


Increasing Database Performance Using the Dell PERC H800 with Solid State Drives

A Dell Technical White Paper

Joe Noyola

Systems Performance Analysis



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.

© 2011 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 and the *DELL* badge are trademarks of Dell Inc. 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 Inc. disclaims any proprietary interest in trademarks and trade names other than its own. TPC and the benchmark name TPC-E are registered trademarks of the Transaction Processing Performance Council. TPC-E results were simulated and do not represent actual TPC benchmark results.

July 2011

Table of Contents

- Executive Summary 4
- Introduction 4
- Key Results 4
- Test Methodology..... 5
- Test Results..... 5
 - Database Performance - Transactions Per Second (tps) 5
 - Database Performance - User Scalability with acceptable average query latency..... 6
 - Storage Performance - I/Os per second (IOPS) 7
 - Price/Performance 7
- Conclusion 8
- Appendix A - System Cost Details (from Dell.com) 9
- Appendix B - Server Hardware Configuration 10
- Appendix C - Storage Hardware Configurations 11
- Appendix D - Database Software / Configuration / Benchmark Factory Configuration Details 13

Tables

- Table 1. Storage Configuration 5

Figures

- Figure 1. Database Performance Comparison 6
- Figure 2. Database Concurrent User Comparison..... 6
- Figure 3. IOPS Scalability Comparison..... 7
- Figure 4. Performance, Price/Performance Comparison..... 8

Executive Summary

Dell's High-Performance SAS Solid State Drives (SSD) significantly improve I/O performance compared to traditional hard drives (HDDs). With the introduction of CacheCade™ functionality, Dell's PERC (PowerEdge™ RAID Controller) H800 with the PowerVault™ MD1220 enclosure now supports two SSD implementation schemes: SSD Caching and SSDs implemented as primary data storage. This white paper evaluates the performance improvements and cost of these SSD usage cases against traditional hard drives using a simulation of the I/O workload of a small-to-medium sized SQL Server TPC-E-like OLTP database.

Introduction

Many IT professionals today are considering SSDs as an option to improve storage performance to mitigate traditional storage I/O bottlenecks. Solid state storage is flash-based and thus does not rely on slow-moving mechanical parts to access data. This makes them ideal for workloads with read-intensive data I/O access patterns. SSDs, while providing maximum performance, do require tradeoffs: given that this is a fairly new technology, they are currently more expensive than traditional HDDs.

Compared to all-SSD implementations that require multiple SSDs per volume, CacheCade with one SSD provides a lower-cost solution which still harnesses the performance of SSDs. CacheCade software (SSD-caching) running on the PERC H800 1GB NV Cache RAID controller improves the performance of hard drive storage by identifying frequently-accessed read operations and caching them on solid-state storage. Application read performance is improved over time, as the CacheCade cache volume is populated and subsequent reads are serviced from SSD cache.

Unlike CacheCade, where I/O performance is achieved as the SSD cache is populated, all-SSD implementations see performance improvements immediately. Applications can utilize the full potential of SSDs, as both read and write operations are accelerated. If maximum storage performance supersedes storage capacity and cost concerns, the PERC H800 combined with multiple SSDs provides a new level of storage I/O performance traditional storage cannot match.

To measure the performance improvements these SSD implementations offer, Quest's Benchmark Factory® was used to build and simulate an industry-standard SQL Server OLTP database. Compared to hard drive storage alone, CacheCade and all-SSD implementations significantly improve OLTP database performance.

Key Results

OLTP Database Performance - Transactions per second (tps)

- Enabling an SSD as a CacheCade volume increased hard drive database performance by 5x¹
- All-SSD storage achieved 15x more database performance than HDD storage alone¹

Database User Scalability Support with Acceptable Response time of 100ms

- CacheCade volume allows hard drive storage an additional 3x more users²
- All-SSD storage supports up to 350 users, an 11x increase in database user support over hard drive storage²

Storage Performance - I/Os per Second (IOPS)

- CacheCade provided an increase of 4x more IOPS than hard drive storage alone³

¹ For detailed performance results see Figure 1. Testing configurations are detailed in the appendices.

² For detailed performance results see Figure 2. Testing configurations are detailed in the appendices.

Increasing Database Performance Using the PERC H800 with Solid State Drives

- All-SSD storage achieved 14x more IOPS than hard drive storage³

Price per IOPS (lower is better)

- Dollars per IOPS of hard drive storage costs \$1.54/IOPS⁴
- An SSD enabled as a CacheCade volume decreases dollars per IOPS to \$0.64/IOPS, a 58 percent savings in price/performance⁴
- All-SSD storage achieved 14x more IOPS than hard drive storage, which drops overall dollar per IOPS to only \$0.31/IOPS, an 79 percent dollars/IOPS savings compared to hard drive storage⁴

Test Methodology

Quest's Benchmark Factory for Databases (BMF) is a database build and workload simulation tool that measures database performance. To compare the performance that CacheCade and all-SSD storage achieve over that of traditional hard drives, a SQL Server 2008 100GB database modeled after the industry-standard TPC-E benchmark was built on a RAID 10 volume made up of fourteen 146GB 10k RPM hard drives. An additional database was built on a 4-drive SAS SSD RAID 10 volume. Table 1 lists the hardware of each storage configuration.

Table 1. Storage Configuration

Database Storage Configuration	RAID Configuration	Drive Type / Speed
Hard Drive baseline	14-drive RAID 10	SAS HDD /10k RPM - 6Gbs
Hard Drive + 1 CacheCade volume	14-drive RAID 10 + 1 SAS SSD	SAS HDD /10k RPM - 6Gbs + 1 SSD
All-SSD baseline	4-drive RAID 10	SAS SSD

Baseline performance was collected from the database built on hard drive storage alone. A single SAS SSD was then enabled as a CacheCade volume and measurements were repeated. Metrics were then collected from the database built on all-SSD storage. During each measurement, BMF simulated users issuing TPC-E like SQL transactions that exercised each database's storage sub-system with a random read and write I/O data access pattern.

BMF simulated an increasing number of virtual users until the average query response time met an acceptable response time of 100 milliseconds. To ensure all I/O requests were serviced from physical storage, the SQL Server 2008 default instance was limited to run with only 1GB of RAM. Limiting the amount of memory available to SQL Server ensures that storage sub-systems are exercised to their full I/O capacity.

Test Results

Database Performance - Transactions Per Second (tps)

A baseline of 40 tps was measured from the hard drive database. Enabling a single SAS SSD as a CacheCade volume increased database performance to over 200 tps, a 5x gain. All-SSD database performance reached a staggering 600 tps; a 15x gain in tps compared to the database on all-HDD storage. Figure 1 shows the achieved database performance CacheCade and all-SSDs provide.

³ For detailed performance results see Figure 3. Testing configurations are detailed in the appendices.

⁴ For detailed performance results see Figure 4. Testing configurations are detailed in the appendices.




Figure 1. Database Performance Comparison

Database Performance - User Scalability with acceptable average query latency

BMF increased the number of virtual users simulating an OLTP workload to each storage configuration while measuring average query latency. Acceptable average query latency was set to 100ms, thus ensuring the simulated users an acceptable quality of service. Hard drive storage supported up to 30 simultaneous users accessing the database while maintaining the acceptable average query response time. Enabling CacheCade allows hard drive storage alone to sustain an additional 3x user load. Database built on solely SSDs supports up to 350 concurrent users, 11x more users than hard drives.

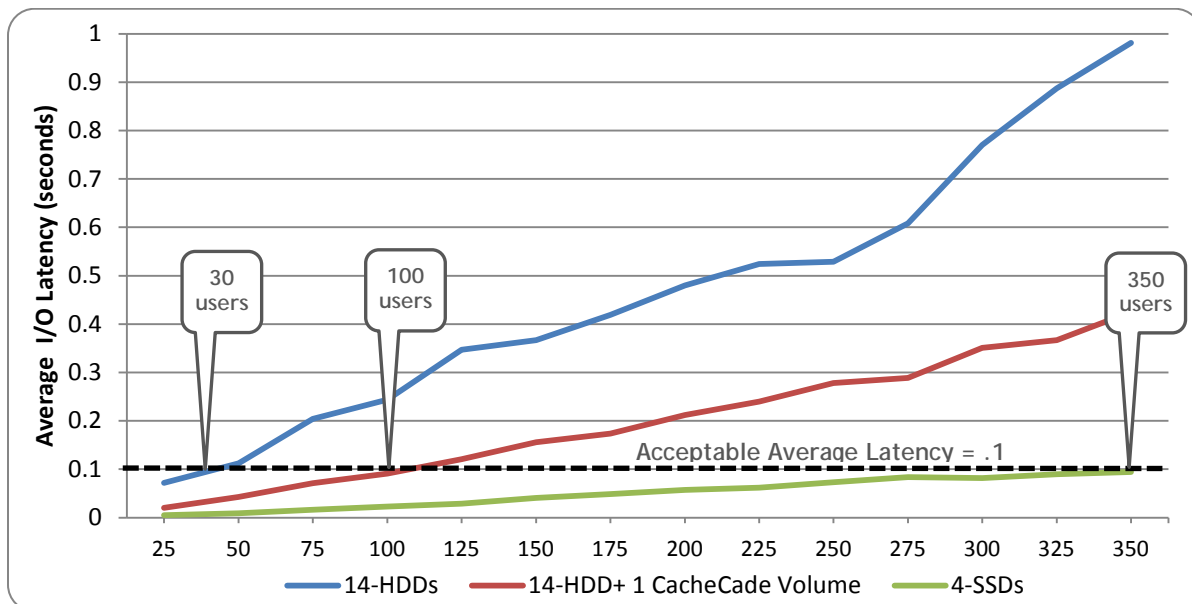


Figure 2. Database Concurrent User Comparison

Storage Performance - I/Os per second (IOPS)

The hard drive storage configuration achieved a maximum of 5,500 IOPS while 30 users concurrently queried the database. Configuring an SSD as a CacheCade volume allowed 100 users to query the database and increased total IOPS of hard drives to 20,450, a 3x improvement. All-SSD storage supported 350 concurrent database users and achieved 76,270 IOPS, 11x more IOPS compared to all-HDD storage. Figure 3 illustrates the IOPS achieved by hard drives, CacheCade, and all-SSD storage.




Figure 3. IOPS Scalability Comparison

Price/Performance

Traditional storage offers a combination of low cost and storage capacities that solid state storage currently cannot match. Due to the electro-mechanical rotating circular design, however, the number of IOPS hard drives can achieve is physically limited. As seen in Figure 4, price per IOPS of a database with hard drive storage is \$1.54, a 58 percent higher cost per IOPS than CacheCade at \$0.64 per IOPS and 79 percent more expensive than \$0.31 per IOPS that all-SSD storage provides.




Figure 4. Performance, Price/Performance Comparison

Conclusion

With the introduction of flash-based Solid State Disks, customers are faced with balancing performance versus cost. CacheCade allows the best of both worlds, where SSDs can be used as a read cache, but traditional HDDs can be used for primary data storage for a lower total cost of ownership. Although carrying the highest purchase price, SSDs implemented at primary data storage differ from CacheCade as both read and write I/O performance is accelerated and available instantly to the application versus SSD-caching.

The results of this Benchmark Factory OLTP database simulation clearly shows that both CacheCade and SSDs used for primary database storage improve database performance while decreasing the overall cost per IOPS. CacheCade provides customers a low-cost SSD solution that improves database performance while continuing to benefit from the best features traditional hard drives offer: storage capacity and cost. If performance is paramount, database applications built on all-SSD storage fully harness the I/O capability of Dell's enterprise SAS SSD and achieve a new era of OLTP database performance above and beyond that of traditional rotating media.

Appendix A - System Cost Details (from Dell.com)

Storage Configuration	Components	Price	Dollars/IOPS
All-HDD storage	1- PERC H800 1GB NVRAM controller 1- PowerVault MD1220 storage enclosure 14 - 10k SAS 146GB drives - data 4 - 10k SAS 146GB drives - log	\$8,478	\$1.54/IOPS
All-HDD + 1-drive CacheCade	All-HDD configuration + 1 - Pliant 149GB SLC SSD as CacheCade	\$13,073	\$0.64/IOPS
All-SSD	1- PERC H800 1GB NVRAM controller 1- PowerVault MD1220 storage enclosure 4 - 149GB SLC SSDs - data 4 - 10k SAS 146GB drives - log	\$23,288	\$0.31/IOPS

Appendix B - Server Hardware Configuration

Server Hardware	Dell PowerEdge R710
Processor	Dual Intel Xeon X5670 2.93Ghz 12MB Cache 6 core processors
Total RAM in system (GB)	16
Vendor and model number	Samsung ST9146852SS
Memory Type	PC3L-10600R
Speed (MHz)	1333
Speed in system as tested	1066
Number of RAM modules	16 x 4 GB
Rank organization	Dual Rank
Vendor and model number	Dell MBE2147RC
Number of disks in system	2
Size (GB)	146
RPM	15,000
Type	SAS 6 Gbps
RAID Type	RAID 1
Controller	PERC Integrated H700 512MB F/W:12.1.0091 / Driver - 4.31.1.64
Name	Microsoft® Windows Server® 2008 R2 Enterprise
File system	NTFS
Language	English
Vendor and model number	Broadcom® BCM5709S NetXtreme® II
Type	Integrated

Appendix C - Storage Hardware Configurations

Storage Configuration: All-HDD Baseline	
PowerVault Storage Enclosure	MD1220
Total Number of Available Drive Slots	24
Number of Drives Installed	14-drive RAID 10 volume for database 4-drive RAID 10 volume for database log
Drive Type	6Gb/s SAS
HDD Speed (RPM)	10k HDD
HDD Capacity (GB)	146GB
Controller	Dell PERC H800 1GB NV Cache
Controller Firmware / Driver	F/W 12.1.0091 / 4.31.1.64
PERC Cache Settings	Write Back Adaptive Read Ahead
Disk Firmware	HT64

Storage Configuration: All-HDD + 1 CacheCade Volume	
PowerVault Storage Enclosure	MD1220
Total Number of Available Drive Slots	24
Number of Drives Installed	14-drive RAID 10 volume for database 4-drive RAID 10 volume for database log 1 Pliant SSD as CacheCade volume
Drive Type	6Gb/s (HDD) / 3Gb/s (SSD)
HDD Speed (RPM)	10k HDD
SSD / HDD Capacity (GB)	149GB (SSD) / 146GB (HDD)
Controller	Dell PERC H800 1GB NV Cache
Controller Firmware / Driver	F/W 12.1.0091 / 4.31.1.64
PERC Cache Settings	Write Back Adaptive Read Ahead
Disk Firmware	D801

Increasing Database Performance Using the PERC H800 with Solid State Drives

Storage Configuration: All-SDD	
PowerVault Storage Enclosure	MD1220
Total Number of Available Drive Slots	24
Number of Drives Installed	4-drive Pliant SSD RAID 10 volume for database 4-drive RAID 10 volume for database log
Drive Type	3Gb/s (SSD) / 6Gb/s (HDD)
HDD Speed (RPM)	10k
SSD / HDD Capacity (GB)	149GB (Database) / 146GB (Log)
Controller	Dell PERC H800 1GB NV Cache
Controller Firmware / Driver	F/W 12.1.0091 / 4.31.1.64
PERC Cache Settings	Write Back Adaptive Read Ahead
Disk Firmware	D011

Appendix D - Database Software / Configuration / Benchmark Factory Configuration Details

Database Software	Microsoft® SQL Server® 2008 R2 Enterprise
Memory Configuration	1GB
Database Container Size	150GB
Log Size	100GB
Database Instance Memory	Limited to 1GB, ensured all I/O be serviced from disk
Primary DB File Location	HDD Database: 14-drive 10K RAID 10 volume SSD Database: 4-drive SAS SSD RAID 10 volume
Benchmark Software	Quest's Benchmark Factory for Databases
Version	6.1.1
Industry Standard Benchmark	TPC-E
Scale Factor	50
Complete Database Size	126GB
Number of Virtual Users	Unlimited
Data Collection Enabled	Yes, with additional logical disk metrics collected by BMF from Windows Performance Monitor