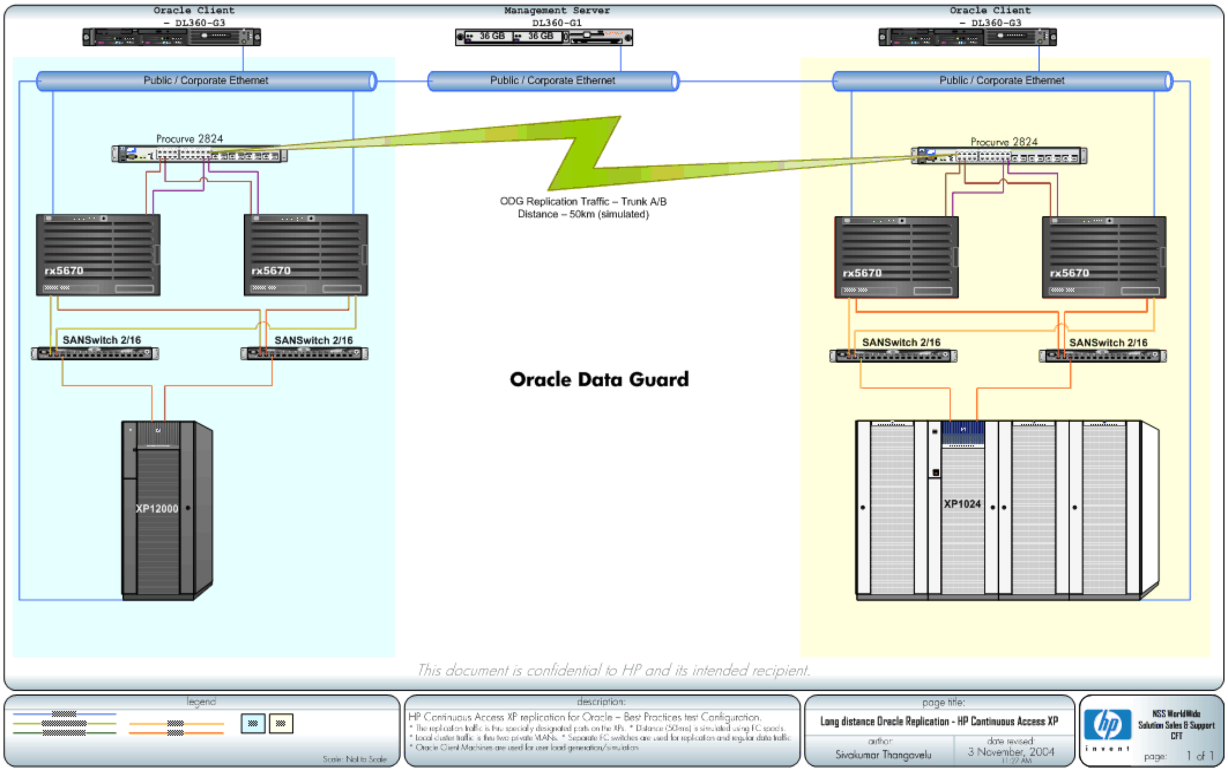

Oracle Client
– DL360-G3
Management Server
DL360-G1
36 GB
36 GB
Oracle Client
– DL360-G3
Public / Corporate Ethernet
Public / Corporate Ethernet
Public / Corporate Ethernet
Procurve 2824
Procurve 2824
ODG Replication Traffic – Trunk A/B
Distance – 50km (simulated)
rx5670
rx5670
rx5670
rx5670
SANSwitch 2/16
SANSwitch 2/16
SANSwitch 2/16
SANSwitch 2/16
Oracle Data Guard
XP12000
XP1024
This document is confidential to HP and its intended recipient.
legend
Scale: Not to Scale
description:
HP Continuous Access XP replication for Oracle – Best Practices test Configuration.
* The replication traffic is thru specially designated ports on the XPs. * Distance (50ms) is simulated using FC spools.
* local cluster traffic is thru two private VLANs. * Separate FC switches are used for replication and regular data traffic
* Oracle Client Machines are used for user load generation/simulation
page title:
Long distance Oracle Replication - HP Continuous Access XP
author:
Sivakumar Thangavelu
date revised:
3 November, 2004
11:27 AM
invent
NSS WorldWide
Solution Sales & Support
CFT
page: 1 of 1

Phase 3: Continuous Access XP and ODG co-existing environments

This system was configured using Continuous Access XP and ODG together to provide a broader capacity for replicating data across a long-distance SAN, as compared to using one or the other alone.

This phase utilized Business Copy XP to access data from the remote site while the primary site was fully functional. In this test environment, the LUNs replicated from the primary site to the remote site were mirrored on the remote site using Business Copy XP. The Continuous Access XP replication was then suspended for a moment, and the fully synchronized Business Copy XP mirror connection broken. The Continuous Access XP link was then re-enabled to allow for the availability of point-in-time data (in the mirror copy) for reporting and backup purposes.

Figure 3. Test configuration for Continuous Access XP and ODG with Oracle 10g RAC

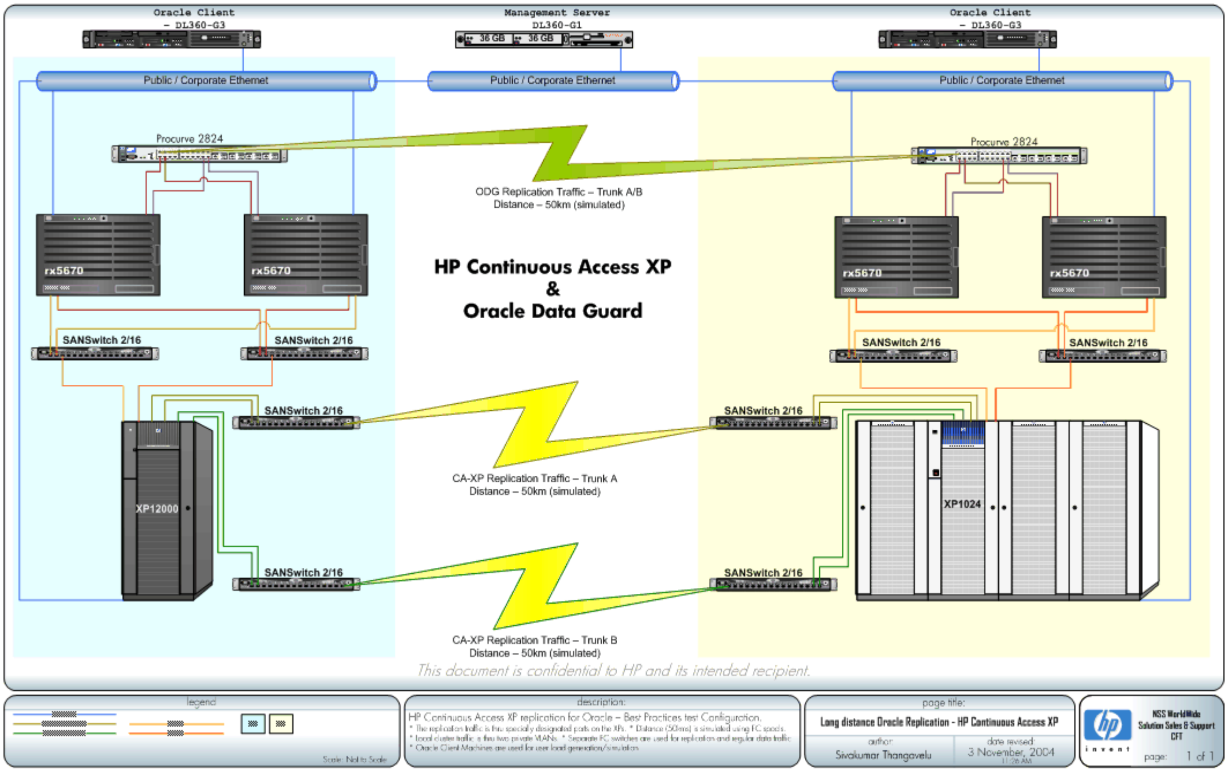
Oracle Client
– DL360-G3
Management Server
DL360-G1
36 GB
36 GB
Oracle Client
– DL360-G3
Public / Corporate Ethernet
Public / Corporate Ethernet
Public / Corporate Ethernet
Procurve 2824
Procurve 2824
ODG Replication Traffic – Trunk A/B
Distance – 50km (simulated)
rx5670
rx5670
HP Continuous Access XP
&
Oracle Data Guard
rx5670
rx5670
SANSwitch 2/16
SANSwitch 2/16
SANSwitch 2/16
SANSwitch 2/16
SANSwitch 2/16
SANSwitch 2/16
XP12000
CA-XP Replication Traffic – Trunk A
Distance – 50km (simulated)
XP1024
SANSwitch 2/16
SANSwitch 2/16
CA-XP Replication Traffic – Trunk B
Distance – 50km (simulated)
This document is confidential to HP and its intended recipient.
legend
Scale: Not to Scale
description:
HP Continuous Access XP replication for Oracle – Best Practices test Configuration.
* The replication traffic is thru specially designated ports on the XPs. * Distance (50km) is simulated using FC spools.
* local cluster traffic is thru two private VLANs. * Separate FC switches are used for replication and regular data traffic
* Oracle Client Machines are used for user load generation/simulation
page title:
Long distance Oracle Replication - HP Continuous Access XP
author:
Sivakumar Thangavelu
date revised:
3 November, 2004
11:26 AM
invent
NSS WorldWide
Solution Sales & Support
CFT
page: 1 of 1

Oracle database load generator

Tests used Benchmark Factory Suite by Quest Software to generate Oracle database loads. Benchmark Factory uses industry-standard benchmarks, such as TPC-C, TPC-D, and ASP3AP, and real workload simulation to provide accurate production database load and performance testing for hardware and software.

For the tests in this report, the TPC-C benchmark suite was selected to meet testing requirements for transaction loads. Benchmark Factory also helped to identify performance-related issues such as production bottlenecks, database migration, and upgrade impact, and to correct the sizing of Oracle 10g RAC node configurations.

Real-world latencies, such as manual keying time and inter-arrival times, were accounted for and included in the test simulations. The workload generation simulated a fully scaleable architecture capable of supporting a very large number of users (maximized for up to 1,000 users) and simulated by using Benchmark Factory agents running on multiple nodes.

Performance monitor tool

The HP CFT team configured Performance Advisor to collect and export data from its database using automated scripts. This allowed for daily or periodic collection of information and report generation. Performance Advisor generated data graphs based on selected metrics; two of these are shown in Figure 4.

For testing the Oracle configurations, certain performance counters were monitored to ensure maximum performance on the system. The metrics measured were:

- Side-file metrics

  When using Continuous Access XP in asynchronous mode, it is important to monitor the side file. The side file uses cache space on the XP array, which can be up to 60 percent of the cache.

- Host port throughput

  The Performance Advisor tool was used to check load balancing across ports, CHIP Fibre port utilization and LDEV reads/writes (MB/sec), load balancing on the CHIP ports, and the number of random and sequential reads/writes.

- Database and log disk I/O utilization

  For these tests, the team measured reads/writes per second at LDEV level and RAID-group level. The number of cache writes pending were monitored to see how quickly data was getting updated to disk. Cache hit utilization was also monitored to determine if enough cache existed for replication processes.

- Back-end metrics

  The tests monitored utilization of ACP processors and reads/writes to and from cache. Because Continuous Access XP and Business Copy XP can increase back-end activity, ongoing monitoring is encouraged to proactively manage the throughput so that it does not negatively impact the primary site host I/O.

- Continuous Access XP link throughput

  To ensure balancing of the workload across multiple ports, Continuous Access Fibre port utilization and LDEV read/write activity was monitored.

Figure 4. Performance Advisor graph of monitored activities

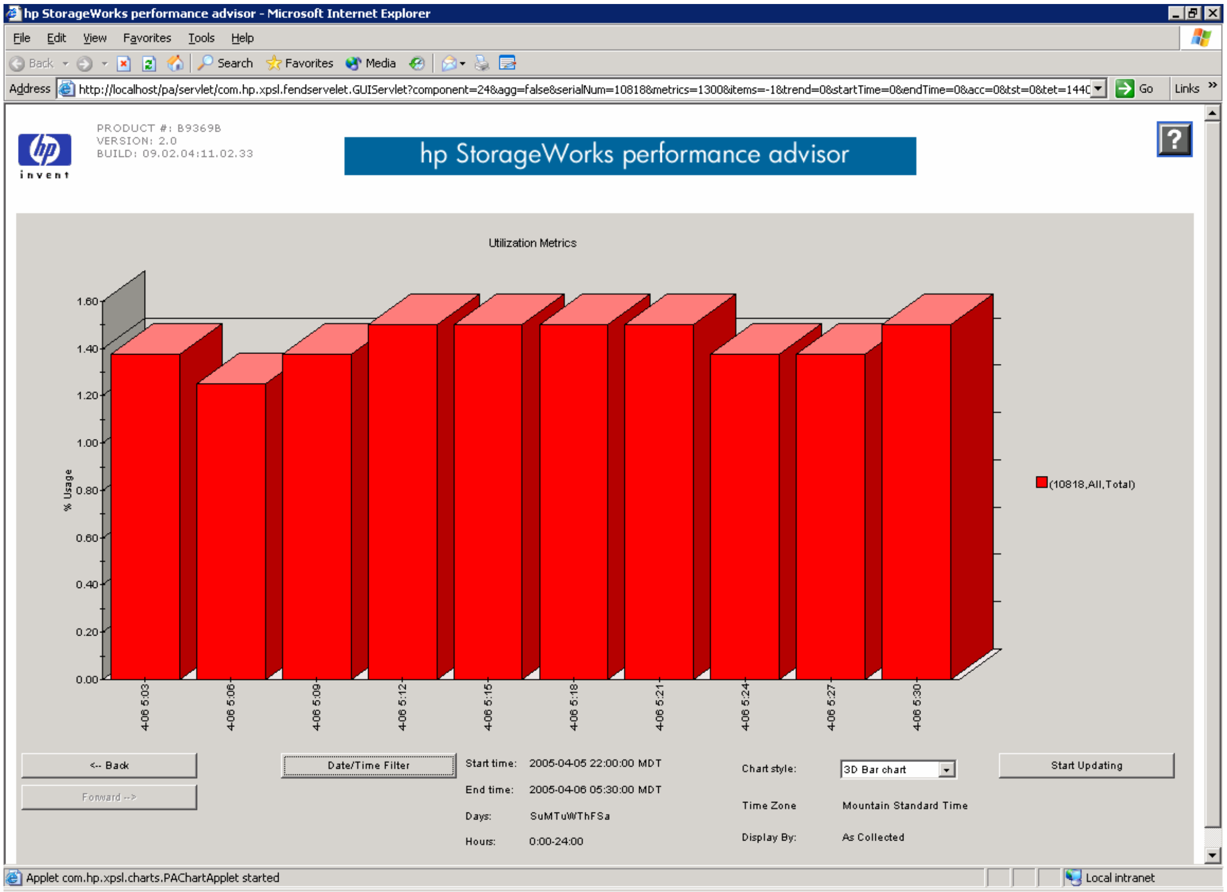

Address: http://localhost/pa/servlet/com.hp.xpsl.fendservelet.GUIServlet?component=24&agg=false&serialNum=10818&metrics=1300&items=-1&trend=0&startTime=0&endTime=0&acc=0&tst=0&tet=144C

PRODUCT #: B9369B
VERSION: 2.0
BUILD: 09.02.04:11.02.33

Start time: 2005-04-05 22:00:00 MDT
End time: 2005-04-06 05:30:00 MDT
Days: SuMTuWThFSa
Hours: 0:00-24:00

Chart style: 3D Bar chart
Time Zone: Mountain Standard Time
Display By: As Collected

Applet com.hp.xpsl.charts.PAChartApplet started

Figure 5. Performance Advisor graph of monitored activities

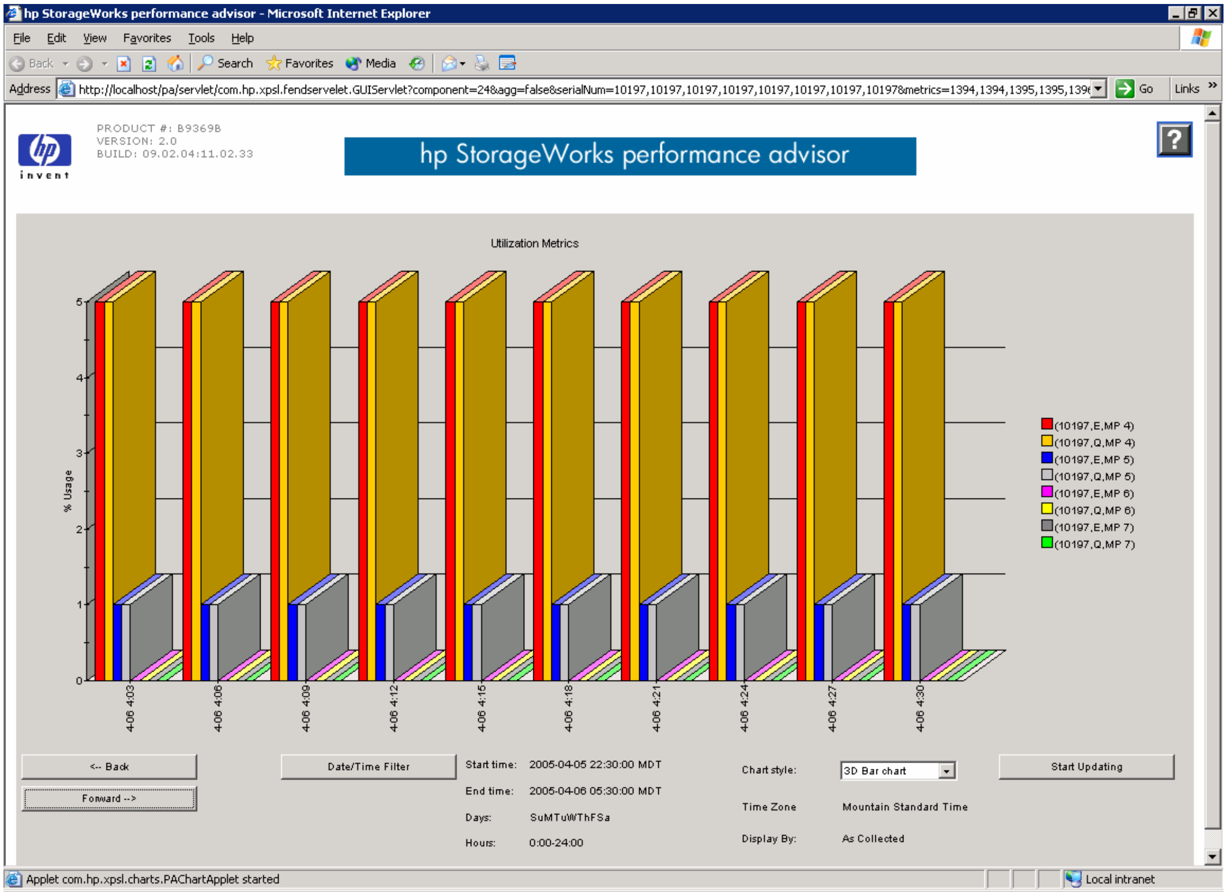
hp StorageWorks performance advisor - Microsoft Internet Explorer

Address http://localhost/pa/servlet/com.hp.xpsl.fendservelet.GUIServlet?component=24&agg=false&serialNum=10197,10197,10197,10197,10197,10197,10197,10197&metrics=1394,1394,1395,1395,139€

PRODUCT #: B9369B
VERSION: 2.0
BUILD: 09.02.04:11.02.33

hp StorageWorks performance advisor

Start time: 2005-04-05 22:30:00 MDT
End time: 2005-04-06 05:30:00 MDT
Days: SuMTuWThFSa
Hours: 0:00-24:00

Chart style: 3D Bar chart
Time Zone Mountain Standard Time
Display By: As Collected

Applet com.hp.xpsl.charts.PAChartApplet started

Figure 4 shows the CHIP port utilization percentage and Figure 5 shows the back-end Microprocessor (MP) usage. This graphed information enabled the testing team to determine which back-end MPs were underutilized and which MPs to use to evenly spread the workload.

Test procedures

The testing procedure consisted of four separate stages:

1. Installation and configuration
2. Validation and performance testing of Continuous Access XP
3. Validation and performance testing of ODG
4. Validation and performance testing of Continuous Access XP and ODG together

The following sections describe the test procedures for these four stages.

Pre-Test: Installation and configuration

During this stage, hardware and software test configurations were developed. The procedure used by the team is described in Table 3 and the resulting configurations are illustrated in Figures 1–3.

Storage systems—XP 12000 Disk Array and XP1024 Disk Array

HP recommends the following guidelines for configuring the primary and remote XP-array storage systems:

- Configure a mix of 7d+1p, 3d+1p, and 2d+2p RAID groups—7d+1p provides for more physical spindles and therefore increases read/write performance; the conventional 3d+1p provides a fine balance of performance and reliability, and 2d+2p provides better protection against disk failures.
- Create Open-Vs for log and data—Carved out of 7d+1p RAID group. This allows a single LUN to span across more physical spindles increasing the write performance.
- Create Open-9/Open-E for Oracle Software, Oracle Cluster Registry, Cluster Services Voting Disk, and Performance Analyzer command device—Carved out of 2d+2d RAID type. This provides better protection against multiple disk failures and gives comparatively good I/O performance.
- Create Open-V for Oracle Recovery Files (flashback, backup, and archive log files)—3d+1p RAID group. Because typically these files have voluminous data, there is a greater need to balance performance and reliability. Open-V spreads the data across all the disks to provide the performance, where as 3d+1p provides good data reliability.
- Create dual redundant paths to each LUN presented to host. This helps to provide access continuity in the event of a path failure.
- Create LUSE volumes comprised of Open-Vs spread across RAID groups. This increases application performance and spindle/track count.
- As you read the following tables and information in the Appendix, notice the way in which the Fibre Channel ports for the XP12000 Disk Array and XP1024 Disk Array are named and numbered.

Red Hat Enterprise Linux Advanced Server 3 Update 3

- Installation of the Linux operating system—Select the “Everything” option for the packages you want installed.
- Default kernel parameters—This test configuration requires changes to the default kernel parameters to accommodate recommendations and requirements for Oracle. The kernel parameter changes can be made permanent by adding them to the `/etc/sysctl.conf` file. See the sample file in Appendix C.
- ssh and/or rsh configuration—The OS user (often “oracle”) that owns the Oracle Software distribution should be set up for user equivalence across all cluster nodes and should be able to log in to other nodes without a passphrase or a password. This is a requirement for proper functioning of Oracle RAC.
- QLogic Failover Driver—For accessing the LUNs and managing the multiple paths available to them, the HP qualified and supported QLogic driver stack was used with the failover option enabled. This allowed for transparent switching of the disk-access paths in the event there was a path failure.
- Shared disks—Due to the clustered nature of the solution, the LUNs that stored the database files, the Oracle cluster registry, and the quorum disk were all required to be shared and visible across all nodes.

- Disk configuration—Due to the inconsistent nature of the “LUN-to-device-name mapping” functionality in Linux, the testing team used the “labels” option of the file systems to ensure correct identification and mounting of the LUNs to the correct directory structure. For information about the `/etc/fstab` file, see Table 10 in Appendix B.
- Directory structure organization for Oracle—Oracle home directories were on local file systems. Because this was not a production environment, both the binaries and shared data were grouped under a single-parent directory structure. For production deployments, HP strongly recommends that you follow recommendations and requirements in the *Oracle Optimal Flexible Architecture* specifications.

In addition to required installation settings, HP recommends that you follow the installation and Linux operating system best practices listed in the Best practices and recommendations section.

Table 3. Installation and Configuration Procedure

Pre-test installation and configuration			
Step	**Task**		**Command Line Input / Comments**
1	SET UP AND CONFIGURE INTEGRATED SAN: Oracle Cluster File System, Oracle 10g RAC, Red Hat Enterprise Linux Advanced Server, and HP StorageWorks XP Arrays		
	1.1	Install Red Hat Linux Advanced Server 3.0 Update 3 using an identical configuration on each server node in the cluster.	

	1.2	Install Linux multi-path driver (QLogic failover driver) identically on each server. In `/etc/modules.conf` file for the QLogic multi-path failover driver, enter…	`options qla2300  ql2xmaxqdepth=16`
			`qlport_down_retry=3`
			`qlogin_retry_count=16`
			`ql2xfailover=1`

	1.3	Give all cluster servers access to the shared vDisks in their respective XP arrays.	See Appendix B, Table 10.
	1.4	Install Oracle Cluster Manager (OCM) and Oracle Cluster File System (OCFS) on each server.	Installations can be done in parallel: - Use multiple tabbed sessions within one ”konsole” terminal session. - Enable the ”send input to multiple sessions” option (CTRL+SHIFT+1). - Store and install the binaries/rpms from the Network File System (NFS) share.
	1.5	Install Oracle 10g and with cluster enabled.	Oracle Universal Installer automatically detects if the CRS is installed and running. It also provides options to provision the RAC components.
2	SMOKE TEST THE QLOGIC MULTI-PATH DRIVER BEHAVIOR		
	2.1	With database active, disconnect a Fibre Channel cable between the host and Fibre Channel switch, and perform steps 2.2–2.4.	
	2.2	Check the database status to verify that it has not been disrupted.	
	2.3	Verify that the active path has failed and automatically switched over to the alternate path.	

	2.4	Restore cable and verify that the path status is restored to normal.	View the `var/log/messages` file for relevant information. Note: Another method is to install and use QLogic SANSurfer utility to view this information.
	2.5	With database active, power off one Fibre Channel switch and perform steps 2.6–2.8.	
	2.6	Check the database status to verify that it has not been disrupted.	Ensure that the connections between the "switch–to" server and "switch-to" storage conform to recommendations in the *SAN Design Guide*.
	2.7	Verify that the path status is marked as "Failed."	
	2.8	Restore the cable and verify that the path status has been restored to normal.	
3	TEST ORACLE 10G RAC BEHAVIOR		
	3.1	Simulate a crash of one server node in primary site.	
	3.2	Verify database was active on other nodes in cluster (same-site).	
	3.3	Fail over to the remote site and bring up the database there.	
	3.4	Verify database is active on the nodes in cluster at the remote-site.	

Phase 1: Functional replication tests with Continuous Access XP

The following test procedure is designed to validate the crash recovery capabilities of Continuous Access XP in an Oracle 10g environment and to gather performance test data to establish best practices for long-distance replication of Oracle SANs.

Phase 1 tested the performance of Continuous Access XP as part of an integrated platform with the Oracle Cluster File System (OCFS) and Oracle 10g RAC running on Red Hat Enterprise Linux Advanced Server 3. Servers at the remote standby site were configured with QLogic multi-path firmware, the same as the primary site servers.

For this test, Continuous Access XP was configured to replicate Oracle 10g databases over a distance of 35 km between the primary site (XP12000 Disk Array) and the remote site (XP1024 Disk Array) using a long-link Fibre Channel spool. After the underlying database and network services were created with the Oracle Database Configuration Assistant Utility, the client node was configured. This server ran Benchmark Factory, which communicated with the database and built the database size to test requirements. In this way, a 300-GB baseline Oracle 10g database was configured using the Benchmark Factory load generation tool to simulate 40 percent of CPU load, with the workload utilizing 40 percent writes and 60 percent reads.

When the data loading activity and creation of the required rows and data structure was complete, the tablespace was taken offline and a backup copy was created (by means of a transportable tablespace) by exporting meta-data using the Oracle data pump utilities. This created a baseline that could then be used to create the baseline for the ODG databases.

Table 4. Phase 1 Continuous Access XP test procedure

Replication with Continuous Access XP			
Step	**Task**		**Command Line Input / Comments**
1	FAIL FC INTER-SITE LINK		
	1.1	Disconnect 50-km link between primary and remote SAN sites.	See Appendix B, Table 11 for the failure procedure.
	1.2	Continue to run application for 24 hours.	
	1.3	Reconnect primary site to remote site.	See Appendix B, Table 12 for the failover procedure.
	1.4	Initiate resynchronization of primary and remote site databases.	See Appendix B, Table 13 for the resynchronization and failback procedure.
2	FAIL NETWORK CLUSTER INTERCONNECT		
	2.1	Simulate the failure of the network interconnect on the primary site cluster by logging in to the network switch and disabling the port or manually unplugging the network cable	Database active and accessible through the primary node.
	2.2	Verify continuity of service.	
3	SITE OUTAGE AT PRIMARY (PRODUCTION) SITE		
		Manual failover procedure	See Appendix B, Table 12.
		Verify	See Appendix B, Table 12.

Phase 2: Functional replication tests with ODG

The Phase 2 test procedure validated the replication capabilities of ODG with the Oracle Cluster File System (OCFS) and Oracle 10g RAC on Red Hat Enterprise Linux Advanced Server 3.0 Update 3.

For this test, ODG was configured to replicate Oracle 10g databases over a distance of 10 km between the primary site (XP12000 Disk Array) and the remote site (XP1024 Disk Array) using a long-link Fibre Channel and TCP/IP as the transmission protocol. A 50-GB baseline Oracle 10g database was configured using the Benchmark Factory load generation tool to simulate 40 percent of the CPU load, with the workload utilizing 40 percent writes and 60 percent reads. Servers at the remote standby site were configured with the QLogic multi-pathing firmware, the same as the primary site servers.

Table 5. Phase 2 Oracle Data Guard Test Procedure

Replication with ODG			
Step	**Task**		**Command Line Input / Comments**
1	CREATE ORACLE PHYSICAL STANDBY		
	1.1	Create physical standby database for replication, transfer log files, and start the Managed Recovery Process (MRP).	See Appendix B, Table 14 for detailed steps.
2	SWITCH OVER WITH PHYSICAL STANDBY (PLANNED)		
	2.1	Convert primary database into physical standby.	See Appendix B, Table 15.
	22.	Convert the old standby database into the primary.	See Appendix B, Table 15.
3	FAIL OVER WITH PHYSICAL STANDBY (UNPLANNED)		
	3.1	Perform Terminal Recovery on the physical standby using MRP, and convert the standby database to primary.	See Appendix B, Table 16.

Phase 3: Functional replication tests with Continuous Access XP, Business Copy XP, and ODG together

Phase 3 testing verified the replication of Oracle databases using Continuous Access XP, Business Copy XP, and ODG with OCFS and Oracle 10g RAC on Red Hat Enterprise Linux Advanced Server 3.0 update 3.

The final configuration included a Continuous Access–based database (named CAXP) and an ODG-based replicated database (named ODG) sharing resources from the same servers (see Table 6). This configuration required dual redundant paths between the primary and remote site to support two-way data replication, creating the need for additional network bandwidth, server bandwidth, and SAN infrastructure. The size and the emulation type for the primary volume and secondary volume in both Continuous Access XP and Business Copy XP configurations needed to be exactly the same (block to block match) on both XP arrays. This is a requirement to successfully configure Continuous Access XP and Business Copy XP.

Table 6. Database configuration used in Phase 3 testing

Database Instance	On Node
CAXP (DB unique name)	Primary/Remote Site
CAXP1 (instance 1)	cftsrv107 / cftsrv111
CAXP2 (instance 2)	cftsrv109 / cftsrv113
ODG1 (DB unique name)	Primary Site
ODG11	cftsrv107
ODG12	cftsrv109
ODG5 (DB unique name)	Remote Site
ODG51	cftsrv111
ODG52	cftsrv113

The fact that two basic replication methods—Continuous Access XP/Business Copy XP and ODG—must be taken into account impacts the configuration of shared memory resources. In this phase, there are two databases and each database must have its own shared memory resources on each server. This adds a level of complexity that does not exist in a Continuous Access XP–only or ODG-only configuration.

For example, an administrator or DBA must consider the I/O demands on particular disk volumes. When both replication methods are used, it is important to not create too much I/O demand on one volume. In certain cases, parallel operations that use both replication methods simultaneously can create I/O contention.

Note that both ODG and Business Copy XP support have an offline backup or reporting process at the remote site. Each method has its own advantages and disadvantages.

Table 7. Phase 3 Continuous Access XP and ODG test procedure

Step	Description	
1	DETERMINE USER IMPACT OF USING BOTH CONTINUOUS ACCESS XP and ODG	
2	CAPTURE BASELINE ORACLE PERFORMANCE DATA	Measured as transactions per second (TPS) and other representative counters.
3	CAPTURE BASELINE SERVER PERFORMANCE DATA	Measured as queue depth and CPU utilization under load.
4	CAPTURE ORACLE AND SERVER PERFORMANCE WITH CONTINUOUS ACCESS XP RUNNING	Oracle workload generated with Benchmark Factory.
5	CAPTURE ORACLE AND SERVER PERFORMANCE WITH ODG RUNNING	Oracle workload generated with Benchmark Factory.
6	CAPTURE RECOVERY POINT AND RECOVERY TIME WITH CONTINUOUS ACCESS XP RUNNING	
7	CAPTURE RECOVERY POINT AND RECOVERY TIME WITH ODG RUNNING	

Test results

The following test results correspond to the test objectives described in the Test Scenarios section.

Validation of bi-directional replication of Oracle database clusters

In Phase 1 with Continuous Access XP, tests verified the bi-directional replication of the Oracle 10g database (contained within the disk volumes replicated) over a distance of 10 km/50 km in synchronous mode and over 50 km in asynchronous mode using Continuous Access XP Extensions. These tests used RAID Manager XP for control and management of replication and XP Command View for centralized administration.

With a reference database size of 50 GB and 300 GB, a workload of 1,000 virtual users based on a TPC-C Benchmark load, and 10-ms–200-ms randomly variable delays (all generated by Benchmark Factory), the metrics obtained showed that there was very little or no impact on the Oracle application when Continuous Access XP was run in asynchronous mode. When run in synchronous mode, very minimal performance impact was noticed. Tests resynchronized only those changes made after the failure and restore of data.

In Phase 2, with ODG, tests verified the bi-directional replication of an Oracle 10g database cluster over a distance of 10 km, simulated using Fibre Channel spools connected through HP ProCurve 2824 network switches (using GigE interconnects) with ODG using a physical standby database and the log transport and log apply services.

Oracle database recovery from a single node failure

Continuous Access XP and ODG (Phase 3) tests verified the ability of Oracle 10g databases on a two-node Oracle 10g RAC running on top of OCFS to recover from a single node failure. With Transparent Application Failover (TAF) turned on and configured within Oracle Net and the Cluster Ready Services fully functional, during the failure of one of the two nodes in the database cluster, the system was able to seamlessly retain access to the database.

Operating system and Oracle 10g RAC recovery from data path failure (intra-site)

For both Continuous Access XP and ODG (Phase 3), tests showed that the intra-site data path (that is, the path from the XP array to the server) uses the failover option to fail over to another available path and maintains access to the database.

Operating system and Oracle 10g RAC recovery from data path failure (inter-site)

For inter-site data path or link failure, the replication software handled the situation differently in each test phase.

For Continuous Access XP, the inter-site data path was the Fibre Channel link used for replication of data from the primary to the remote site. Failure of this link put the volume pair (consisting of the primary volume "`p-vol`" and secondary volume "`s-vol`") into suspended mode. After the link was restored, the suspended volume was resynchronized either manually or programmatically. In this case, the tests were performed using RAID Manager XP, which allowed fault detection and management mechanisms to be incorporated into scripts.

For ODG, the inter-site data path was typically the network interconnect that carried the replication traffic (log and control data). When the network failed, "dead connection detection" was used on the primary site. Depending on the process, the MAX_FAILURE, NET_TIMEOUT, and REOPEN parameter values in the ODG definition determined how the logs were sent across the network. If the log writer process (LGWR) detected a dead connection, it waited for the operating system TCP stack to clear out the session. On the remote site, the Remote File Server (RFS) process terminated itself when the OS network software detected a failure, proceeding as normal until this happened.

Application recovery time from network failure (cluster interconnect)

In all test phases, the replication software successfully triggered a restart on the secondary node.

Application recovery time from inter-site connection failure

In Phase 1, there was no impact on the Oracle database or the Oracle clients upon the failure of ISL. Continuous Access XP recognized the failure, put the volume pairs in suspend mode, and started accumulating the transmit data to the side file. After the link was reestablished and resynchronization initiated, the volume pair resumed the replication as normal. In the event that the side file exceeded the limit, a full pairing operation was automatically initiated.

In Phase 2, where the ISL was the optical link between the two GigE network switches, the Oracle media recovery process terminated after the operating system network stack timed out. This required the testing team to restart the process on the remote site after the link was reestablished.

Recovery period after site outage (loss of power to production site)

In Phase 1 with Continuous Access XP, control was transferred over to the remote volumes by issuing the `horctakeover` command available with RAID Manager XP. This forced the `s-vol` to become the `p-vol`. The next step was to mount the LUNs on to the hosts and bring up the database. While starting the database, the remote site recognized an improper shutdown of the instance on the primary site, and the crash recovery process for the node instance then started up and recovered to a known stable state.

In Phase 2 with ODG, the media recovery process stopped when the network connection to the primary database timed out. The failover method was manually initiated, during which the role of the primary database was transferred to the remote database. The remote database then was mounted and opened as the primary database. Possibility of data loss was dependent on the type of crash, the amount of data in the transmission link, and the type of ODG configuration.

In Phase 3, results were the same as the combination of the preceding two phases.

Capture of data accessibility, information, and statistics at the remote site for Continuous Access XP and ODG

For both replication technologies, data on the remote site was accessible only when some form of control transfer had occurred from the primary database to the remote. The referenced configuration used manual intervention to transfer control from the primary site to the remote site.

Depending on the amount of data that had to be recovered, the recovery time for both Continuous Access XP and ODG varied. With Continuous Access XP, even with a manual process, the timeframe was well within five minutes when recovering a database instance from a crash, and this included the time required to issue the `horctakeover` command, mount the LUNs, and bring up the database. With ODG, the time period was similar, but was increased when the log files that had not been applied until that point were applied. While the reverse synchronization or failback option was straight forward for Continuous Access XP, for ODG it required that the Oracle parameter file to be set up beforehand to support the failback option.

When considering availability of data for reporting and backup purposes on the remote side, tests showed that Business Copy XP provided an alternative point-in-time data accessibility on an XP-based configuration. While this method required twice the space on the remote array, it also provided an uninterrupted method of replication. Whereas with ODG, although the physical standby database could be mounted for read-only access and enabled for reporting or shutdown for a cold backup, it could not be used at the same time for replication or running the log application service. In other words, with ODG and physical standby databases, replication and access to data on a remote standby database were exclusive to each other.

Best practices and recommendations

Replication across less than 100 km

For short to medium distances (for example, 5 km to 50 km), synchronous replication using Continuous Access XP is a very good option. Continuous Access XP offers high-level data consistency and the advantage of updating changes to remote volumes through real-time replication. This strategy also provides consistent recovery from a crash situation that is determined by the sequence of completed I/O at the target side.

The ODG physical standby mode provides a log-based replication and recovery strategy that requires a robust network connection to maintain synchronized data. All updates to data are accomplished using log files. Since the bandwidth requirements are comparatively lower (anything from a T1 or T3 to GigE), the connection can use an existing IP-based data network and is a less expensive alternative. However, this setup involves more time to recover. The recovery time period is determined by the amount of log data that needs to be applied to bring the database to a known consistent state.

Replication across more than 100 km

For long distances, asynchronous replication using Continuous Access XP provides the best choice. Continuous Access XP uses side-file technology to maintain performance on the primary site. The rate of updates must be managed based on the latency of the long distance link.

Installation

If possible, when planning an Oracle 10g RAC implementation along with a SAN, plan on using the same hardware for all servers, as this will make deployment easier.

If you are running tests on identical configurations of hardware and software, you may want to consider using mirrors of a clean install from one server (to deploy to other servers) or setting up a deployment server. For example, you could use an NFS server to hold the common master configuration files for all the servers, set the configuration parameters once, and then deploy the files to other servers. The Linux Network Information Service (NIS/NIS+) is an option that would also work and is preferable to NFS.

General Oracle configuration

All Oracle configurations for tests described in this paper use OCFS. This is a recommended method because of the OCFS ability to permit shared files and shared LUNs with very minimal OS overhead and without impacting performance. OCFS can also support virtually all types of Oracle files, including data files, log files, control files, and recovery (flash recovery area) files.

In general, use shared disks for data files, log files, archive logs, control files, and server parameter files, and use private or local disks for Oracle software distribution (Oracle Home), dump files, and Cluster Registry Service binaries (CRS Home). These can also be made available on a shared disk when OCFS version 2 is released with support for shared home and context dependant symbolic links (CDSL), such as in Tru clusters on the Tru64 operating system.

When attempting to balance user loads on Oracle, HP recommends the use of shared servers for network connections where large numbers of users are running short transactions (typically OLTP Load). It is best to use dedicated servers for a few users who initiate long-running queries (typically DSS Load). Using a combination of both dedicated and shared server modes and spreading these across instances (nodes) for load balancing is recommended for better utilization of available resources and takes advantage of Oracle load balancing and transparent application failover features.

Configure at least one network service with a transparent basic failover option to enable the switchover of connections from one node instance to others to provide a seamless connectivity during a failure to the oracle client applications.

For better data protection and performance, as a rule, separate the data files, redo logs, and archive logs on separate disk groups or mount points as recommended by the *Oracle Optimal Flexible Architecture* definitions.

As a matter of good practice, always ensure that kernel parameters for shared memory, semaphores, and so forth are compatible with the database startup parameters and vice versa.

While it is possible to configure Oracle listeners to use multiple listening ports, if you want to increase performance and availability, configure dedicated listeners for each database.

Oracle RAC

RAC components depend on the private link established between the node instances for cluster heartbeat, block transfer, and cache fusion. Our recommendation is to have at least a one Gig-E minimum for inter-node connectivity with possible redundancy. Redundancy can be obtained by “bonding” or “teaming” techniques. Due to cache fusion technology used in Oracle 10g RAC, modified data blocks traverse between instances through the interconnects. The larger the number of nodes in a cluster, the greater the need for higher bandwidth in the interconnects.

For the Database Grid control to function properly, at least one network service should be set up for each instance and one common for the database.

OCFS

When you configure the file system, also configure the “label” option to ensure that you can mount a volume based on label instead of the device file name. This is important as it improves the device name mapping and LUN persistence on Linux.

To improve recovery from intra-site path failure, size the ”timeout” value of the ”hang check timer” module so that it does not automatically reboot the server while the multi-path driver is still processing the failover options. Typically, this can happen when a system is heavily loaded with a large number of paths, and when the device driver must traverse all of them. In a typical configuration, the time-out value is set to 180 seconds of tolerance and a 30-second polling interval.

Continuous Access XP and Oracle 10g

In Continuous Access XP Oracle configurations, be sure to create identical databases on both primary and remote sites. This includes creating the same directory structure, file names, and `init.ora` parameters.

This is very important when using Continuous Access XP and Business Copy XP together in the same XP array. Upon initiation of replication, the shared database structures on the remote site are overwritten with the data from the primary site, but the configuration details that are stored locally in the Oracle Home and elsewhere are still retained. This helps to duplicate the Oracle instance while failing over because the remote site thinks it is accessing the same data files, even though they are actually stored on the primary site and are replicated through the Continuous Access XP or Business Copy XP operations.

When using Continuous Access XP in asynchronous mode, to ensure consistency and data integrity, group data files and log files together in the same consistency group. Also, be sure to monitor the side file and link regularly to ensure that the system is not falling behind in the copying of data and thus overfilling the cache memory. Whether these conditions occur depends on bandwidth, replication distance, and data change rate. Note that the system will suspend the pair if the side file overflows, and this may require manual intervention to restart and resynchronize the replication, resulting in a complete initial copy operation.

ODG and Oracle 10g

In ODG and Oracle configurations, use a common Server Parameter File (`spfile`) on the shared storage that contains instance specific values. This reduces human errors and removes the necessity to maintain two separate parameter files. You can create a `pfile` from the `spfile` to edit and add ODG-specific values. Be sure to convert it back to `spfile` and place it back into the shared storage for access by all instances (nodes) in the configuration. See Appendix C for examples of `pfile–init.ora` files for both primary and standby sites.

When you create the primary database in the RAC cluster, use normal methods (dbca recommended).

Be sure to configure the network services with one common service for the database and at least one instance-specific service for each of the node instances to allow for proper functioning of the DB grid control as well as for setting up the TAF option within Oracle Net.

Continuous Access XP and Business Copy XP

When you configure a port as an initiator or a target for use with Continuous Access XP, it also is automatically set to the same type by the XP arrays.

Use of scripts

In general, use scripts to perform redundant operations as this helps to eliminate human error and improves failover and failback times.

Conclusion

In large enterprise environments with demanding requirements for replication of multiple business-applications, Continuous Access XP and Business Copy XP are a suitable technology. As a hardware/array-based solution, they utilize block-level transfers between XP disk arrays to provide excellent data consistency without using network resources. This reduces the time it takes to transfer and recover data, getting your business system back online faster in the face of a disaster or data loss.

In array-based replication, the process of copying is performed on the disk array, so that replication occurs independent of the server. This provides several advantages. For example, it enables administrators to use remote servers for local purposes and to take them offline for non-critical use until they are needed for data recovery. This helps your business achieve a higher ROI by more fully utilizing your existing hardware and services.

XP arrays are highly scaleable and allow connectivity to multiple heterogeneous servers simultaneously; they also allow aggregation of the ISLs to allow for better resilience and bandwidth. The array-based solution further offers the advantage of low-level hardware integrity checks, minimizing the incurred load overheads of such activities.

Continuous Access XP and Business Copy XP together provide a common management process for all replicated applications and data, and can be configured to copy data in both directions. With asynchronous replication, side-file operations allow for efficient transfer of data over long distances while still maintaining the I/O write order. These advantages are further enhanced by thoughtful configuration and a replication SAN structure that supports the distance and throughput demands of applications. Tests showed consistently that crash recovery occurred faster when using Continuous Access XP and Business Copy XP, as compared to other solutions tested.

In general, the Continuous Access XP and Business Copy XP replication solution is an excellent choice for environments that require high data availability and sustained business continuity. Both integrate well with cluster extension software and perform highly reliable replication of data over long-distance SANs without extensive operating system configuration. For ease-of-use, they can be used with HP RAID Manager XP to generate scripts that are created once and used often for dependable copy operations and status checks.

As a SAN hardware-based solution, however, there is a required investment in resources that does not occur for an "Oracle data only" solution; nor is it advisable to try to minimize SAN hardware. If an array is configured with minimal hardware resources, an extra I/O load is added to the back-end of the XP array, which affects replication performance. Table 8 in the Appendix lists all equipment required for the tested configurations. Separate licenses are required for Continuous Access XP and Business Copy XP.

ODG is a good choice where SAN resources do not currently exist and the only applications that require replication use Oracle data. Because ODG only replicates Oracle databases, it requires an integrated solution, such as the one shown in the Phase 3 configuration, to replicate **any** data used by business applications. However, by being Oracle only, this also means that the replication solution does not require investment in a SAN infrastructure; in fact, it does not require a SAN at all.

The ODG product comes pre-packaged with Oracle 10g databases. These are typically configured as one primary and one or more standby databases. Each additional standby database requires more server and network resources compared to Continuous Access XP, which does not require these. Log replication and recovery further require a considerable amount of network bandwidth and disk space for the archive logs, and compared to Continuous Access XP, is not as fast.

To control transaction consistency and data availability levels during data recovery, ODG offers three data protection modes—maximum protection, availability, and performance—including a zero data loss (maximum protection and availability) option. You can select asynchronous operations only by selecting the maximum performance mode. Thus, configuring for transactionally consistent recovery and zero-data loss will always be at the expense of performance, since these preset modes require synchronous operations.

For minimizing application downtime during planned upgrades or unplanned disasters and outages, it is easy to use the ODG failover and failback options to do role reversal between the primary and standby databases. It is not possible however to take the remote site servers offline for other uses, because the Oracle processes must be running on the operating system at all times.

In general, the ODG replication solution is good for an administrator who is an Oracle DBA where the only application to replicate is an Oracle database. For anyone who is not an experienced DBA, getting the system up and running may mean more research and more work because ODG requires more manual configuration compared to the Continuous Access XP solution. In reverse, an Oracle DBA may or may not be familiar with the process of setting up hardware replication and integrating Oracle with Continuous Access XP. Oracle software is not aware of hardware replication.

It is the purpose of this document to bridge this technical gap by providing a comprehensive description and technical explanation of three different replication solutions—(1) using Continuous Access XP only, (2) using ODG only, and (3) using Continuous Access XP with Business Copy XP and ODG together. This paper provides detailed information about configurations and test procedures to give IT administrators and Oracle DBAs the technical guidance and best practice recommendations required to develop a robust and reliable replication system.

Appendix A

List of software and hardware

The following table lists the software and hardware used to run the replication, validation, and performance tests described in this document. For descriptions and diagrams of test configurations, refer to the Test configurations section.

Table 8. Software and hardware configuration details

Configuration components	Quantity	Comments
Linux software		
Red Hat Enterprise Linux Advanced Server 3.0 update 3	1 license per server	4 servers www.redhat.com
QLogic Failover Driver with MPIO enabled	1 kit per server	4 servers Download from www.hp.com
HP software		
HP Storage Works Continuous Access XP (5 TB)	1 per XP array	For 2 arrays
Continuous Access XP Extension	1 per XP array	For 2 arrays
Business Copy XP (5TB)	1 per XP array	For 2 arrays
Command View XP 2.0	1	On the Management Station
Performance Advisor XP 2.0	1	On the Management Station
RAID Manager XP for Replication Control	1 per server	Total of four servers
Oracle software		
ODG for 10g		Part of Oracle 10g Enterprise Edition suite v10.1.0.3
Oracle 10g RAC for Linux		Part of Oracle 10g Enterprise Edition suite v10.1.0.3
Oracle 10g Database Enterprise Edition		Version 10.1.0.3
Oracle Cluster Ready Services CRS for Linux		
Oracle File System (OCFS) for Linux	1 per server per 4 servers	Version 1.0.13-1
Other		
Benchmark Factory 3.3.5	1 Virtual Command Console	1,000 virtual users per TPC-C and TPC-D loads Unlimited virtual users available. See www.quest.com.
Oracle Database Servers		
HP Integrity Servers rx5670 •4-way IA64—Itanium® II 1.5GHz processors •8-GB memory, 2 x 36-GB 15k HDD •2 x A6826A Dual ported Fibre Channel HBA •3 x GbE Network Interface Cards (1 public, 2 private) •4 x IP addresses (1 public IP, 1 public Virtual IP, 2 private IPs)	2 per site (clustered)	4 servers total

Oracle Client Machines		
HP ProLiant DL360-G3 server •Single or Dual Intel® Pentium® Xeon™ 2.80 GHz Processors •2-GB memory, 2 x 36-GB HDD •1 x Dual port Fibre Channel HBA •2 x GbE Network Interface Cards (1 x public, 1x private)	1 per site	2 total Runs this software: •Bench Mark Factory 3.3.5 •Oracle Client Components
Storage Management Station		
HP ProLiant DL360-G1 server •Pentium III 800 MHz processor •1.25-GB memory, 2 x 36-GB HDD •1 x Dual port Fibre Channel HBA •2 x 10/100 Network Interface Cards	1 per site (recommended)	2 total Runs this software: •Command View XP 2.0 •Performance Advisor XP 2.0
Primary Site Storage		
HP StorageWorks XP12000 Disk Array •5.6 TB in usable volume capacity •16-GB cache •4-GB shared memory •2 x 16HS fibre 8ch4mp CHIP boards •120 x 72-GB DKS2C-K72FC disk drives •8 x 72-GB DKS2C-K72FC spares •Primary and Redundant SVPs		Located at Primary site.
Remote Storage		
HP StorageWorks XP1042 Disk Array •5.5 TB in usable volume capacity •32-GB cache •1.5-GB shared memory •2 pairs of 8 port CHIP* cards •(8ch-4mp) •116 x 72-GB DKS2C-K72FC disks •4 x 36-GB DKS2C-K36FC disks •7 x 72-GB DKS2C-K72FC spares •1 x 36-GB DKS2C-K36FC spares		Located at Remote Stand-by site.
Infrastructure Network		
HP ProCurve Networking 2824 Manageable Switches •3 VLANs (1 Public, 2 Private) •10-km long-distance network SFPs used for ISL between sites •1 switch per site •Netsim4 10-km Fibre Channel Spool—for distance		

SAN		
HP StorageWorks 2/16 (16x2GB ports) Fibre Channel switches •B-Series switches (Brocade) •35-km Extended Long Distance SFPs used for ISL between sites •4 switches per site (redundant fabrics for host-to-storage and storage-to-storage connectivity) •Netsim8 10.5/20.5/50km Fibre Channel spools for distance		

Table 9: List of equipment

Description	Quantity	Part number
Third-party software		
Red Hat Enterprise Linux Advanced Server 3 Update 3	1	See www.redhat.com.
QLogic Failover Driver 7.01.01 (or higher) for RHEL3/IA64	1	Download from www.hp.com.
Oracle software Oracle Cluster Manager (OCM) for Linux Oracle Cluster File System (OCFS) for Linux Oracle 10g Database Enterprise Edition Oracle 10g RAC for Linux Oracle Data Guard for 10g	1	See www.oracle.com.
Benchmark Factory Suite 3.3.5	1	
Empirix GbE network simulator or Net8 Fibre Channel Spools (for simulation only)	1	Replace with long-distance connection.
Hardware solutions		
HP rx5670 Itanium2 1.5-GHz CPU Solution Includes one Itanium™2 1.5-GHz CPU, system board, core I/O, and two power supplies. No memory or disk. Rack ready.	4 (2 per site)	A6838B
Itanium2 1.5-GHz CPU for HP rx5670	3	A9810A
4-GB DDR memory quad for HP rx5670 (4x1-GB RAM)	1	A6834A
4-GB DDR memory quad for HP rx5670 (4x1-GB RAM)	1	A6834A
Memory carrier board for HP rx5670	1	A6747A
36-GB 15K Hot Plug Ultra320 disk, rx5670 Half height 15,000 RPM drive	2	A7049A
DVD ROM drive for HP Svr rp54X0, rx5670	1	A5557B
Win/Linux 1000Base-T Gigabit Eth Adpt PCI 1-port 10/100/1000Base-T Gigabit Ethernet Copper LAN Adapter, with short form factor, for IA-64 systems running Microsoft® Windows® and Linux.	2	A7061A
Factory Rack Kit, slides, install	1	A5581A

Hot Swap power supply for HP rx5670 Additional Hot Swap DC power supply for N + 1 redundancy	1	A7093A
HP Linux Enablement Kit LTU, Integrity Installation, configuration and recovery tools for Linux on HP servers rx2600 and rx4670. Media kit.	1	T2387AA
PCI-X Dual Channel 2GB Fibre Channel HBA PCI-X dual channel, 2GB Fibre Channel Adapter, PCI-X (64-bit, 133 MHz) with 2 LC connectors, auto negotiates 2-GB or 1-GB transfer mode.	2	A6826A
HP StorageWorks XP12000 Disk Array SSP Solution	1	AE001A
XP12000 Disk Control Frame (DKC)	1	AE002A
3-Phase 30A/60Hz for XP12000 DKC	1	AE002A 001
XP12000 73-GB 15k rpm Array Group-4 disks	32	AE050A
XP12000 73-GB 15k rpm Spare Disk*	4	AE050AS
XP12000 Disk Array Frame (DKU)	1	AE045A
3-Phase 30 A/60 Hz for XP12000 DKU	1	AE045A 001
XP12000 Cable set for DKU R1-Basic	1	AE040A
XP12000 Standard Performance ACP Pair	1	AE034A
XP12000 16-Port 1-2-GB/sec FC SW CHIP pr	1	AE006A
Fibre Channel Host Cable	16	A5750A
16M LC/LC FC cable, 50/125 micron, multi	16	A5750A 004
50M LC/LC FC cable,50/125 micron, multi	0	A5750A 005
200M LC/LC FC cable, 50/125 micron, multi	0	A5750A 006
XP12000 4-GB Cache Memory Module	3	AE025A
XP12000 1-GB Shared Memory Module	4	AE030A
XP12000 DKC Power Supply	1	AE024A
XP12000 DKC-DKU Battery	1	AE028A
RAID Manager XP LTU	1	T1610A
HP Command View XP LTU	1	B9357AJ
HP LUN Config/Sec Mgr 1-TB LTU (7-15TB)	4	T1714AC
HP LUN Config/Sec Mgr 1-TB LTU (2-6TB)	5	T1714AB
HP LUN Config/Sec Mgr 1-TB LTU (0-1TB)	1	T1714AA
HP StorageWorks XP1024 Disk Array Solution		A7905A
XP1024 Disk Control Frame (DKC)	1	A7906A
3 Phase 50 or 60 Hz for XP1024	0	A7906A 001
Single Phase 50 Hz for XP1024	0	A7906A 004

Single Phase 60 Hz for XP1024	0	A7906A 003
3 Phase 50 Hz/30 A for XP1024	0	A7906A 006
3 Phase 60 Hz/30 A for XP1024	One (1)	A7906A 005
XP1024 73-GB 15k rpm FC (4 disks)	Thirty-two (32)	A7931A
XP1024 73-GB 15k rpm FC spare disk drive	Four (4)	A7931S
XP1024/128 8-port 1–2-GB/sec FC Enh Chip	One (1)	A7912B
Fibre Channel Host Cable	Eight (8)	A5750A
16M LC/LC FC cable, 50/125 micron, multi	8	A5750A 004
50M LC/LC FC cable,50/125 micron, multi	0	A5750A 005
200M LC/LC FC cable, 50/125 micron, multi	0	A5750A 006
2M LC/SC FC cable, 50/125 micron, multi	0	A5750A 007
16M LC/SC FC cable, 50/125 micron, multi	0	A5750A 008
2M LC Male adapter kit	0	A5750A 010
XP1024 Disk Array Frame	One (1)	A7925A
3 Phase 60 Hz/30 A for XP1024 DKU	One (1)	A7925A 005
Single Phase 50 Hz for XP1024 DKU	0	A7925A 004
Single Phase 60 Hz for XP1024 DKU	0	A7925A 003
3 Phase 50 Hz/30 A for XP1024 DKU	0	A7925A 006
XP1024 Cache Platform Board	One (1)	A7919A
XP1024/128 2-GB Cache Memory Module	Four (4)	A7918A
XP1024/128 512-MB Shared Memory Module	One (1)	A7921A
XP 1024/128 1-GB Shared Memory Module	One (1)	A7935A
RAID Manager XP LTU	One (1)	T1610A
HP Command View XP LTU	One (1)	B9357AJ
Continuous Access XP 1-TB LTU (2–6 TB)	Five (5)	T1611AB
Continuous Access XP 1-TB LTU (up to 1 TB)	One (1)	T1611AA
LUN Conf/Sec Mgr XP 1-TB LTU (7–15 TB)	Two (2)	T1614AC
LUN Conf/Sec Mgr XP 1-TB LTU (2-6 TB)	Five (5)	T1614AB
LUN Conf/Sec Mgr XP 1-TB LTU (up to 1 TB)	One (1)	T1614AA

* Designated from the previous list if only DKU-R0.

See Table 8 for a complete list of third-party software used in the test configurations.

Appendix B

Procedures

This section provides detailed steps and command line instructions for procedures described in the Test procedures section. These instructions are provided here for reference only. They are samples, and, as such, are without technical support, and therefore, should be used accordingly.

Access to shared vDisks in HP StorageWorks XP Array

This procedure gives all cluster servers, including those in the remote SAN site, access to the shared vDisks in their respective XP arrays.

Table 10. Procedure for shared vDisk access

Shared vDisk access procedure			
Step	**Task**		**Command Line**
1	From the master `/etc/fstab` file, perform the following steps:		

	1.1	NFS Server to store common or master configuration files	`tiger1:/users/software  /software`
			`nfs   _netdev  0 0`

	1.2	Local/private file system for db dump, Oracle software binaries and inventory	`LABEL=ORABASE`
			`/u01/orabase  ext3  defaults 0 0`

	1.3	Local/private file system for Oracle Cluster Ready Services (OCRS)	`LABEL=CRSHOME`
			`/u01/crshome  ext3  defaults 0 0`

	1.4	Shared OCFS volume for Oracle cluster registry and voting disk	`LABEL=OCR            /u01/ocr`
			`ocfs  _netdev  0 0`

	1.5	Shared OCFS volume for flash-back area, archive logs, and backup files	`LABEL=ORF            /u01/orf`
			`ocfs  _netdev  0 0`

	1.6	Shared OCFS volume for storing the Continuous Access XP data, control and server parameter files	`LABEL=CAXPDATA`
			`/u01/caxpdata ocfs  _netdev  0 0`

	1.7	Shared OCFS volume for storing the Continuous Access XP database online log files	`LABEL=CAXPLOG`
			`/u01/caxplog  ocfs  _netdev  0 0`

	1.8	Shared OCFS volume for storing the ODG data, control and server parameter files	`LABEL=ODGDATA`
			`/u01/odgdata  ocfs  _netdev  0 0`

	1.9	Shared OCFS volume for storing the ODG database online log files	`LABEL=ODGLOG`
			`/u01/odglog   ocfs  _netdev  0 0`

	1.10	Shared OCFS volume for storing the baseline Benchmark Factory tablespaces and other miscellaneous database files	`LABEL=BACKUP`
			`/u01/backup   ocfs  _netdev  0 0`

Continuous Access XP: Failure on primary site

This section contains the failure test procedure used in the simulated network connection link plus site failure scenario for the primary database site.

Table 11. Test procedure for primary site failure

FAILURE PROCEDURE ON PRIMARY SITE—CONTINUOUS ACCESS XP			
Step	**Task**		**Command Line**
1	SHOW ENVIRONMENT		
	1.1	Show status of Oracle database. Start database, if required.	`srvctl status db -d CAXP`
	1.2	Show status of the replication pair.	`pairdisplay -g caxp -fcx`
	1.3	List mounted file systems with size.	`df -kh`
2	PERFORM A DDL/DML		
	2.1	Log in as user `caxp`.	`su - caxp`
	2.2	Log in to `sqlplus` as user `scott`.	`sqlplus scott/tiger@caxp`
	2.3	Create a test table.	`Create table test01 as select * from emp;`
	2.4	Show the list of tables for the schema.	`select * from tab;`
	2.5	Display the contents of table "`test01`."	`Select * from test01;`
	2.6	Commit the changes to the database.	`Commit;`
3	VERIFY ISL STATUS AND SIMULATE ISL FAILURE		
	3.1	Telnet to the two switches bass111 and bass112.	
	3.2	List the current status of switch.	`switchshow`
	3.3	Disable the ISL connections (simulates inter-site link failure).	`portdisable 15`
4	VERIFY THE REPLICATION PAIR STATUS		
	4.1	Show status of replication pair. Status should be PSUE.	`pairdisplay -g caxp -fcx`
5	SHUT DOWN ORACLE DATABASE, UNMOUNT LUNS, and SHUT DOWN SERVERS		
	5.1	Shut down the Continuous Access XP database instances (only executed once from one server).	`srvctl stop db -d CAXP`
	5.2	Unmount Data volume (both servers).	`umount /u01/caxpdata`
	5.3	Unmount Log volume (both servers).	`umount /u01/caxplog`
	5.4	Shut down the server (both servers).	`shutdown`

Continuous Access XP: Failover on remote site

This section contains the failure test procedure used in the simulated network connection link plus site failure scenario for the remote standby site.

Table 12. Test procedure for remote site failover

FAILOVER PROCEDURE ON REMOTE SITE—CONTINUOUS ACCESS XP			
Step	**Task**		**Command Line**
1	CHECK REPLICATION PAIR STATUS		
	1.1	Show status of replication pair. Display shows the status PSUE for P-Vol.	`pairdisplay -g caxp -fcx`
2	NOW (FORCE) MOVE CONTROL TO THE S-VOL		
	2.1	Take over the control at remote site.	`horctakeover -g caxp`
3	CHECK REPLICATION PAIR STATUS		
	3.1	Show status of replication pair. Display shows P-Vol → PSUE and S-Vol → SSWS.	`pairdisplay -g caxp -fcx`
4	MOUNT LUNS (S-VOL) ON SITE-B		
	4.1	Mount replicated Data volume.	`mount -t ocfs /dev/sdi1 /u01/caxpdata`
	4.2	Mount replicated Log volume.	`mount -t ocfs /dev/sdj1 /u01/caxplog`
5	START THE DATABASE		
	5.1	Log in as the user `caxp`.	`su - caxp`
	5.2	Perform crash recovery and bring up the database.	`svrctl start db -d CAXP`
6	PROOF OF REPLICATED DATA and DDL FOR FAILBACK		
	6.1	Connect to database as user `scott`.	`sqlplus scott/tiger@caxp`
	6.2	List the existing tables including "`test01`" created on the Primary site.	`select * from tab`
	6.3	List all rows from table "`test01`."	`select * from test01`
	6.4	Create table "`test02`."	`create table test02 as select * from test01;`
	6.5	List tables including "`test02`."	`select * from tab`

Continuous Access XP: Resynchronization and failback of primary site

This section contains the resynchronization test procedure used in the simulated failback scenario for the primary database site.

Table 13. Test procedure for resynchronization and failback on primary site

RESYNCHRONIZATION AND FAILBACK ON PRIMARY SITE—CONTINUOUS ACCESS XP			
Step	**Task**		**Command Line**
1	SIMULATE ISL LINKS RESTORATION		
	1.1	Telnet to the two switches bass111 and bass112.	
	1.2	List the current status of switch.	switchshow
	1.3	Enable ISL connections (simulates inter-site link restoration).	portenable 15
2	REVERSE SYNC DATA – S-VOL to P-VOL		
	2.1	Use "swaps" to change resync direction.	pairresync -g caxp -swaps
	2.2	Display the status of resync operation.	pairdisplay -g caxp -fcx
3	SHUT DOWN ORACLE DATABASE, UNMOUNT LUNS, and SHUT DOWN SERVERS		
	3.1	Shut down the Continuous Access XP database instances (only executed once from one server).	srvctl stop db -d CAXP
	3.2	Unmount Data volume (both servers).	umount /u01/caxpdata
	3.3	Unmount Log volume (both servers).	umount /u01/caxplog
	3.4	Shut down the server (both servers).	shutdown
4	VERIFY REPLICATION STATUS		
	4.1	Show status of replication. Display should show "pair" status; else wait.	pairdisplay -g caxp -fcx
5	MOVE THE P-VOL back to PRIMARY SITE		
		Run on the primary site.	horctakeover -g caxp
		Display status to verify successful transfer.	pairdisplay -g CAXP -fcx
6	MOUNT LUNS (P-VOL)ON SITE-A		
	6.1	Resynchronize P-VOL—data. (This must be done on both the servers—cftsrv107, cftsrv109.)	mount -t ocfs /dev/sdi1 /u01/caxpdata
	6.2	Resynchronize P-VOL—log. (This must be done on both the servers—cftsrv107, cftsrv109.)	mount -t ocfs /dev/sdj1 /u01/caxplog
7	START THE DATABASE		
		Log in as the user caxp.	su – caxp
		Perform crash recovery and bring up the database.	svrctl start db -d CAXP
8	PROOF OF RESYNCHRONIZED DATA		
		Connect to database as user scott.	sqlplus scott/tiger@caxp
		List the existing tables including "test01" and "test02."	select * from tab
		List all rows from table "test01."	select * from test02

ODG: Creation of Oracle physical standby database

This section describes the procedure for creating an Oracle physical standby database.

Table 14. Test Procedure for Oracle Physical Standby Database

CREATION OF PHYSICAL STANDBY DATABASE —ODG			
Step	**Task**		**Command Line**
1	ENABLE ARCHIVING		
	1.1	Shut down database.	shutdown immediate;
	1.2	Start database in ”mount” mode.	startup mount;
	1.3	Enable archive logging.	alter database archivelog;
	1.4	Change database state to “open.”	alter database open
2	CREATE A PASSWORD FILE		
	2.1	Change directory to ”dbs.”	$cd $ORACLE_HOME/dbs
	2.2	Create the password file.	$orapwd file=orapwODG1 password=oracle
	2.3	Shut down database.	shutdown immediate;
	2.4	Start database in “nomount” mode.	startup nomount;
	2.5	Set the remote password file option to exclusive in the “spfile.”	alter system set remote_login_passwordfile=exclusive scope=spfile
	2.6	Change database state to ”open.”	alter database open;
3	ENABLE FORCE LOGGING		
		Enable “force logging” of the database.	alter database force logging
4	CREATE STANDBY REDO LOGS FOR EACH THREAD (count = online redos +1)		
	4.1	Create standby “redo01” log file and set file size.	alter database add standby logfile thread 1 (‘/u01/odglog/ODG1/stby_redo01.log’) size 30M;
	4.2	Create standby “redo02” log file and set file size.	alter database add standby logfile thread 1 (‘/u01/odglog/ODG1/stby_redo02.log’) size 30M;
	4.3	Create standby ‘redo03’ log file and set file size.	alter database add standby logfile thread 1 (‘/u01/odglog/ODG1/stby_redo03.log’) size 30M;
	4.4	Create standby ‘redo04’ log file and set file size.	alter database add standby logfile thread 2 (‘/u01/odglog/ODG1/stby_redo04.log’) size 30M;
	4.5	Create standby ‘redo05’ log file and set file size.	alter database add standby logfile thread 2 (‘/u01/odglog/ODG1/stby_redo05.log’) size 30M;
	4.6	Create standby ‘redo06’ log file and set file size.	alter database add standby logfile thread 2 (‘/u01/odglog/ODG1/stby_redo06.log’) size 30M;
5	CONFIGURE THE PRIMARY INITIALIZATION PARAMETERS		
	5.1	See Appendix C.	

6	CREATE A BACKUP OF THE PRIMARY DATABASE and TRANSFER TO THE STANDBY HOST		
		Shut down.	`shutdown immediate;`
		Create a backup copy of the Data volume on the primary site.	`scp -r /u01/odgdata/ODG1/* root@cftsrv111:/u01/odgdata/ODG5/.`
		Create a backup copy of the Log volume on the primary site.	`scp -r /u01/odglog/ODG1/* root@cftsrv111:/u01/odglog/ODG5/.`
		Transfer to standby host.	`scp -r /u01/orf/ODG1/* root@cftsrv111:/u01/orf/ODG5/.`
7	CREATE A STANDBY CONTROL FILE		
		Create standby control file with the database in "`open`" or "`mount`" state.	`alter database create standby controlfile as '/u01/odgdata/ODG1/stby_ctrl01.ctl';`
8	CREATE THE INITIALIZATION PARAMETER FILE FOR THE STANDBY DATABASE		
	8.1	Create parameter file from the "`spfile`" using the sql cmd.	`Create pfile='/u01/odgdata/ODG1/initODG11.ora' from spfile;`
	8.2	Remote/Standby initialization file.	See Appendix C.
9	PREPARE THE STANDBY HOST		
		For Linux, set up environment variables.	Follow Oracle installation requirements.
10	CREATE THE STANDBY PASSWORD FILE		
	10.1	Change directory to "`dbs`" under Oracle Home.	`$cd $ORACLE_HOME/dbs`
	10.2	Create the password file.	`$orapwd file=orapwODG5 password=oracle`
11	CONFIGURE ORACLE NET COMPONENTS		
	11.1	For ease of use, create a common `tnsnames.ora` file that includes all instance and databases.	Use the'netca' utility or manual method to create the 'tnsnames.ora' file. For sample file, see Appendix E.
	11.2	Create at least one listener each for each host (both on primary and remote standby sites).	Use the '`netca`' utility or manual method to create the 'listener.ora' file. For sample file, see Appendix E.
12	CREATE SPFILE ON STANDBY SITE.		
	12.1	Create a shared "`spfile`" for use on the remote standby site	`create spfile from pfile;`
13	START THE REMOTE STANDBY DATABASE		
	13.1	Start database in "`mount`" mode.	`startup mount;`
14	BEGIN SHIPPING LOGS TO STANDBY HOST		
		Ship logs to remote standby host.	
15	VERIFY LOG TRANSPORT SERVICES and DATA TRANSMISSIONS		
	15.1	Connect as user `scott.`	`connect scott/tiger@ODG1`
	15.2	Create a table "`test01`" under `scott.`	`create table test01 as select * from DEPT;`
	15.3		`conn sys/oracle@ODG1`
	15.4	Perform a log switch on the primary site.	`alter system switch logfile;`

	15.5		select * from v$archive_dest where dest_id=2;
	15.6	To verify a successful transmission, display the log file at the standby archive log destination (/u01/orf/ODG5/).	
16	APPLY THE ARCHIVE LOG FILES TO STANDBY DATABASE		
		Apply archive log files transmitted to standby database.	alter database recover managed standby database disconnect;
17	STOP THE MANAGED RECOVERY PROCESS (MRP)		
	17.1	Cancel the MRP on the remote standby database.	alter database recover managed standby database cancel;
18	START THE MANAGED RECOVERY PROCESS (MRP)		
	18.1	Start MRP with a delay (in minutes)	alter database recover managed standby database delay 60;

ODG: Planned switchover with physical standby

This section describes the steps for performing a planned switchover of Oracle databases: (1) primary to new standby, and (2) standby to new primary.

Table 15. Test procedure for planned switchover of primary and physical standby databases

SWITCHOVER WITH PHYSICAL STANDBY (PLANNED)—ODG			
Step	**Task**		**Command Line**
1	VERIFY SWITCHOVER_STATUS ON THE PRIMARY DATABASE		
	1.1	Verify switchover status on primary site.	select switchover_status from v$database;
	1.2	If the returned result is: • TO STANDBY, proceed to step 2.1. • TO LOGICAL STANDBY, proceed to step 2.2. • SESSIONS ACTIVE, proceed to step 2.3	
2	CONVERT PRIMARY DATABASE INTO PHYSICAL STANDBY		
	2.1	If status was TO STANDBY, use...	alter database commit to switchover to physical standby;
	2.2	If status was TO LOGICAL STANDBY, use...	
	2.3	If status was SESSIONS ACTIVE, use...	alter database commit to switchover to physical standby with session shutdown;
3	SHUT DOWN and RESTART THE OLD PRIMARY AS THE NEW STANDBY		
	3.1	Shut down primary database.	shutdown immediate;
	3.2	Restart the database in "mount" mode.	startup mount;

4	VERIFY THE SWITCHOVER STATUS		
	4.1	Query the switchover status on standby database.	`select switchover_status from v$database;`
	4.2	If returned result was: • TO_PRIMARY, proceed to step 5.1. • SESSIONS ACTIVE, proceed to step 5.2. • NOT ALLOWED, proceed to step 5.3.	
5	CONVERT THE OLD STANDBY TO PRIMARY DATABASE		
	5.1	If result was TO_PRIMARY, use…	`alter database commit to switchover to primary;`
	5.2	If result was SESSIONS ACTIVE, use…	`alter database commit to switchover to primary with session shutdown;`
	5.3	If result was NOT ALLOWED, validate the settings. Switchover process cannot proceed further.	
6	SHUT DOWN and RESTART THE NEWLY CONVERTED PRIMARY		
	6.1	Shut down the new primary database.	`shutdown immediate;`
	6.2	Restart it.	`startup;`
7	REPEAT STEPS TO FAILBACK		
	7.1	Repeat steps 1–6 on the new Primary and Standby to fail back to original setup.	

ODG: Unplanned failover with physical standby

Here are a few things to note before initiating a failover:

- Try to get all unapplied log files from the primary to standby site. (Server, storage, and site failures cannot guarantee this, but it is possible with an ISL failure.)
- Ensure that temporary tablespaces on the standby are populated with temp files.
- In a RAC configuration, ensure that all databases except the primary standby instance are down.
- Remove delay settings for recovery of redo from the primary.
- Change the protection mode of the standby database to maximum performance.

Table 16: Test procedure for unplanned failover

FAILOVER WITH PHYSICAL STANDBY (UNPLANNED)—ODG			
Step	**Task**		**Command Line**
1	RESOLVE GAPS ON THE STANDBY		
	1.1	Resolve any gaps that may exist on the standby.	select * from V$ARCHIVE_GAP
	1.2	Try copying archive logs from primary to standby with those that have sequence numbers higher than the last one to arrive at the standby.	<use your favorite OS level file copy utility>
2	COPY ARCHIVE LOGS AND REGISTER		
	2.1	Copy any archive logs that need to be registered.	alter database register physical logfile '*<logfile location>*';
3	PERFORM TERMINAL RECOVERY ON THE STANDBY BY STARTING MANAGED RECOVERY WITH "FINISH" KEYWORD		
	3.1	With active standby redo logs, use…	alter database recover managed standby database finish;
	3.2	With non-existent or non-active standby redo logs, use…	alter database recover managed standby database finish skip standby logfile;
4	CONVERT THE STANDBY DATABASE TO PRIMARY		
	4.1	After recovery completes, switch over to primary.	alter database commit to switchover to primary;
5	SHUT DOWN AND RESTART THE NEW PRIMARY DATABASE		
	5.1	Shut down primary database.	shutdown immediate;
	5,2	Restart database.	startup;

Appendix C

This section lists the contents of the database server parameter files for the primary and secondary Oracle Data Guard sites and the operating system control files. Note that these are applicable only for the referenced configuration and differ for other implementations. These files are not endorsed and are provided here for reference only. They are samples, and, as such, are without technical support, and therefore, should be used accordingly. This section contains examples of the following files:

`DB pfile-init.ora` (Oracle primary site)
`DB pfile-init.ora` (Oracle standby site)
`OS sysctl.conf` (Linux configuration file)
`OS hosts` (hosts file -- for resolving host names to IP addresses)

Database server parameter files

pfile-init.ora (ODG): primary site

```
ODG12.__db_cache_size=754974720
ODG11.__db_cache_size=704643072
ODG12.__java_pool_size=16777216
ODG11.__java_pool_size=16777216
ODG12.__large_pool_size=33554432
ODG11.__large_pool_size=33554432
ODG12.__shared_pool_size=603979776
ODG11.__shared_pool_size=654311424
*.background_dump_dest='/u01/orabase/admin/ODG1/bdump'
*.cluster_database_instances=2
*.cluster_database=true
*.compatible='10.1.0.2.0'
*.control_files='/u01/odgdata/ODG1/control01.ctl','/u01/odgdata/ODG1/control02.ct
l','/u01/odgdata/ODG1/control03.ctl'
*.core_dump_dest='/u01/orabase/admin/ODG1/cdump'
*.db_block_size=8192
*.db_domain=''
*.db_file_multiblock_read_count=16
*.db_file_name_convert='/u01/odgdata/ODG5/','/u01/odgdata/ODG1/'
*.db_name='ODG1'
*.db_unique_name='ODG1'
*.dispatchers='(protocol=TCP)(listener=listeners_ODG1)'
*.fal_client='ODG51'
*.fal_server='ODG51','ODG52'
ODG12.instance_number=2
ODG11.instance_number=1
*.job_queue_processes=10
*.log_archive_config='DG_CONFIG=(ODG1,ODG5)'
*.log_archive_dest_1='LOCATION=/u01/orf/ODG1/ VALID_FOR=(ALL_LOGFILES,ALL_ROLES)
DB_UNIQUE_NAME=ODG1'
*.log_archive_dest_2='SERVICE=ODG51 VALID_FOR=(ONLINE_LOGFILES,PRIMARY_ROLE)
DB_UNIQUE_NAME=ODG5'
*.log_archive_dest_state_1='enable'
*.log_archive_dest_state_2='ENABLE'
*.log_archive_format='%t_%s_%r.dbf'
*.log_file_name_convert='/u01/odglog/ODG5/','/u01/odglog/ODG1/'
*.open_cursors=300
*.pga_aggregate_target=473956352
*.processes=5000
*.remote_listener='LISTENERS_ODG1'
*.remote_login_passwordfile='exclusive'
ODG11.service_names='ODG1','ODG11'
ODG12.service_names='ODG1','ODG12'
*.sessions=5505
*.sga_target=1423966208
*.standby_file_management='auto'
ODG12.thread=2
ODG11.thread=1
```

```
*.undo_management='AUTO'
ODG11.undo_tablespace='UNDOTBS1'
ODG12.undo_tablespace='UNDOTBS2'
*.user_dump_dest='/u01/orabase/admin/ODG1/udump'
```

pfile-init.ora (ODG): standby site

```
ODG52.__db_cache_size=872415232
ODG51.__db_cache_size=872415232
ODG52.__java_pool_size=16777216
ODG51.__java_pool_size=16777216
ODG52.__large_pool_size=16777216
ODG51.__large_pool_size=16777216
ODG52.__shared_pool_size=503316480
ODG51.__shared_pool_size=503316480
*.background_dump_dest='/u01/orabase/admin/ODG5/bdump'
*.cluster_database_instances=2
*.cluster_database=true
*.compatible='10.1.0.2.0'
*.control_files='/u01/odgdata/ODG5/stby_ctrl01.ctl','/u01/odgdata/ODG5/stby_ctrl0
2.ctl','/u01/odgdata/ODG5/stby_ctrl03.ctl'
*.core_dump_dest='/u01/orabase/admin/ODG5/cdump'
*.db_block_size=8192
*.db_domain=''
*.db_file_multiblock_read_count=16
*.db_file_name_convert='/u01/odgdata/ODG1/','/u01/odgdata/ODG5/'
*.db_name='ODG1'
*.db_unique_name='ODG5'
*.dispatchers='(protocol=TCP)(listener=listeners_ODG1)'
*.fal_client='ODG11'
*.fal_server='ODG11','ODG12'
ODG52.instance_number=2
ODG51.instance_number=1
*.job_queue_processes=10
*.log_archive_config='DG_CONFIG=(ODG1,ODG5)'
*.log_archive_dest_1='LOCATION=/u01/orf/ODG5/ VALID_FOR=(ALL_LOGFILES,ALL_ROLES)
DB_UNIQUE_NAME=ODG5'
*.log_archive_dest_2='SERVICE=ODG11 VALID_FOR=(ONLINE_LOGFILES,PRIMARY_ROLE)
DB_UNIQUE_NAME=ODG1'
*.log_archive_dest_state_1='enable'
*.log_archive_dest_state_2='enable'
*.log_archive_format='%t_%s_%r.dbf'
*.log_file_name_convert='/u01/odglog/ODG1/','/u01/odglog/ODG5/'
*.open_cursors=300
*.pga_aggregate_target=473956352
*.processes=5000
*.remote_listener='LISTENERS_ODG5'
*.remote_login_passwordfile='exclusive'
ODG51.service_names='ODG5','ODG51'
ODG52.service_names='ODG5','ODG52'
*.sessions=5505
*.sga_target=1423966208
*.standby_file_management='auto'
ODG52.thread=2
ODG51.thread=1
*.undo_management='AUTO'
ODG51.undo_tablespace='UNDOTBS1'
ODG52.undo_tablespace='UNDOTBS2'
*.user_dump_dest='/u01/orabase/admin/ODG5/udump'
```

Operating system control files

/etc/sysctl.conf

```
# Kernel sysctl configuration file for Red Hat Linux
# For binary values, 0 is disabled, 1 is enabled.  See sysctl(8) and
# sysctl.conf(5) for more details.
# Controls IP packet forwarding
```

```
net.ipv4.ip_forward = 0
# Controls source route verification
net.ipv4.conf.default.rp_filter = 1
# Controls the System Request debugging functionality of the kernel
kernel.sysrq = 0
# Controls whether core dumps will append the PID to the core filename.
# Useful for debugging multi-threaded applications.
kernel.core_uses_pid = 1
# Increase the default and maximum send and receive buffer size to 256kb
# This is to support Oracle10g RAC implementation
net.core.rmem_default = 262144
net.core.wmem_default = 262144
net.core.rmem_max = 262144
net.core.wmem_max = 262144
# Kernel Paramaters changed for Oracle10g RAC setup
# ---------------------------------------------
# kernel.sem = 250 32000 32 128
# kernel.shmmni = 4096
# kernel.shmall = 524288
# kernel.shmmax = 33554432
# kernel.shmmax = 2099748864
# s.file-max = 65536
kernel.sem = 250 32000 100 128
kernel.shmmni = 4096
kernel.shmall = 2097152
kernel.shmmax = 2147483648
fs.file-max = 416646
# net.ipv4.ip_local_port_range = 32768    61000
net.ipv4.ip_local_port_range = 1024    65000
```

/etc/hosts

```
# Do not remove the following line, or various programs
# that require network functionality will fail.
# Host 1
10.10.14.107   cftsrv107.cftlabs.hpq.net    cftsrv107  #Public IP
10.10.14.96    cftsrv096.cftlabs.hpq.net    cftsrv096  #Public Virtual IP
192.168.0.107   cftsrv107-red               s1-pip1    #Private IP -1
192.168.1.107   cftsrv107-blue              s1-pip2    #Private IP -2
# Host 2
10.10.14.109   cftsrv109.cftlabs.hpq.net    cftsrv109  #Public IP
10.10.14.97    cftsrv097.cftlabs.hpq.net    cftsrv097  #Public Virtual IP
192.168.0.109   cftsrv109-red               s2-pip1    #Private IP -1
192.168.1.109   cftsrv109-blue              s2-pip2    #Private IP -2
# Host 3
10.10.14.111   cftsrv111.cftlabs.hpq.net      cftsrv111       #Public IP
10.10.14.98    cftsrv098.cftlabs.hpq.net      cftsrv098       #Public Virtual IP
192.168.0.111   cftsrv111-red                 s3-pip1         #Private IP -1
192.168.1.111   cftsrv111-blue                s3-pip2         #Private IP -2
# Host 4
10.10.14.113   cftsrv113.cftlabs.hpq.net      cftsrv113       #Public IP
10.10.14.99    cftsrv099.cftlabs.hpq.net      cftsrv099       #Public Virtual IP
192.168.0.113   cftsrv113-red                 s4-pip1         #Private IP -1
192.168.1.113   cftsrv113-blue                s4-pip2         #Private IP -2
# NFS servers
10.10.1.11     tiger1.cftlabs.hpq.net         tiger1          # NFS Server
# Oracle Client servers
# Used for Load Generation
# Runs W2k3 and Benchmark Factory
10.10.14.105   cftsrv105.cftlabs.hpq.net      cftsrv105          # Client 1
# <dhcp>        cftsrv105ilo.americas.cpqcorp.net cftsrv105ilo   # Client 1 - iLO
10.10.14.106   cftsrv106.cftlabs.hpq.net      cftsrv106          # Client 2
# <dhcp>        cftsrv106ilo.americas.cpqcorp.net cftsrv106ilo   # Client 2 - iLO
```

Appendix D

RAID Manager XP configuration

Environment variables

```
      On cftsrv107      HORCMINST=1
      On cftsrv109      HORCMINST=2
      On cftsrv111      HORCMINST=3
      On cftsrv113      HORCMINST=4
```

SCSI targets and LUN numbers

All SCSI target IDs and LUN numbers **must** be obtained from your configuration. Do **not** use the examples. Use the raidscan command to obtain your SCSI targets and LUN numbers.

Following is an example output of the raidscan command. The information with which you are concerned is: the SCSI Target, the LUN #, and the LDEV #.

Be sure that the RAID Manager XP setup and configuration are executed with root privileges (or with root equivalence).

```
 [root@cftsrv107 root]# raidscan -p CL1-A -CLI -fx
PORT# /ALPA/C TID#  LU# Seq#   Num  LDEV# P/S  Status Fence  P-Seq# P-LDEV#
CL1-A    ef  0   0     0 10197     1   200 SMPL     -      -       -       -
CL1-A    ef  0   0     1 10197     1   220 SMPL     -      -       -       -
CL1-A    ef  0   0     2 10197     1   222 SMPL     -      -       -       -
CL1-A    ef  0   0     3 10197    10   100 SMPL     -      -       -       -
CL1-A    ef  0   0     4 10197    10     0 P-VOL PAIR  NEVER   10818       0
CL1-A    ef  0   0     5 10197    10    20 P-VOL PAIR  NEVER   10818      20
CL1-A    ef  0   0     6 10197    10    2a SMPL     -      -       -       -
CL1-A    ef  0   0     7 10197     1   201 SMPL     -      -       -       -
CL1-A    ef  0   0     8 10197     1   221 SMPL     -      -       -       -
CL1-A    ef  0   0     9 10197     1   222 SMPL     -      -       -       -
CL1-A    ef  0   0    10 10197    10   100 SMPL     -      -       -       -
CL1-A    ef  0   0    11 10197    10     0 P-VOL PAIR  NEVER   10818       0
CL1-A    ef  0   0    12 10197    10    20 P-VOL PAIR  NEVER   10818      20
CL1-A    ef  0   0    13 10197    10     a SMPL     -      -       -       -
CL1-A    ef  0   0    14 10197    10     a SMPL     -      -       -       -
CL1-A    ef  0   0    15 10197    10    2a SMPL     -      -       -       -
```

RAID Manager configuration files for Continuous Access XP

In this configuration, the two primary servers that communicate to perform the Continuous Access XP operations are cftsrv107 and cftsrv111. This equates to HORCM configuration files `horcm1.conf` and `horcm3.conf`. One server resides on the primary site (SITE A), and the other server resides on the remote site (SITE B). The two remaining servers act as a backup configuration, but they must be used together. In other words, if cftsrv107 or cftsrv111 are not accessible for use by RAID Manager, then cftsrv109 and cftsrv113 must be used together to enable the backup configuration.

CFTSRV107 HORCM configuration file: `horcm1.conf`

```
#ip_address      service      poll(10ms)  timeout(10ms)
10.10.14.107     horcm1        1000            3000

HORCM_CMD
#dev_name          dev_name       dev_name
/dev/sdm

HORCM_DEV
#dev_group       dev_name    port#   TargetID       LU#     MU#
caxp            caxpdata    CL1-A   0              4
caxp            caxplog     CL1-A   0              5
```

```
HORCM_INST
#dev_group           ip_address          service
caxp                 10.10.14.111        horcm3
```

CFTSRV109 HORCM configuration file: `horcm2.conf`

```
HORCM_MON
#ip_address     service     poll(10ms)       timeout(10ms)
10.10.14.109    horcm2      1000             3000

HORCM_CMD
#dev_name       dev_name       dev_name
/dev/sdm

HORCM_DEV
#dev_group      dev_name       port#   TargetID       LU#
caxp            caxpdata       CL1-A      0            4
caxp            caxplog        CL1-A      0            5

HORCM_INST
#dev_group           ip_address          service
caxp                 10.10.14.113        horcm4
```

CFTSRV111 HORCM configuration file: `horcm3.conf`

```
HORCM_MON
#ip_address     service     poll(10ms)       timeout(10ms)
10.10.14.111    horcm3      1000             3000

HORCM_CMD
#dev_name        dev_name       dev_name
/dev/sdm

HORCM_DEV
#dev_group       dev_name     port#    TargetID       LU#
caxp             caxpdata     CL1-E       0            8
caxp             caxplog      CL1-E       0           10

HORCM_INST
#dev_group           ip_address          service
caxp                 10.10.14.107        horcm1
```

CFTSRV113 HORCM configuration file: `horcm4.conf`

```
HORCM_MON
#ip_address     service   poll(10ms)     timeout(10ms)
10.10.14.113    horcm4      1000            3000

HORCM_CMD
#dev_name       dev_name       dev_name
/dev/sdm

HORCM_DEV
#dev_group      dev_name     port#    TargetID       LU#
caxp            caxpdata     CL1-E      0            8
caxp            caxplog      CL1-E      0           10

HORCM_INST
#dev_group      ip_address          service
caxp            10.10.14.109        horcm2
```

Startup and shutdown commands

To start the processes for RAID Manager Continuous Access XP, you must successfully run the following command on cftsrv107 and cftsrv111:

```
[root@cftsrv107 root]# horcmstart.sh
```

The following message displays if there is a successful **startup** execution of the command `horcmstart.sh`:

```
starting HORCM inst 1
HORCM inst 1 starts successfully.
```

This message displays when there is a successful **shutdown** execution of the command: `horcmshutdown.sh:`

```
[root@cftsrv107 root]# horcmshutdown.sh
HORCM Shutdown inst 1 !!!
```

Root .bash_profile

The following is an example of how a typical root `.bash_profile` looks:

```
# .bash_profile

# Get the aliases and functions
if [ -f ~/.bashrc ]; then
        . ~/.bashrc
fi

# User specific environment and startup programs
PATH=$PATH:$HOME/bin:/software/etc
BASH_ENV=$HOME/.bashrc
USERNAME="root"
HORCMINST=1
export HORCMINST DISPLAY USERNAME BASH_ENV PATH
```

Pair display command

The following is an example of the pair display command : ‘`pairdisplay –g caxp –fcx`’:

```
[root@cftsrv107 root]# pairdisplay -g caxp -fcx
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,Fence,   %,P-LDEV# M
caxp    caxpdata(L) (CL1-A , 0,   4)10197     0.P-VOL PAIR NEVER ,  100     0 -
caxp    caxpdata(R) (CL1-E , 0,   8)10818     0.S-VOL PAIR NEVER ,  100     0 -
caxp    caxplog(L)  (CL1-A , 0,   5)10197    20.P-VOL PAIR NEVER ,  100    20 -
caxp    caxplog(R)  (CL1-E , 0,  10)10818    20.S-VOL PAIR NEVER ,  100    20 -
```

RAID Manager configuration files for Business Copy XP

Business Copy XP was run on server cftsrv111, but it also could be configured to cftsrv113.

The process for performing Business Copy XP replication allows the Continuous Access XP link to be active while the copy is taking place. However, before splitting the Business Copy XP pair, the Continuous Access XP link must be temporarily suspended, and all changes pending in the Continuous Access XP link must be completely synchronized.

Sample information for configuring Business Copy XP for the Continuous Access XP replicated database and logs on SITE B (XP 1024) follows.

Environment variables

```
On cftsrv111        HORCMINST=5
HORCC_MRCF=1 (this must be set     for BC operations)
```

CFTSRV111 HORCM BC configuration files: `horcm5.conf`

```
HORCM_MON
#ip_address   service  poll(10ms)       timeout(10ms)
10.10.14.111  horcm5   1000              3000

HORCM_CMD
#dev_name      dev_name       dev_name
/dev/sdm

HORCM_DEV
#dev_group     dev_name    port#  TargetID   LU#  MU#
bcxp           bcxplog     CL1-E   0         10    0
bcxp           bcxpdata    CL1-E   0         8     0

HORCM_INST
#dev_group     ip_address       service
bcxp           10.10.14.111     horcm6
```

CFTSRV111 HORCM BC configuration files: `horcm6.conf`

```
HORCM_MON
#ip_address   service    poll(10ms)    timeout(10ms)
10.10.14.111  horcm6      1000           3000

HORCM_CMD
#dev_name      dev_name       dev_name
/dev/sdm

HORCM_DEV
#dev_group     dev_name    port#  TargetID  LU#   MU#
bcxp           bcxplog     CL1-E    0       19     0
bcxp           bcxpdata    CL1-E    0       21     0

HORCM_INST
#dev_group           ip_address          service
bcxp                 10.10.14.111        horcm5
```

Business Copy XP commands

`paircreate -g bcxp -vl`	Initiate a pair create
`Pairsplit -s -g bcxp`	Suspend the pair but keep the pair relationship
`Pairsplit -S -g bcxp`	Split the pair relationship
`pairdisplay -f bcxp -fcx`	Check the status of the pair
`pairresync -g bcxp`	Resynchronize a suspended pair

Appendix E

This appendix contains the sample files used to set up the Linux environment variables and Oracle network configurations. Note that the only difference between the user profiles (“caxp” and “odg”) is the ORACLE_SID values. The ”listener.ora” file varies by the different values given for the server addresses. The “tnsnames.ora” file is the same across all the servers.

Oracle environment variables for Linux

CA-XP: .bash_profile

```
# .bash_profile
# Get the aliases and functions
if [ -f ~/.bashrc ]; then
        . ~/.bashrc
fi
# User specific environment and startup programs
PATH=$PATH:$HOME/bin:/software/etc:/software/scripts
export PATH
unset USERNAME
########################################
## Oracle specific environment variables
###------------------------------------
## Used for CA-XP configurations
## by user 'caxp'
########################################
ORACLE_BASE=/u01/orabase
export ORACLE_BASE
ORACLE_HOME=/u01/orabase/ora10g
export ORACLE_HOME
CRS_HOME=/u01/crshome
export CRS_HOME
PATH=$PATH:$ORACLE_HOME/bin
export PATH
CLASSPATH=$CLASSPATH:$ORACLE_HOME/JRE:$ORACLE_HOME/jlib
CLASSPATH=$CLASSPATH:$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib
export CLASSPATH
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$ORACLE_HOME/lib:$ORACLE_HOME/oracm/lib
export LD_LIBRARY_PATH

HOSTNUMBER=`hostname|cut -c6-8`
case "$HOSTNUMBER" in
        107)    ORACLE_SID=CAXP1
                ;;
        109)    ORACLE_SID=CAXP2
                ;;
        111)    ORACLE_SID=CAXP1
                ;;
        113)    ORACLE_SID=CAXP2
                ;;
        *)      ORACLE_SID=""
                ;;
esac
export ORACLE_SID
export EDITOR=vi
umask 0022
```

ODG: .bash_profile

```
# .bash_profile
# Get the aliases and functions
if [ -f ~/.bashrc ]; then
        . ~/.bashrc
fi
# User specific environment and startup programs
PATH=$PATH:$HOME/bin:/software/etc:/software/scripts
```

```
export PATH
unset USERNAME
########################################
## Oracle specific environment variables
###-----------------------------------
## Used for ORACLE DATA Guard
## by user 'odg'
########################################
ORACLE_BASE=/u01/orabase
export ORACLE_BASE
ORACLE_HOME=/u01/orabase/ora10g
export ORACLE_HOME
CRS_HOME=/u01/crshome
export CRS_HOME
PATH=$ORACLE_HOME/bin:$PATH
export PATH
CLASSPATH=$CLASSPATH:$ORACLE_HOME/JRE:$ORACLE_HOME/jlib
CLASSPATH=$CLASSPATH:$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib
export CLASSPATH
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$ORACLE_HOME/lib:$ORACLE_HOME/oracm/lib
export LD_LIBRARY_PATH
HOSTNUMBER=`hostname|cut -c6-8`
case "$HOSTNUMBER" in
        107)    ORACLE_SID=ODG11
                ;;
        109)    ORACLE_SID=ODG12
                ;;
        111)    ORACLE_SID=ODG51
                ;;
        113)    ORACLE_SID=ODG52
                ;;
        *)      ORACLE_SID=""
                ;;
esac
export ORACLE_SID
export EDITOR=vi
umask 0022
```

Oracle network configuration files

listener.ora

```
# listener.ora.cftsrv111 Network Configuration File:
/u01/orabase/ora10g/network/admin/listener.ora.cftsrv111
# Generated by Oracle configuration tools.

SID_LIST_LISTENER_CFTSRV111 =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = PLSExtProc)
      (ORACLE_HOME = /u01/orabase/ora10g)
      (PROGRAM = extproc)
    )
  )
LISTENER_CFTSRV111 =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC))
      )
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv098.cftlabs.hpq.net)(PORT =
1521)(IP = FIRST))
      )
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = TCP)(HOST = 10.10.14.111)(PORT = 1521)(IP =
FIRST))
      )
    )
  )
```

tnsnames.ora

```
# tnsnames.ora Network Configuration File:
/u01/orabase/ora10g/network/admin/tnsnames.ora
# Generated by Oracle configuration tools.

LISTENERS_CAXP =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv098.cftlabs.hpq.net)(PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv099.cftlabs.hpq.net)(PORT = 1521))
  )
CAXP =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv098.cftlabs.hpq.net)(PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv099.cftlabs.hpq.net)(PORT = 1521))
    (LOAD_BALANCE = yes)
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = CAXP)
      (FAILOVER_MODE =
        (TYPE = SELECT)
        (METHOD = BASIC)
        (RETRIES = 180)
        (DELAY = 5)
      )
    )
  )
CAXP1 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv098.cftlabs.hpq.net)(PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = CAXP)
      (INSTANCE_NAME = CAXP1)
    )
  )
CAXP2 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv099.cftlabs.hpq.net)(PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = CAXP)
      (INSTANCE_NAME = CAXP2)
    )
  )
EXTPROC_CONNECTION_DATA =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC))
    )
    (CONNECT_DATA =
      (SID = PLSExtProc)
      (PRESENTATION = RO)
    )
  )
LISTENERS_ODG1 =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv096.cftlabs.hpq.net)(PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv097.cftlabs.hpq.net)(PORT = 1521))
  )
LISTENERS_ODG5 =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv098.cftlabs.hpq.net)(PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv099.cftlabs.hpq.net)(PORT = 1521))
  )
ODG1 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv096.cftlabs.hpq.net)(PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv097.cftlabs.hpq.net)(PORT = 1521))
    (LOAD_BALANCE = yes)
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = ODG1)
      (FAILOVER_MODE =
        (TYPE = SELECT)
        (METHOD = BASIC)
        (RETRIES = 180)
        (DELAY = 5)
      )
    )
  )
ODG11 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv096.cftlabs.hpq.net)(PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = ODG1)
      (INSTANCE_NAME = ODG11)
    )
  )
ODG12 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv097.cftlabs.hpq.net)(PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = ODG1)
      (INSTANCE_NAME = ODG12)
    )
  )
ODG5 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv098.cftlabs.hpq.net)(PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv099.cftlabs.hpq.net)(PORT = 1521))
    (LOAD_BALANCE = yes)
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = ODG5)
      (FAILOVER_MODE =
        (TYPE = SELECT)
        (METHOD = BASIC)
        (RETRIES = 180)
        (DELAY = 5)
      )
    )
  )
ODG51 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv098.cftlabs.hpq.net)(PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = ODG5)
      (INSTANCE_NAME = ODG51)
    )
  )
ODG52 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = cftsrv099.cftlabs.hpq.net)(PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = ODG5)
      (INSTANCE_NAME = ODG52)
    )
  )
```

For more information

HP StorageWorks

For the latest HP StorageWorks product information:

http://www.hp.com/country/us/eng/prodserv/storage.html

The following links provide information about:

- HP StorageWorks Continuous Access XP
  http://www.hp.com/products1/storage/products/disk_arrays/xpstoragesw/continuousaccess/
- HP StorageWorks XP arrays
  http://h18006.www1.hp.com/products/storageworks/enterprise/index.html
- HP StorageWorks Business Copy XP
  http://www.hp.com/products1/storage/products/disk_arrays/xpstoragesw/business/
- XP12000 hardware
  http://h200002.www2.hp.com/bc/docs/support/SupportManual/c00318013/c00318013.pdf

Oracle

For the latest Oracle product information:

http://www.oracle.com/products/index.html

Technical guides

You can access these technical guides online:

- Continuous Access Configuration Manual
  http://h200007.www2.hp.com/bc/docs/support/SupportManual/c00098520/c00098520.pdf
- Business Copy Configuration Manual
  http://h200007.www2.hp.com/bc/docs/support/SupportManual/c00098533/c00098533.pdf
- RAID Manager Manual
  http://h200002.www2.hp.com/bc/docs/support/SupportManual/c00313867/c00313867.pdf
- Performance Advisor Manual
  http://h200002.www2.hp.com/bc/docs/support/SupportManual/c00312112/c00312112.pdf
- Command View User Manual
  http://h200001.www2.hp.com/bc/docs/support/SupportManual/c00312425/c00312425.pdf
- Red Hat Linux Documentation
  http://www.redhat.com/docs/
- Oracle 10g Documentation Library
  http://www.oracle.com/technology/documentation/database10g.html

4AA0-1650ENW, August 2005