

Best Practices for Deploying  
SSDs in a Microsoft® SQL  
Server™ 2008 OLTP  
Environment with Dell  
EqualLogic™ PS-Series Arrays

Database Solutions  
Engineering

By

Murali Krishnan.K

Dell Product Group

October 2009



### Executive Summary

This whitepaper discusses the performance improvement on using Dell EqualLogic™ SSD arrays to deploy Microsoft® SQL Server™ 2008 database. This paper details the experimental analysis for the following scenarios, based on a standard OLTP workload using Dell EqualLogic arrays.

- Deployment of database in non-tiered Storage.
- Database Deployment in tiered storage.

This whitepaper also discusses some of the best practices for deploying SSDs in an OLTP environment using Dell EqualLogic PS series arrays.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2009 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

*Dell*, the *DELL* logo, *PowerEdge*, and *PowerVault* are trademarks of Dell Inc. Quest is a registered trademark of Quest Software, Inc. *Microsoft* and *SQL Server* are registered trademarks of Microsoft in the United States and/or other countries.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

## Contents

- Executive Summary..... 2
- Introduction ..... 4
- Dell EqualLogic PS-Series Arrays ..... 4
  - SSDs in DELL Equal Logic PS Series Storage Architecture ..... 5
- Test Methodology ..... 6
  - Test Configuration..... 6
  - SQL Server 2008 – Database Setup..... 8
  - Performance Counters ..... 8
- Test Strategy and Results..... 9
  - Deployment of Whole Database in Non-tiered Storage with PS Array ..... 9
  - Test Result..... 9
  - Analysis ..... 11
  - Database Deployment in Tiered Storage with PS Array ..... 11
  - Test Result..... 12
  - Analysis ..... 14
- Conclusion..... 15
- Figures..... 15
- Tables ..... 16
- References ..... 16

### Introduction

Many modern business applications demand very fast response times. The processor and the conventional storage devices have been improved in terms of capacity and capabilities, but most of the real world application workloads still need to have huge IO with very minimal response time. The business application databases may be either OLAP or OLTP. The OLAP workload could be characterized by large, sequential IOs, for example, long and complex queries involving full table scans. Whereas OLTP is small, random IO, for example, DML operations in a transaction block. With OLTP the response time is crucial for the business continuity to handle many transactions for a huge user workload.

With high server processing capabilities, IT expects high storage performance to meet the low disk latency needs. But the conventional storage disks may not be fast enough to meet the growing performance needs of the industry. The storage performance is a major factor in deciding the response time of the database. If the backend storage is not fast enough to process the request from the client, the whole database performance gets affected. To meet the high processing needs of the industry, Dell EqualLogic™ has storage offerings with the SSD (Solid State Device) support. The Dell EqualLogic PS6000S array provides high performance and low latency with the SSD drives. The SSD technology helps access the data with near zero seek time and with no rotational delay.

This paper discusses the capabilities and performance improvement achieved on deploying Microsoft® SQL Server™ 2008 Database with Dell EqualLogic PS6000S arrays on the backend. Several experiments were executed to compare the Dell EqualLogic PS6000S with SSD drives with Dell EqualLogic PS6000XV arrays with conventional SAS HDDs in SQL server 2008 environment, using a standard OLTP workload.

### Dell EqualLogic PS-Series Arrays

Dell EqualLogic PS-Series arrays are high performing iSCSI storage systems, which are built on patented peer storage architecture. The PS arrays support SSD, SAS and SATA disks and provide multiple high-performance network interfaces with automatic load balancing capability. The redundant storage controllers are fully mirrored and are equipped with battery-backed caches. PS series arrays support RAID configurations of 10, 50, 5 and 6 along with two disks as dedicated hot spares.

The Major terminologies related to Dell EqualLogic PS arrays are as follows:

**Member:** A single PS series array.

**PS Series group:** An iSCSI SAN comprised of one or more PS Series storage arrays (members) connected to an IP network and managed as a single system.

**Storage Pool:** In a PS Series group, storage can be organized into a multiple storage pools consisting of all or a subset of the group members. Volume data is automatically spread across the pool members and load balanced, as needed.

## Deploying Microsoft SQL Server 2008 with Dell EqualLogic PS6000S iSCSI Storage Arrays

**Volume:** A LUN or virtual disk that represents a subset of the capacity of a pool. Each volume is assigned a size, storage pool, and access controls.

A PS series group is a top level entity with single or multiple arrays working in it. The group is considered a single-shared storage for application and end users. The PS series array is a PS series storage pool and its RAID configuration. A PS Series array can have different RAID storage pools. A volume is created from a PS series storage pool.

The figure-1 represents the Dell EqualLogic peer storage architecture. In figure-1, a single volume is carved out of a storage pool. The storage pool here consists of four PS member arrays.

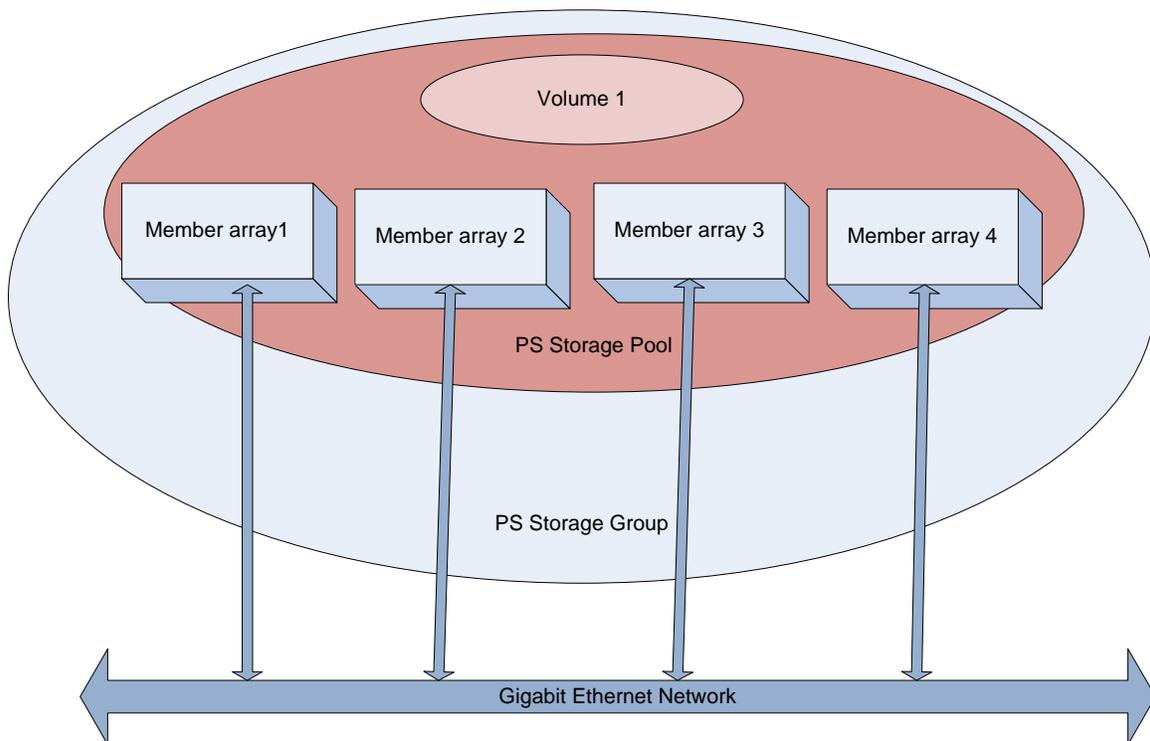


Figure 1: EqualLogic Scale-out Peer Storage

### SSDs in Dell Equal Logic PS Series Storage Architecture

The peer storage architecture is noted for its scalability features with sustained performance. In general, traditional scale-up methods make member's disks share the array controller. Dell Equal Logic PS series storage will have a dedicated controller and network interfaces. The peer architecture provides a way to leverage the SSD performance with existing PS series arrays. Using peer architecture, the storage pool can be extended, both in capacity and performance, by including a member to the pool in the online mode.

## Deploying Microsoft SQL Server 2008 with Dell EqualLogic PS6000S iSCSI Storage Arrays

The Dell EqualLogic PS6000S array may also be configured in a separate pool, with other PS-Series arrays, within the same SAN as a Tiered Storage (Figure-2) to scale out performance and capacity. The PS6000S array uses SSD drives. Each individual solid state drive (SSD) offers a capacity of 50 GB. The total array capacity is 800 GB (16 X 50GB).

In Figure 2, the database has been configured with two volumes, one from PS6000XV and the other from PS6000S.

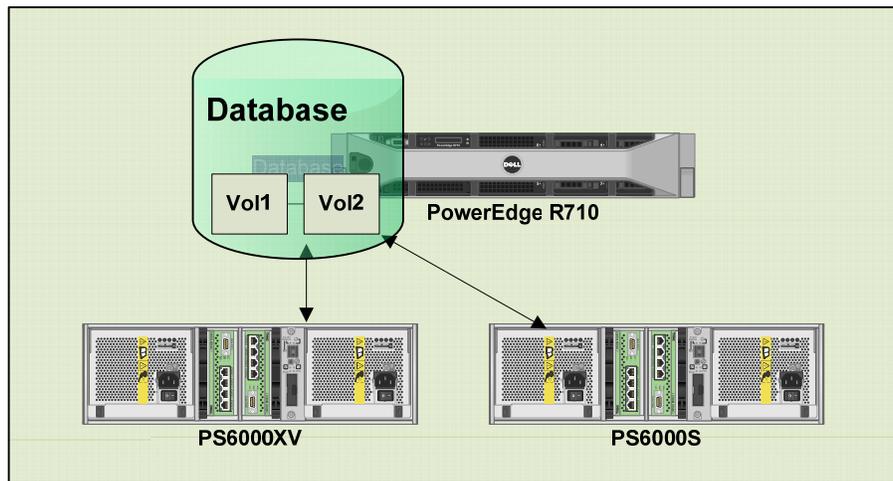


Figure 2: Tiered Storage

Several experiments were carried out in Dell Labs to analyze the performance improvement of Dell EqualLogic PS6000S arrays over the conventional SAS arrays. The latter sections of this paper discuss the performed test methodology and performance analysis.

## Test Methodology

This section provides the PS6000S and PS6000XV performance characterization with Microsoft SQL Server 2008 for a standard OLTP Workload.

## Test Configuration

To do the characterization of Dell EqualLogic PS6000S iSCSI array with SQL server 2008, a dedicated Dell PowerEdge R710 Server was set up with SQL server 2008 Enterprise Edition database. The database was stressed with a standard OLTP workload using Quest® Benchmark Factory™ tool from a client system. The detailed test configuration is given in Table 1 and the Test tools are given in Table2.

## Deploying Microsoft SQL Server 2008 with Dell EqualLogic PS6000S iSCSI Storage Arrays

Table 1: Test Configuration

HARDWARE AND SOFTWARE CONFIGURATUON FOR BENCHMARK FACTORY TEST	
Server	One Dell PowerEdge R710 Server with: <ul style="list-style-type: none"> <li>• Two Intel Xeon quad-core 2.40 GHz CPU</li> <li>• 12 GB of RAM</li> <li>• 2 Gigabit Broadcom BCM5709C NIC ports for iSCSI traffic</li> </ul>
External Storage	Dell EqualLogic PS6000XV OR PS6000S RAID 10 with 2 spare disks in each member. PS6000XV: 300GB SAS disks , 15k RPM PS6000S : 50GB SSD disks Firmware : Storage Array Firmware V 4.1.4(R91505)
Database Volume	300 GB
OS and Device Driver	Microsoft Windows 2008 Servers X64 R2 Microsoft iSCSI initiator EqualLogic Multiple IOPS Device Specification Module (DSM) Version 3.2.0.4722
Storage Network	2 stacked Dell PowerConnect 5448 Gigabit Ethernet switches for iSCSI SAN
Test Software	Quest Benchmark Factory 6.0

Table 2: Test Tools

Tools and Usage	
Benchmark Factory	Benchmark Factory for database tool, from Quest software, was used to simulate users and OLTP transactions on the SQL server database.
SQL Server Profiler	Microsoft SQL server profiler is a graphical user interface to SQL Trace for monitoring an instance of the Database Engine.  In this paper, SQL server profiler was used to identify the read/ Write intensive database object in the OLTP work Load. The SQL trace was run when the benchmark Factory was performing the OLTP workload and the database objects events are captured for the analysis.
Windows Performance Monitor	The Windows Performance Monitor (perform) is a tool provided by Microsoft along with Windows to track and monitor the server activities. In this paper, perform was used to capture the server performance counters like Processor - %idle time, Physical Disk - Avg.Disk Queue Length, Avg. Disk sec/Read, Avg. Disk sec/Write , Disk Reads/sec and Disk Writes/sec , when the database was under load.

### SQL Server 2008 – Database Setup

A single instance SQL server 2008 Enterprise Edition with SP1 was installed on a Dell PowerEdge Server. OLTP Database was created with approximately 105 GB data by Quest Benchmark Factory Tool. The SQL server instance “Maximum Server Memory” parameter was configured with 1 GB RAM. This was done to decrease the amount of cached data at the SQL server database, to enable the database to have more IOs to the storage system.

The standard OLTP workload was performed with 25 to 200 users (in 25 user increments) for each test scenario. The benchmark factory tool simulated the typical OLTP applications in which most IO operations are random reads and writes. During the experiments, the storage IOPS and response time were captured and analyzed.

### Performance Counters

When the workload was executed, the key performance counters from the server were captured to measure the storage performance. Several performance variables were observed to make sure that no server side bottlenecks exist for the setup.

Window Performance Monitor tool was used to capture the following performance counters when OLTP work load was executed by the Benchmark Factory tool.

Table 3: Performance Counters

Performance Counter	Description
Processor - %Processor Time	This is a percentage of time the processor spent executing a non-idle thread. This indicates that there are no CPU bottlenecks in the server.
Physical Disk - Avg.Disk Queue Length	Average number of requests that waited for the read and write requests in the particular disk.
Avg. Disk sec/Read	The number indicates the average time, in seconds, of a read of data from the disk. This was used to calculate volume response time.
Avg. Disk sec/Write	The number indicates the average time, in seconds, of a write of data from the disk. This was used to calculate volume response time.
Disk Reads/sec	The number shows the average time, in seconds, of a read of data from the disk. This was used to calculate volume IOPS.
Disk Writes/sec	The number shows the average time in seconds that it takes to write data to the disk. This was used to calculate volume IOPS.
SQL Server:General Statistics - User Connections	The number of users currently connected to the SQL Server.

## Deploying Microsoft SQL Server 2008 with Dell EqualLogic PS6000S iSCSI Storage Arrays

The database storage volume was configured as RAID10. The IOPS was calculated by the following formula:

$$\text{IOPS} = 1 \text{ Disk Reads/sec} + (2 * \text{Disk Writes/sec})$$

The response time was calculated by using the formula:

$$\text{Response Time} = \text{Avg. Disk sec/Read} + \text{Avg. Disk sec/Write}$$

## Test Strategy and Results

To analyze the effect on database performance using SSD drives, the following database configurations were setup:

- Database deployment in non-tiered storage with PS array.
- Database deployment in tiered storage with PS array.

The performance counters were captured to calculate IOPS and response time for each test configuration.

### Deployment of Whole Database in Non-tiered Storage with PS Array

This section discusses the effect on database performance on implementing the whole database on SSD drives. For this analysis, the whole database is deployed in one PS6000S array. The Quest Benchmark Factory tool was used to stimulate a 25 to 200 user (in 25 user increments) workload for the test scenario. The performance counters were captured while Benchmark Factory performed the workload transactions.

The same test was repeated with PS6000XV array and its performance was compared with PS6000S performance.

### Test Result

With the specific OLTP workload, PS6000S was delivering a maximum IOPS of approximately 10117, the average disk high queue length was 1095 and the maximum response time was approximately 260 msec.

PS6000XV was delivering a maximum IOPS of approximately 7907, the average disk high queue length was 1273 and the maximum response time was approximately 792 msec.

## Deploying Microsoft SQL Server 2008 with Dell EqualLogic PS6000S iSCSI Storage Arrays

Table 4: Test Results - Non-Tiered Storage Environment

Setup details	Max .IOPS (per array)	Max. Response Time (msec)	Max. Avg. Disk Queue Length
Complete DB on PS6000S (14 drives)	10117.75	260.45	1095.00
Complete DB on PS6000XV (14 drives)	7907.01	792.99	1273.99

Figure-3 shows the difference in IOPS and Response Time (msec) for the particular workload. In the figure, the X- axis represents the number of users in the workload. In the first graph, the Y-axis shows the corresponding IOPS achieved in the test and in the second graph, Y-axis depicts the response time in milliseconds.

For the specific work load, the IOPS was averaged to approximately 8908.94085 for PS6000S and 5306.884188 for PS6000XV. The average response time in milliseconds was approximately 194.8856831 for PS6000S and 450.8124066 for PS6000XV.

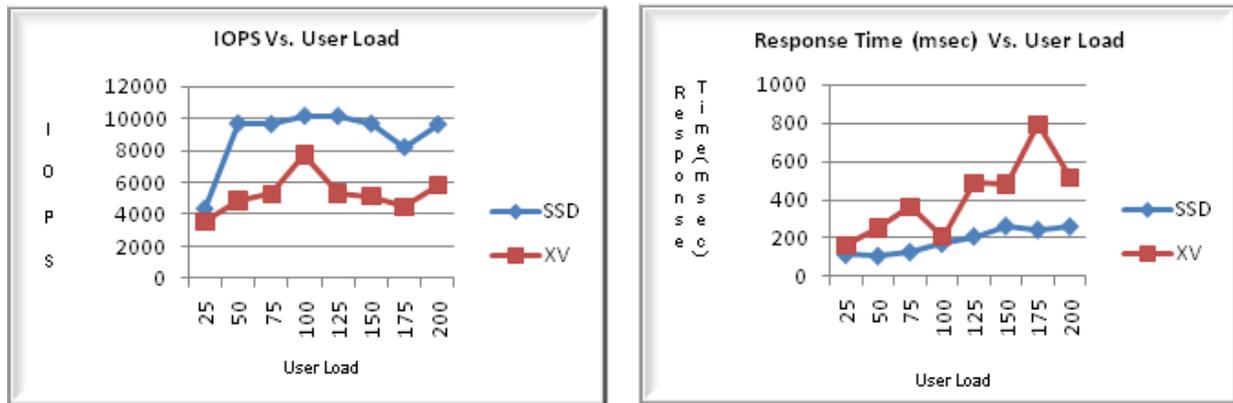


Figure 3: IOPS and Response Time (PS6000XV vs. PS6000S) from OLTP workload

Figure-4 shows the CPU Utilization and the disk queue length during the workload. In the figure, X- axis represents the number of users in the workload. In the first and second graph, the Y-axis shows the CPU utilization and average disk queue length respectively.

For the particular workload, the average CPU utilization percentage was approximately 10.41912833 for PS6000S and 7.831456582 for PS6000XV. The Average Disk Queue length was approximately 729.0438552 for PS6000S and 857.5919734 for PS6000XV.

## Deploying Microsoft SQL Server 2008 with Dell EqualLogic PS6000S iSCSI Storage Arrays

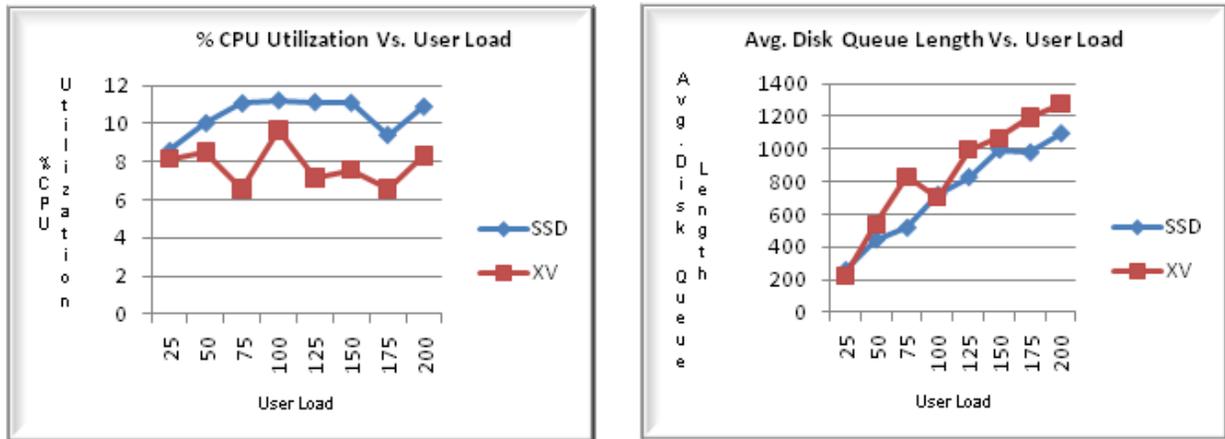


Figure 4: CPU Utilization and Avg. Disk Queue Length (PS6000XV vs. PS6000S) from OLTP workload

### Analysis

The performance and overall scalability of the storage system are measured based on IOPS and response time.

Figure-3 shows that the database performance improved significantly on using the SSD array as the backend. The experiment results shows that the PS6000S array gave 23% to 49% more IOPS, compared to PS6000XV array for the given workload. Figure-3 shows that response time was also drastically reduced with PS6000S array. The test result gave 21% to 229% improvement with PS6000S. This is due to the neat, zero-seek time and no rotational delay for the SSD drives. In addition, it is observed that the response time remarkably increased in PS6000XV on increasing the user load. But the PS6000S array shows a relatively constant response time on increasing user load. This is because of the greater request processing capability of the SSD drives.

The server CPU utilizations were continuously monitored to make sure that there are no server side bottlenecks for the workload. The CPU utilization for the workloads remained within 12% throughout the test period.

The heavy IO queues observed at the database disk was due to the imposed memory constraint of 1GB, at the SQL Server instance. This was done to direct the maximum IOs to the storage. In PS6000S, the logical disk queue length was 7% to 21 % lesser than the PS6000XV. This is because the PS6000S SSD disks were able to process the request faster than the PS6000XV SAS disks.

### Database Deployment in Tiered Storage with PS Array

This section explains the database performance difference caused by moving the IO intensive database objects to the SSD array.

One of the traditional approaches to improve SQL server performance is to move the most READ/ WRITE database objects to a different logical drive. This section examines how PS6000S can be beneficial when

## Deploying Microsoft SQL Server 2008 with Dell EqualLogic PS6000S iSCSI Storage Arrays

it is deployed as part of a tiered storage database. PS6000S can be used to accommodate mission critical database objects (tables, Indexes and indexed views).

Initially the database was setup on a single volume housed in a single PS6000XV array. Using a sample workload, SQL server Profiler was used to find the IO intensive objects in the database. The Top 3 groups of IO intensive Database Objects are displayed in Table 5 below. (The details of write intensive are ignored since there were few writes in the given workload.)

Table 5: I/O-Intensive Objects in the Workload

Object Name	Object Type	Physical Reads	% Read
E_TRADE_HISTORY,E_TRADE, E_HOLDING,E_CUSTOMER	Table	31324700	45%
E_BROKER ,E_TRADE_REQUEST, E_SECUIRT ,E_COMPANY,E_INDUSTRY E_SECTORY	Table	1573100	2%
E_WATCH_ITEM,E_WATCH_LIST	Table	795600	1%

The database objects in Table 6 were then moved to a second PS6000XV array. These identified database objects summed up to approximately 13.32 GB of data, about 12% of the whole database size. Then the standard OLTP workload was performed and the server side performance counters were captured.

Table 6: Selected Database Objects

Object Name	Object Type	Size(MB)
E_CUSTOMER	Table	3.26
E_TRADE_HISTORY	Table	13567.11
E_WATCH_ITEM	Table	70.03

The second PS6000XV was then replaced with one PS6000S array. The database was again stressed with the same OLTP workload and the performance counters were captured.

## Test Result

Table 7 shows the results for the test scenario described above. With the specific OLTP workload, the configuration 'A' was delivered a maximum IOPS of approximately 6516, the average disk queue length was 1372 and the maximum response time approached 1372 msec.

The configuration B was delivered a maximum IOPS of approximately 6677, the average disk high queue length was 1056 and the maximum response time came close to 485 msec.

## Deploying Microsoft SQL Server 2008 with Dell EqualLogic PS6000S iSCSI Storage Arrays

Table 7: Test Results - Tiered Storage Environment

Configuration	Setup details	Max .IOPS (per array)	Max.Response Time (msec)	Max. Avg.Disk Queue Length
<b>A</b>	Database on PS6000XV + PS6000XV	6516.36	1372.66	1739.85
<b>B</b>	Database on PS6000XV +PS6000S	6677.19	485.98	1056.60

Figure-5 shows difference in IOPS and Response time (msec) for the test scenario. In the figure, X-axis represents the number of users in the workload. In the first graph, the Y- axis shows the corresponding IOPS achieved in the test and in the second graph, Y-axis depicts the response time in milliseconds.

For the specific work load, the IOPS was averaged to approximately 5559.069589 for configuration 'A' and 6099.19337 for configuration 'B'. The average response time in milliseconds was approximately 634.172878 for confirmation 'A' and 359.3008568 for configuration 'B.'

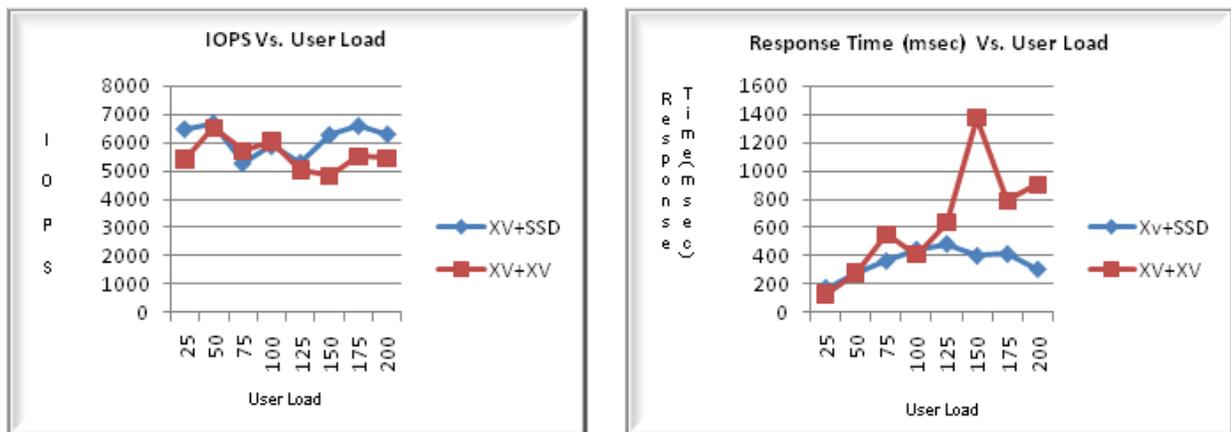


Figure 5: IOPS and Response Time (Configuration B vs. Configuration A) from OLTP workload

Figure-6 shows the CPU Utilization and the disk queue length during the workload. In the figure, X- axis represents the number of users in the workload. In the first and second graph, the Y-axis shows the CPU utilization and average disk queue length.

For the particular workload, the average CPU utilization percentage was approximately 7.665852479 for configuration 'A' and 9.773133304 for PS6000XV. The Avg. Disk Queue length was approximately 1187.035142 for configuration 'A' and 813.627881 for configuration 'B.'

## Deploying Microsoft SQL Server 2008 with Dell EqualLogic PS6000S iSCSI Storage Arrays

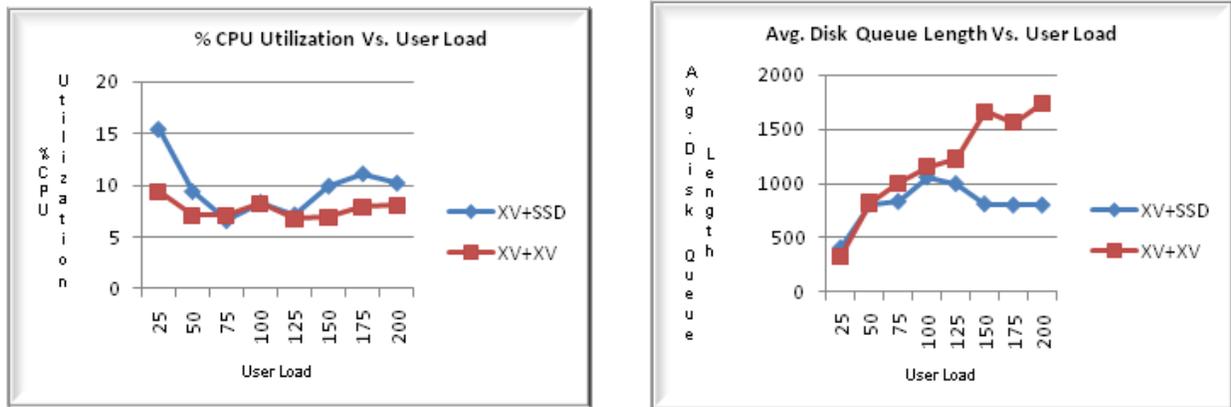


Figure 6: CPU Utilization and Avg. Disk Queue Length (Configuration B vs. Configuration A) from OLTP workload

### Analysis

Figure-5 shows that the database performance improvement on introducing the PS6000S array as part of tiered storage. The experiment results show that configuration 'B' gave 7% to 23% more IOPS, then configuration 'A' for this given workload. Figure-5 also shows that response time was observed to be reduced with configuration 'B.' The test result gave 23% to 194% improvement with configuration 'B.' This is because in configuration 'B' the IO intensive operations were processed using the fast SSD drives in the tiered architecture.

Also, CPU Utilizations for the workloads remained within 39% throughout the test period. This confirms that no server CPU bottleneck exists.

It was observed that in configuration 'B' the database disk queue length was 19% to 117 % lesser than configuration 'A.' The high processing capability of the SSD array processes the database requests faster, which in turn reduces the volume queue length.

## Conclusion

Based on the experimental analysis, we observed that the SQL Server Database deployed on Dell EqualLogic PS6000S iSCSI array with SSD drives may provide up to 229% faster response time than the Dell EqualLogic PS6000XV array with SAS HDDs, depending on the workload. We also observed that the database IOPS was improved up to 49%, for the particular workload. As a single PS6000S array provides an overall capacity of approximately 700GB, small databases may be accommodated entirely within the PS6000S array for high performance.

In a tiered storage configuration, a combination of Dell EqualLogic PS6000XV and Dell EqualLogic PS6000S arrays may provide significant performance improvement. On moving the IO intensive objects of the workload to the SSD array, we noticed that the database response time may be improved up to 194% and the overall database IOPS may be improved up to 23%, depending on the workload. It is recommended to move the IO intensive components of the database to the SSD array to improve the database performance.

The SSD drives come with a higher cost than the conventional SAS HDDs. In data center environments, where multiple databases are deployed, the frequently used IO intensive objects from multiple databases may be moved to a shared PS6000S array. This will ensure that PS6000S array is effectively utilized and will provide a cost effective high performing deployment.

Based on the analysis discussed in this paper, it may be concluded that SQL server 2008 database deployments on Dell EqualLogic PS6000S iSCSI arrays provides better IOPS and improved response time. Hence, Dell EqualLogic PS6000S iSCSI Array is an excellent solution for high performance – latency intensive applications.

## Figures

Figure 1: EqualLogic Scale-out Peer Storage.....	5
Figure 2: Tiered Storage .....	6
Figure 3: IOPS and Response Time (PS6000XV vs. PS6000S) from OLTP workload .....	10
Figure 4: CPU Utilization and Avg. Disk Queue Length (PS6000XV vs. PS6000S) from OLTP workload .....	11
Figure 5: IOPS and Response Time (Configuration B vs. Configuration A) from OLTP workload .....	13
Figure 6: CPU Utilization and Avg. Disk Queue Length (Configuration B vs. Configuration A) from OLTP workload .....	14

## Tables

Table 1: Test Configuration..... 7  
Table 2: Test Tools ..... 7  
Table 3: Performance Counters ..... 8  
Table 4: Test Results - Non-Tiered Storage Environment..... 10  
Table 5: I/O-Intensive Objects in the Workload ..... 12  
Table 6: Selected Database Objects..... 12  
Table 7: Test Results - Tiered Storage Environment..... 13

## References

1. Oracle - Best practices for implementing SSDs with EqualLogicTM PS6000 storage running an OLTP workload.  
[http://www.dell.com/downloads/global/solutions/best\\_practices\\_EqualLogic\\_PS6000\\_storage.pdf?c=us&cs=555&l=en&s=biz](http://www.dell.com/downloads/global/solutions/best_practices_EqualLogic_PS6000_storage.pdf?c=us&cs=555&l=en&s=biz)
2. Faster SQL Server Access with Solid State Disks by Texas Memory Systems.  
<http://www.texmemsys.com/files/f000174.pdf>
3. Solid State Storage for SQL Server  
<http://www.sqlmag.com/Articles/Index.cfm?ArticleID=98704&DisplayTab=Article>