Accelerating Oracle OLTP with Compellent Flash Optimized Array

Dell Technical Solutions
November 2013

A Dell Reference Architecture

Revisions

Date	Description
November 2013	Initial release

Table of contents

Acknowledgements

Rajendraprasad B.Hosamani, Dell Product Marketing

Wes Vaske, Dell Global Solutions Engineering

Kai Yu, Dell Global Solutions Engineering

Mark Tomczik , Dell Technical Solutions Engineering

Executive summary

Online Transaction Processing (OLTP) information systems enable transaction oriented applications. Such systems are characterized by the need to store and retrieve a very large number of transactions. Naturally, a high performing and responsive storage system is a necessity for such systems. As the last mechanical component in the enterprise data center stack, Hard Disk Drive (HDD) based systems simply cannot meet the performance and latency requirements of the current and future OLTP systems. Traditionally, system architects have used methods such as disk overprovisioning, short stroking and manual performance tuning to keep up with the high performance requirements demanded by the OLTP information systems. Flash technology promises to eliminate the overprovisioning and the high-touch manual tuning by providing a storage medium with extreme performance characteristics, but at a high cost. In addition, most available solutions fail to unlock the true potential of flash because they treat flash simply as a single expensive storage medium with high performance characteristics.

The Dell Compellent Flash-optimized array is a radically new storage platform that leverages flash in an innovative way. The Compellent flexible storage architecture acknowledges that not all flash is created equal thereby leveraging the unique strengths of different types of flash media appropriately. As a result, the Dell Compellent Flash-optimized array redefines the economics of high performance storage by delivering a truly high performing storage platform that is more cost effective than pure HDD systems. The solution discussed in this paper describes and demonstrates a reference architecture that delivers flash performance for Oracle OLTP workloads at the price of mechanical disk. The testing compares a Dell Compellent Flash-Optimized solution to a similarly priced solution consisting of 15k SAS HDD's.

In this paper the benefits of the Compellent Flash-optimized array are presented by comparing its transactional performance under simulated (ORION based) and real world (TPC-C style) OLTP workloads to a similarly priced HDD based Compellent array. The key findings based on the testing conducted in Dell labs are:

- Improved Performance and reduced latency
  - Over 300% improvement in raw IOPS compared to HDD configuration with similar cost
  - Over 300% improvement in real world (TPC-C style) transactions per second (TPS)
  - Over 80% reduction in latency
- Higher Capacity
  - 14% more raw capacity
- Improved Power and Space efficiency
  - 45% reduction in power consumption
  - 40% reduction in Datacenter space requirements
- Enterprise Grade Features
  - Full enterprise functionality such as snapshots, replication and live volume
  - Enterprise grade monitoring and management functionality
  - Flexibility to tailor system configuration to cost effectively balance performance and capacity requirements

1 Introduction

The purpose of this white paper is to introduce the Dell Compellent Flash-optimized storage array and demonstrate how it can be leveraged to accelerate Oracle OLTP infrastructures.

1.1 Scope

The scope of this paper includes:

- Design of a high performance storage array
- Attributes of flash storage as it pertains to its deployment in enterprise grade storage systems
- Barriers to adoption of flash in the enterprise
- Dell Compellent Flash-optimized array and how it eliminates the barriers to flash adoption.
- Performance of the Dell Compellent Flash-optimized array for OLTP workloads and the value it delivers
- Architecture, design and components used in the solution

1.2 Assumptions

This document assumes familiarity with storage and database concepts. The reader should have a high level familiarity with concepts such as RAID, Fibre Channel, Serial Attached SCSI (SAS), LUN and general IP networking concepts. Familiarity with the Dell Compellent Storage Platform and Oracle Real Application Clusters, although not needed, will be beneficial in understanding the solution presented in this document.

This document is not intended to be a step by step configuration guide. Instead the focus of this document is to demonstrate the performance and the economic value the solution brings to Oracle OLTP deployments.

1.3 Audience

This document is intended for Oracle DBAs, database architects, storage administrators, architects and customers seeking to deploy cost effective Oracle OLTP infrastructures. The solution presented in this paper provides the audience with a more cost effective alternative to address the demanding performance requirements of Oracle OLTP infrastructures.

2 High performance storage and flash

Over the past decade, SAN-based shared storage has established itself as the staple technology to satisfy the ever growing appetite for data storage. As the scale of information stored on such devices grows, the storage industry has continued to innovate and provide technologies that help customers manage the ever-increasing storage requirements. While satisfying the capacity requirements has been simpler, the performance requirements are not; especially when they have to be cost effective.

Some common approaches storage vendors have taken to meet the performance requirements are:

- Use of expensive, high performance storage media such as solid state drives
- Implementing caching mechanisms
- Aggregating drive performance through carefully chosen RAID schemes
- Automatically tiering across different storage media based on access characteristics

Automated tiered storage solutions have emerged as one of the leading contenders in delivering high performance cost effectively. Unfortunately, the tiering solutions in the market treat flash as a single medium. As a result, current flash storage solutions fail to fully exploit the cost saving potential of flash. The high cost of flash storage solutions remains the biggest barrier to its wide spread adoption.

2.1 Understanding flash Storage

There are two main aspects of flash storage impacting its adoption. Flash is a fundamentally different medium and not all flash is created equal. Flash storage is composed of electronic NAND gates whereas HDDs are composed of rotating magnetic media. This results in flash storage having very different performance, cost and data retention characteristics.

The core component of an SSD is the NAND flash—of which there are two basic types—SLC and MLC. Unlike magnetic media on HDD storage, data stored on flash needs to be erased before new data can be written or programmed; —this is known as the Program-Erase Cycle (PE/C). The maximum number of PE/Cs for NAND is dependent on the technology (SLC or MLC). This is typically on the order of a few thousand cycles per NAND cell, after which, the performance and reliability of the flash storage cannot be guaranteed. This characteristic of flash technology limits the number of write operations that can be performed on a flash drive.

Each cell in SLC NAND is capable of storing a single bit of data. MLC NAND, on the other hand, can store multiple bits of data in a single cell. This enables SLC drives to write faster and achieve high cell endurance while making the drives more expensive than their MLC siblings. Cell endurance is defined as the number of times the media can be rewritten (erased and programmed). MLC NAND has significantly higher memory density, thereby reducing cost. This comes at the expense of slower write speeds and significantly lower cell endurance. However, MLC NAND still possesses exceptional random read performance.

As with HDDs, SSDs are typically developed and sold for two distinct markets: enterprise and personal storage. Enterprise-grade SSDs typically have features not found on consumer products such as nonvolatile write cache, significant amounts of NAND over-provisioning, more write channels, built in advanced error detection/correction and a dual-ported SAS interface. All of these features are important for data integrity, high availability and enterprise-grade performance.

2.2 Flash Tiering: An untapped opportunity

The two classifications for enterprise-grade SAS SSDs that Compellent uses are write-intensive (WI) and read-intensive (RI). The term write-intensive SSD refers to an SLC SSD, and the term read-intensive SSD refers to an MLC SSD. The main distinctions between these drive types are their endurance specifications, capacities and cost. Compellent has been using write-intensive SSDs for more than five years, and has found that the endurance characteristics of these drives make them a great choice for deployment in the data center. Adding read-intensive SSD's opens up possibilities for additional optimizations. Advanced firmware intelligence in the storage controller software is required to unlock this potential. Figure 1 below summarizes the differences between WI and RI SSDs. Note that the RI SSDs can deliver exceptional read performance and higher capacity for significantly lower cost.

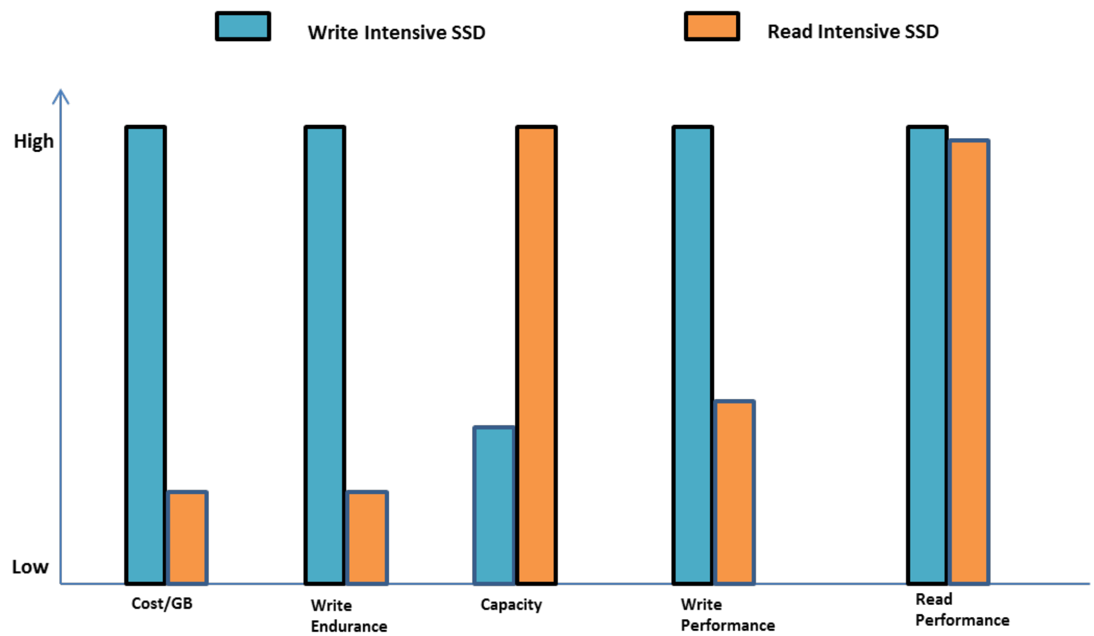

Figure 1 Write-intensive versus read-intensive SSDs

Traditional tiering solutions fail to take advantage of the differences between WI and RI drives. By treating Flash as a single type of storage medium, traditional tiering solutions are unable to leverage the performance, cost and capacity attributes of different type of SSD drives. As a result, solutions that use only write-intensive SSDs are expensive, while solutions that use only read-intensive drives (MLC) do not meet the demanding reliability requirements for deployment in the enterprise datacenter. Opportunity exists to intelligently combine the unique characteristics of WI and RI SSDs to lower the cost of high performance storage. A fundamentally flexible storage architecture is required to leverage the unique cost, capacity and performance characteristics of RI and WI SSDs.

2.3 Compellent Flash Optimization

The Compellent Storage Center 6.4 (SC6.4) release builds on the robust storage virtualization foundation of the Dell Compellent Fluid Data Architecture. The core features that enable Dell Compellent to deliver cost effective high performance are:

- **Flash aware data progression**: Whereas the traditional data progression algorithms would run only once each day to optimize RAID levels and data tiering, Flash-optimized data progression has the capability to move data across tiers throughout the day. As an example, data that has recently been frozen in a replay (such as a snapshot) is now read-only and a great candidate for being moved to a read-intensive tier of storage. This is all handled in the background with little to no impact to the host. By moving the data to a RAID level and storage media that has exceptional read I/O performance characteristics (for example space efficient RAID 5 and inexpensive MLC SSDs), the cost of storing the read only data is significantly reduced.

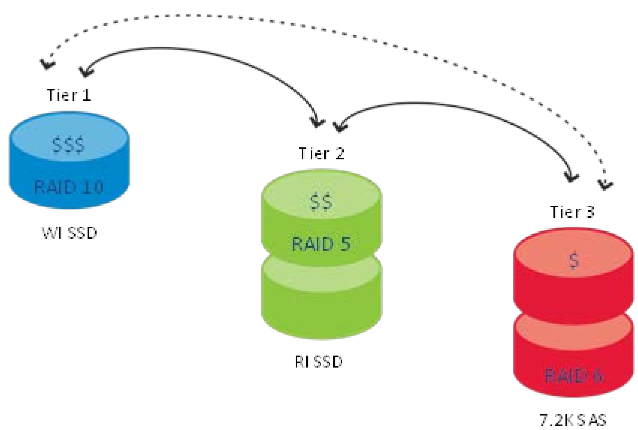

Figure 2 Compellent data progression

- **Flash aware storage profiles**: Storage profiles are a feature of the Dell Compellent storage platform that radically simplifies tiering policy management. By assigning a storage profile to a volume, the volume inherits the tiering attributes defined in the storage profile. SC6.4 has added two new profiles in order to place data on the most appropriate types of drives, Flash-optimized and cost-optimized storage profiles. A Flash-optimized storage profile directs all writes to high-performance Tier 1 storage (SLC SSDs in RAID 10,). On-demand data progression automatically converts replay and read-intensive data to RAID 5 and places it on Tier 2 read-intensive flash storage that has exceptional read performance characteristics. A cost-optimized storage profile allows the administrator to create a volume for cost-optimized applications such as backup, archive and low priority applications. The new storage profiles simplify tiering management which enables storage administrators to assign flash storage to the most active data while retaining the flexibility of storing cold data in a more cost effective tier.
- **Proactive WI SSD space management:** The Compellent architecture directs all writes to high performance Tier 1 storage. It is very important that the Tier 1 is used judiciously to ensure that sufficient space in Tier 1 is always available for incoming host writes. Enhancements in SC6.4

proactively create space management replays to move data to Tier 2; these are automatically triggered when Tier 1 space consumption exceeds predefined thresholds. This ensures that sufficient Tier 1 space is available to service new writes. Tier 1 space management happens in the background without administrator involvement.

- **Management and monitoring:** Compellent Storage Center monitoring tools have been enhanced with the ability to monitor wear on any SSD in the array. Additionally, since the flash storage will yield dramatic reduction in I/O latencies, Dell Compellent Enterprise Manager and Storage Center user interfaces now report sub-millisecond latencies to provide a more precise view of flash storage performance.

The combined effect of all these changes is illustrated in Figure 3. Compared to traditional tiered storage architectures, the enhancements in SC6.4 significantly improve the average performance of the system. The low cost and high densities of the read-intensive SSDs contribute to increasing the usable SSD capacity without increasing the cost of the solution.

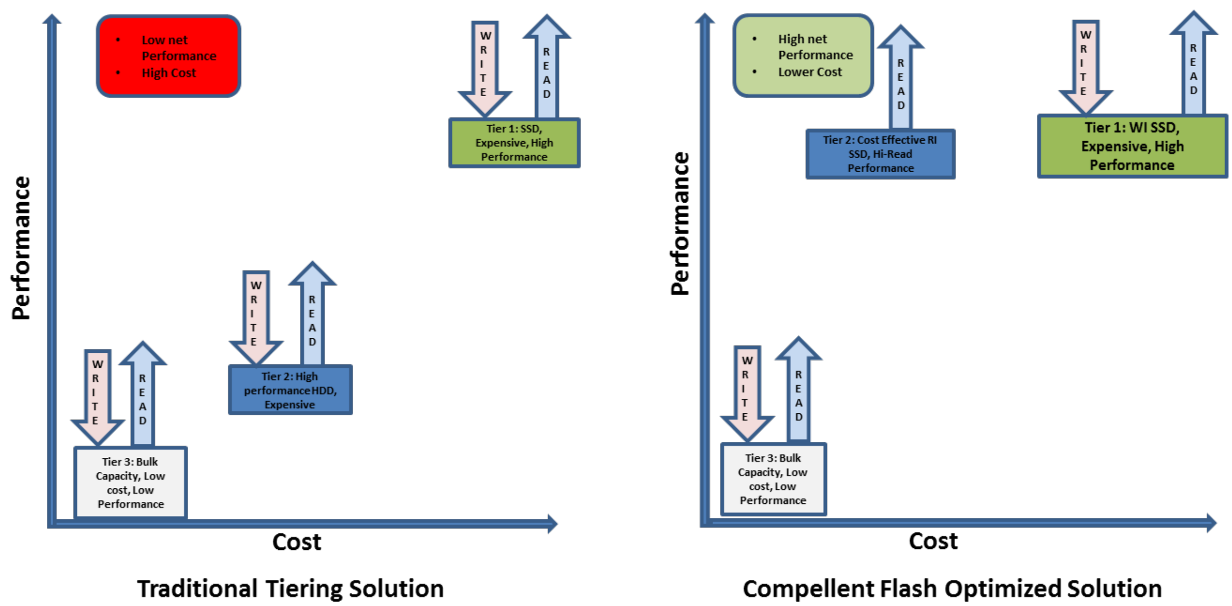

Figure 3 Traditional tiering versus the Compellent Flash-optimized solution

3 Developing an Oracle OLTP reference architecture

Database applications can broadly be classified into two types based on the nature of their data access. OLTP represents applications that are transactional in nature where database applications read and write in small block sizes. These systems are actively creating new data and or updating existing data. Typical examples of OLTP systems include order processing systems and financial transaction systems. Online Application Processing (OLAP) systems are a class deployed predominantly for developing Business Intelligence (BI) and Decision Support. The data access made by these applications differs from the OLTP systems. Since OLAP systems are typically analyzing existing data, they read large amounts of data, in large block sizes, to be processed by powerful servers.

This paper discusses the development and characterization of OLTP reference architectures in order to highlight the benefits of the Compellent Flash-optimized array when compared to a Compellent array with high performance HDDs.

3.1 Design Goals

Enhancements in SC6.4 seek to change the economics of high performance storage. Through careful selection of SSD components and innovative software enhancements, Dell Compellent delivers industry leading performance without compromising the cost or capacity offered by the storage system. This provides a new, cost-effective alternative to address the demanding performance and capacity requirements of OLTP deployments.

Common methods to address the high IOPS and low latency requirements of OLTP deployments are:

- **Performance driven overprovisioning**: High performance HDDs, such as 15k RPM drives, are the leading choice in addressing the performance and latency requirements of OLTP applications. The performance of the 15k SAS drives is limited by the mechanical nature of the HDD. As a result, this approach invariably leads to capacity overprovisioning. In other words, customers will (in most cases) be required to buy more capacity than needed in order to meet the performance requirements.
- **Implementing all flash storage**: There are several flash solutions available in the market. These solutions, while delivering the needed performance, are typically very expensive and may not satisfy the capacity requirements.
- **Hybrid storage solutions**: Another common approach is to implement hybrid storage solutions which implement a flash based cache or flash as a tier within the storage system. These solutions treat flash as a single medium and typically deploy higher cost SLC based SSDs.

The design goal for this reference architecture is to demonstrate that the Compellent Flash-optimized solution eliminates the pain points associated with traditional approaches. This paper demonstrates that the Compellent Flash-optimized solution:

- Delivers high performance and low latency
- Eliminates the need for overprovisioning
- Delivers cost effective performance compared to other flash offerings

3.2 Test methodology

The test methodology used compared the performance, latency, power, and space characteristics of two Compellent based storage solutions with similar cost characteristics. One solution deployed high performance 15k HDDs which will be compared to a flash optimized configuration. The testing began with a flash optimized configuration consisting of six 400G Write-intensive SSDs plus six 1.6TB Read-intensive SSDs. This configuration was cost equivalent to a homogeneous HDD configuration consisting of 72 146G 15k RPM SAS drives.

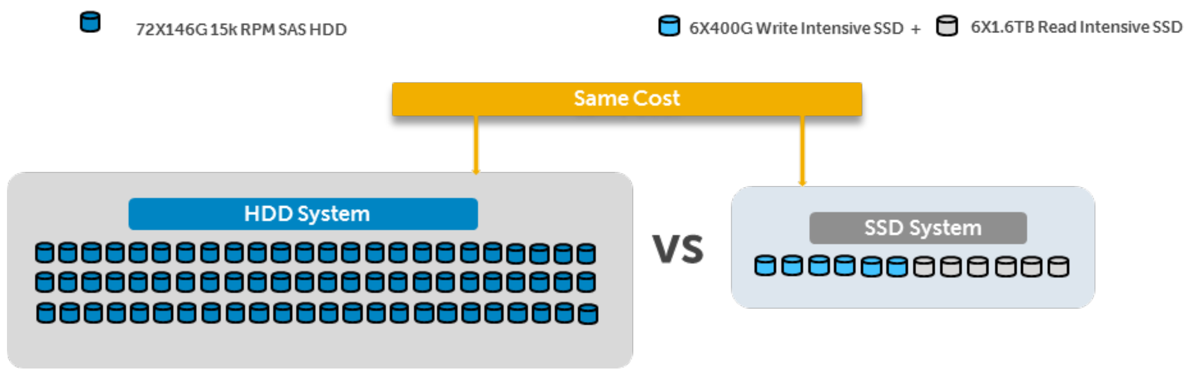

Figure 4 HDD versus an SSD system

The testing was performed in two phases:

- **ORION I/O profiling**: In the first phase, a series of tests were executed using the Oracle I/O Calibration Tool (ORION)[1]. ORION uses the Oracle I/O software stack to generate simulated I/O workloads without having to create an Oracle database, load applications, and simulate users. ORION is commonly used for benchmarking the I/O performance characteristics of Oracle database storage systems. These tests were focused on measuring the storage I/O performance and system limits for pure OLTP workloads.
- **Real world TPC-C style transactional workload**: In the second phase, Benchmark Factory[2] for Databases was used to simulate a real TPC-C style workload. The workload was tested on a two node Oracle 11gR2 RAC database to characterize the maximum TPS capability of the two configurations for an end-to-end transactional response time of less than a second. The testing simulates a real world 500GB OLTP deployment.

The suggested reference architecture is based on the testing conducted with TPC-C style transactional workload. Detailed storage and database configuration information of the reference architecture is presented in Appendix A.

---

[1] http://www.oracle.com/technetwork/topics/index-089595.html
[2] http://www.quest.com/benchmark-factory/

4 Results and analysis

4.1 ORION I/O profiling

The ORION tool was used to simulate OLTP database transactions. The goal of the test was to measure the maximum IOPS and latency capability of the HDD based configuration and the Flash-optimized configuration.

Table 1 ORION test parameters

Parameter	Purpose
-oltp	Simulates random 8k I/O
-write	Set to 0 for 100% read workload Set to 30 to create a workload with 70/30 R/W ratio

The IOPS and latency observed on the two configurations are shown below.

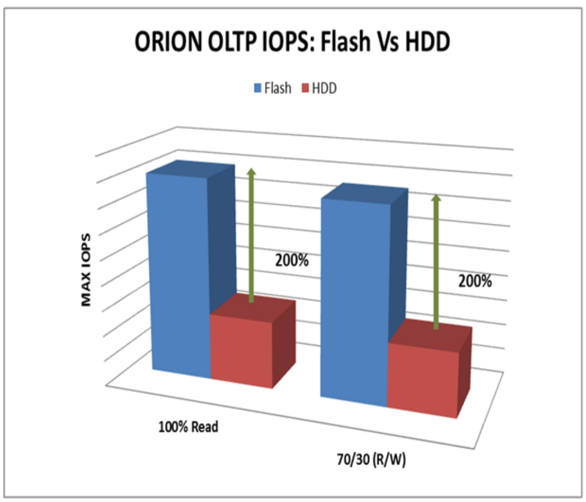

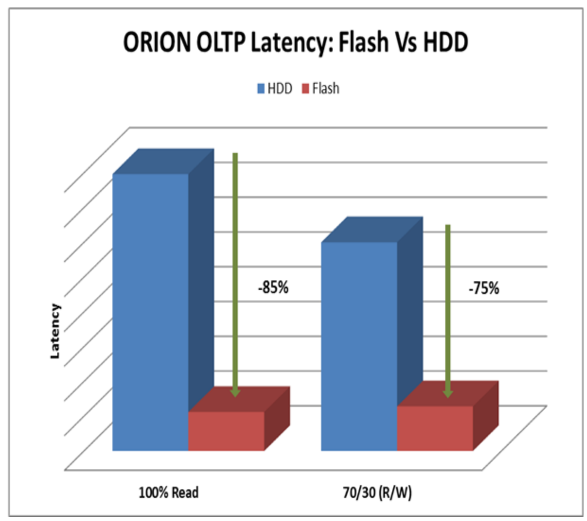

Figure 5 Flash Vs HDD: Comparing OLTP IOPS and Latency

4.1.1 Key takeaways

- The flash optimized configuration delivered up to **three times the raw IOPS** compared to the pure HDD configuration for both read-only and mixed workloads.
- The flash optimized configuration delivered three times the raw IOPS at **85% lower latency for read-only** workloads and **75% lower latency for mixed workloads**.

Why this matters:
The Compellent Flash-optimized array delivers up to three times higher performance at one forth the latency making it a compelling choice for accelerating Oracle OLTP applications.

4.2 TPC-C style workload analysis

A TPC-C style workload was chosen to simulate a real-world OLTP workload. The TPC-C style workload simulates an order entry deployment where orders from multiple concurrent users are processed by the Oracle database. This testing was used to understand the impact of the performance enhancing features of the Compellent Flash-optimized array on a real-world workload. In this test, maximum sustained TPS and the Application Response Time at maximum TPS were monitored. Additionally, the user load at max sustained TPS was also monitored. These metrics provide insight into the business value associated with the performance enhancements.

Table 2 TPC-C workload test results

Configuration	Max Sustained TPS	Application Response Time at max TPS (in sec)	User Load at Max TPS
HDD Configuration	985	0.60	2600
Flash Configuration	3400	0.58	9000

4.2.1 Key takeaways

- The Compellent Flash optimized array delivers up to **3.5 times the TPS** compared to the HDD configuration that is similar in cost.
- This translates directly to a **70% reduction in cost per TPS**
- The Compellent Flash-optimized array can support up to **3.5 times more concurrent users without compromising application SLAs** (maintains low-latency at max TPS)

Why this matters:
The Compellent Flash-optimized array can deliver a 3.5 times increase in concurrent user load while reducing cost per TPS by 70%. The increased concurrent user load coupled with reduced cost per transaction opens up opportunities for database server and storage consolidation which can result in significant hardware and Oracle licensing savings.

4.3 Comparing Capacity, Power and Space consumption

In designing high performance storage, several criteria have to be taken into account. Flash storage, while delivering high performance, has traditionally not been able to match the capacity offered by HDD's. Furthermore, space and power consumption in the data center are also major factors that system architects have to balance. Achieving an overall balance has been one of the major design criteria for the

Compellent Storage system. Figure 6 compares the capacity, power consumption and space consumption of the two configurations tested as part of this effort.

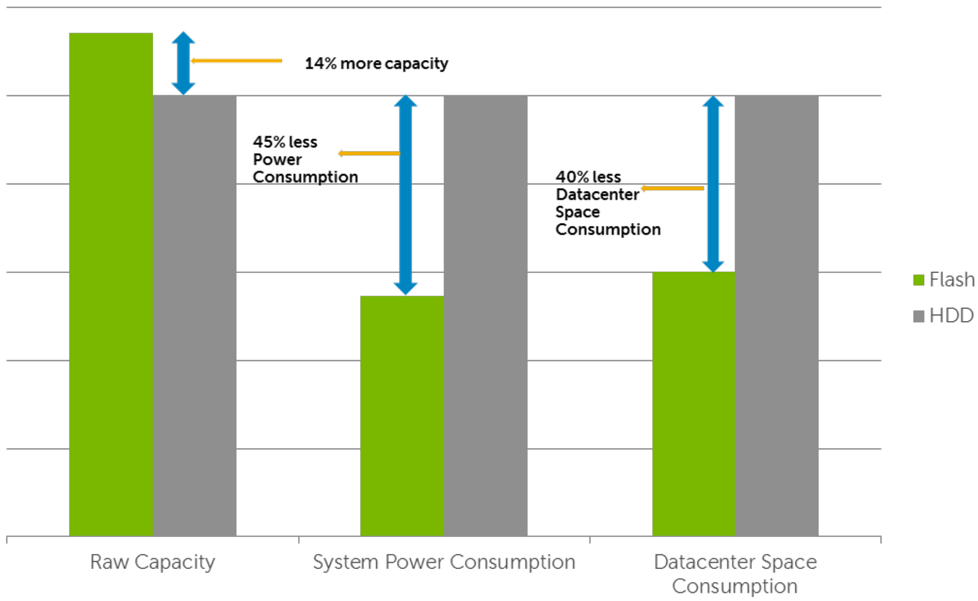

Figure 6 Flash versus HDD: Capacity, power and datacenter space

4.3.1 Key takeaways

- The Compellent Flash-optimized array delivers **14% more high-performance capacity** (12TB raw flash capacity) for the same cost compared to the HDD configuration (10.5TB raw).
- The Compellent Flash-optimized array **consumes 45% less power** than the HDD based configuration.
- The Compellent Flash-optimized array **consumes (6U) 40% less datacenter space** compared to the HDD based configuration (10U).

Why this matters?
Deploying the Compellent Flash-optimized array can reduce the environmental foot print and result in cost savings associated with power/cooling and data center space.

5 Conclusion

The Dell Compellent Storage platform differentiates itself through its unique approach to storage virtualization and tiering. The SC6.4 platform extends its leadership in designing well balanced enterprise-grade storage systems to include flash storage. Flash storage holds the potential to radically change the performance characteristics of the storage systems used in the datacenter. Historically, cost and capacity concerns have held back flash storage from reaching its true potential. The Dynamic Block Architecture foundation of the Dell Compellent platform enables building storage platforms that transcend traditional limitations.

With SC6.4, Dell Compellent has eliminated the last barrier to wide spread adoption of flash storage in the enterprise. By combining robust and flexible storage architecture with practical components and innovative software features, the Dell Compellent Flash-optimized storage platform produces high performance flash technology without significant CAPEX increases. This redefines the economics of high performance storage and empowers administrators to do more. The high performance storage enables customers to exceed end-user SLAs, accommodate more concurrent users, and explore database consolidation opportunities to reduce environmental footprint and reduce other CAPEX line items. Dell Compellent is a truly balanced, high performance, storage platform that delivers the most cost effective storage performance in the industry.

A Compellent Flash-optimized reference architecture for Oracle OLTP

The SSD based solution presented in this section is a cost effective alternative to deploying a HDD based infrastructure. In comparison to a Compellent based solution consisting of mechanical HDDs, this reference architecture is capable of delivering:

- 12TB raw storage capacity; up to 14% more than the HDD based configuration
- A sustained throughput of 3,400 TPS , or 204,000 transactions per minute, with a sub second response time under a real world TPC-C style workload
- Support for up to 9,000 concurrent users while maintaining a sub second response time
- Up to 300% more raw IOPS and 75% lower latency
- A 67% reduction in cost per transaction while supporting 350% more concurrent users
- Up to 45% reduction in power consumption and 40% reduction in space consumption

The hardware resources used in this reference architecture are listed in Table 3.

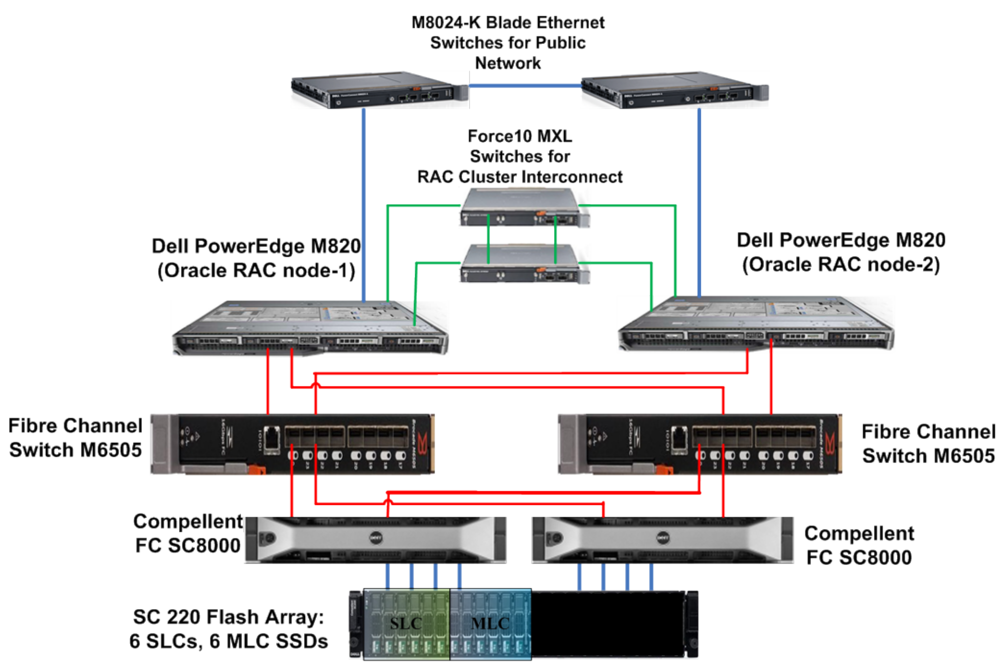

Figure 7 Reference architecture

A.1 Hardware and configuration

A.1.1 Physical server hardware

The servers implemented for this reference architecture were Dell PowerEdge M820 Blade servers (12th Generation). These servers have Intel Xeon E5-4600 series processors offering up to 32 cores per server and 1.5TB of memory. All network interfaces used two dual-port cards to provide switch and HBA redundancy. The half-height M620 Blade server provides the same CPU and memory density but does not allow for HBA redundancy—a feature specific to full-height blades.

Table 3 Server hardware

Component	Description
CPU	Quad Intel Xeon E5-4650 per server (8-cores, 2.7GHz, 20M cache)
Memory	512GB per server (32 x16GB dimms at 1600MHz)
Operating System	Oracle Enterprise Linux 6.4
UEK Revision	kernel-uek-2.6.39-400.109.4.el6uek.x86_64
Fibre Channel Interfaces	Two dual-port Emulex 16 Gb FibreChannel HBAs per server • Model: LPe16002-D • Firmware: 1.1.21.4 • Driver: 8.3.7.10.4p
Cluster Interconnect Interfaces	Two dual-port Broadcom NetXtreme II 10 Gb NICS per server • Model: BCM57810) • Firmware: 7.4.8 • Driver: 1.76.54
Client IP Interfaces	Two dual-port Broadcom NetXtreme II 10 Gb NICS per server • Model: BCM57810) • Firmware: 7.4.8 • Driver: 1.76.54
OS Disk	4x 146GB 15k SAS Disks in RAID10 Dell PERC H310 Mini RAID Controller (Embedded) • Firmware: 20.12.0-0004 • Driver: 06.505.02.00

A.1.2 Storage hardware

Storage was provided by a dual-controller Compellent array connected to a single SC220. Front-end connectivity was provided by a dual-port 16Gb Fibre Channel HBA in each controller. HBA redundancy was supported by adding one or more dual-port 16Gb Fibre Channel HBA per controller. Back-end connectivity was provided by one quad-port 6Gb SAS per controller.

Component	Description
Storage Controller	2x SC8000
Storage firmware revision	Storage Center 6.4
Enclosure	1x SC220 (24x 2.5" Drives)
Write-intensive drives	6x 400GB WI SSD (SLC)
Read-intensive drives	6x 1.6TB RI SSD (MLC)
Backend connectivity	6Gb SAS
Frontend connectivity	One dual-port 16Gb FC HBA per controller

A.1.3 Network switching hardware

The PowerEdge m1000e blade chassis supports three redundant fabrics (labeled: Fabric A, Fabric B and Fabric C). In this architecture the Public Interface (Client IP) used Dell PowerConnect M8024-k 10 GbE switches that connected to the lab network via dual-port LAG (Fabric A). The Private Interconnect for Cluster Services used Dell Force10 MXL 10/40GbEW switches that used two 40 GbE ports in a LAG to connect the two switches (Fabric B). The Fibre Channel network used redundant, isolated Brocade M6505 16Gb switches (Fabric C).

Table 4

Component	Description
Fibre Channel switch	Brocade M6505 Gen5 16Gb FC blade switch
Cluster Interconnect switch	Dell Force 10 MXL 10/40GbE blade switch
Client IP switch	Dell PowerConnect M8024-K 10GbE blade switch

A.2 Oracle RAC configuration

This architecture was designed for and tested with Oracle Database 11gR2 (11.2.0.3) in a two node RAC configuration. The tested database was 1TB with space for growth to 2TB. The load was generated by Benchmark Factory using the TPC-C benchmark. Oracle System Global Area (SGA) was artificially limited to 16GB per node to increase the load on the storage system. In practice, Dell recommends using Oracle Automatic Memory Management to fully exploit the RAM available.

Table 5 Oracle Database configuration

Component	Description
Database type	Oracle 11gR2 (11.2.0.3.0) 2-node RAC
Database size	1TB
ASM	Yes—with external redundancy
Oracle System Global Area (SGA)	16GB per node[3]
Load generator	BenchMark Factory TPC-C style workload
User scaling	100 – 11,900 in increments of 200 users

A.2.1 Redo log tuning

The default starting configuration for Oracle Database 11gR2 provides two redo log files per server (50MB each). This causes severe performance degradation if not tuned for the specific workload. For the tests presented, ten redo log files per server (20 total) of 250MB each were created. The number and size of redo logs are specific to individual environments.

A.2.2 Processes and sessions

The processes and sessions parameters need to be increased from the default values to prevent limited performance. For these tests, the processes parameter was set to 8000 and the sessions parameter to 8805.

[3] The SGA Target in this test was set far below the available memory to reduce caching and increase the load on the storage subsystem. In a production environment, we recommend using Automatic Memory Management.

A.3 Oracle storage layout

The logical storage layout used in this reference architecture is shown in Figure 7. Oracle RAC was configured with ASM (external redundancy). A dedicated disk group was created for DATA, FRA, OCR and REDO logs.

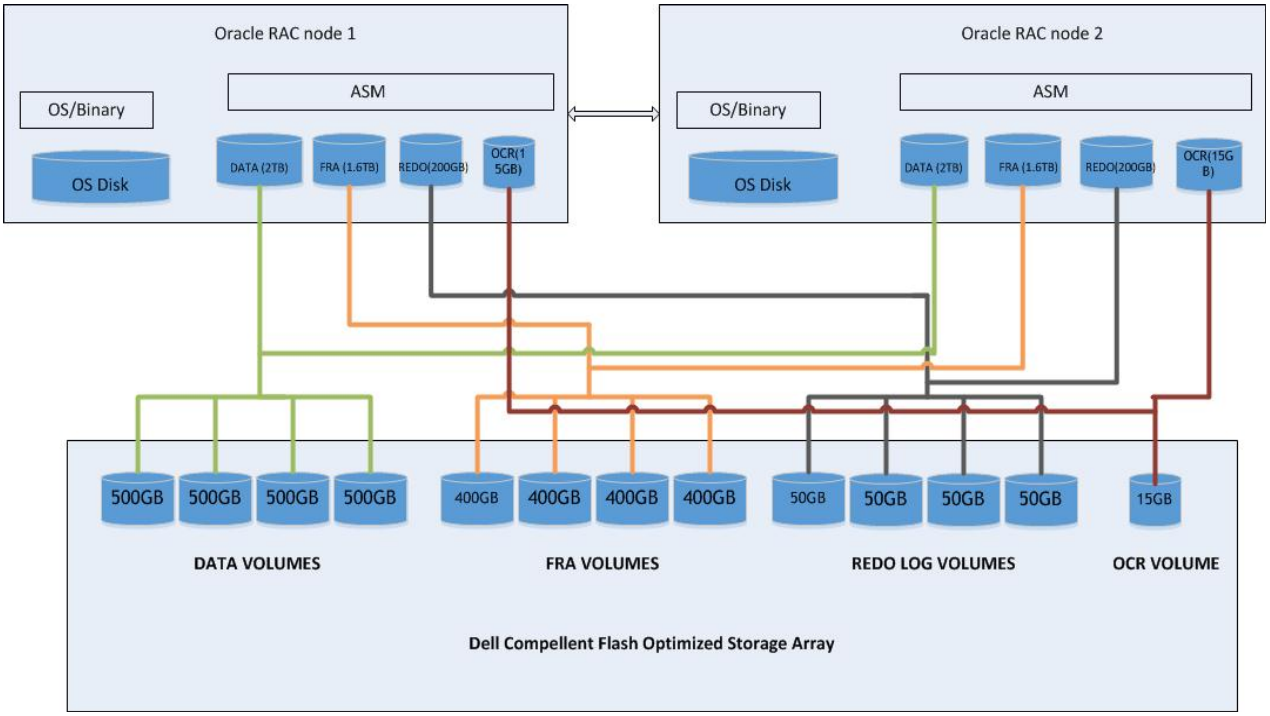

Figure 8 Oracle storage logical Layout

Table 6 Storage layout

Disk Group	Storage Profile	Number of Luns	Total Disk Group Size
DATA, temp, index	Flash Optimized	4X500GB	2TB
FRA[4]	Flash Optimized	4X400GB	1.6TB
REDO Logs	Flash Optimized	4X50GB	200GB
OCR	Flash Optimized	1X15G	15GB

[4] A small FRA Disk Group sufficed for the testing conducted as part of this effort since the test plan did not involve archiving use cases. Dell recommends following the Compellent Oracle Deployment best practices recommendations referenced in Appendix C to size the FRA disk group appropriately

A.4 Benchmark Factory table layout

The following script was used to create the tables for the performance tests. This differs slightly from the standard TPC-C design by using a clustered C_DISTRICT table. The data is initially created in SCHEMA1 and loaded into the testing schema via CTAS.

Table 7 SQL for creating tables

```
CREATE TABLE C_WAREHOUSE
NOLOGGING NOCACHE NOMONITORING
PARALLEL (DEGREE 8) NOCOMPRESS
AS
SELECT * FROM SCHEMA1.C_WAREHOUSE;

create cluster c_district_cluster (
  d_id number,
  d_w_id number
)
single table
size 4096
hashkeys 500;

CREATE TABLE C_DISTRICT
cluster c_district_cluster (d_id, d_w_id)
NOMONITORING
NOCOMPRESS
AS
SELECT * FROM SCHEMA1.C_DISTRICT;

CREATE TABLE C_CUSTOMER
NOLOGGING NOCACHE NOMONITORING
PARALLEL (DEGREE 8) NOCOMPRESS
AS
SELECT * FROM SCHEMA1.C_CUSTOMER;

CREATE TABLE C_ORDER
NOLOGGING NOCACHE NOMONITORING
PARALLEL (DEGREE 8) NOCOMPRESS
AS
SELECT * FROM SCHEMA1.C_ORDER;
```

```
CREATE TABLE C_HISTORY
NOLOGGING NOCACHE NOMONITORING
PARALLEL (DEGREE 8) NOCOMPRESS
AS
SELECT * FROM SCHEMA1.C_HISTORY;

CREATE TABLE C_NEW_ORDER
NOLOGGING NOCACHE NOMONITORING
PARALLEL (DEGREE 8) NOCOMPRESS
AS
SELECT * FROM SCHEMA1.C_NEW_ORDER;

CREATE TABLE C_ORDER_LINE
NOLOGGING NOCACHE NOMONITORING
PARALLEL (DEGREE 8) NOCOMPRESS
AS
SELECT * FROM SCHEMA1.C_ORDER_LINE;

CREATE TABLE C_ITEM
NOLOGGING NOCACHE NOMONITORING
PARALLEL (DEGREE 8) NOCOMPRESS
AS
SELECT * FROM SCHEMA1.C_ITEM;

CREATE TABLE C_STOCK
NOLOGGING NOCACHE NOMONITORING
PARALLEL (DEGREE 8) NOCOMPRESS
AS
SELECT * FROM SCHEMA1.C_STOCK;
```

Table 8 SQL for creating indexes

CREATE UNIQUE INDEX C_WAREHOUSE_I1 ON C_WAREHOUSE (W_ID) NOLOGGING PARALLEL (DEGREE 8) INITRANS 8; CREATE UNIQUE INDEX C_DISTRICT_I1 ON C_DISTRICT (D_W_ID, D_ID) NOLOGGING PARALLEL (DEGREE 8) INITRANS 8; CREATE UNIQUE INDEX C_CUSTOMER_I1 ON C_CUSTOMER (C_W_ID, C_D_ID, C_ID) NOLOGGING PARALLEL (DEGREE 8) INITRANS 8; CREATE INDEX C_CUSTOMER_I2 ON C_CUSTOMER (C_W_ID, C_D_ID, C_LAST, C_FIRST) NOLOGGING PARALLEL (DEGREE 8) INITRANS 8; CREATE UNIQUE INDEX C_ORDER_I1 ON C_ORDER (O_W_ID, O_D_ID, O_C_ID, O_ID) NOLOGGING PARALLEL (DEGREE 8) INITRANS 8; CREATE UNIQUE INDEX C_NEW_ORDER_I1 ON C_NEW_ORDER (NO_W_ID, NO_D_ID, NO_O_ID) NOLOGGING PARALLEL (DEGREE 8) INITRANS 8; create index c_order_i2 on c_order(o_w_id) NOLOGGING PARALLEL (DEGREE 8) INITRANS 8; create index c_order_i3 on c_order(o_d_id) NOLOGGING PARALLEL (DEGREE 8) INITRANS 8; create index c_order_i4 on c_order(o_c_id) NOLOGGING PARALLEL (DEGREE 8) INITRANS 8; CREATE UNIQUE INDEX C_ORDER_LINE_I1 ON C_ORDER_LINE (OL_W_ID, OL_D_ID, OL_O_ID, OL_NUMBER) NOLOGGING PARALLEL (DEGREE 8) INITRANS 8; CREATE UNIQUE INDEX C_ITEM_I1 ON C_ITEM (I_ID) NOLOGGING PARALLEL (DEGREE 8) INITRANS 8; CREATE UNIQUE INDEX C_STOCK_I1 ON C_STOCK (S_I_ID, S_W_ID) NOLOGGING PARALLEL (DEGREE 8) INITRANS 8;	ALTER TABLE C_WAREHOUSE NOPARALLEL; ALTER TABLE C_STOCK NOPARALLEL; ALTER TABLE C_ITEM NOPARALLEL; ALTER TABLE C_DISTRICT NOPARALLEL; ALTER TABLE C_CUSTOMER NOPARALLEL; ALTER TABLE C_ORDER NOPARALLEL; ALTER TABLE C_ORDER_LINE NOPARALLEL; ALTER TABLE C_NEW_ORDER NOPARALLEL; ALTER TABLE C_HISTORY NOPARALLEL; ALTER INDEX C_WAREHOUSE_I1 MONITORING USAGE NOPARALLEL; ALTER INDEX C_DISTRICT_I1 MONITORING USAGE NOPARALLEL; ALTER INDEX C_CUSTOMER_I1 MONITORING USAGE NOPARALLEL; ALTER INDEX C_CUSTOMER_I2 MONITORING USAGE NOPARALLEL; ALTER INDEX C_ORDER_I1 MONITORING USAGE NOPARALLEL; ALTER INDEX C_ORDER_I2 MONITORING USAGE NOPARALLEL; ALTER INDEX C_ORDER_I3 MONITORING USAGE NOPARALLEL; ALTER INDEX C_ORDER_I4 MONITORING USAGE NOPARALLEL; ALTER INDEX C_NEW_ORDER_I1 MONITORING USAGE NOPARALLEL; ALTER INDEX C_ORDER_LINE_I1 MONITORING USAGE NOPARALLEL; ALTER INDEX C_ITEM_I1 MONITORING USAGE NOPARALLEL; ALTER INDEX C_STOCK_I1 MONITORING USAGE NOPARALLEL; exit

Additional resources

Support.dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com is an IT Community where you can connect with Dell Customers and Dell employees for the purpose of sharing knowledge, best practices, and information about Dell products and installations.

Referenced or recommended Dell publications:

- How to deploy Oracle 11gR2 on RHEL6/Oracle Linux 6:
  http://en.community.dell.com/techcenter/enterprise-solutions/m/oracle_db_gallery/20212999.aspx

- Oracle Best Practices on Compellent Storage Center
  http://www.dellstorage.com/WorkArea/DownloadAsset.aspx?id=3055

- New economies of storage with the Compellent Flash-optimized solutions
  http://partnerdirect.dell.com/sites/channel/Documents/Dell-Compellent-New-Economies-of-Storage-Flash-Optimized-Solutions-Business-White-Paper.pdf

Referenced or recommended Oracle publications:

- Oracle 11g Database Concepts:
- http://download.oracle.com/docs/cd/B28359_01/server.111/b28318/toc.htm
-
- Oracle Database Storage Administrators Guide
- http://download.oracle.com/docs/cd/B28359_01/server.111/b31107.pdf