Hewlett Packard Enterprise

HPE 3PAR All-Flash Acceleration for Oracle ASM Preferred Reads

HPE 3PAR StoreServ Storage All-Flash Arrays

Technical white paper

Contents

Executive summary

Oracle databases sit at the heart of a customer's business, demanding the highest levels of performance, availability, and data services. Many Oracle database environments are overloaded because they are utilizing aging storage. Costs are too high and constant storage tuning is required to keep up with service-level agreements (SLAs). Change can introduce risk, so many customers feel their only alternatives are to increase capacity or upgrade their legacy storage.

Hewlett Packard Enterprise 3PAR StoreServ All-Flash Arrays (AFAs) can help address this challenge. The addition of an HPE 3PAR StoreServ AFA as a preferred read device with Oracle ASM increases database performance significantly at a fraction of the cost of a legacy storage upgrade. This solution provides a low-impact, low-risk alternative to upgrading a legacy storage environment.

Introduction

Maintaining Oracle database service levels can be challenging due to the number of factors that affect overall performance. These include host hardware (CPU, memory, I/O), operating system (OS), SAN infrastructure, Oracle database configuration settings, and the storage system. Poorly written applications or application code merged over many years may also have a significant impact on Oracle database performance.

Database administrators (DBAs) will often use Oracle Automatic Workload Repository (AWR) reports to locate and pinpoint problem areas with Oracle applications. In a large number of cases the storage system is found to be the bottleneck in high-performance, mission-critical databases.

The problems found center around using older spinning hard disk technologies, outdated storage array architectures, and unforeseen growth in database utilization. This is caused typically by applications and software updates outpacing hardware refresh cycles in the data center. In order to mitigate challenges with storage system performance, constant storage tuning is required.

The obvious solution, to increase capacity or upgrade the current disk array ecosystem, requires a substantial financial investment, followed by the risk of a full data migration that may take months. This paper describes an alternative path: inserting an HPE 3PAR StoreServ AFA alongside the legacy array as a preferred read device under Oracle Automatic Storage Management (ASM) to enhance the overall performance of the existing storage subsystem.

This solution can be implemented with minimal disruption, at a fraction of the cost of a legacy storage upgrade. Additionally, this solution has been extensively tested and validated by HPE.

Technology overview

Prerequisites

The concepts explained in this paper have been verified for Oracle versions 11g R2 and 12c in a single-instance and Real Application Clusters (RAC) setup. The use of Oracle ASM is mandatory.

HPE 3PAR StoreServ All-Flash Array

HPE 3PAR StoreServ Arrays deliver the agility and efficiency required by modern data centers. Instant scalability and fast, just-in-time provisioning meets the changing needs of the business quickly and cost-effectively. Hardware-assisted thin provisioning combined with inline deduplication reduces the footprint of storage, saving power, space, and cooling in the data center. Autonomic provisioning and secure multi-tenant support achieve higher service levels for more users and applications with less infrastructure and administrator effort. An extended portfolio of software titles for replication, data placement optimization, quality of service (QoS), and more builds an enterprise-class, Tier-1 ecosystem with unmatched performance, agility, and resilience. The HPE 3PAR StoreServ All-Flash storage system offers extreme performance with the same benefits as spinning disk systems. The synergy of an Oracle database and an HPE 3PAR StoreServ AFA was demonstrated in a white paper[1] by HPE.

Note: The HPE 3PAR StoreServ AFA needs HPE 3PAR OS 3.2.1 or higher.

[1] prime31.sharepoint.hp.com/teams/hpesaleslibrary/PublicDocuments1/4aa5-7457enw.pdf

Oracle ASM

Oracle ASM is a volume manager and a file system for Oracle database files, which supports single-instance and Oracle RAC configurations. ASM is Oracle's recommended storage management solution that provides an alternative to conventional volume managers, file systems, and raw devices.

Oracle ASM makes use of disk groups to store database files. An ASM disk group is a collection of disks that ASM manages as a unit. Within a disk group, ASM exposes a file system interface for Oracle database files. The ASM stores these files evenly distributed or striped across the physical disks in the disk group to eliminate hot spots and provide uniform performance. Its performance is comparable to the performance of raw devices, an involved alternative to ASM.

One of the most beneficial features of Oracle ASM is the support of failure groups. Failure groups ensure that data and its redundant copy do not both reside on disks that are likely to fail together. The disk hardware for a failure group may be close to a node running an Oracle instance or farther away from it. It may be more efficient for ASM to read from a disk block (extent) on a node-local failure group, even if that extent is a secondary one. ASM can be configured to preferentially read from a secondary extent if that is 'closer' to the node requesting the data. This concept is most useful in geographically extended clusters.

In this paper we will show how preferred reads on selected failure groups placed on an HPE 3PAR StoreServ AFA accelerate Oracle databases.

Oracle enterprise database challenges

Keeping Oracle databases running at peak performance in an enterprise organization can be very challenging. Most large organizations have online transaction processing (OLTP) databases to capture and process customer orders and manage finances. Figure 1 shows an example of a real-time business intelligence application. Data entered by a user or captured by telemetry is stored in an Oracle OLTP database. That information is copied, transformed using an extract/transform/load (ETL) tool, and stored in data warehouse databases. Next, an Oracle database analytics engine processes the data. This data may be used internally or sold to external customers via data marts.

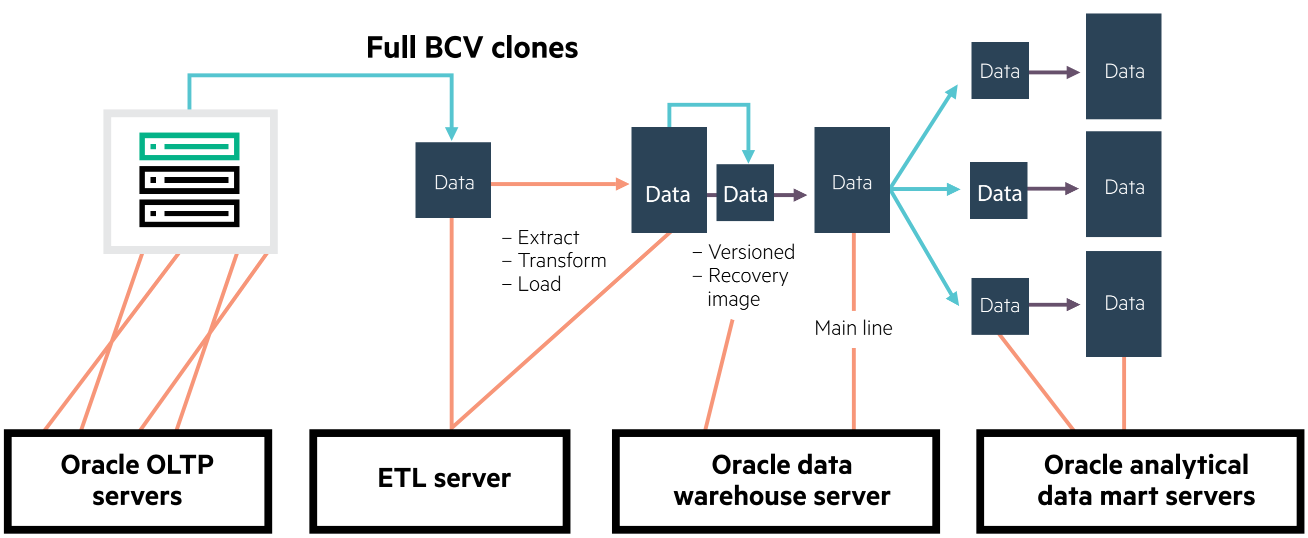

Figure 1: Real-time business intelligence application

Running an enterprise with OLTP, ETL, data warehouse, and analytical database servers is often very difficult and challenging. DBAs may spend hours monitoring performance by checking Oracle Enterprise Manager and reviewing AWR reports. From these AWR reports, DBAs learn how well the databases operate and what may cause their slow performance. To solve Oracle performance issues related to storage, DBAs and storage administrators check the physical read/write numbers and the top 10 key wait states found in AWR reports.

Note: AWR reports should not cover more than two hours. Any longer can dilute statistics and hide the real problem.

When physical reads and writes have a latency greater than 10 ms, upgrading the array by adding faster drives, more cache, and sub-LUN-tiering are the primary tactics used to offset poor performance. The next option available is a full array upgrade.

As shown in figure 1, enterprises often have multiple large legacy storage arrays running their Oracle databases. These legacy arrays typically have limitations on cache size and connectivity to the SANs, usually found to be 4 or 8 GB per second with Fibre Channel. Some legacy storage arrays can add flash SSD drives, but internal design limitations prevent them from taking advantage of the significant increase in IOPS and throughput of SSDs. Replacing just one of the multiple legacy arrays with a new one in figure 1 may not solve the performance problem, meaning a large investment is needed to rebuild the data handling chain for the next three to four years. Financial executives and DBAs would like to find a way to augment the performance of the Oracle environment without a major investment in capital and migration work. This is where the HPE StoreServ All-Flash solution fits.

HPE solution to accelerate legacy storage environments

HPE's solution to address the performance issue of an Oracle database on a legacy storage hardware adds an HPE 3PAR StoreServ AFA alongside the primary array. A new Oracle ASM failure group is added to the environment pointing to the 3PAR AFA. Implementing failure group Preferred Reads ensures that reads will come from the fastest storage array and will boost the performance of the Oracle database significantly: with read operations completion times in sub-milliseconds and writes in 1 ms or less. The setup for this is shown in figure 2. The orange arrows are for read operations.

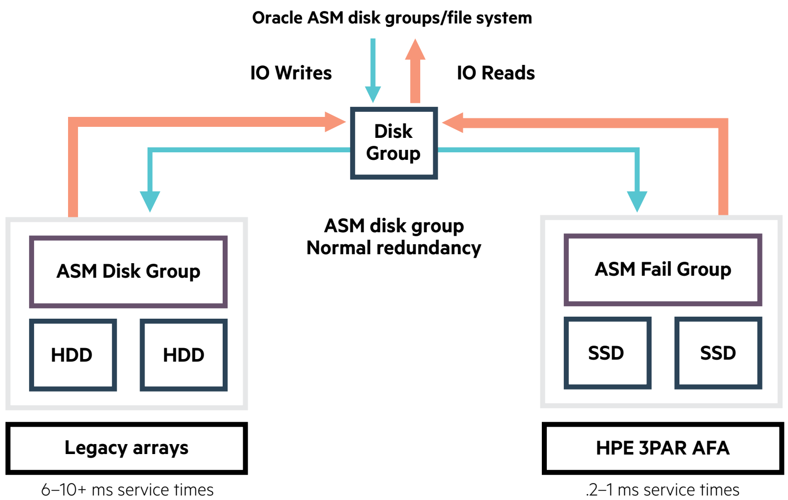

Figure 2: Adding an HPE 3PAR StoreServ All-Flash Array for Oracle ASM Preferred Reads

It is recommended to use the native MPIO drivers on the Oracle database hosts. In this solution, using HPE 3PAR with EMC products together, HPE recommends using the OS standard in-box drivers.

In the configuration described in figure 2, nearly all physical reads from the database will be coming from the HPE 3PAR StoreServ AFA. Writes go to both arrays, but the HPE 3PAR StoreServ AFA will acknowledge writes faster, speeding up the database as seen by HPE AWR reports. Using or adding a failure group helps minimize any risks in production environments. If one of the arrays fails, the Oracle database will be run from the array that is still online.

To prove the concept, HPE tested several third-party legacy storage arrays in the environment described in figure 2. Appendix A of this document demonstrates the setup and testing using an HPE 3PAR StoreServ 7450c AFA added as a failure group to a legacy EMC VMAX in an Oracle RAC environment with ASM.

Test results

To benchmark the performance gain in the setup of Figure 2, the Silly Oracle Benchmark (SLOB) was used. It is widely used by DBAs and enterprises to test their Oracle environments. SLOB supports the following features:

- Oracle logical reads (SGA buffer gets) scaling
- Physical random single-block reads
- Physical random single-block writes

SLOB avoids any application contention and uses simple SQL/PL commands to run. The best setting when using SLOB in Oracle ASM Preferred Read testing was to use logical Read/Write and scalability settings for the benchmark. The tests conducted using SLOB ran for two hours with maximum physical I/O to the database during the benchmarks.

Three sets of different benchmark tests were conducted to test Oracle ASM Preferred Reads. The following lists the benchmark tests:

- Heavily constrained legacy storage array with 6 ms latency service times to Oracle RAC
  - Baseline test using EMC VMAX with latency to Oracle RAC and running SLOB for two hours
  - Added HPE 3PAR StoreServ 7450c AFA as part of an Oracle failure group for preferred reads, and running SLOB for two hours
- Moderately constrained legacy storage array with 4 ms latency service times to Oracle RAC
  - Baseline test using EMC VMAX with latency to Oracle RAC and running SLOB for two hours
  - Added 3PAR 7450c AFA as part of an Oracle failure group for preferred reads, and running SLOB for two hours

Table 1 and 2 show the AWR test results of the benchmarks, with the physical reads and writes in MB per second, plus the service times for each test. Table 3 has data from HPE 3PAR System Reporter capturing data from SLOB during the two hour tests.

From the tables 1, 2, and 3 we see significant performance increases while using the HPE 3PAR StoreServ AFA with Oracle ASM Preferred Reads. The latency of storage access times drops to under 1 ms. In the first test with 3PAR AFA, the service database latency time is lowered to 250 microseconds. The Oracle ASM write latency drops from 6 ms to 1 ms, increasing overall storage performance.

By adding the HPE 3PAR StoreServ AFA into the Oracle RAC ASM failure group for Preferred Reads, the latency drops to 200 microseconds or 80 percent gain with reads, 76 percent with writes, and a combined 79 percent read/write overall performance increase.

Table 1: Performance results

USE CASE	READ IO (MB PER SECOND)	WRITE IO (MB PER SECOND)	USER IO AVG WAIT (MS)	TOTAL PHYSICAL READ IMPROVEMENT	PHYSICAL WRITES PER SECOND	PHYSICAL WRITES PER SECOND	TOTAL PHYSICAL WRITE IMPROVEMENT	TOTAL PHYSICAL R/WS	TOTAL PHYSICAL R/W IMPROVEMENT
VMAX heavily constrained (6 ms latency)	18.6	9.8	6 ms	2,376	Base	1,256	Base	3,632	Base
VMAX heavily constrained (6 ms latency with 3PAR AFA)	32.5	17.5	< 1 ms	4,160	75%	2,236	78%	6,369	76%
VMAX moderately constrained (4 ms latency)	31.1	16.8	4 ms	3,974	Base	2,155	Base	6,129	Base
VMAX moderately latency with 3PAR AFA)	66.9	29.7	<1 ms	7,153	80%	3,803	76%	10,956	79%

Note: The AWR report shows latency that appears to be the lowest number it can record. 3PAR System Reporter and CLI data shows actual latencies to be less than 500 microsecond response times.

The Oracle AWR results and the HPE 3PAR StoreServ System Reporter output show that HPE 3PAR StoreServ AFA can help significantly accelerate Oracle RAC database environment. The full AWR reports are in Appendix B.

Table 2: HPE 3PAR System Reporter data captured during two hour SLOB tests

TEST	IOPS	KBs PER SECOND	READ SERVICE TIMES (MS)	WRITE SERVICE TIMES (MS)
VMAX with heavily constrained service times (6 ms) with 3PAR AFA	5,800	63,000	.25	1
VMAX with moderately constrained service times (4 ms) with 3PAR AFA	10,500	100,000	.2	1

HPE 3PAR StoreServ System Reporter also captured the 3PAR CLI statvlun data to record VLUN performance while running failure groups with Preferred Reads from the volumes during the two hour SLOB tests. Table 3 shows the increase in performance from HPE 3PAR StoreServ AFA during the tests seen in the AWR reports and also highlighted in table 2 with System Reporter data being captured. The 3PAR AFA achieved 100,775 IOPS, with a latency less than .27 ms as captured during the Oracle RAC database constrained tests with VMAX.

Table 3: HPE 3PAR CLI Stat Virtual LUN data captured during two-hour SLOB tests

TEST	IOPS	KBs PER SECOND	READ SERVICE TIMES (MS)	WRITE SERVICE TIMES (MS)
VMAX with heavily constrained (6 ms) service times with 3PAR AFA	5,800	86,000	.24	.97
VMAX with moderately constrained service times (4 ms) with 3PAR AFA	10,500	100,775	.27	1.1

Table 4 shows the increase in write performance from HPE 3PAR StoreServ AFA during the tests seen in the AWR reports. While using 3PAR AFA writes were accelerated by 1,119 IOPS with constrained test and 1,820 IOPS with constrained environments. Along with read acceleration the writes also improved significantly allowing the database running SLOB benchmark to run faster.

Table 4: AWR Report of physical write acceleration comparison from baseline to AFA acceleration results

AWR REPORT	Physical writes baseline	Physical writes AFA acceleration
VMAX with heavily constrained service times (6 ms) with 3PAR AFA	1,256.4	2,375.6
VMAX with moderately constrained service times (4 ms) with 3PAR AFA	2,154.0	3,974.0

All above testing and validation was performed by Hewlett Packard Enterprise, providing customers with the assurance that complex solutions have been tested and architected by HPE for interoperability assurance and best practice configurations.

TCO improvement

Cost is a major factor for all Oracle database deployments. While SLAs must be met, costs must also be contained. Figure 3 shows the results of a TCO analysis conducted by Wikibon, comparing the cost of a VMAX upgrade to the addition of HPE 3PAR StoreServ AFA for read acceleration. Savings come from lower storage and server costs, plus a reduction in Oracle licensing costs. The three-year TCO analysis assumed a database size of 50 TB. The VMAX upgrade was from four to eight controllers with 100K IOPS required. The 3PAR AFA is a 4-node 8450 with 62 TB useable in RAID5 with no compression or dedupe. Also included in the 3PAR AFA figures were costs for additional host bus adapters and consulting services.

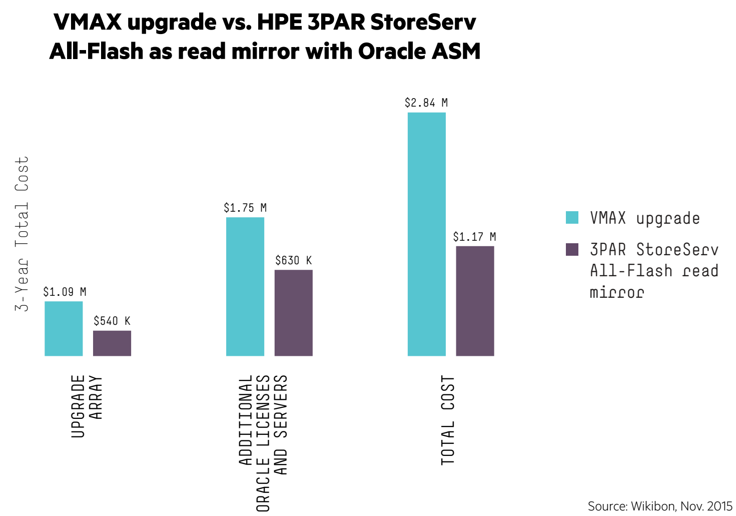

Figure 3: VMAX upgrade costs compared with use of HPE 3PAR StoreServ AFA for read acceleration co- located with existing VMAX deployment

In addition to savings from the cost of the array and Oracle license and server savings, the 3PAR AFA solution affords significant innovation as well as savings in the area of application development. Developers are demanding full database copies and read database copy access for DevOps (particularly for mobile development), and also high performance copies, like those from an All-Flash array instead of a legacy array.

The HPE 3PAR StoreServ All-Flash solution solves this challenge with read capacity on the 3PAR AFA mirror. Copies off of the 3PAR AFA are full, up-to-date copies with All-Flash performance. Additionally, copies off of the 3PAR AFA are created with 3PAR Virtual Copy snapshots with copy-on-first-write technology. These 3PAR snapshots are highly efficient compared to full copies/clones and also significantly reduce the storage capacity consumed by traditional copy proliferation. The net-net of all these benefits is a doubling of productivity for the application development community. Using the same three-year TCO model provided by Wikibon and described above, for a community of 10 application developers the additional cost savings to the business in productivity is $1.6 million USD.[2]

Table 5: Supported Configurations

OPERATING SYSTEM	VERSION SUPPORTED
Red Hat Enterprise Linux	6.x
Oracle Enterprise LINUX (OEL)	6.x
Microsoft Windows	2008, 2012
Standard inbox multipathing OS drivers supported	

[2] Wikibon 2015. Three-year view assuming 50 percent boost in productivity with application developer salary of $140,000 USD; benefits begin to kick in six months into the three year period

Host bus adapters

Separate host bus adapters (HBAs) are required for use with EMC and 3PAR StoreServ storage arrays to prevent array co-existence on the same HBA. Co-existence is not supported. This solution is restricted to using different HBA vendors for EMC and 3PAR StoreServ storage arrays to prohibit HBA Driver/BIOS/Firmware incompatibility. A single HBA vendor for EMC and 3PAR StoreServ storage arrays is not supported.

3PAR StoreServ Operating System/HBA/Array support can be found on SPOCK: https://spock.corp.hp.com/SPOCK/index.aspx

EMC Operating System/HBA/Array support can be found on EMC support website: https://support.emc.com

Summary

From the tests with an Oracle RAC database using ASM disk groups, we were able to conclude that placing an ASM failure group on an HPE 3PAR StoreServ AFA disk array will significantly accelerate Oracle database performance. This increase in performance was demonstrated with the 3PAR StoreServ AFA configured alongside a legacy storage array (EMC VMAX) without requiring a major upgrade or replacement.

Using the SLOB benchmark, an increase of 75 percent in performance by using the HPE 3PAR StoreServ AFA unit was reported by Oracle RAC AWR reports, HPE 3PAR System Reporter and the HPE 3PAR CLI statvlun command. All three reporting tools showed a database read access time of below 300 microseconds, a 6X improvement over the baselines of the EMC VMAX with the additional failure groups on the HPE 3PAR StoreServ AFA.

The Oracle RAC write performance benefited as well from the presence of the HPE 3PAR StoreServ AFA: physical write service times were lowered to giving 4X better from Oracle RAC baseline with EMC VMAX and a 6X improvement over the Oracle RAC baseline.

The HPE 3PAR StoreServ AFA was only marginally used in the SLOB testing: only 100,775 IOPS were used during the tests while this type of array is capable of many times these IOPS with Oracle databases. See the “Realize extreme performance and high availability with 3PAR StoreServ 7450 All-Flash for Oracle” white paper in the “For more information” section of this white paper for more details.

An HPE 3PAR StoreServ AFA has a number of great built-in technology features that work well with Oracle databases, like HPE 3PAR Priority Optimization or quality of service (QoS). This allows the user to set performance limits for IOPS or bandwidth, and a latency goal. QoS is often used on a consolidating HPE 3PAR StoreServ system where the dev/test group is allowed to use 10 percent of the resources on the array granting the other 90 percent to the production database. HPE 3PAR StoreServ AFAs also support Virtual Copy snapshots and Full Copy of volumes. As described earlier, the TCO analysis showed 3PAR All-Flash Virtual Copy snapshots provide a 50 percent boost in application developer productivity compared to slower-performing and less easily accessible copies off of a VMAX array. Using these features will enhance database operations significantly without impacting performance.

The conclusion of the research in this paper is that when mission-critical applications running on Oracle databases with ASM request more performance, adding an HPE 3PAR StoreServ AFA to store failure groups with Preferred Reads, major upgrades on legacy storage arrays or the acquisition of a new array can be avoided. The failure group added to spinning disk Oracle ASM environments significantly increases the performance of Oracle non-RAC or Oracle RAC databases. The HPE 3PAR StoreServ AFA used for this purpose can be added easily to the environment using standard OS MPIO drivers, and it is fully supported by HPE.

The addition of HPE 3PAR StoreServ AFA can offer great benefits for the Oracle development community. The AFA offers the ability to serve up near-instant copies of the full database for application developers using 3PAR Virtual Copy snapshots. These copies are full copies that perform at extremely high speed compared to those of a legacy tiered array. Also, copies do not need to be shared due to the ample capacity of the 3PAR AFA. The result is significantly improved productivity for the application development community.

Lastly, the addition of the HPE 3PAR StoreServ can be done at a fraction of the cost of a legacy storage upgrade.

For more information

- Realize extreme performance and high availability with the HPE Universal Database Solution with 3PAR StoreServ 7450 All-Flash for Oracle: http://www8.hp.com/h20195/V2/GetDocument.aspx?docname=4AA5-7457ENW&cc=us&lc=en
- HPE 3PAR Command Line Interface Reference: h20565.www2.hpe.com/hpsc/doc/public/display?docId=emr_na-c03618134-3
- HPE 3PAR Recovery Manager Software for Oracle (RM-O): h20392.www2.hpe.com/portal/swdepot/displayProductInfo.do?productNumber=HP3PARRMO
- h20565.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=5383441&docId=emr_na-c03817217&docLocale=en_US
- HPE 3PAR Virtual Copy: www8.hp.com/us/en/products/storage-software/product-detail.html?oid=5044626

To help us improve our documents, please provide feedback at hp.com/solutions/feedback

Appendix A: Solution architecture overview

Platform details

- EMC VMAX
- Oracle 11g R2
- Oracle ASM
- HPE 3PAR StoreServ 7450c AFA
- Red Hat Enterprise Linux 6.6
- HPE ProLiant DL580 Server host, 32 GB DRAM

SLOB benchmark test configuration

3PAR AFA OS version 3.2.1

Test tool: SLOB

Oracle Parameter changed so that SLOB run will generate most physical IOs.

```
alter system set "_db_block_prefetch_limit"=0 scope=spfile sid='*'; alter system set "_db_block_prefetch_quota"=0 scope=spfile sid='*';

alter system set "_db_file_noncontig_mblock_read_count"=0 scope=spfile sid='*'; alter system set "cpu_count"=1 scope=spfile sid='*';

alter system set "db_cache_size"=4m scope=spfile sid='*'; alter system set "shared_pool_size"=500m scope=spfile sid='*'; alter system set "sga_target"=0 scope=spfile sid='*';
```

Primary database setup

Primary DB setup involves setting up the storage volumes and configuration of disk multipath software, marking storage volumes for ASM usage, and creating an ASM disk group followed by creating a sample Oracle database. It's assumed that Grid/Oracle binary setup has already been done on the source and target systems.

Appendix B: Oracle RAC database test results

The following AWL report data shows test results from running SLOB benchmark utility. The tests were conducted over a two-hour test run and monitored on an Oracle RAC Database.

Oracle RAC with 3PAR AFA constrained with VMAX

The following test results show increased benefits of using Hewlett Packard Enterprise 3PAR AFA with simulated workload on VMAX constrained for both arrays.

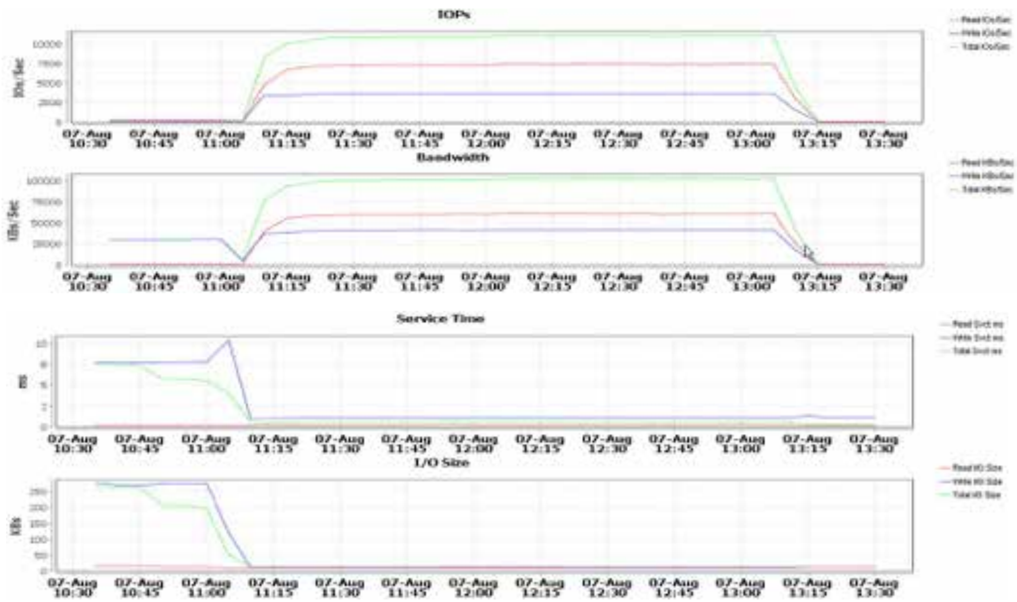

Figure 4: System reporter performance reports VLUNs

Table 6: Oracle automatic workload repository report

DB NAME	DB ID	INSTANCE	INST NUM	STARTUP TIME	RELEASE	RAC
TestDB	260599568	testdb1	1	05-Aug 15:08	11.2.0.4.0	YES

HOST NAME	PLATFORM	CPUS	CORES	SOCKETS	MEMORY (GB)
dl380-rac1.hp.superna.net	Linux x86 64-bit	32	16	2	220.68

	SNAP ID	SNAP TIME	SESSIONS	CURSORS/SESSIONS	INSTANCES
Begin Snap:	1918	07-Aug-15 11:05:11	76	2.3	2
End Snap:	1921	07-Aug-15 13:07:04	49	2.2	2
Elapsed:		121.90 (mins)			
DB Time:		1,953.08 (mins)			

	PER SECOND	PER TRANSACTION	PER EXEC	PER CALL
DB Times (s):	16.0	2.4	0.52	1.34
DB CPU(s):	1.0	0.2	0.03	0.08
Redo size (bytes):	9,345,994.7	1,397,768.6		
Logical read (blocks):	9,971.2	1,491.3		
Block changes:	5,618.2	840.3		
Physical read (blocks):	7,152.9	1,069.8		
Physical write (blocks):	3,803.0	568.8		
Read IO requests:	7,152.9	1,069.8		
Write IO requests:	3,737.2	558.9		
Read IO (MB):	55.9	8.4		
Write IO (MB):	29.7	4.4		
Global Cache blocks received:	1.1	0.2		
Global Cache blocks served:	0.4	0.1		
User calls:	12.0	1.8		
Parses (SQL):	3.2	0.5		
Hard parses (SQL):	0.0	0.0		
SQL Work Area (MB):	0.8	0.1		
Logons:	1.6	0.2		
Executes (SQL):	31.1	4.7		
Rollbacks:	0.0	0.0		
Transactions:	6.7			

Report summary

Table 7: Load Profile Instance efficiency percentages (Target 100%)

Buffer nowait%:	99.91	**RedoNoWait%:**	99.67
Buffer Hit%:	28.27	**In-memory Sort%:**	100.00
Library Hit%:	99.68	**SoftParse%:**	98.65
Execute to Parse%:	89.63	**Latch Hit%:**	98.48
Parse CPU to Parse Elapsed%:	37.03	**% Non-Parse CPU:**	99.97

Table 8: Top 10 foreground events by total wait time

EVENT	WAITS	TOTAL WAIT TIME (SEC)	WAIT AVG (MS)	% DB TIME	WAIT CLASS
db file sequential read	52,254,698	78.2K	1	66.7	User I/O
free buffer waits	9,845,598	27.9K	3	6.1	Configuration
DB CPU		7095.8		6.1	
write complete waits	121,270	5258.4	43	4.5	Configuration
log file switch completion	8,774	864.3	99	.7	Configuration
gc cr grant 2-way	323,108	110.8	0	.1	Cluster
latch: cache buffers lru chain	1,493,307	70.5	0	.1	Other
read by other session	51,506	50	1	.0	User I/O
waitlist latch free	29,862	32.4	1	.0	Other
Disk file Mirror Read	3,341	32.1	10	.0	User I/O

Table 9: Wait classes by total wait time

EVENT	WAITS	TOTAL WAIT TIME (SEC)	WAIT AVG (MS)	% DB TIME	AVG ACTIVE SESSIONS
User I/O	52,378,106	78,382	1	66.9	10.7
Configuration	10,014,393	34.055	3	29.1	4.7
System I/O	6,347,638	12,314	2	10.5	1.7
DB CPU		7,096		6.1	1.0
Other	2,906,636	229	0	.2	0.0
Cluster	355,190	124	0	.1	0.0
Concurrency	440,372	24	0	.0	0.0
Commit	443	10	22	.0	0.0
Application	2,894	1	0	.0	0.0
Network	33,883	0	0	.0	0.0

Table 10: Host CPU

CPUS	CORES	SOCKETS	LOAD AVERAGE BEGIN	LOAD AVERAGE END	%USER
32	16	2	0.06	0.34	2.8

Table 11: Instance CPU

% TOTAL CPU	% BUSY CPU	% DB TIME WAITING FOR CPU (RESOURCE MANAGER)
4.0	77.9	0.0

	READ+WRITE PER SECOND	READ PER SECOND	WRITE PER SECOND
Total Requests:	10,960.3	7,162.4	3,797.9
Database Requests:	10,890.1	1,152.9	3,373.2
Optimized Requests:	0.0	0.0	0.0
Redo Requests:	55.1	0.4	54.7
Total (MB):	95.0	56.0	39.0
Database (MB):	85.6	55.9	29.7
Optimized Total (MB):	0.0	0.0	0.0
Redo (MB):	9.2	0.0	9.2
Database (blocks):	10,955.9	7,152.9	3,803.0
Via Buffer Cache (blocks):	10,952.4	7,149.5	3,802.9
Direct (blocks):	3.5	3.4	0.1

IO Profile

Table 12: Memory Statistics

	BEGIN	END
HOST MEM (MB):	225,976.8	225,976.8
SGA use (MB):	513.7	513.7
PGA use (MB):	292.5	236.2
% Host Mem used for SGA + PGA:	0.36	0.33

Table 13: Cache sizes

	BEGIN	END		
Buffer Cache:	4M	4M	**Std Block Size:**	8K
Shared Pool Size:	389M	391M	**Log Buffer:**	3,656K

Table 14: Shared pool statistics

	HOST NAME	PLATFORM	CPUS	CORES	SOCKETS	MEMORY (GB)
	dl380-rac1.hp.superna.net		32	16	2	220.68
Memory Usage %:		91.58		91.03		
% SQL with executions >1:		94.97		94.98		
% Memory for SQL w/exec >1:		83.78		83.96		

Oracle RAC with VMAX constrained with 3PAR AFA

The following test results show increased benefits of using HPE 3PAR AFA with simulated 6 ms workload on VMAX for both arrays.

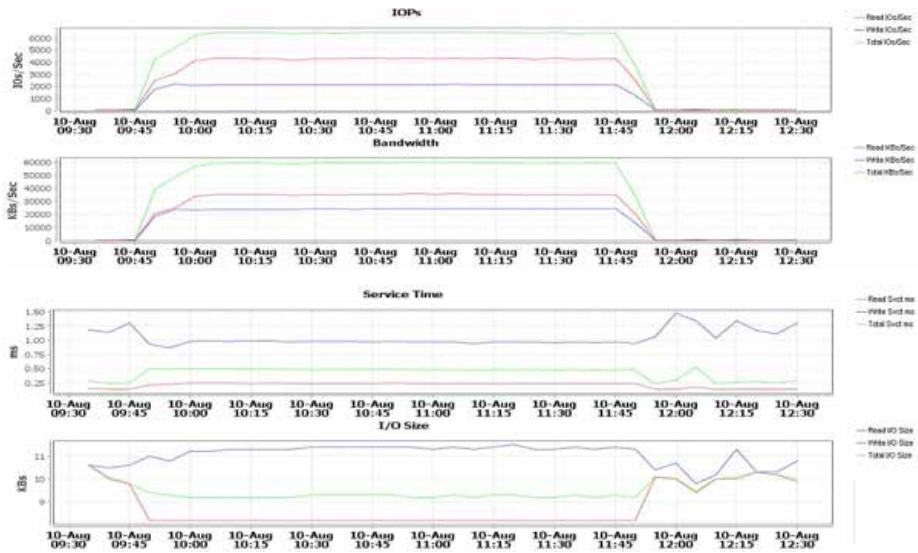

Figure 5: System reporter performance reports VLUNs

Table 15: Workload repository report

DB NAME	DB ID	INSTANCE	INST NUM	STARTUP TIME	RELEASE
TEST DB	2660599568	Testdb1	1	05-Aug-15 15:08	11.2.0.4

	SNAP ID	SNAP TIME	SESSIONS	CURSORS/SESSIONS	INSTANCES
Begin Snap:	1992	10-Aug-15 09:46:00	62	1.9	2
End Snap:	1994	10-Aug-15 11:47:46	48	1.8	2
Elapsed:		121.77 (mins)			
DB Time:		1,955.70 (mins)			

Report Summary

Table 16: Load Profile

	PER SECOND	PER TRANSACTION	PER EXEC	PER CALL
DB TIME(s):	16.1	4.1	0.76	1.51
DB CPU(s):	0.8	0.2	0.04	0.07
Redo size (bytes):	5,454,509.4	1,379,031.9		
Logical read (blocks):	5,904.0	1,492.7		
Block changes:	3,321.3	839.7		

Physical Read (blocks):	4,159.7	1,051.7
Physical write (blocks):	2,236.5	565.4
Read IO requests:	4,159.7	1,051.7
Write IO requests:	2,206.0	557.7
Read IO (MB):	32.5	8.2
Write IO (MB):	17.5	4.4
Global Cache blocks received:	2.4	0.6
Global Cache blocks served:	0.3	0.1
User calls:	10.7	2.7
Parses (SQL):	4.0	1.0
Hard parses (SQL):	0.1	0.0
SQL Work Area (MB):	0.8	0.2
Logons:	1.4	0.4
Executes (SQL):	21.0	5.3
Rollbacks:	0.0	0.0
Transactions:	4.0	

Table 17: Instance efficiency percentages (target 100%)

Buffer Nowait %:	99.95	**RedoNoWait %:**	99.61
Buffer Hit %:	29.59	**In-memory Sort %:**	100.00
Library Hit %:	99.02	**Soft Parse %:**	98.31
Execute to Parse %:	81.08	**Latch Hit %”**	98.08
Parse CPU to Parse Elapsed %:	40.26	**% Non-Parse CPU:**	99.90

Table 18: Top 10 foreground events by total wait time

EVENT	WAITS	TOTAL WAIT TIME (SEC)	WAIT AVG (MS)	% DB TIME	WAIT CLASS
free buffer waits	17,177,870	66.2K	4	56.4	Configuration
db file sequential read	30,305,341	32.6K	1	27.8	User I/O
write complete waits	185,379	14.2K	76	12.1	Configuration
DB CPU		5450.7		4.6	
log file switch completion	3,635	624.1	172	.5	Configuration
gc cr grant2-way	195,578	68.6	0	.1	Cluster
latch: cache buffers lru chain	1,283,757	57.8	0	.0	Other
Disk file Mirror Read	2,705	50.1	19	.0	Other
log file sync	855	40.9	48	.0	Commit
waitlist latch free	32,023	34.7	1	.0	Other

WAIT CLASS	WAITS	TOTAL WAIT TIME (SEC)	WAIT AVG (MS)	% DB TIME	AVG ACTIVE SESSIONS
Configuration	17,388,928	81,099	5	69.1	11.1
User I/O	30,398,359	32,786	1	27.9	4.5
System I/O	3,458,363	13,258	4	11.3	1.8
DB CPU		5,451		4.6	0.7
Other	2,592,130	286	0	.2	0.0
Cluster	228,583	85	0	.1	0.0
Commit	881	42	48	.0	0.0
Concurrency	343,017	31	0	.0	0.0
Application	2,925	1	0	.0	0.0
Network	30,150	0	0	.0	0.0

Wait classes by total wait time

Table 19: Load Profile

CPUS	CORES	SOCKETS	LOAD AVERAGE BEGIN	LOAD AVERAGE END	% USER	% SYSTEM
32	16	2	0.11	0.19	2.3	1.5

Table 20: Instance CPU

% TOTAL CPU	% BUSY CPU	% DB TIME WAITING FOR CPU (RESOURCE MANAGER)
3.0	72.9	0.0

Table 21: IO Profile

	READ + WRITE PER SECOND	READ PER SECOND	WRITE PER SECOND
Total Requests:	6,406.6	4,167.3	2,239.2
Database Requests:	6,365.7	4,159.7	2,206.0
Optimized Requests:	0.0	0.0	0.0
Redo Requests:	29.7	0.2	29.4
Total (MB):	55.5	32.6	22.9
Database (MB):	50.0	32.5	17.5
Optimized Total (MB):	0.0	0.0	0.0
Redo (MB):	5.4	0.0	5.4
Database (blocks):	6,396.2	4,159.7	2,236.5
Via Buffer Cache (blocks):	6,387.0	4,150.6	2,236.4
Direct (blocks):	9.1	9.1	0.1

Table 22: Memory statistics

	BEGIN	END
Host Mem (MB):	225,976.8	225,976.8
SGA use (MB):	513.7	513.7
PGA use (MB):	236.8	227.3
% Host Mem used for SGA+PGA:	0.33	0.33

11:04:22 08/10/2015		r/w	I/O per second			KBytes per sec			Svt ms		IOSz KB		
Port	D/C		Cur	Avg	Max	Cur	Avg	Max	Cur	Avg	Cur	Avg	Qlen
2:2:2	Data	r	4409	3215	9833	36176	26427	80619	0.24	0.24	8.2	8.2	-
2:2:2	Data	w	2185	1646	2457	23866	18574	27055	0.87	0.97	10.9	11.3	-
2:2:2	Data	t	6594	4861	11625	60041	45001	96303	0.45	0.49	9.1	9.3	1
1	Data	r	4409	3215		36176	26427		0.24	0.24	8.2	8.2	-
1	Data	w	2185	1646		23866	18574		0.87	0.97	10.9	11.3	-
1	Data	t	6594	4861		60041	45001		0.45	0.49	9.1	9.3	1

Figure 6: 3PAR CLI STAT VLUN Command taken with 6ms constrained legacy array helping with preferred reads

11:40:33 08/07/2015		r/w	I/O per second			KBytes per sec			Svt ms		IOSz KB		
Port	D/C		Cur	Avg	Max	Cur	Avg	Max	Cur	Avg	Cur	Avg	Qlen
2:2:2	Data	r	7285	5799	10237	59784	47596	83907	0.27	0.27	8.2	8.2	-
2:2:2	Data	w	3607	3032	4021	40970	35388	44465	1.15	1.12	11.4	11.7	-
2:2:2	Data	t	10892	8831	12975	100755	82984	113713	0.56	0.56	9.2	9.4	5
1	Data	r	7285	5799		59784	47596		0.27	0.27	8.2	8.2	-
1	Data	w	3607	3032		40970	35388		1.15	1.12	11.4	11.7	-
1	Data	t	10892	8831		100755	82984		0.56	0.56	9.2	9.4	5

Figure 7: 3PAR CLI STAT VLUN Command taken 4ms with constrained legacy array helping with preferred reads

Learn more at
hpe.com/info/3par

Sign up for updates

★ Rate this document

Hewlett Packard Enterprise

4AA6-3375ENW, December 2015