Dell EqualLogic Best Practices Series

Performance Baseline for Deploying Oracle® 11*g* Release 2 Based Decision Support Systems using Dell™ EqualLogic™ PS6110XV Arrays

A Supplemental Test Report

Base White paper:
Sizing and Best Practices for Deploying Oracle 11g Release 2 Based Decision Support Systems using Dell EqualLogic PS6110XV Arrays
A Dell Technical White paper
http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/20059328/download.aspx

Storage Infrastructure and Solutions Engineering
Dell Product Group
May 2012

Table of Contents

Acknowledgements

This whitepaper was produced by the PG Storage Infrastructure and Solutions of Dell Inc.

The team that created this whitepaper:

Chidambara Shashikiran, Chhandomay Mandal, and Camille Daily

We would like to thank the following Dell team members for providing significant support during development and review:

Ananda Sankaran and Suresh Jasrasaria

Feedback

We encourage readers of this publication to provide feedback on the quality and usefulness of this information by sending an email to SISfeedback@Dell.com.

SISfeedback@Dell.com

1 Introduction

This paper provides a summary of results for tests executed using EqualLogic PS6110XV arrays. These tests were set up the same as and run in addition to the set of test scenarios described in the Dell white paper titled, *Sizing and Best Practices for Deploying Oracle 11g Release 2 Based Decision Support Systems with Dell EqualLogic 10GbE iSCSI SAN*. All tests executed as a part of that paper were executed using PS6010XV arrays. The key objectives addressed were:

- Determine the best practices and sizing guidelines for deploying Oracle based Decision Support Systems (DSS) applications.
- Determine the I/O throughput of the storage using different RAID configurations and array models via I/O simulation.
- Determine the scalability of the I/O throughput via I/O simulation as storage arrays are added.
- Determine the scalability of the I/O throughput via DSS application simulation as storage arrays are added, while ensuring that the overall configuration is balanced with no resource bottlenecks existing on server.

A series of tests were conducted and based on those test results, best practices were validated and scalability guidelines have been noted for designing storage architectures that are best suited to support the storage demands of DSS applications running on Oracle 11g Release 2 databases.

The complete paper covering these initial test results, sizing guidelines and best practices is located on Dell TechCenter at http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/20059328/download.aspx

The key objectives of the additional tests presented in this paper are:

- Determine performance baseline for Oracle 11g Decision Support Systems (DSS) using PS6110XV arrays.
- Compare the performance of Oracle 11g DSS workload on PS6110XV arrays compared to PS6010XV.
- Evaluate the performance scalability for DSS type of workload when additional PS6110XV array members are added.
- Compare RAID 10 versus RAID50 performance using PS6110XV arrays.
- Evaluate benefits of PS6110XV arrays compared to PS6010XV arrays.

1.1 Audience

This paper is for solution architects, storage network engineers, system administrators, and IT managers who are interested in deploying the Oracle 11*g* Release 2 based DSS solution using EqualLogic PS6110XV arrays.

Readers should first refer to the *Sizing and Best Practices for Deploying Oracle 11g Release 2 Based Decision Support Systems with Dell EqualLogic 10GbE iSCSI SAN* white paper mentioned above, and then refer to this paper for additional test results using EqualLogic PS6110XV arrays.

2 DSS I/O simulation

2.1 Test methodology

A subset of I/O simulation tests were executed with large sized I/O block requests simulating a DSS database I/O to understand the read and write I/O throughput rate offered by the EqualLogic PS6110XV storage arrays.

2.2 Test scenarios

The main objectives of these tests were to determine the difference in performance of PS6110XV arrays compared to the older generation PS6010XV arrays. The performance numbers for PS6010XV arrays are taken directly from the Oracle 11g DSS white paper mentioned in the Introduction. The following test scenarios were completed using the ORION (Oracle I/O Calibration Tool) on a PS6110XV configuration.

One PS6110XV and RAID 10

- Sequential read I/O in concat mode and 256KB I/O block size
- Sequential read I/O in RAID 0 mode and 1MB I/O block size
- Sequential write I/O in RAID 0 mode and 1MB I/O block size
- Large random read I/O in RAID 0 mode and 1MB I/O block size

One PS6110XV and RAID 50

- Sequential read I/O in concat mode and 256KB I/O block size
- Sequential read I/O in RAID 0 mode and 1MB I/O block size
- Sequential write I/O in RAID 0 mode and 1MB I/O block size
- Large random read I/O in RAID 0 mode and 1MB I/O block size

Two PS6110XV and RAID 50

- Sequential read I/O in concat mode and 256KB I/O block size
- Sequential read I/O in RAID 0 mode and 1MB I/O block size
- Sequential write I/O in RAID 0 mode and 1MB I/O block size
- Large random read I/O in RAID 0 mode and 1MB I/O block size

2.3 Test studies and results

This section describes the ORION tests conducted and the results observed. The same test parameters used for the ORION tests on PS6010XV arrays were used for executing the ORION tests on PS6110XV arrays.

All tests were executed using RAID 0 first to simulate Oracle ASM behavior. The tests were repeated again using the default concat option to determine the performance impact due to ASM striping.

2.3.1 PS6010XV and PS6110XV – RAID 10

The four ORION I/O workload scenarios described earlier were executed on a PS6110XV array and the results were compared with a PS6010XV array. The storage arrays were configured with RAID 10 for all these scenarios.

The following chart compares the throughput produced by PS60110XV and PS6110XV arrays in RAID 10 configuration.

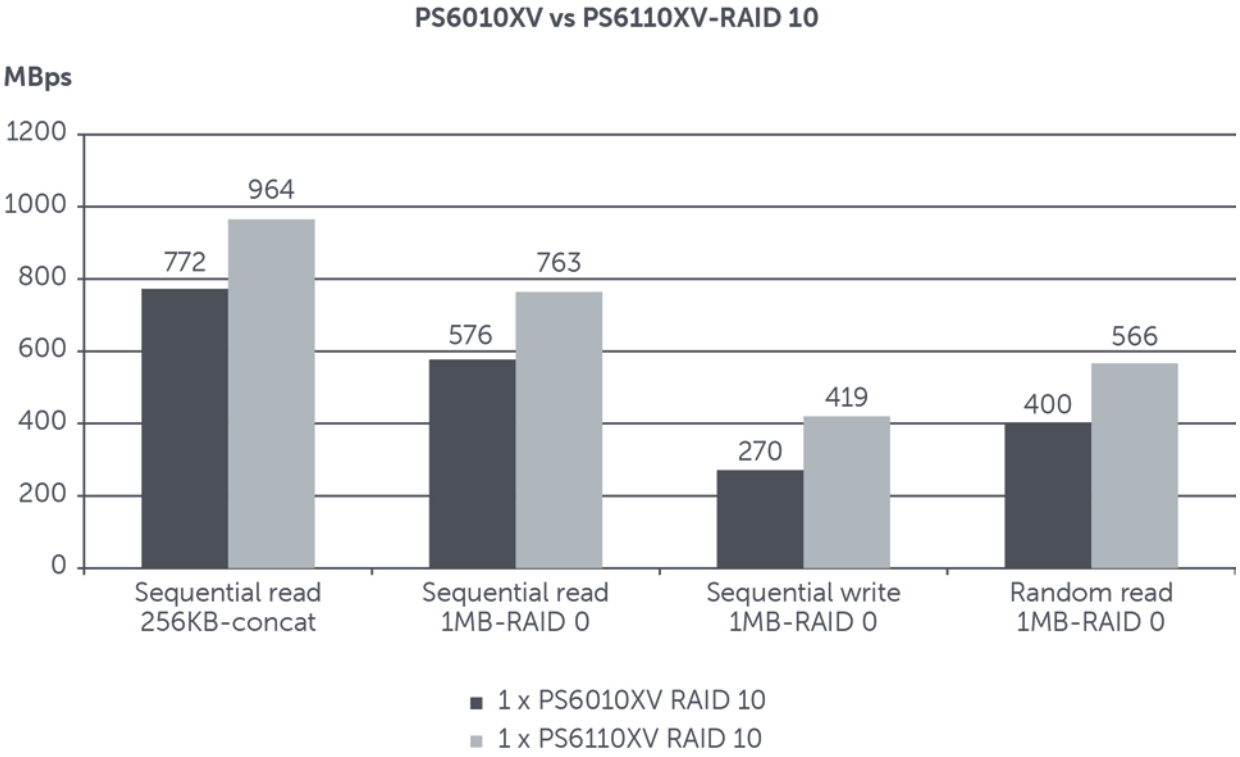

Figure 1 ORION I/O throughput in RAID10 configuration: PS6010XV versus PS6110XV

PS6110XV arrays provided better throughput than the PS6010XV arrays as shown in the above chart.

The behavior can primarily be attributed to the increased number of disks hosting active data on PS6110XV arrays. PS6110XV arrays have 24 disks (22 active and 2 hot spares) while PS6010XV arrays have only 16 disks (14 active and 2 hot spares).

The following table shows the percentage of increase in throughput for different I/O workloads in RAID 10 configuration.

I/O Workload	1 x PS6010XV RAID 10	1 x PS6110XV RAID 10	% Increase
Sequential read: 256KB - concat	772.36	964.92	24.93
Sequential read: 1MB - RAID 0	576.86	763.93	32.42
Sequential write: 1MB -RAID 0	270.31	419.45	55.17
Random read: 1MB - RAID 0	400.6	566.56	41.42

2.3.2 PS6010XV and PS6110XV – RAID 50

The four ORION I/O workload scenarios described earlier were executed on a PS6110XV array and the results were compared with a PS6010XV array. The storage arrays were configured with RAID 50 for all scenarios.

The following chart compares the throughput produced by PS60110XV and PS6110XV arrays in RAID 50 configuration.

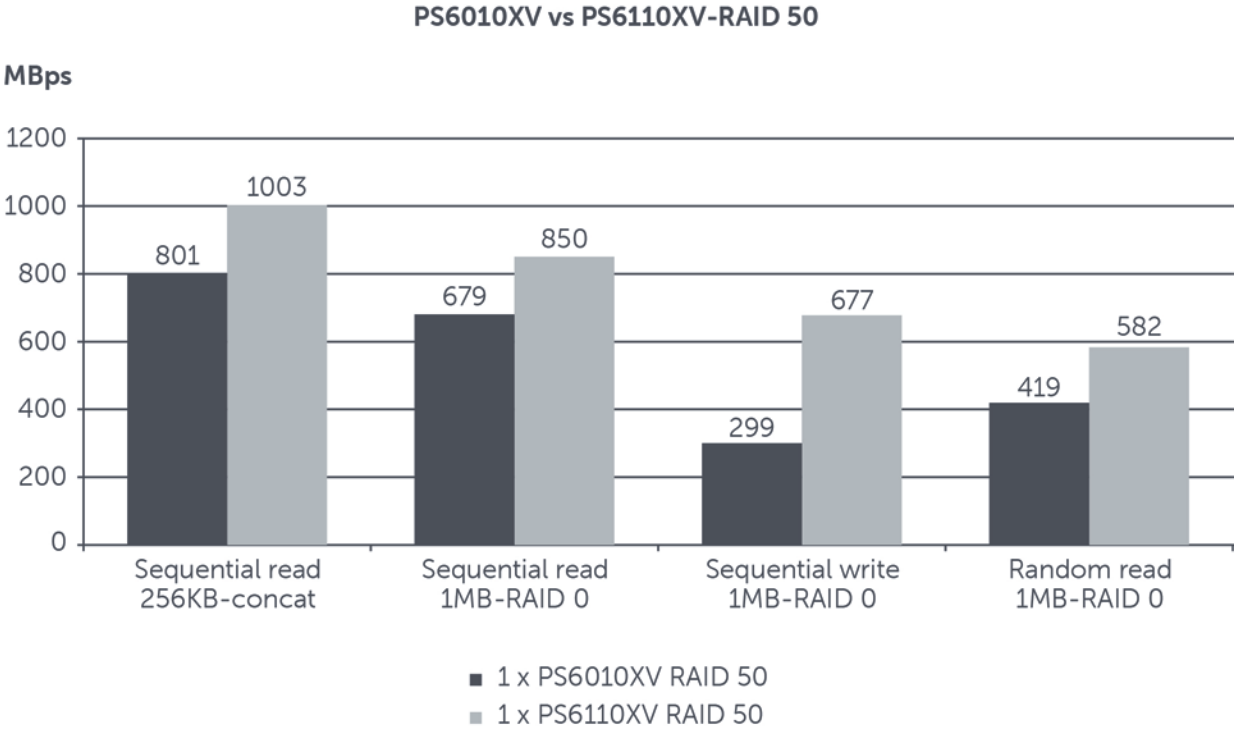

Figure 2 ORION I/O throughput in RAID 50 configuration: PS6010XV versus PS6110XV

A significant increase in throughput (125%) was observed for the 100% sequential write I/O workload. The behavior can primarily be attributed to the increased number of disks hosting active data on PS6110XV arrays. The sequential write I/O operations got significant performance benefits due to the additional eight disks hosting active data on PS6110XV arrays.

The following table shows the percentage of increase in throughput for different I/O workloads in RAID 50 configuration.

I/O Workload	1 x PS6010XV RAID 50	1 x PS6110XV RAID 50	% Increase
Sequential read: 256KB - concat	801.82	1003.35	24.13
Sequential read: 1MB - RAID 0	679.97	850.17	25.03
Sequential write: 1MB -RAID 0	299.75	677.41	125.99
Random read: 1MB - RAID 0	419.3	582.81	38.99

2.3.3 PS6110XV – RAID 10 versus RAID 50

The four ORION workloads described earlier were executed on a PS6110XV array in RAID 10 and in RAID 50. The I/O throughout reported by ORION for both configurations are summarized in the chart below.

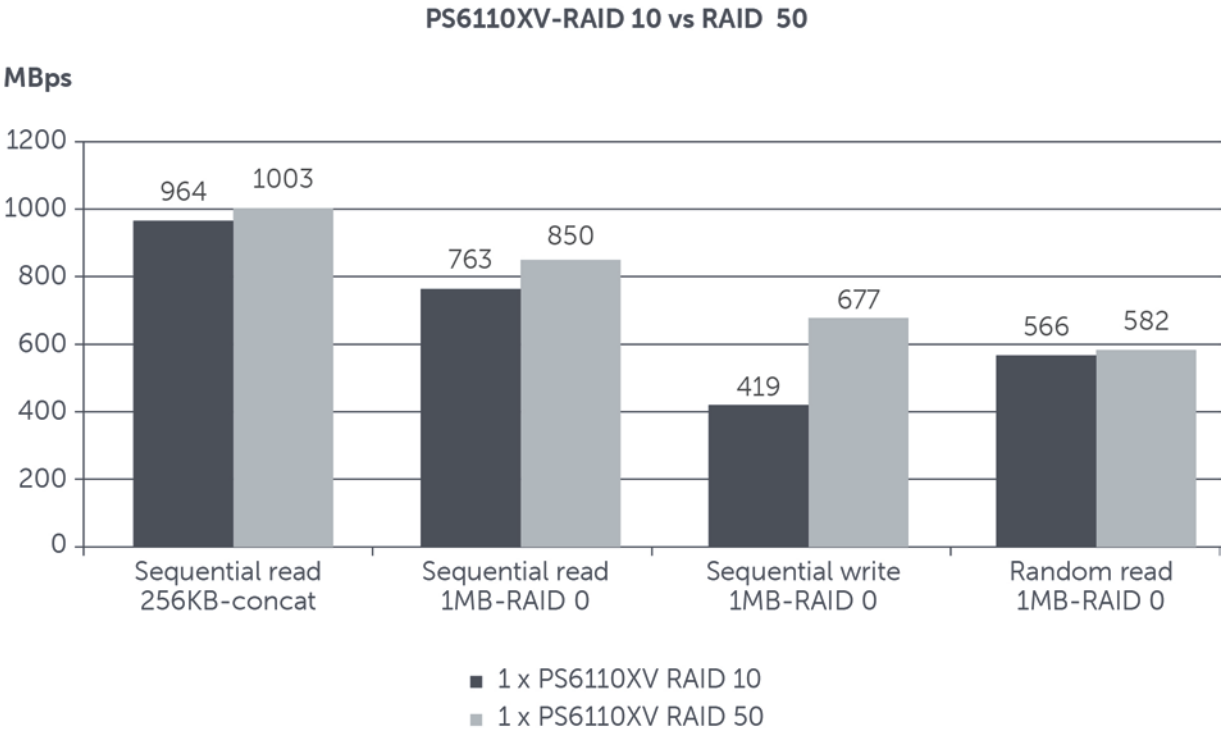

Figure 3 ORION I/O throughput for PS6110XV: RAID 10 versus RAID 50

The RAID 50 configuration provided a slightly better throughput than the RAID 10 configuration. This behavior can be primarily attributed to respective RAID implementation differences, such as more disks hosting active data than redundant data in RAID 50, which can improve read I/O performance.

This test shows an increase of more than 60% in throughput for 100% sequential write operations.

The following table shows the percentage of increase in throughput for different I/O workloads in RAID 50 as compared to RAID 10.

I/O Workload	1 x PS6110XV RAID 10	1 x PS6110XV RAID 50	% Increase
Sequential read: 256KB - concat	964.92	1003.35	3.8
Sequential read: 1MB - RAID 0	763.93	850.17	11.28
Sequential write: 1MB -RAID 0	419.45	677.41	61.49
Random read: 1MB - RAID 0	566.56	582.81	0.03

2.3.4 PS6110XV scaling studies

The four ORION workloads described in section 2.2 were executed on PS6110XV arrays using RAID 50. The storage pool was configured with one array initially and then scaled to two arrays in the same storage pool. Test volumes were recreated and mounted on a server for each scaled configuration. The I/O throughput reported by ORION across the scaled configurations is shown in Figure 4.

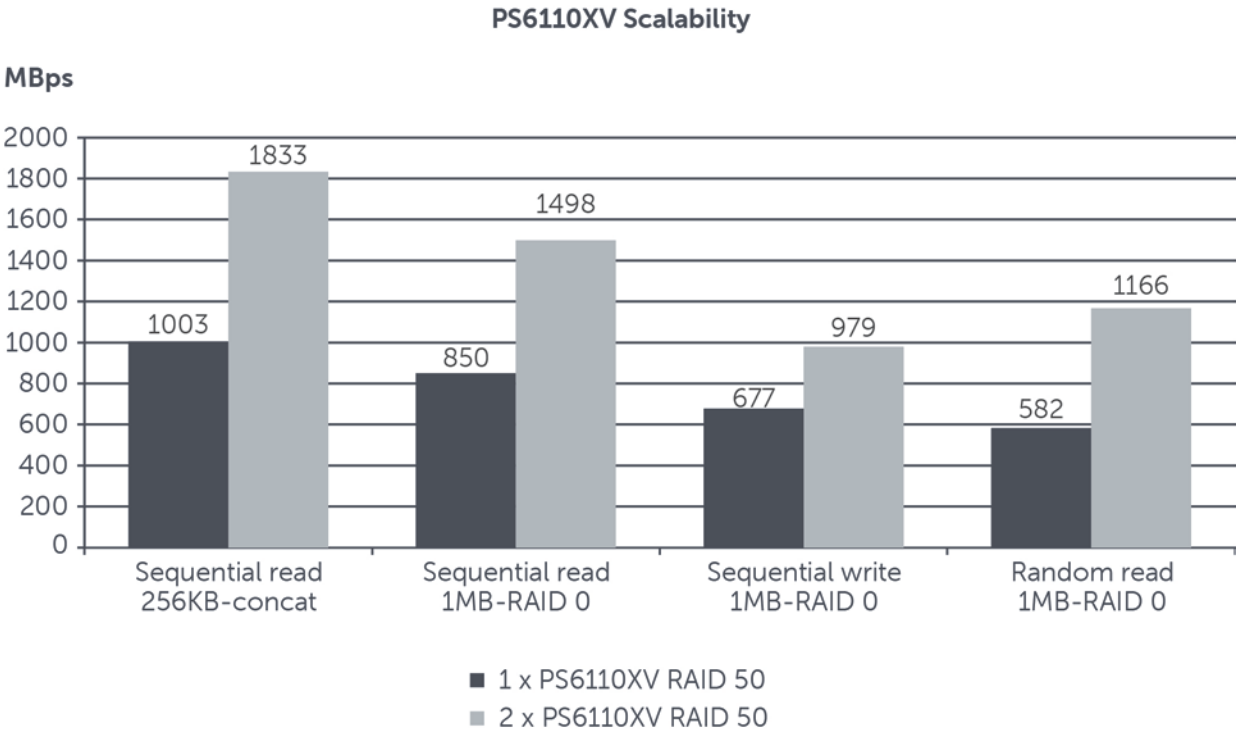

Figure 4 ORION I/O throughput: PS6110XV Scalability

The following table shows the percentage of increase in throughput for different I/O workloads.

I/O Workload	1 x PS6110XV RAID 50	2 x PS6110XV RAID 50	% Increase
Sequential read: 256KB - concat	1003.65	1833.43	82.67
Sequential read: 1MB - RAID 0	850.17	1498	76.20
Sequential write: 1MB -RAID 0	677.41	979.3	44.56
Random read: 1MB - RAID 0	582.81	1166.86	100.21

With sequential reads and sequential writes in RAID 0, the throughput did not increase linearly as array members were added. This is because the ORION tool RAID 0 option introduced additional randomness to the sequential I/O workloads while simulating Oracle ASM striping behavior. As a result, the throughput did not scale linearly. But the large random read I/O operations using 1MB block size scaled as expected and more than 100% scaling in the performance was observed. This shows that introduction of randomness to the already random I/O workload did not have any impact to the performance.

To confirm that the ORION tool RAID 0 option is introducing randomness while simulating ASM behavior and resulting in a decrease in performance, the same set of tests was run a second time using the ORION tool with the default

concat option. Executing the ORION tool with this option is like running any other I/O workload simulator like VDBench or IOMeter.

The following chart summarizes the performance scalability results using the ORION tool concat option.

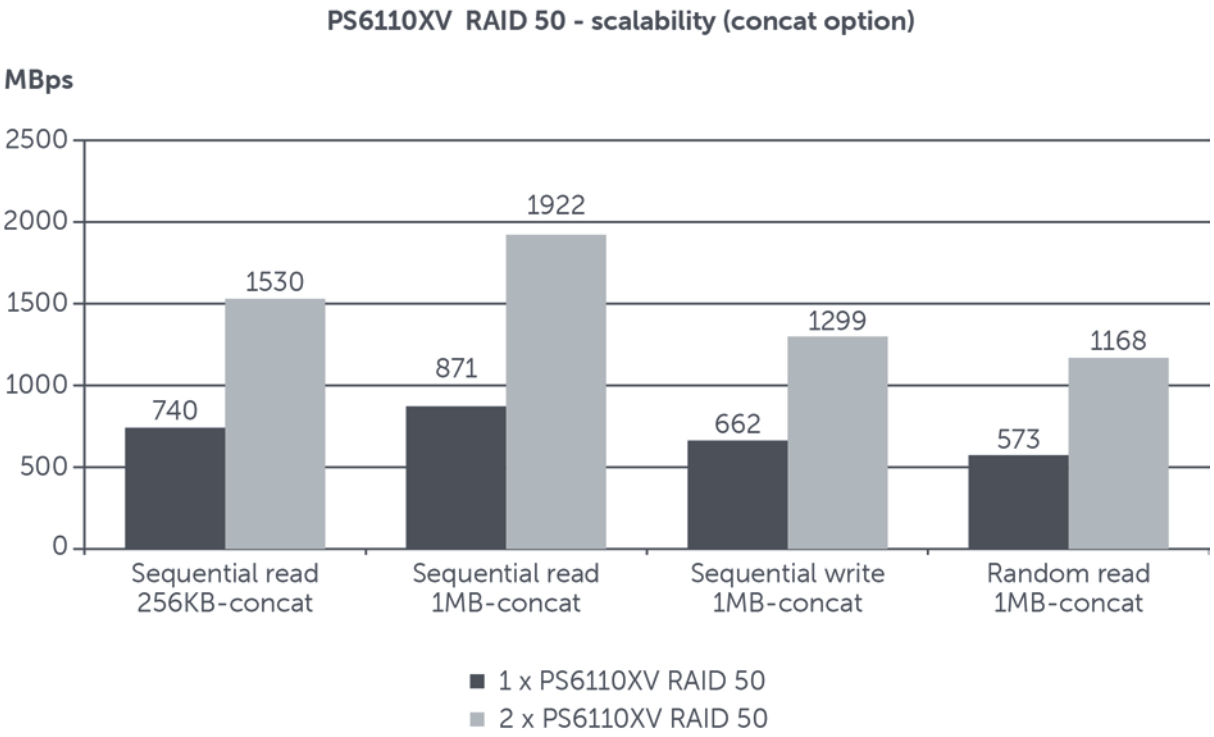

Figure 5 ORION I/O throughput on PS6110XV Scalability: Using ORION concat option

The following table shows the percentage of increase in throughput for different I/O workloads when ORION was executed using the default concat option.

I/O Workload	1 x PS6110XV RAID 50	2 x PS6110XV RAID 50	% Increase
Sequential read: 256KB – concat	1003.65	1833.43	82.67
Sequential read: 1MB - concat	871.92	1922.09	120.44
Sequential write: 1MB -concat	662.89	1299.70	96.06
Random read: 1MB - concat	573.11	1168.9	103.95

As expected, an almost linear scaling in performance was observed when the default concat option was used with ORION. The random read throughput can be considered the I/O baseline for DSS application tests conducted with an Oracle database installed.

3 DSS application simulation

3.1 Simulation tool and workload definition

Additional DSS tests were conducted to understand the throughput behavior of the PS6110XV arrays when the Oracle Real Application Cluster (RAC) database executed DSS queries.

The following additional DSS query simulation tests (TPC-H), using an Oracle database, were executed on the storage configurations listed below.

- One RAC node, one PS6110XV array, and RAID 50
- One RAC node, two PS6110XV arrays, and RAID 50

> **Note:** Additional tests listed below were executed using a two RAC node configuration as well to evaluate the performance impact when additional database nodes were configured.

- Two RAC nodes, one PS6110XV array, and RAID 50
- Two RAC nodes, two PS6110XV arrays, and RAID 50

3.2 Test studies and results

3.2.1 PS6110XV scaling studies for Oracle RAC DSS Database

First, the DSS query simulation was executed on one array in a single storage pool. The number of arrays was then increased to two within the same storage pool and the same DSS test was rerun. The tests were executed using both one node and two node Oracle RAC configurations with one and two arrays.

3.2.1.1 Storage I/O throughput

The storage I/O throughput measured during the DSS tests with one to two arrays is shown in Figure 6 below. The storage array throughput from the PS6110XV arrays is compared with PS6010XV arrays in the same chart.

The I/O profile of the database during DSS tests was large block sized I/O read requests which were random in nature. Both the maximum and average throughput measured is shown. The throughput was measured as the rate of data sent by the array NICs in MBps.

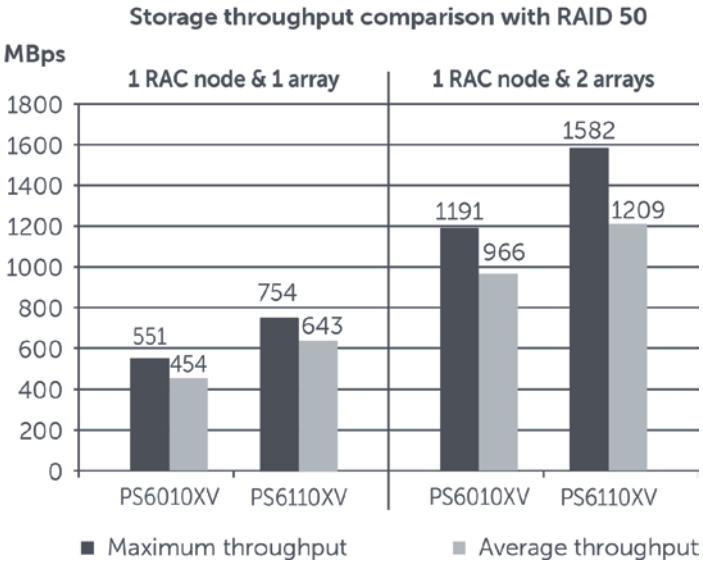

Figure 6 Storage I/O throughput – PS6010XV versus PS6110XV scalability

As seen in the bar graph in Figure 6, both the average and maximum throughput scaled linearly with the number of arrays. The throughput from PS6110XV arrays was significantly higher compared to the PS6010XV results. The behavior can primarily be attributed to the increased number of disks hosting active data on PS6110XV arrays. The linear scaling was possible because of EqualLogic's scale-out architecture; an array adds not only disks but also controllers and NIC ports.

The following table shows the percentage of increase in throughput for PS6110XV scalability tests. The throughput increased almost linearly as additional array members were added within the same pool. A single RAC node was able to saturate two PS6110XV arrays with sufficient system resources available. But the two RAC node configurations were also tested to evaluate the benefits of horizontal scaling. The test results from two RAC node configurations are also included in the tables below.

DSS Application throughput	1 x PS6110XV RAID 50	2 x PS6110XV RAID 50	% Increase
MAX Throughput			
one RAC node	754.76	1582.89	109
two RAC nodes	896.36	1647.34	84
AVG Throughput			
one RAC node	643.20	1209.62	88
two RAC nodes	672.14	1331.31	98

The following table compares the performance of PS6010XV and PS6110XV arrays while executing DSS simulation tests on the actual Oracle database.

DSS Application throughput	1 x PS6010XV RAID 50	1 x PS6110XV RAID 50	% Increase
MAX Throughput			
one RAC node	551.04	754.76	37
two RAC nodes	709.74	896.36	26
AVG Throughput			
one RAC node	454.94	643.20	42
two RAC nodes	446.40	672.14	51

The I/O throughput measured during the DSS test runs using the actual Oracle database matched well with the random read I/O test conducted with ORION as shown in Figure 7.

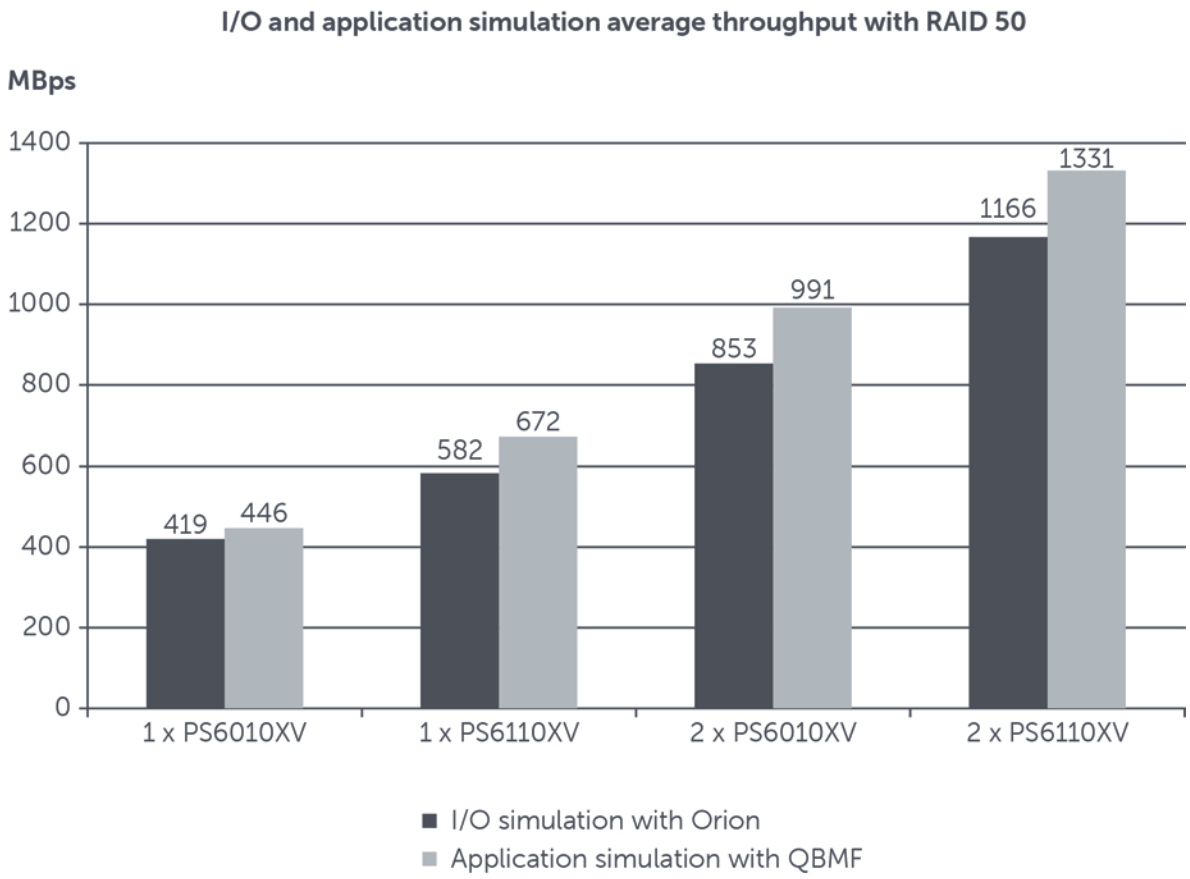

Figure 7 Storage I/O throughput – I/O (ORION) and application simulation

The following table shows the percentage of increase in throughput on the PS6110XV array configuration when DSS simulation tests were executed using the actual Oracle database as compared to I/O simulation.

I/O Workload	1 x PS6010XV RAID 50	1 x PS6110XV RAID 50	% Increase
TPCH Workload using 5 concurrent streams.	446.40	672.14	50.56
I/O Workload	**2 x PS6010XV RAID 50**	**2 x PS6110XV RAID 50**	**% Increase**
TPCH Workload using 5 concurrent streams.	991.81	1331.31	34.23
I/O Workload	**1 x PS6110XV RAID 50**	**2 x PS6110XV RAID 50**	**% Increase**
TPCH Workload using 5 concurrent streams.	672.14	1331.31	98.07

As shown in the above chart, the I/O throughput scaled almost linearly when additional array members were added within the same pool and also increased with the PS6110XV array due to the additional disk drives.

The I/O throughput across storage array members, as reported by EqualLogic SAN HQ in the two-array configuration, is shown in Figure 8.

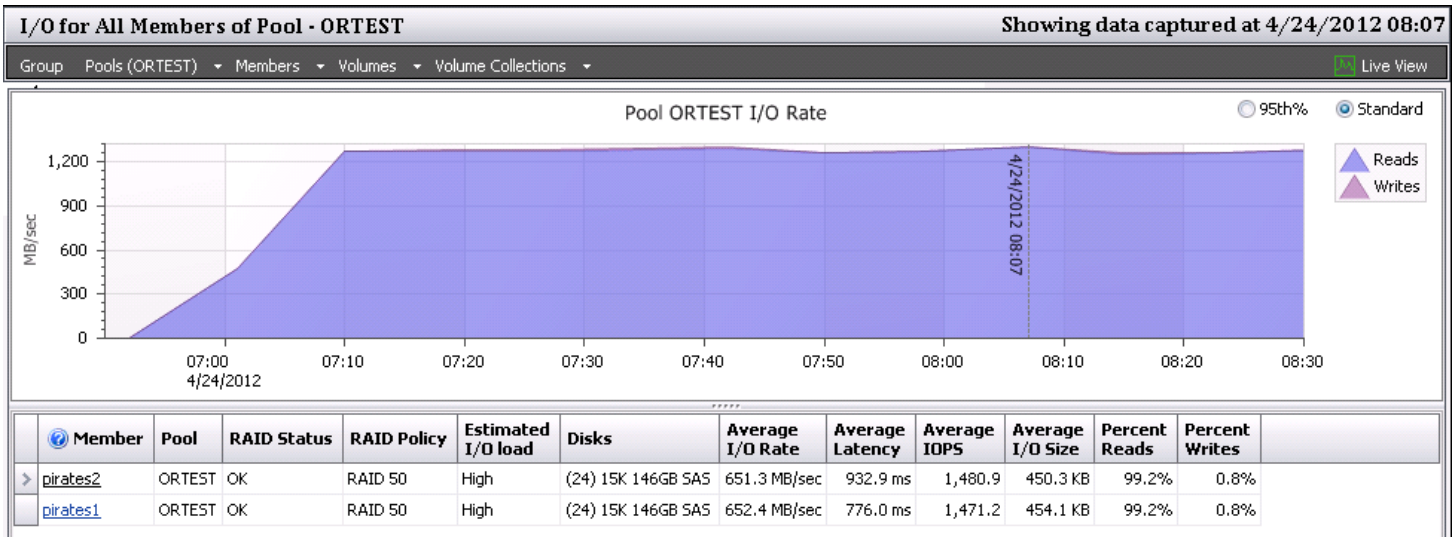

Member	Pool	RAID Status	RAID Policy	Estimated I/O load	Disks	Average I/O Rate	Average Latency	Average IOPS	Average I/O Size	Percent Reads	Percent Writes
pirates2	ORTEST	OK	RAID 50	High	(24) 15K 146GB SAS	651.3 MB/sec	932.9 ms	1,480.9	450.3 KB	99.2%	0.8%
pirates1	ORTEST	OK	RAID 50	High	(24) 15K 146GB SAS	652.4 MB/sec	776.0 ms	1,471.2	454.1 KB	99.2%	0.8%

Figure 8 Storage I/O throughput – SAN HeadQuarters

The initial ramp up in Figure 8 happened during the start of the test when all five users login and start submitting queries. As the users login, the database server commences query processing and generates increasing I/O requests resulting in a ramp up of I/O throughput. After this initial ramp up period, the I/O throughput achieved stabilizes at a constant level based on I/O throughput available from storage.

3.2.1.2 CPU Statistics on RAC nodes

The average CPU utilization percentage and percent I/O wait time were measured during the DSS tests are shown in Figure 9.

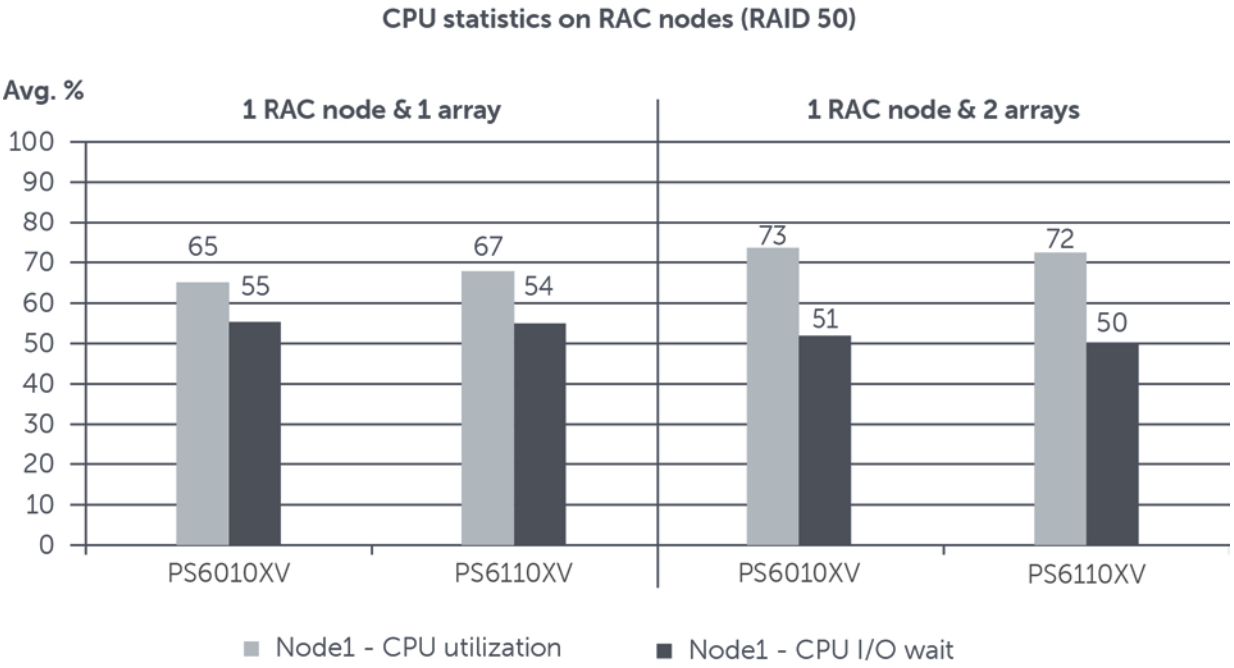

Figure 9 CPU statistics on RAC nodes

The average CPU utilization was under 75% across all tests. The percentage of I/O wait time occupied a significant portion of the CPU utilization. I/O wait time was at 70% to 85% of the CPU utilization. This indicated that the database server reached I/O limits with sufficient CPU resources available.

3.2.1.3 Query Completion Time

The end-to-end completion time for all 22 queries submitted by the five simulated users is shown in Figure 10 below. The duration is between the time when the first query was submitted to the database and when the last query was completed. This graph also displays a comparison between PS6010XV and PS6110XV arrays with respect to query completion time.

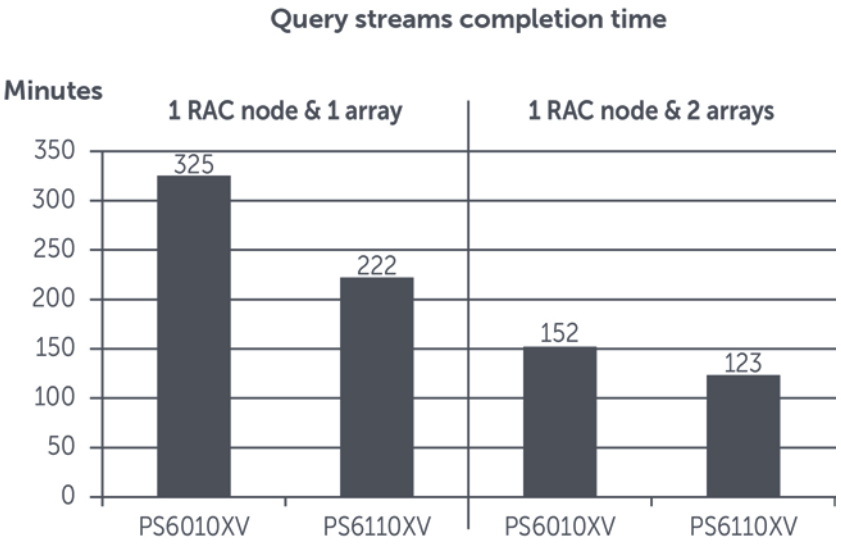

Figure 10 Query streams completion time: PS6110XV versus PS6010XV

Due to the increased storage throughput available as the number of arrays increased, the database was able to fetch data at a much faster rate. The queries were completed faster with an increased number of arrays, resulting in decreased completion times.

The following table shows the percentage of decrease in query completion time for different configurations.

Configuration	1 x PS6010XV RAID 50	1 x PS6110XV RAID 50	% Decrease
one RAC node	325.18	222.14	31.68
Configuration	**2 x PS6010XV RAID 50**	**2 x PS6110XV RAID 50**	**% Decrease**
one RAC node	152.20	123.35	18.95
Configuration	**1 x PS6110XV RAID 50**	**2 x PS6110XV RAID 50**	**% Decrease**
one RAC node	222.14	123.35	44.47

The query completion time was significantly reduced with PS6110XV arrays compared to PS6010XV arrays. With two PS6110XV arrays, the completion time was 44.47% less than with a single array. The nonlinear decrease in query completion time for 1-array and 2-array configurations can be attributed to other database processing factors including workload characteristics and database size.

4 Conclusion

EqualLogic 10 GbE iSCSI PS Series arrays provide high levels of I/O throughput warranted by DSS. By adding arrays, PS Series scale not only capacity but also I/O throughput as well. This is due to the scale-out architecture where all array resources, including controllers and NICs, are being scaled proportionately. Optimal operation of DSS was achieved when the applicable best practices laid out in the reference white paper are adhered to. It must be ensured that the entire ecosystem, including server resources and network switches, is sized and configured appropriately to meet the workload performance requirements. It must also be ensured that the operating system and database are configured with the optimal settings and design recommendations mentioned in this paper.

Some of the key conclusions from the tests are listed below.

- Significant increase in throughput observed for EqualLogic PS6110XV arrays compared to PS6010XV arrays in both RAID 10 and RAID 50 configurations. Almost 125% increase in throughput for large (1 MB) 100% sequential write I/O workload was observed.

- The RAID 50 configuration provided slightly better throughput than the RAID 10 configuration. More than 60% increase in throughput for 100% sequential write operations were observed.

- Large sequential read and write throughput did not scale linearly when additional array members were added when the ORION tool RAID 0 option was used to simulate Oracle ASM striping behavior. The throughput was primarily affected due to the introduction of additional randomness to the I/O while simulating ASM striping behavior. The tests were repeated with the concat option and verified that the throughput scaled linearly.

- Horizontal and vertical scaling in performance was observed when additional array members were added and additional servers added to the RAC database configuration.

- The I/O throughput measured during DSS test runs using the actual Oracle database matched well with the random read I/O test conducted with ORION.

- The I/O throughput measured during DSS test runs using actual Oracle database on PS6110XV arrays were significantly higher than PS6010XV configuration.

- Almost linear scaling even with DSS simulation while executing TPCH workload on 1-array and 2-array configurations was observed.

- The average CPU utilization was under 75% across all tests. The percentage of I/O wait time occupied a significant portion of the CPU utilization which indicated that the database server reached I/O limits with sufficient CPU resources available.

- With two PS6110XV arrays, the end-to-end completion time for all 22 queries submitted by the five simulated DSS users was 44.47% less than with a single array. It was also observed that there was more than a 30% decrease in query completion time when PS6110XV arrays were used in place of PS6010XV arrays.