Microsoft SQL Server 2012 Fast Track Reference Configuration using PowerEdge R720 and PowerVault MD3620i

This whitepaper describes the Dell Microsoft SQL Server Fast Track reference architecture configuration and performance results

Jisha J

Dell Database Solutions Engineering

This document 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.

July 2012| Rev 1.0

Contents

Tables

Figures

Executive summary

Dell™ and Microsoft®, in cooperation, provide guidelines and design principles to assist customers in designing and implementing a balanced configuration for Microsoft SQL Server® data warehouse workloads to achieve "out-of-the-box" scalable performance. These database reference architectures enable each of the components in the database stack to provide optimal throughput to match the database capabilities of the specific setup. Dell 12th generation servers, along with the robust storage backends, form efficient candidates for a high performing data warehouse solution.

This white paper describes the architectural design principles and guidelines used to achieve an optimally balanced SQL Server Data Warehouse solution using Dell PowerEdge R720 and PowerVault MD3620i iSCSI SAN storage along with Microsoft SQL Server 2012. The performance numbers achieved for the reference configuration are also discussed in detail.

Introduction to Microsoft SQL Server Fast Track reference architectures

The objective behind the database Fast Track reference configurations is to provide guidance and resources to help customers identify the right architecture and configuration for data warehouse solutions. The following are tangible benefits of following these recommended configuration best practices and guidelines:

- Accelerated data warehouse projects with pre-tested hardware configurations
- Reduced hardware and maintenance costs by purchasing the right balanced hardware solution and optimizing it for a data warehouse workload.
- Reduced planning and setup costs leveraging the Certified Reference Architecture configurations.
- Predictable performance by configuring the system correctly, and taking advantage of the tuning directions.

The Fast Track reference architectures avoid the risk of improperly designed and configured systems. These guidelines ensure that the capability and throughput for the entire system is maximally utilized. Configurations are designed specifically for the data warehouse or BI (Business Intelligence) systems.

Visit msdn.microsoft.com/en-us/library/dd459146%28v=sql.100%29.aspx for more information on Fast Track Data Warehouse architectures.

What's new with Dell Microsoft Fast Track reference architectures?

Dell and Microsoft have refreshed the Fast Track reference architecture offerings with the latest technology advancements. Microsoft SQL Server 2012 and Dell PowerEdge™ 12th generation servers are the latest additions to the list of reference architecture components.

Microsoft SQL Server 2012 comes with several exciting features which directly benefit database environments. Column Store Index is one of these important features, which impacts the data warehouse database configurations. This feature enables storing data in columnar fashion, in contrast

to the traditional row-based approach. This technology enables better compression rates within the database, which is very beneficial for data warehouses because of the huge amount of data handled. Column store indexes also benefit common data warehousing queries such as filtering, aggregating, grouping, and star-join queries. (See msdn.microsoft.com/en-us/library/gg492088.aspx for more information on Column Store Indexes.) Microsoft Fast Track Data Warehouse 4.0 guidelines incorporate the benefit of Column Store Indexes for improved query performance.

Dell's 12th generation server platforms, featuring enhanced onboard memory, storage, and processor speeds, have been optimized for better data warehouse performance. The latest Intel® Xeon® E5 series processors, larger memory capacities, higher memory speed, and third generation PCI Express slots on the newer PowerEdge platforms ensure faster database throughput.

The Dell Microsoft Fast Track reference architectures are engineered jointly by Dell and Microsoft. The hardware-based optimizations and best practices are proposed by Dell and the software-based optimizations are proposed by Microsoft. This approach presents "Faster time-to-value" using integrated, balanced, and verified architectures.

Dell Microsoft SQL Server 2012 Fast Track reference architectures using PowerEdge R720 and PowerVault MD3620i

We are proposing two different fast track reference architectures comprised of SQL Server 2012, PowerEdge R720, and PowerVault™ MD3620i. Table 1 lists the Proposed Reference architectures along with the assigned Solution IDs.

Table 1. Recommended reference architectures with Solution IDs

Reference Architecture	Solution ID
Single Server Reference Architecture	2405036
Highly Available Reference Architecture	2405110

The following sections describe the hardware and software configuration details for the recommended reference architectures.

Single server reference architecture

Figure 1 depicts the single server reference architecture. Table 2 lists the details of the configuration.

Figure 1. Single server reference architecture

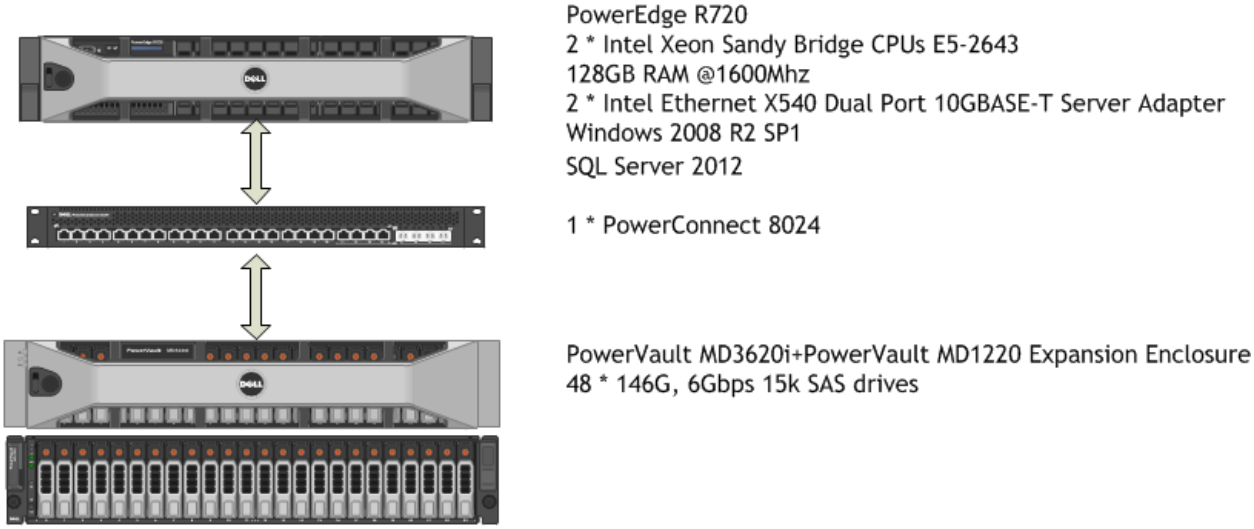

Table 2. Single server reference architecture details

Hardware Components	Details
Server	PowerEdge R720
Processors	2* Intel Xeon Sandy Bridge CPUs E5-2643 @ 3.3Ghz
Total Cores per Socket	4
Total Logical Processors (HT enabled)	16
Total Installed Memory	128GB @ 1600Mhz
Network Adapters	2* Intel Ethernet X540 Dual Port 10GBASE-T Server Adapter
Multipathing Software	Dell MD Series DSM
Multipathing Policy	Least Queue Depth(Default)
External Storage	1 * PowerVault MD3620i with a MD1220 Expansion Enclosure
Disks	48 * 146G, 6Gbps 15k SAS drives
Network Switch	1 * PowerConnect 8024
Operating System	Windows 2008 R2 SP1 Enterprise Edition
Database Software	SQL Server 2012 Enterprise Edition

Highly available reference architecture

For achieving high availability for the database, we recommend using Microsoft Database Clustering. Using Microsoft clustering services, one database server is configured as the primary (active) server and the second server is configured as secondary (passive). The secondary server should have exactly the same configuration as the primary server. Since the database is active on a single server at any point of time, the performance of the database on the primary server (active) is comparable to the single server configuration (discussed earlier).

Figure 2 depicts the highly available fast track reference architecture, and Table 3 lists the configuration details.

Figure 2. Highly available reference architecture

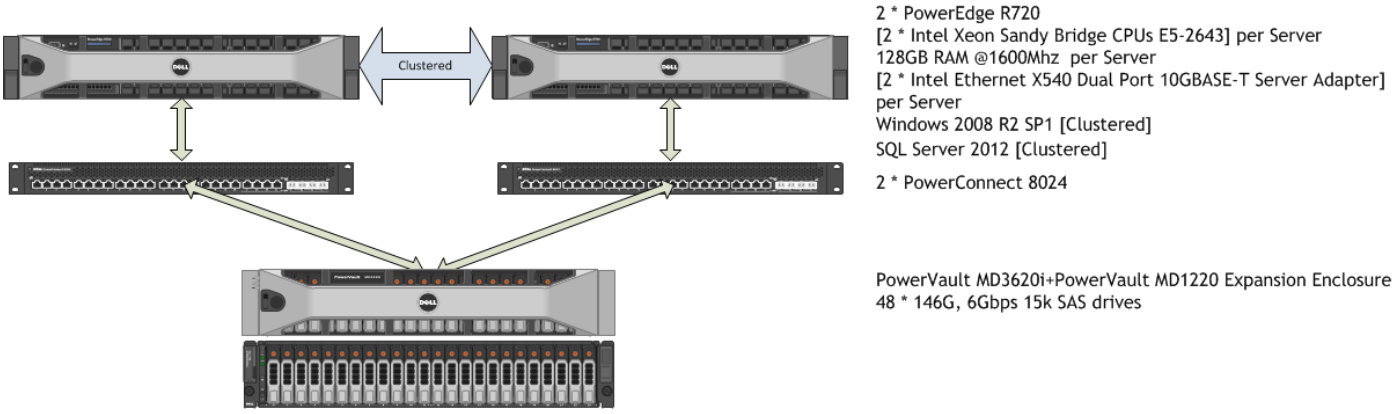

Table 3. Highly available reference architecture details

Hardware Components	Details
Server	2 * PowerEdge R720
Processors	[2* Intel Xeon Sandy Bridge CPUs E5-2643 @ 3.3Ghz] per Server
Total Cores per Socket	4
Total Logical Processors (HT Enabled)	16 per Server
Total Installed Memory	128GB @ 1600Mhz per Server
Network Adapters	[2* Intel Ethernet X540 Dual Port 10GBASE-T Server Adapter] per Server
Multipathing Software	Dell MD Series DSM
Multipathing Policy	Least Queue Depth(Default)
External Storage	1 * PowerVault MD3620i with a MD1220 Expansion Enclosure

Disks	48* 146G, 6Gbps 15k SAS drives
Network Switch	2 * PowerConnect 8024
Operating System	Windows 2008 R2 SP1 Enterprise Edition
Clustering Software	Microsoft Windows Clustering
Database Software	SQL Server 2012 Enterprise Edition

Reference architecture configuration hardware component details

This section describes the hardware details used for the specific reference configurations.

Dell PowerEdge R720 server

Dell launched the 12th generation PowerEdge servers with support for Intel Xeon E5 series processors. The PowerEdge R720 is a 2-socket 2U rack server that offers simplified management, purposeful design, and energy efficiency with support of Intel E5 Series Sandy-Bridge processors and ECC DDR3 RDIMMs @1600MHz with a maximum capacity of 768GB Memory. PowerEdge R720 supports five x8 PCIe Gen3 & two x16 PCIe Gen3 slots.

Visit dell.com/us/enterprise/p/poweredge-r720/pd for more information on PowerEdge R720 Servers.

Intel Ethernet X540 DP 10GBASE-T server adapter

The proposed reference architectures use two dual port Intel Ethernet X540 DP 10GBASE-T Server Adapter cards. This new dual port adapter has RJ-45 copper interfaces and is available both in low profile and full height interfaces.

Visit intel.com/content/www/us/en/network-adapters/10-gigabit-network-adapters/ethernet-10gigabit-adapters.html for more information on Intel Ethernet X540 DP 10GBASE-T Server Adapters.

Dell PowerConnect 8024 Ethernet switch

Dell PowerConnect 8024 switch is a rack mountable 10Gb Ethernet switch with 24 ports. The last four ports (21-24) are combo ports which have the option to use SFP+ slots. These allow the use of fiber optic connection uplinks to core switches.

Visit dell.com/us/enterprise/p/switch-powerconnect for more information on PowerConnect 8024 switches.

PowerVault MD3620i iSCSI storage arrays

PowerVault MD3620i storage arrays support 2.5 inch 6Gbps SAS disks, which enable better backend performance. MD3620i supports up to two controllers. Each controller has two 10GbE host side ports and one SAS out port, which is used to connect to any additional expansion enclosures. Each RAID controller module contains 2GB of cache that is mirrored with the other controller's cache for high availability and protected by a battery-powered cache offload mechanism.

Dell PowerVault MD3620i storage comes with its own Device Specific Module (DSM) software, included in the MD3620i's Resource DVD, which is fully integrated with Microsoft MPIO and helps you configure multipath solutions.

Figure 3 depicts how the cables are connected in the configuration.

Figure 3. Cabling diagram for single server configuration

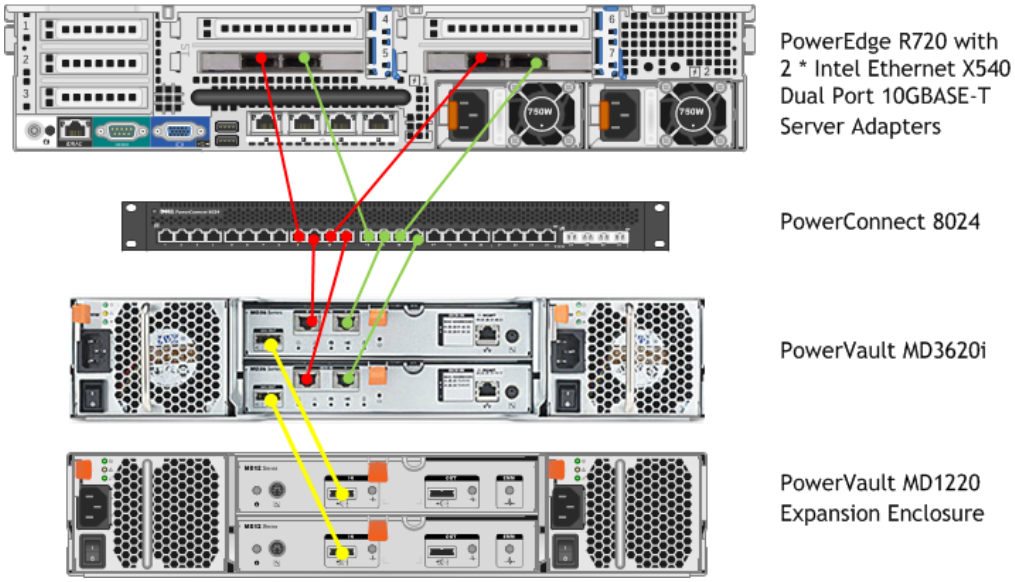

Visit dell.com/us/enterprise/p/powervault-iscsi-arrays for more information on PowerVault MD3620i storage arrays.

Reference architecture configuration tuning details

This section explains best practices and tuning guidelines to achieve the best performance using the discussed reference architectures.

PowerVault MD3620i storage tuning

The following modifications are needed to optimize the storage for maximum sequential throughput:

- Enable Jumbo Frames on all storage iSCSI host ports
- Enable High Performance Tier feature on MD3620i
- Set the Storage Cache Block Size to 32KB
- Set the virtual disk segment size to 256KB

Enable Jumbo Frames on all storage iSCSI host ports

Jumbo frames need to be enabled on all four storage iSCSI host ports and the MTU size should be set to 9000 bytes/frame. This can be done using the **Advanced Host Port Settings** window (Figure 4), accessible from the Configure iSCSI host ports link of the PowerVault Modular Storage Manager (MDSM) Setup tab.

Figure 4. Enabling Jumbo frames on storage iSCSI host ports

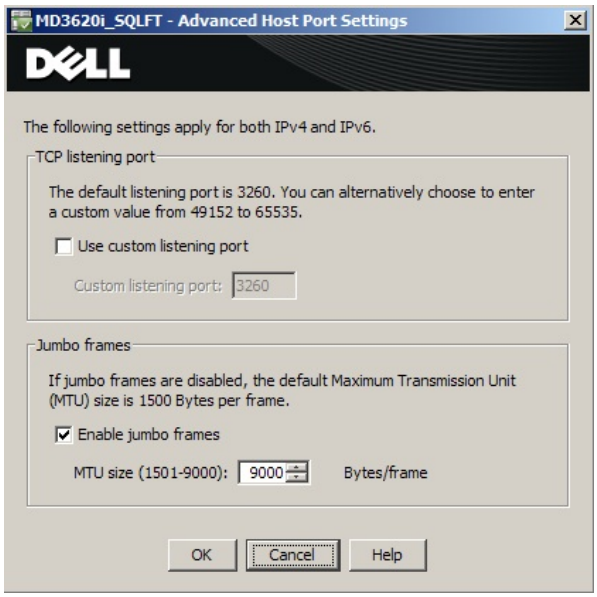

Enable High Performance Tier feature on MD3620i

High performance Tier (turbo mode), one of the premium licensed features of the MD3620i array, enables the storage to deliver increased performance. For the proposed fast track architecture, you should enable the HPT feature to achieve improved throughput from the array. Other features may be enabled based on specific customer requirements. The premium licensed features of the MD3620i storage array may be enabled or disabled from the **Premium Features and Features Pack** information window on the MD Storage Management GUI (Figure 5).

Figure 5. Enabling storage High Performance Tier(Turbo Mode)

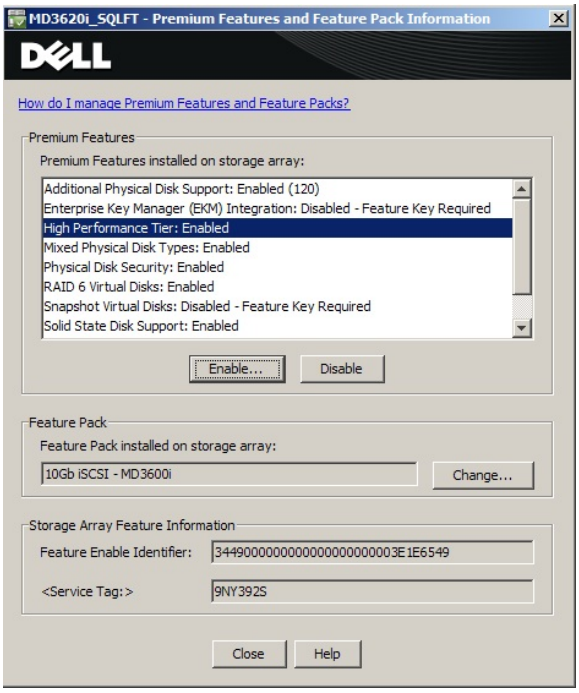

Set the Storage Cache Block Size to 32KB

In this specific Fast Track configuration, Dell recommends setting the storage cache block size to 32KB for maximum array throughput. The **Cache block size** is set in the Change **Cache Settings** window (Figure 6).

Figure 6. Modifying storage cache block size

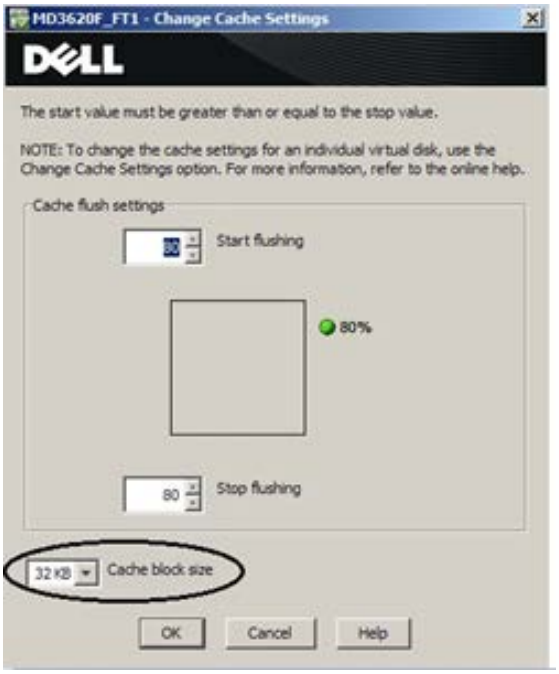

Set the virtual disk segment size to 256KB

For the recommended Fast Track configurations, Dell recommends setting the virtual disk segment size to 256KB. This benefits the large block sequential workloads. The segment size is set in the **Customize Advanced Virtual Disk Parameters** window (Figure 7).

Figure 7. Modifying the virtual disk segment size

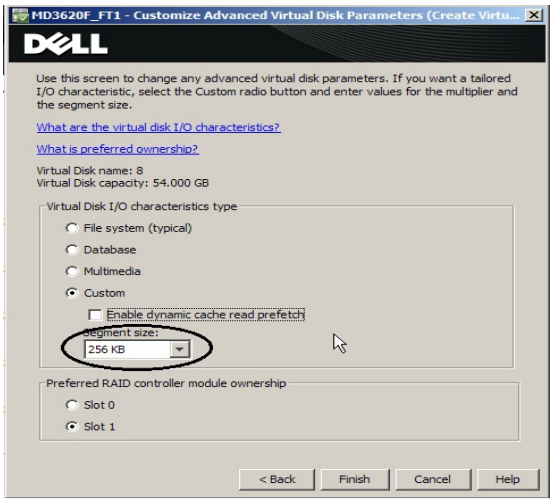

Intel X540 10GbE NIC tuning

For the specific reference configurations, Jumbo Packets were enabled on all NIC ports and set to a packet size of 9014 bytes. All other parameters were set to the default settings. The jumbo packet settings can be changed using the **Advanced** tab of the Adapter properties window (Figure 8).

Figure 8. Enabling Jumbo packets on network adapter ports

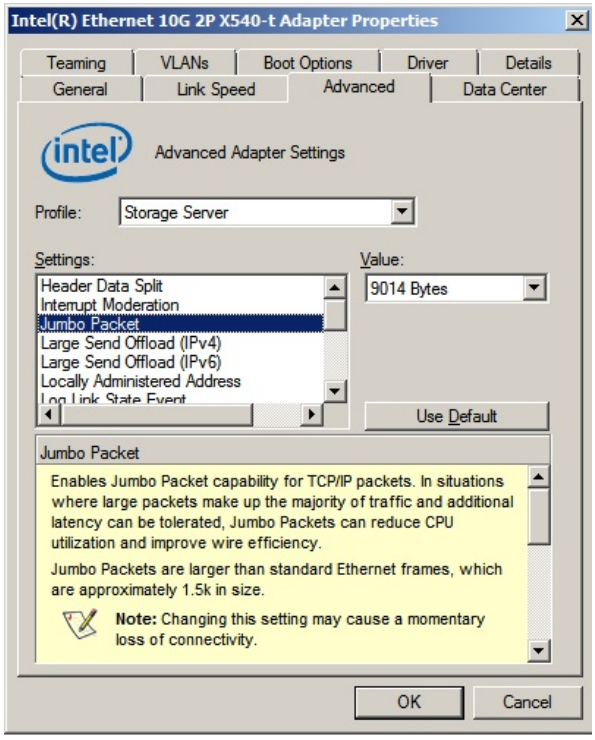

PowerConnect 8024 Switch Tuning

For the reference configuration, Dell proposes to have Jumbo frames enabled on all participating iSCSI ports of the configuration. The MTU size of all the concerned ports should be set to 9216. Figure 9 shows how to enable jumbo frames for all the switch ports from the serial console window.

Figure 9. Enabling Jumbo frames on switch ports

```
console>enable

console#configure

console(config)#interface range Tengigabitethernet All


console(config-if)#mtu 9216

console(config-if)#exit

console(config)#exit

console#copy running-config startup-config

This operation may take a few minutes.
Management interfaces will not be available during this time.

Are you sure you want to save? (y/n) y

Configuration Saved!
console#exit

console>
```

Windows tuning

The allocation unit size for all the database hard drives was set to 64KB.

SQL server tuning

The following startup options were added to the SQL Server Startup options:

- **-E**: This parameter increases the number of contiguous extends in each file that are allocated to a database table as it grows. This option is beneficial because it improves sequential access.
- **-T1117**: This trace flag ensures the even growth of all files in a file group when auto growth is enabled. Note that the Fast Track reference guidelines recommend to pre-allocate the data file space rather than to depend on auto grow.
- **SQL Server Maximum Memory**: For SQL Server 2012, FTDW 4.0 guidelines suggest allocating no more than 92% of total server RAM to SQL Server. If additional applications will share the server, the amount of RAM left available to the OS should be adjusted accordingly. For this reference architecture the maximum server memory was set at 117.76GB.

  For the specific recommended fast track configuration, the resource governor (disabled by default) and the 'max degree of parallelism' settings (0 by default) were set to the default settings.

Storage system configuration

The Fast Track Reference Architecture guidelines define three primary layers of storage configuration:

- Physical disk array (RAID Groups for Data and Logs)
- Operating system volume assignment (LUN)
- Databases: User, System Temp, Log

On each storage array, for the primary user and system temp database files:

- Four RAID 5 disk groups were created, each consisting of eleven disks (two disk groups each from the base and the expansion enclosure)
- One virtual disk (LUN) was created out of each disk group. These LUNs were evenly distributed across the storage controllers.

Therefore, there were four LUNs (virtual disks) for the primary user data. Each storage controller was assigned two of the LUNs.

For the primary user and system temp database transaction logs, one RAID1 disk group was created of 2 disks. This RAID group was dedicated to host the database transaction log files.

The remaining 2 disks on the storage array were assigned as the storage hot spares (global).

Figure 10 shows the pictorial representation of the LUN layout in the storage array.

Figure 10. Storage LUN layout

For Fast Track architectures, Dell recommends using mount points for the storage access rather than using drive letters. It is also important to assign the appropriate LUN/volume and mount point names to the configuration in order to simplify troubleshooting and performance analysis. The mount point names should be assigned in such a way that the logical file system reflects the underlying physical storage enclosure mapping. Table 4 shows the LUN and mount point names used for the specific reference configuration and the appropriate storage layer mapping. All the logical volumes were mounted to the **C:\FT** folder.

Table 4. Mount point naming and storage enclosure mapping

Storage Array	Storage Enclosure	Storage Processor	Disk Group	Logical LUN (volume)	Logical Volume Label	Logical Array Label	Full Volume Path
SA1	SE1 (Base)	1	1	1	SA1-SE1-SP1-DG1-v1	PRI	C:\FT\PRI\ SA1-SE1-SP1-DG1-v1
SA1	SE1 (Base)	2	2	2	SA1-SE1-SP2-DG2-v2	PRI	C:\FT\PRI\ SA1-SE1-SP2-DG2-v2
SA1	SE2 (Expansion)	1	3	3	SA1-SE2-SP1-DG3-v3	PRI	C:\FT\PRI\ SA1-SE2-SP1-DG3-v3
SA1	SE2 (Expansion)	2	4	4	SA1-SE2-SP2-DG4-v4	PRI	C:\FT\PRI\ SA1-SE2-SP2-DG4-v4
SA1	-	1	5	5	SA1-SE2-SP1-DG5-v5	LOG	C:\FT\LOG\ SA1-SP1-DG5-v5

SA: Storage Array, SE: Storage Enclosure, SP: Storage Processors, DG: Disk Group, V: Volume

Reference architecture performance details

Table 5 shows the performance numbers reported for the recommended reference configuration.

Table 5. Performance metrics

Metric	Value	Description
FTDW Rated Capacity (TB)	10	This capacity rating is based on "up-to" capacity but adjusted to account for Fast Track Rated I/O.
FTDW Rated I/O (MB/s)	1690	Core performance metric for validation. This is the midpoint of Physical and Logical I/O.
Up-to User Data Capacity (TB)	14	Maximum user data capacity. Includes the total disk capacity of all disks allocated to primary data storage and does assume SQL Server page compression ratio of 3.5:1. This metric is not limited by rated bandwidth.
Benchmark Scan Rate Logical (MB/s)	2070	Reflects actual user query throughput which includes reads from RAM/Buffer cache.
Benchmark Scan Rate Physical (MB/s)	1320	Reflects physical I/O read from disk during benchmark.
FTDW Peak I/O (MB/s)	2460	Maximum observed I/O rate.
FTDW Rated CSI (MB/s)	4410	Represents potential throughput using Columnstore Index.

Summary

Dell, in partnership with Microsoft, enables customers to enhance ROI (Return of Investment) on their data warehouse systems using Fast Track Data Warehouse architectures. These uniquely-designed architectures ensure optimally designed and architected BI solutions. The end to end database best practices and recommendations enable the customer to achieve a balanced data warehouse environment with greater performance than traditional data warehouse systems.

The Dell Microsoft Fast Track Architecture provides the following benefits to customers:

- Delivers a tested and validated configuration with proven methodology and performance behavior.
- Achieves a balanced and optimized system at all the levels of the stack by following the best practices of hardware and software components.
- Avoids over-provisioning of hardware resources.
- Offers high availability at all levels of setup (host, switches, and storage).
- Helps customers avoid the pitfalls of an improperly designed and configured system.
- Reduces future support costs by limiting solution re-architect efforts because of scalability challenges.