HP NAS Data Center with iSCSI support for Oracle white paper

Executive summary

The HP NAS Data Center with Internet SCSI (iSCSI) support for Oracle® is an affordable, easy-to-manage enterprise storage solution. By combining the HP ProLiant Storage Server iSCSI Feature Pack for Oracle database and the NAS Data Center solution with HP OpenView Storage Mirroring, this white paper shows how small to medium businesses can reap the benefits of reliable data protection without the cost of a complex storage system. The performance tests described herein utilize the NAS Data Center solution environment—one central location for data collection and multiple satellite offices with Oracle databases connected to the central site using iSCSI. They further show the effectiveness of using Storage Mirroring to perform remote copy over an IP LAN and demonstrate how this application provides a very cost-effective disaster recovery alternative from one site to another within a LAN or direct attached storage (DAS) system.

This white paper focuses on Oracle 10g performance test results using two test scenarios that support continued system operation if there is a (1) disaster at the central site, (2) loss of the LAN, or (3) loss of local storage data at a satellite site. To assist HP customers in the development of Oracle environments within a NAS-iSCSI infrastructure, recommended practices and conclusions based on data result comparisons are also presented.

Introduction

The performance tests described in this white paper use the HP NAS Data Center solution as the model for the design of the test environment and configuration. One central location, called the central site, was configured for dedicated data collection and housed an HP ProLiant DL380 (NAS) Storage Server device. Two low-cost satellite sites, site A and site B, were created to simulate distant offices and were configured differently based on the required test scenarios. Site A ran HP OpenView Storage Mirroring and sent mirrored data from a local Oracle server to the central site. Site B used iSCSI to connect Oracle databases directly to Oracle data storage at the central site. NAS Storage Servers at site A provided local storage to the local office. The central location stored the mirrored data on an HP StorageWorks Enterprise Virtual Array 5000 (EVA5000) device (for site A) or on a ProLiant DL380 Storage Server (for site B). By consolidating all storage at the central location, it was easy to recover one of the satellite offices in case of a disaster. Also, there was no need for storage management at the satellite site since all of the data was stored at the central site.

In the tests, the HP OpenView Storage Mirroring software was used to maintain a continuously updated and immediately accessible copy of site A data on the central site. It was also used to provide data protection and disaster recovery during test-simulated site disaster and LAN failures.

The HP ProLiant Storage Server (NAS) iSCSI Feature Pack was installed on site B to provide a single storage platform for delivering file, print, and application storage (block) services. iSCSI is a cost-effective solution as it enables transmission of SCSI block commands and data over existing IP networks using the TCP/IP protocol.

All connections were made over a Gigabit Ethernet (GigE) using industry-standard hardware and software to demonstrate the low-cost aspect of utilizing existing infrastructures.

The following sections present the test results, recommended practices, and observations of the HP Customer Focus Team (CFT) related to the proof of concept tests for the HP NAS Data Center solution with iSCSI support for Oracle and Storage Mirroring.

Test scenarios

The HP NAS Data Center with iSCSI support for Oracle test configuration consisted of two satellite sites (A and B) and a central site.

Site A is an Oracle 10g installation that has all of its working data stored locally on site. It is a completely stand-alone configuration and is able to operate even if the central site fails or the LAN connection is disrupted. Site A provides local NAS or iSCSI storage to local clients. The Oracle data on site A is mirrored to the central site using HP OpenView Storage Mirroring. The Storage Mirroring provides data protection by backing up site A's Oracle data in real time, and also provides for disaster recovery if the satellite site A is destroyed. These tests will show that recovery from a disaster at site A was easily managed and relatively quick.

Site B is also an Oracle 10g installation. The site's Oracle storage is provided by way of an iSCSI connection from local clients. The Oracle path to the data stored on the central site is completely iSCSI. The distance between site B and the central site is simulated to be 80 km using a network delay tool called "DummyNet."

The central site provides "data center" capabilities including a data repository, data management (including backup), and disaster recovery.

The tests were separated into two basic scenarios: (1) testing of Oracle 10g performance and (2) testing of failure and recovery. Configurations were designed for easy recovery from two different disaster points—failure of the LAN backbone and failure of the local storage device on site A. Test scenarios are presented in Table 1.

Table 1. Test scenarios

Testing performance of Oracle 10g performance
• Performance of satellite site A on Oracle 10g
• Performance of satellite site B on Oracle 10g
• Performance of sites A and B running simultaneously
Failure testing
• Failure of the LAN backbone
• Failure of local storage on satellite site A
• Recovery of local storage on satellite site A

Test configuration

The test configuration for each site follows and is illustrated in Figure 1.

Satellite site A

- (1) HP ProLiant DL380 G4 Storage Server with a 120-GB RAID 1+0 file share and HP iSCSI Feature Pack
- (1) HP ProLiant DL580 Oracle Server running Microsoft® Windows® 2003 Enterprise, with Microsoft iSCSI Initiator Package loaded
- (1) HP ProLiant DL360 client running Benchmark Factory as the Oracle load generator

Satellite site B

- (1) HP ProLiant DL580 server configured as an Oracle 10g server, with Microsoft iSCSI Initiator Package loaded
- (1) HP ProLiant DL360 client running Benchmark Factory as the Oracle load generator

Central site

- (2) HP ProCurve Switch 6108 GigE LAN switch that provided:
  - Connectivity of the satellite sites to the central site
  - Component connectivity within the central site
- (1) HP ProLiant DL360 server. This server also ran HP OpenView Storage Mirroring to create and maintain the mirrored data from site A on an EVA5000.
- (1) HP ProLiant DL380 G4 Storage Server with HP ISCSI Feature Pack installed. The server received block-level commands and data from the Oracle server at site B over an iSCSI connection. This data was stored on the local (internal disks) of the storage server.
- (1) HP StorageWorks SAN Switch 2/16T 2-Gb Fibre Channel switch
- (1) EVA5000 2C2D disk array with 72-GB disk drives, providing 962 GB of "user" storage (RAID 5, double sparing) for storing HP OpenView Storage Mirroring data from site A
- (1) HP OpenView Storage Management Appliance supporting EVA storage management

LAN infrastructure

- (1) HP ProLiant DL320 system running FreeBSD and "DummyNet" was used for the network delay.

The LAN provided component connectivity within a satellite site as well as connectivity from the satellite sites to the central site. All LAN switches and NICs are 1 GigE.

The two HP ProCurve switches connected the satellite sites to the central site. This infrastructure also provided network interconnectivity for components within the satellite sites. The LAN contained a Linux system running "DummyNet" to simulate an 80-km distance between the satellite sites and the central site. All IP traffic from the satellite sites was delayed by 20 ms through the "DummyNet" system.

For test equipment and software specifications, refer to Appendix A.

Figure 1. Logical configuration diagram for all three sites

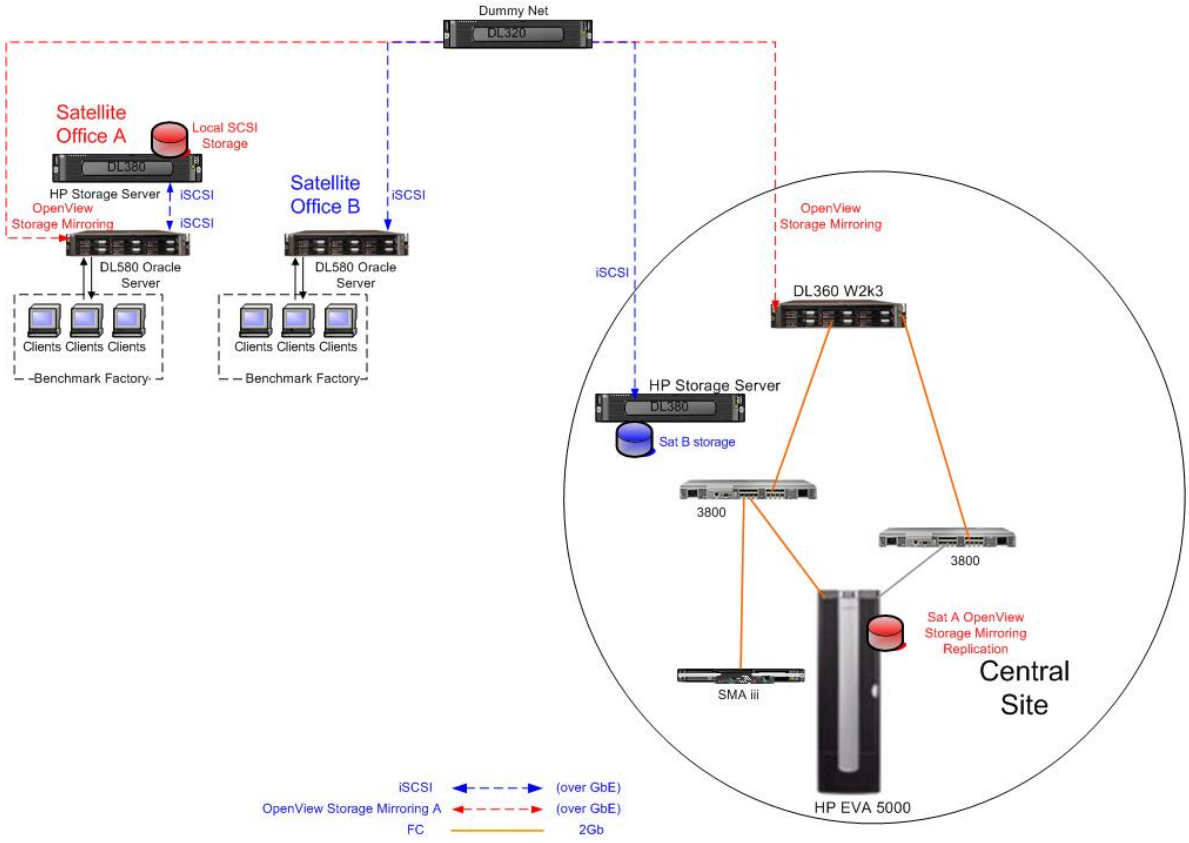

Oracle load generator

The Benchmark Factory load generator by Quest Software was installed. All testing used a load profile based on the TPC-B specification (www.tpc.org). An initial load profile provided data that showed how the configuration performed as the number of users changed. Based on this data, all subsequent tests were run with a user load of 50. At this load level, the Oracle server CPUs were approximately 25% busy on average with Oracle-related processes. The disks were heavily utilized, delivering nearly 1,400 IO/sec at peak performance.

Configuration of HP OpenView Storage Mirroring

Because HP OpenView Storage Mirroring is a major component in this proof of concept testing, it was important to properly configure and test it before starting the NAS Data Center tests. This preliminary testing involved configuration of the software and running multiple trials of high-speed data generation to evaluate the HP OpenView Storage Mirroring mirror for data integrity and performance.

Test procedures

The purpose of these tests is to demonstrate the performance and reliability of Oracle in an iSCSI environment and show that the systems are able to recover easily from a catastrophic event. The following sections describe the measured performance results for Oracle on each site and the observed results of each site failure test. In all test cases, the database size was set to 200 GB. Tests used Benchmark Factory to initially create and populate an Oracle database with data and then to generate the Oracle workload for each test scenario.

Test 1: Oracle performance on site A

This test evaluated the typical performance of the site A while running Benchmark Factory. All tests were performed with HP OpenView Storage Mirroring operating in continuous mirroring mode. The following steps describe the procedure for this test:

1. Simulate user loads on site A.
    a. Configure the Benchmark Factory load generator to force the Oracle server's response time 150 ms (relatively high response time).
2. Take Oracle performance measurements.
    a. Measure the Oracle transactions per second (TPS) using the Benchmark Factory's data reporting tools on the Benchmark Factory server.
    b. Measure the following on the Oracle server at site A: Read/s, Write/s, MB/s, and the reported LAN traffic from HP OpenView Storage Mirroring to its mirror site at the central office.
    c. Measure the CPU utilization, Read/s, Write/s, MB/s, and network traffic on the DL380 Storage Server in the satellite site A.
    d. Measure the CPU utilization, Read/s, Write/s, MB/s, and network traffic on the DL380 Storage Server in the central site.

Test 2: Oracle performance on site B

This test evaluated the performance of site B using the Benchmark Factory load generator. All tests were performed with HP OpenView Storage Mirroring operating in continuous mirroring mode. The following are the procedure steps for this test:

1. Simulate user loads on site B.
    a. Configure the Benchmark Factory load generator to force the Oracle server's response time 150 ms (response time limit).
2. Take Oracle performance measurements.
    a. Measure the Oracle TPS using the Benchmark Factory data reporting tools on the Oracle server. In this case the IO block-level traffic happens over the iSCSI link, so measure only network utilization on the Oracle server at site B.
    b. Measure the CPU utilization, Read/s, Write/s, MB/s, and network traffic on the iSCSI Storage Server at the central site.

Test 3: Oracle performance on sites A and B

In this test, the two satellite sites ran Oracle and sent data to the central site. Site A sent HP OpenView Storage Mirroring data to the central site, and site B sent block data over an iSCSI link.

1. Simulate user loads on site A and site B.
    a. Configure the Benchmark Factory load generator to force the Oracle server's response time 150 ms on sites A and B.
2. Take Oracle performance measurements.
    a. Measure the Oracle TPS using Benchmark Factory data reporting tools for all Oracle servers on sites A and B.
    b. Measure the following on the Oracle server at site A: Read/s, Write/s, MB/s, and the reported traffic from HP OpenView Storage Mirroring to its mirror site at the central site.
    c. Measure the network utilization on the Oracle server at site B.
    d. Measure the CPU utilization, Read/s, Write/s, MB/s and network traffic on both DL380 Storage Servers in the central site.

Test 4: Fail LAN backbone

This test demonstrated that HP OpenView Storage Mirroring was able to recover from a 1-hour LAN outage to the central site, while the satellite site A continued to process Oracle data (generated by the Benchmark Factory). It also showed HP OpenView Storage Mirroring capability to recover after the LAN was restored and then to successfully re-synchronize to mirror data.

1. Simulate user loads on site A and site B.
    a. Configure the Benchmark Factory load generator to force the Oracle server's response time 150 ms on sites A and B.
    b. Check that the two satellite site systems are stable and issue the expected amount of IO traffic.
2. Simulate 1-hour outage to central site.
    a. Break the GigE LAN link connecting the satellite sites and the central site.
    b. Verify that Oracle continues to run on site A.

c. Allow 1 hour to elapse.
d. Re-connect the GigE LAN link between the satellite sites and the central site.

3. Verify the HP OpenView Storage Mirroring recovery and data re-synchronization.
    a. Measure the time required for HP OpenView Storage Mirroring to re-synchronize the mirror on the central site.
    b. Verify that Oracle can again start on site B.

Test 5: Fail and recover local storage on site A

This test shows that when data stored at a satellite site is completely destroyed, it can be recovered from the mirror on the central site, and Oracle can be mounted again without corrupting the database. This test was performed on satellite site A and required the following steps:

1. Simulate user loads on site A.
    a. Configure the Benchmark Factory load generator to force the Oracle server's response time 150 ms on the Oracle server on site A.
2. Simulate the complete destruction of site A.
    a. Power down the HP ProLiant DL380 G4 Storage Server and the Oracle server.
    b. Remove the Oracle data disks from the HP ProLiant DL380 G4 Storage Server.
    c. Insert four blank data disk drives into the HP ProLiant DL380 G4 Storage Server and leave the root/application disks in the server.

3. Recover local storage on site A.
    a. Reboot the HP ProLiant DL380 G4 Storage Server.
    b. Request HP OpenView Storage Mirroring to initiate a re-sync of the mirror at the central site to the site A Storage Server.
    c. Wait until the re-sync completes before going on to step 4.

4. Remount Oracle without corrupting the database.
    a. Reboot the Oracle server on site A.
    b. Start Oracle and attempt to mount the database, that is, the four previously blank data disk drives on the HP ProLiant DL380 G4 Storage Server. These disks now have the restored HP OpenView Storage Mirroring data on them.
    c. Re-start Benchmark Factory and check that the performance is similar to the original configuration.

Test results

Oracle 10g performance results

To run performance tests on sites A or B, the response time limit had to be determined. Multiple Benchmark Factory tests were executed that incremented the number of users until a limit was found. In the case of sites A and B, the performance was limited by disk IO performance. The graph in Figure 2 plots the test response time against Oracle TPS. This graph clearly shows what is termed "the response time wall." At a certain load, additional Oracle users did not significantly increase the TPS value but only increased the response time (the time users wait for their data). This is usually the point at which administrators realize that additional hardware is required.

Figure 2. Response time curve

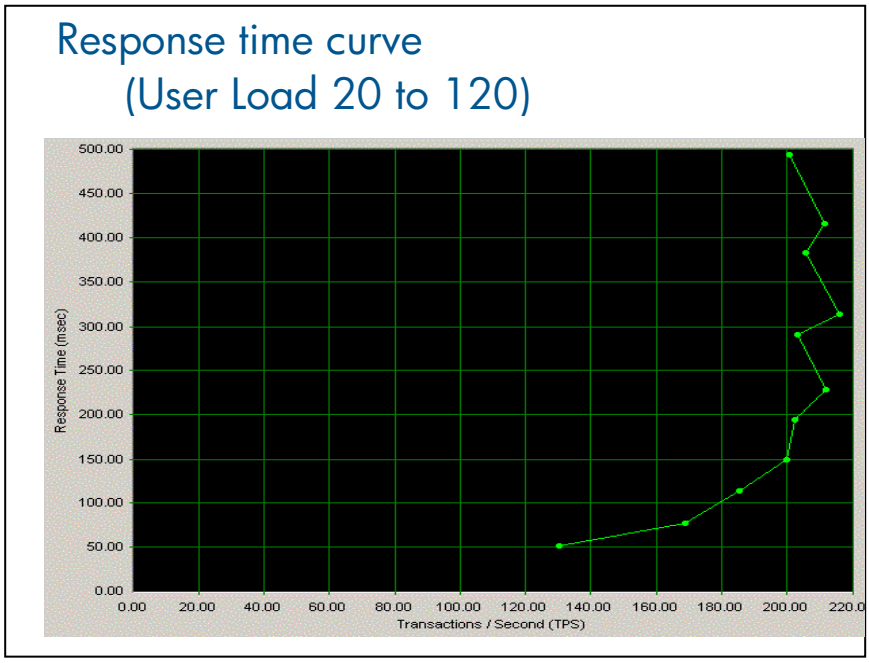

Another curve that also shows this value is generated when TPS is compared to the user load (see Figure 3). The plots showed that the maximum user load for each of the satellite sites should be approximately 50 users.

Figure 3. Oracle TPS

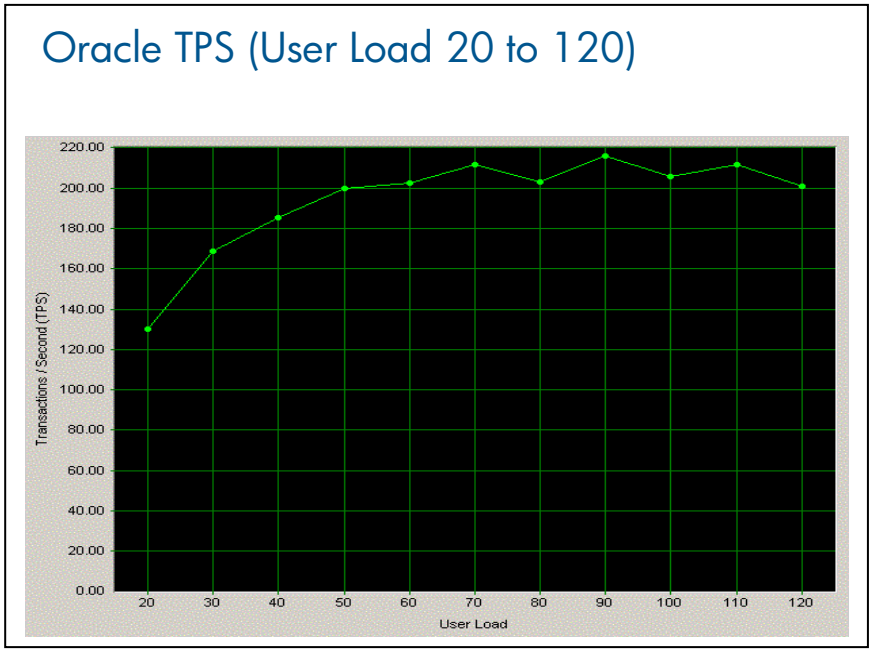

Performance of site A with Oracle 10g

Site A simulated a low-cost, scalable satellite office that was configured to be easily recoverable from a catastrophe at the site. To this end, HP OpenView Storage Mirroring was run to continuously mirror the Oracle data to a distant (80-km) central site. A network delay between the sites was set to 2 ms, which is equivalent to an 80-km distance with four switch hops. All of the iSCSI and HP OpenView Storage Mirroring mirror data went out of the Oracle server through a single 1-Gb/s NIC. For this reason, it was important to measure both the traffic to the storage and the traffic to the distance mirror. The Benchmark Factory generated a test workload of 50 TCP-B users.

This performance test was run on site A alone to ensure that the site did not overload the storage or the IP network. The following figures show the performance results for site A. They show the output measurements from the Microsoft Windows utility Perfmon. The data was obtained by setting Perfmon to take a snapshot of the site's performance near its peak value. Because these snapshots were taken at different times and on different computers, data comparisons between computers can only be approximate.

Figure 4. Oracle 10g performance results for site A

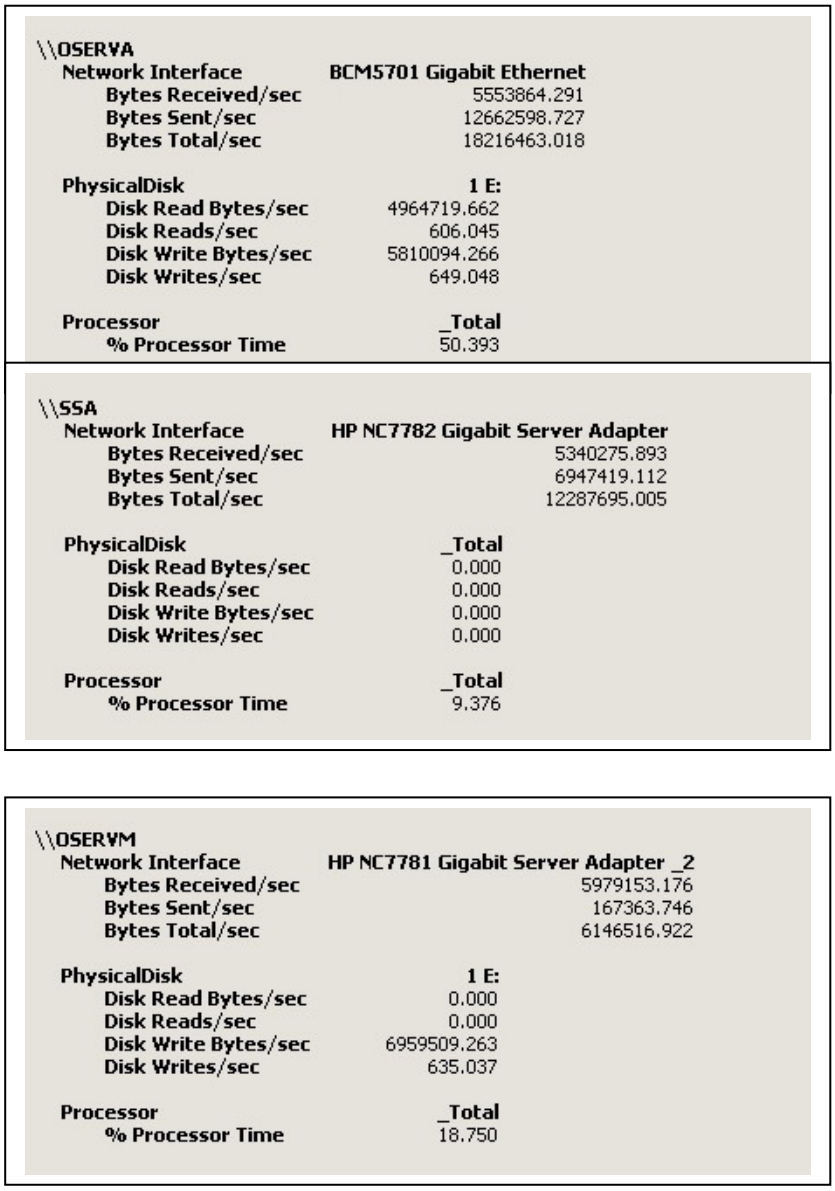

The output from the Benchmark Factory is shown in the following table.

Table 2. Benchmark summary statistics

Userload	TPS	Rows	Bytes	Avg. time	Min. time	Max. time	Errors
50	196.030	117585	470340	0.154	0.002	3.175	0

The key value is the number of Oracle TPS. For site A, this value was 196 TPS.

A better way to visualize the data flow is to incorporate the preceding data into a diagram showing where the data is flowing. Figure 5 shows site A and its data flows.

Figure 5. Data flow diagram for site A

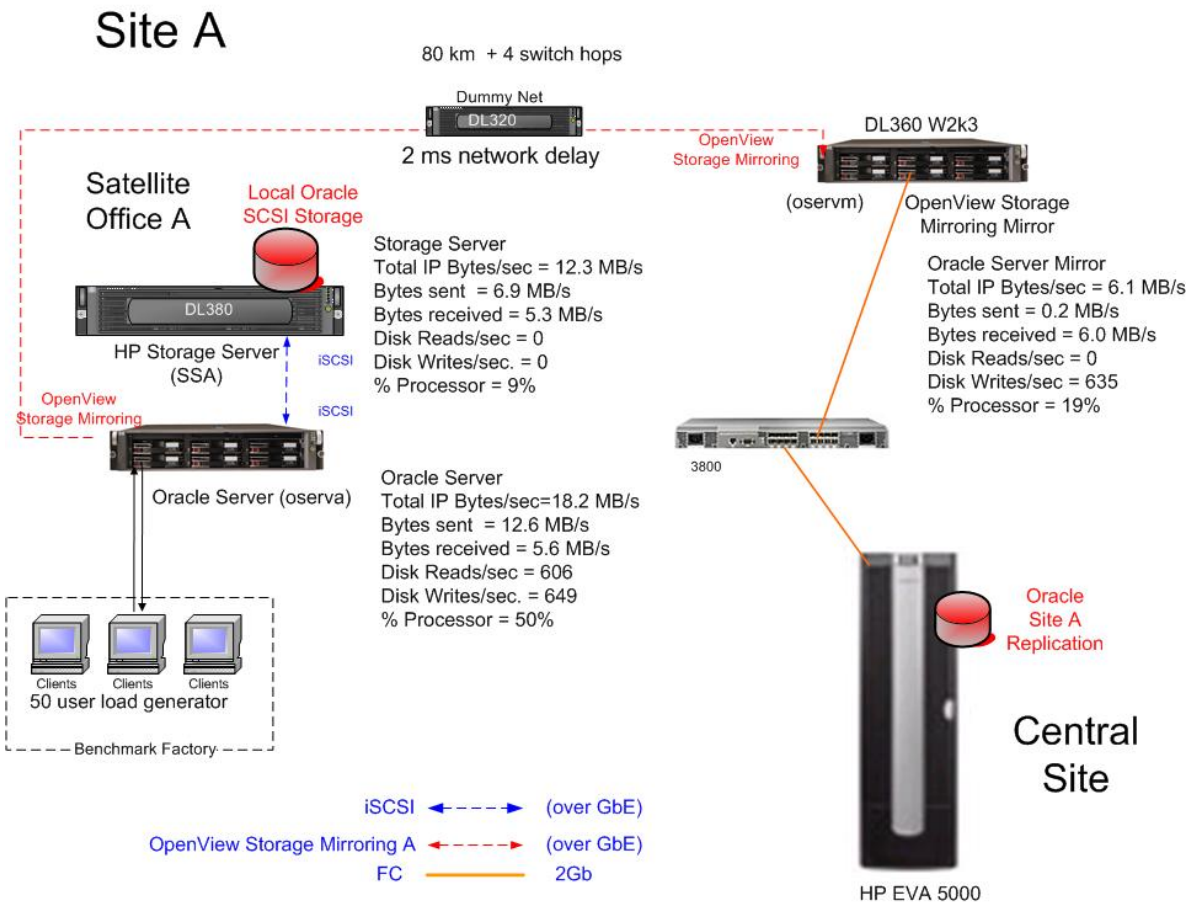

Some important patterns in the data flow should be pointed out.

With a user load of 50, the total number of writes by the Oracle server OSERVA is about 649 Writes/sec (5.8 MB/s). The number of Reads/sec is 606. This shows that the read and the write data rates are approximately the same.

The total traffic on the IP network is 18 MB/s. This amount is composed of the reads and writes to the iSCSI target from Oracle, plus the additional write traffic going from the Oracle server (OSERVA) over the IP network to the HP OpenView Storage Mirroring mirror OSERVM.

The Storage Server A (SSA) shows IP traffic as bytes sent and bytes received, but these only add up to about 12 MB/s. These reads and writes are the true iSCSI block transfers from Oracle to the storage on the storage server, and they are approximately 6 MB/s each.

The physical disk reads and writes on the storage server are zero (0). This is because this storage is not presented to the SSA host, but rather it is presented through the iSCSI link to the Oracle server OSERVA.

All of the traffic on the system OSERVM (the mirror) is “write” traffic. This is expected because a mirror will be receiving only writes from the primary side of the mirror (SSA).

CPU utilization for the Oracle server is high because it was performing both Oracle operations and the HP OpenView Storage Mirroring mirroring operations. By comparing OSERVA to a system that is running Oracle but not running mirroring software, the Oracle component is then about 25% and the HP OpenView Storage Mirroring component is also about 25%.

The IP network utilization never exceeded 20%, when compared to the theoretical maximum bandwidth of a 1 GigE network of 100 MB/sec.

Performance of site B with Oracle 10g

Site B was configured to show that Oracle would successfully run when storage was attached with an iSCSI link and located at some distance from the Oracle server. The test simulated a distance of 80 km by creating a 20-ms delay. The Windows utility Perfmon was used to record data and was set up to capture the peak volume of total IP traffic.

Figure 6. Oracle 10g performance results for site B

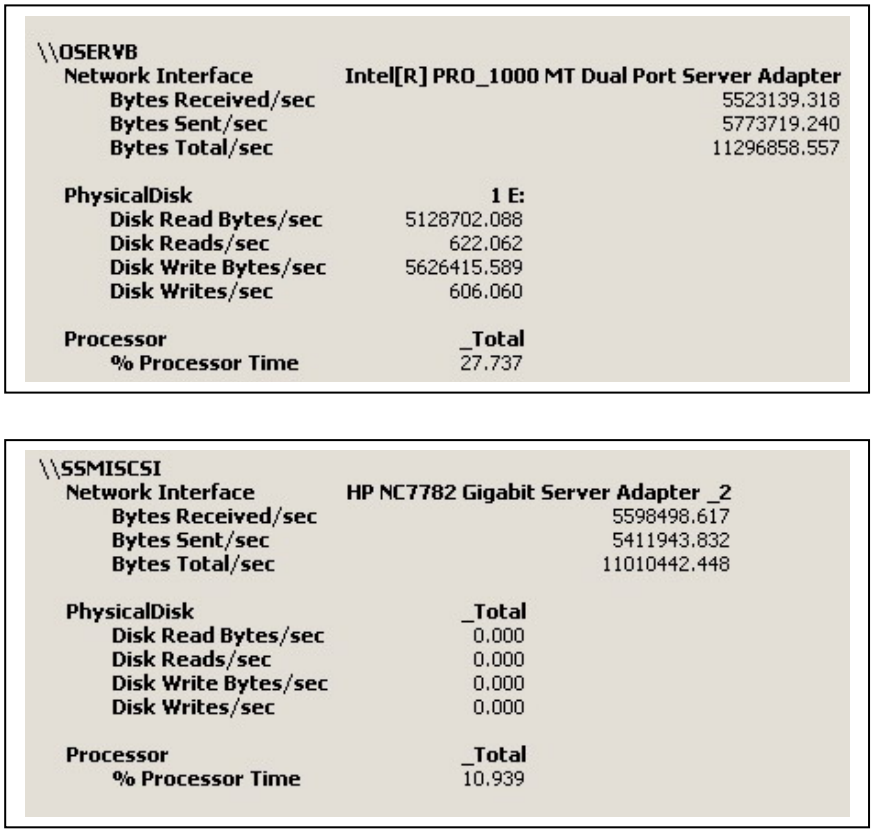

The output from the Benchmark Factory is shown in the following table.

Table 3. Benchmark summary statistics

Userload	TPS	Rows	Bytes	Avg. time	Min. time	Max. time	Errors
50	187.070	112243	448972	0.166	0.002	2.649	0

The key performance indicator is the number of Oracle TPS of 187. This value is slightly lower than it was for site A. This is because all of the Oracle data on site B went through the IP network delay of 2 ms.

Another way to visualize this data is to incorporate it into a data flow diagram. Figure 7 shows the data flow for site B.

Figure 7. Data flow diagram for site B

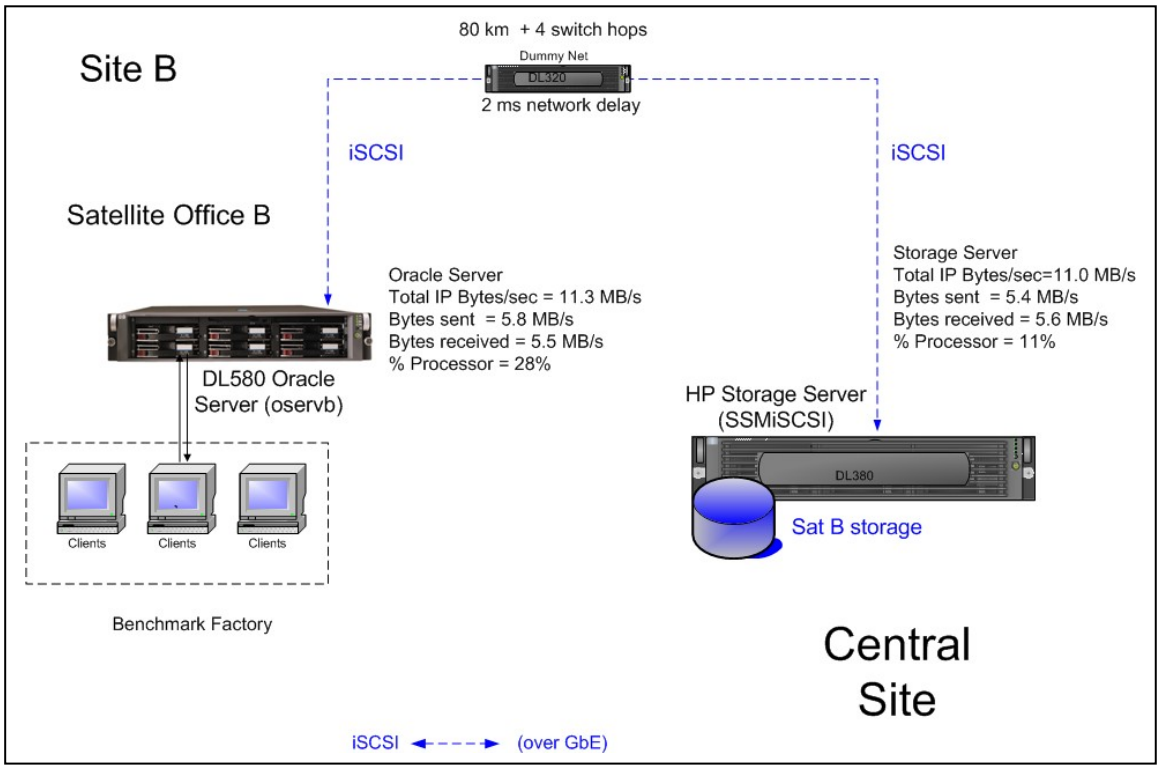

- The data transfers are roughly the same between the Oracle server and the Storage Server iSCSI disk drives.
- The network utilization is still very low, that is, it is under 15% in all cases.
- The iSCSI traffic is similar to site A: roughly 6 MB/s read, 6 MB/s write, 12 MB/s total.
- The CPU utilization is much lower on site B compared to site A because the HP OpenView Storage Mirroring mirroring is not running on site B.
- Oracle ran flawlessly over the 80-km distance to its storage (2.0-ms delay).

- Note that the iSCSI Storage Server cannot see its internal iSCSI drives. Therefore, as can be seen in Figure 6, the reads and writes on SSMISCSI are zero (0). These drives were configured with the iSCSI Feature Pack to be invisible to the storage server and to be presented as block-level SCSI devices to the Oracle server 80 km away.

Performance of sites running simultaneously

This test ran a workload of 50 users on both sites A and B. The test was designed to show that there was no interference between sites over the IP network. Since the servers, storage, and mirrors were independent, the only places where a bottleneck could occur were on the Benchmark Factory server or on the IP network. After repeated trials, tests showed nearly identical performance numbers to those achieved when running the sites separately. This is not surprising since the Oracle storage in both cases was separate and independent and the Oracle performance was clearly limited by the internal disks on the storage servers.

Failure testing of sites A and B

Several failure tests were performed to demonstrate the data recovery capabilities of the NAS Data Center solution. These tests included:

- Failure testing of the LAN backbone
- Failure of the local storage on site A
- Recovery of the local storage on site A

Failure testing of the LAN backbone

Failure tests were performed to determine how well Oracle would recover in an iSCSI configuration when the LAN that connected Oracle to its data storage was disrupted for a significant period of time. Also, the tests were important in determining the effects of removing site A's access to its mirror. This could be thought of as a catastrophic event at the central site.

The test resulted in a successful recovery from the failure. To start the test, the Benchmark Factory was run at a user load of 50. The iSCSI LAN cables were then suddenly removed from the IP network switch connecting the satellite sites to the central site. The response was different for each of the two sites.

The response at site A was:

1. Oracle continued to process data using the local storage on the storage server. This showed that site A could continue processing Oracle transactions even if the central site disappeared for an extended period of time.
2. The HP OpenView Storage Mirroring mirror running on site A's Oracle server quickly realized that the mirror on the main site had disappeared. It then began to buffer all of the Oracle write operations that should go to the mirror.
3. Since the HP OpenView Storage Mirroring service immediately began buffering data that was destined to be sent to the mirror, it was very important that sufficient storage space be made available to Storage Mirroring to store this data. Note that if HP OpenView Storage Mirroring does not have sufficient space to store all Oracle write operations, then the HP OpenView Storage Mirroring application must be stopped. This, however, leaves site A unprotected until the connection to the central site is restored. After the central site is restored, the mirror will resynchronize itself and normal operations will resume.
4. After the LAN connection was re-established, HP OpenView Storage Mirroring began to restore the mirror located at the central site. This operation occurred quickly because all of the write operations had been buffered on the Oracle server. These were then copied to the central site

mirror. If the LAN had stayed down long enough to require a full restore of the mirror, then it would have taken 7 hours to copy the full 200 GB to the mirror.

As previously mentioned, the response at site B was different.

1. Within 10 seconds the Windows Plug and Play functionality noticed that the LAN connection from the Oracle server had disappeared. This LAN connection was equivalent to a block-level SCSI connection to the Oracle storage (drive E: on the Oracle server). This caused Windows to declare that the drive formerly located at E: was no longer available.
2. Losing access to its data storage caused the Oracle instance to crash.
3. After the LAN connection was restored, Oracle was re-started. It easily recovered to the last known good check point.

HP OpenView Storage Mirroring was not running on site B because there was no local storage at the satellite site to be mirrored to the central site.

Failure and recovery of local storage on site A

The purpose of this test was to ensure that if a catastrophe occurred at site A, the site's Oracle data would be restored successfully and quickly from the mirror located at the central site. This test simulated a complete loss of site A. The Oracle load was forced to its maximum of 200 TPS. Then, all power cords were pulled simultaneously from the Oracle server (OSERVA) and the storage server (SSA). Next while the servers were down, all of the Oracle data disks were removed from the storage server and blank disks were inserted. In this way, none of the original Oracle data was available to the Oracle server when it booted up. Next, the HP OpenView Storage Mirroring mirror functions on the mirror server at the central site were turned off. The mirror was then prepared to perform a full restore of the Oracle data over the IP network to site A.

The response to the failure test procedure was as follows:

1. The Oracle server crashed in the middle of a relatively high volume of TPS (about 200 TPS).
2. The storage server also crashed and went down. While it was down, the Oracle storage disks were replaced with blank disk drives. This was drive E: on the Oracle server.
3. The Oracle server was then booted and the Oracle instance was not started. HP OpenView Storage Mirroring was started and a full restore from the mirror at the remote site was initiated. This restore took 7 hours. The long restore period was caused by the 2-ms network delay to the mirrored storage. When this delay was set to zero (0), the restore occurred in 3.3 hours.
4. The Oracle application was started, and it attempted to mount and open the database. Oracle immediately discovered that a catastrophic event had occurred and began its data recovery routine. This resulted in Oracle recovering the database to its last successful check point before the "event."

The HP OpenView Storage Mirroring restore was easy and relatively quick. The mirror operation also worked exactly as described in the documentation.

HP recommended practices

The HP Customer Focused Testing (CFT) team offers these recommendations for the installation and configuration of a NAS Data Center with iSCSI support for Oracle.

Ethernet

A GigE infrastructure greatly reduces the possibility of bottlenecks when connecting to storage on the LAN. HP recommends that you use at minimum a 1-Gb/sec Ethernet infrastructure for the iSCSI link to storage devices. GigE gives you additional performance over a 100 BaseT configuration, and it also

has the advantage of hardware-implemented TCP CRC offload. This means that all CRC calculations are done with dedicated hardware and do not use valuable CPU bandwidth.

WAN versus LAN

HP recommends that you avoid sending storage data over the WAN. This recommendation is for security reasons as well as performance. When data goes over the WAN, no one can predict where it goes or who may have access to it. Also, it is possible that the data may be routed through a very slow infrastructure and cause significant performance degradation.

System configuration

- Add an additional CPU in all iSCSI servers. This gives you additional dedicated CPU power to perform the TCP-IP overhead for iSCSI more efficiently.
- Add an additional 1 GB of RAM to allow sufficient buffering for improved performance.
- Use separate RAID groups for iSCSI Logical Unit Numbers (LUNs). This will minimize contention for data disk spindles.
- Use 15k RPM disk drives for the iSCSI storage when possible.
- Upgrade all Windows computers with the latest current patches and security fixes.
- For additional performance, configure Oracle log files with a separate LUN. This will allow increased parallel loading on the storage.

Conclusions

The performance tests described in this white paper focus on Oracle 10g performance using several test scenarios that support continued system operation if there is a (1) disaster at the central site, (2) loss of the LAN, or (3) loss of local storage data at a satellite site.

Within these testing constraints, each test was designed to demonstrate the value and flexibility of HP NAS Data Center solutions when combined with iSCSI support for Oracle and HP OpenView Storage Mirroring.

The test configurations were based on the NAS Data Center model. A central site provided a centralized repository of stored data from two satellite sites—A and B.

- Site A simulated a low-cost, easily scalable satellite office configured with HP OpenView Storage Mirroring to provide data protection and resiliency to a catastrophe at the site.
- Site B was attached to storage using an iSCSI link and was located at a simulated 80-km distance from the Oracle server. This site was tested to observe how Oracle would operate with its data storage located at a long distance and connected with iSCSI.

Storage Mirroring was installed to mirror Oracle data at the central site, 80 km away from the satellite office. This solution worked very well. After tests crashed site A, HP OpenView Storage Mirroring was able to easily and reliably restore Oracle to the remote site, bring Oracle online, and successfully restore a 200-GB Oracle 10g database to the last known checkpoint over the IP network without resorting to backup tapes. In general, the tests demonstrated how the failover and data replication capabilities of HP OpenView Storage Mirroring minimize the impact of a NAS server failure at a satellite site. The HP OpenView Storage Mirroring failover function reduced the downtime to minutes and provided automatic recovery of all key data.

The Oracle 10g storage performance tests with iSCSI were also highly reliable even with the data located at a simulated distance of 80 km from the site. The HP ProLiant Storage Servers (NAS) with iSCSI Feature Pack provided a fast, easy-to-manage, and cost-effective solution for attaching iSCSI storage to Oracle 10g servers. The iSCSI protocol over a GigE LAN was demonstrated to be robust, reliable, and in every way a solid performer. Test results showed iSCSI utilized no more than 20% of the available network bandwidth at any point during the tests.

These Oracle 10g performance tests further demonstrated how easy it was to add a new satellite to the central site, requiring only an HP ProLiant Storage Server with iSCSI support for Oracle and/or HP OpenView Storage Mirroring.

Appendix

The following table lists the equipment used to run the performance tests described in this white paper.

Table 4. CFT NAS test equipment and software specifications

Oracle and iSCSI for NAS		
NAS Storage and Software	**Quantity**	**Comment**
HP ProLiant DL380 G4 Storage Server	4	
HP OpenView Storage Mirroring	1	Software licenses
HP ProLiant iSCSI Feature Pack	1	
Oracle Clients		
HP ProLiant DL360 server	1	
Storage		
HP StorageWorks Modular Smart Array 1000 (MSA1000)	1	
EVA5000—2C2D	1	EVA5000 storage unit with 2.0-TB raw storage
FC HDD INT CTO 72G 15K	28	
VCS HSV110 Media Kit 3.0b	1 per array	
Management Appliance and SAN components		
HP OpenView Storage Management Appliance III	1	Manages EVA5000 (includes required software)
Fibre Channel Switch 2/16	2	SAN switches to connect storage and NAS
Oracle servers		
Oracle Server—DL580 2-GB memory	3	Only Windows 2003 is supported with Oracle 10g
HP ProLiant DL360 server	1	
Oracle 10g licenses	4	One license per OLTP server
Benchmark Factory	1	Load generator tool with license
LAN		
HP ProCurve Switch 6108	2	1 GbE LAN switch
HP ProLiant DL320 server running FreeBSD	1	Runs “DummyNet” to simulate 80-km distance between sites

For more information

HP storage

Access http://www.hp.com/country/us/eng/prodserv/storage.html for the latest information on this solution. From this website, select the appropriate product or solution.

These links provide information about solutions or components mentioned in this white paper:

- HP OpenView Storage Mirroring
  http://h18006.www1.hp.com/products/storage/software/sm/index.html
- HP ProLiant Storage Server iSCSI Feature Pack
  http://h18006.www1.hp.com/products/storage/software/inas/
- HP ProLiant Storage Servers (NAS)
  http://h18006.www1.hp.com/storage/networkattached.html
- HP StorageWorks Enterprise Virtual Array 5000
  http://h18006.www1.hp.com/products/storageworks/enterprise/index.html

HP authorized reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518.
- In Canada, call 1-800-263-5868.
- Elsewhere, see the HP website for locations and telephone numbers at http://www.hp.com.

HP technical support

Telephone numbers for worldwide technical support are listed at http://www.hp.com/support/. From this website, select the country of origin.

Note: For continuous quality improvement, calls may be recorded or monitored.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

5983-2074EN, 05/2005