


DELL Reference Configuration for Microsoft® SQL Server® 2008 Fast Track Data Warehouse

A Dell Technical Configuration Guide

Database Solutions Engineering

Dell Product Group

Anthony Fernandez



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.

© 2010 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, *PowerEdge*, and *PowerConnect* are trademarks of Dell Inc. *Microsoft* and *SQL Server* are registered trademarks of Microsoft Corporation in the United States and/or other countries. *Intel* and *Xeon* are registered trademarks of Intel Corporation in the U.S. 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 Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

April 2011

Contents

Executive Summary	3
Introduction	3
Audience and Scope	3
Microsoft Fast Track Data Warehouse Overview	4
Reference Architecture	4
Recommendations and Best Practices	5
Database Availability	5
Internal Storage	8
Internal PCIe I/O Channel	9
Database Settings	9
Storage Settings	9
Storage Cache Configuration.....	11
High Performance Tier (Turbo Mode) Configuration.....	13
Multipathing Options and Settings	14
Microsoft MPIO	14
Performance Characteristics	15
Baseline Hardware Validation	15
Conclusion	18
References.....	19

Tables

Table 1. Reference Architecture List	4
Table 2. PowerEdge R710 Configuration Details	4

Figures

Figure 1. Fast Track Reference Architecture	4
Figure 2. Reference Architecture	6
Figure 3. High Capacity Configuration.....	6
Figure 4. Reference Architecture with Clustering	7
Figure 5. Minimum Memory Configuration with 8GB DIMMs	8
Figure 6. Internal Storage Configuration	8
Figure 7. HBA Slot Placement.....	9
Figure 8. Dell PowerVault MD3220 Chassis.....	10
Figure 9. Disk Group Layout.....	10
Figure 10. Storage Controller Disk Group Assignment	11
Figure 11. Storage Cache Settings	12

Figure 12.	Large Cache Block Size Setting.....	12
Figure 13.	Cabling Diagram Configuration	13
Figure 14.	Cascading MD3220 with Expansion.....	14
Figure 15.	Microsoft MPIO Settings	15
Figure 16.	Synthetic I/O Rate Reads from Cache.....	16
Figure 17.	Synthetic I/O Rate Reads from Hard Disks.....	17

Executive Summary

Data warehouses provide the foundation for business intelligence systems, such as the analysis and reporting services that are critical to enabling an organization to scan through very large amounts of data as efficiently as possible. Data warehouse configurations typically suffer from design principles that were originally intended for online transaction processing systems (OLTP). As the amount of data grows, so do scan times that affect the overall time required to gather critical information.

Dell and Microsoft have collaborated to provide a set of guidelines and design principles called Fast Track Data Warehouse 3.0 Reference Guide to help customers design and implement balanced configurations specifically for Microsoft SQL Server Data Warehouse databases with the goal of providing a hardware-balanced approach and predictable out-of-box performance.

In addition, Dell provides services designed to help customers implement complete Business Intelligence Solutions utilizing reference architectures that have been thoroughly tested and validated by Dell and Microsoft.

This white paper describes the architecture design principles to achieve a balanced configuration for Dell™ PowerEdge™ R710 and Dell PowerVault™ MD3220 Direct Attached 6GBps SAS storage. This configuration is optimized for 4TB of compressed user capacity with capabilities to expand up to 19TB with Dell PowerVault MD1220 arrays. It offers a high availability configuration with Microsoft Cluster Service for environments that require maximum availability.

Introduction

This white paper describes a reference architecture implementation of a Microsoft® SQL Server® Fast Track Data Warehouse 3.0 Reference Guide utilizing Dell PowerEdge R710 servers and Dell PowerVault MD3220 Direct Attached 6GBps SAS storage. The goal of the Microsoft Fast Track program is to define a methodology to build a balanced and optimized hardware and software configuration specifically for SQL Server Data Warehouse deployments.

Utilizing a building block approach, a Microsoft Fast Track solution offers a cost effective and proven platform that has been tested and optimized to offer customers a faster way to deploy and configure a data warehouse infrastructure. A data warehouse is the central component in a business intelligence solution that stores large quantities of data and information. As data continues to grow at a fast pace, optimizing data retrieval is crucial for organizations to maintain service level agreements (SLAs). Microsoft Fast Track provides a framework that allows customers to select a reference architecture that has been designed with those principals in mind that best fit their needs, the amount of data, and their performance requirements.

Audience and Scope

This white paper is intended for customers, partners, solution architects, storage administrators, and database administrators who are evaluating, planning, or deploying a balanced data warehouse solution. The scope is limited to the main repository of data or data warehouse and the balanced configuration thereof. Other systems that utilize the data warehouse as a source of data, such as analysis and reporting services, are not within the scope.

Microsoft Fast Track Data Warehouse Overview

The Microsoft Fast Track Data Warehouse initiative provides a framework to build and deploy a balanced configuration for a SQL Server data warehouse. The initiative provides guidelines and best practices to configure software and hardware to achieve optimal cost and performance. Microsoft Fast Track utilizes a building block approach focused on balanced configurations that have been tested and validated for data warehouse workloads which are large sequential in nature.

Reference Architecture

Table 1 lists recommended hardware that has been tested and selected for a balanced configuration to deploy a Microsoft SQL Server Fast Track Data Warehouse. This reference configuration focuses on a configuration that includes one PowerEdge R710 server with two Dell PowerVault MD3220 DAS SAS storage arrays for a capacity range of 4TB to 8TB of compressed data (3:1 ratio).

Table 1. Reference Architecture List

Server	CPU	CPU Cores	SAN	Data Drive Count	Initial Capacity	Maximum Capacity
PowerEdge R710	(2) Intel® Xeon® X5680 3.33GHz	12	(2) Dell PowerVault MD3220	(40) 146GB 15K 6Gbps	4TB	8TB

Figure 1 shows the hardware components that comprise the data warehouse solution.

Figure 1. Fast Track Reference Architecture



Table 2. PowerEdge R710 Configuration Details

Base Solution ID	1407335.1
Highly Available Solution ID	1407293.1

High Capacity Solution ID	1407392.1
Server	PowerEdge R710
CPU	(2) Six-Core Intel® Xeon® X5680 3.33GHz 12MB cache, 6.4GT/s Intel QPI, turbo
Number of Cores	12
Logical Processors	Enabled. (24 Logical Cores)
Memory	96GB RAM (12x8GB RDIMMS @1333MHz)
PCI-E Slots	Two PCIe Gen2 x8 slots (Slots 3,4)
Internal Storage Controller	PERC6/I or H700 512MB/1GB Cache
Internal Drives	(2) 143GB 6Gpbs SAS 15K HDD (Operating System) (6) <i>Various sizes can be configured for additional storage staging space.</i>
Network Adapters	Two embedded Broadcom 5709C dual-port Gigabit Ethernet adapters (four ports total)
Remote Management	iDRAC6 Enterprise
Network Switch	Dell PowerConnect 55XX Series
Storage Adaptor	(2) Dell SAS6 Host Bus Adaptors
Storage	(2) Dell PowerVault MD3220 (2) Dell PowerVault MD1220 optional expansion unit

Recommendations and Best Practices


This section details the recommendations and best practices for implementing a high-performing Fast Track data warehouse on the latest PowerEdge R710 server and Dell PowerVault MD3220 Direct Attached SAS storage enclosure. In this section, the hardware-based optimizations are given preference compared to the software and database parameters.

Database Availability

Based on business requirements, it may be necessary to have a highly reliable and available database configuration with high performance for data warehousing. Redundancy is implemented at every layer to provide network and I/O path redundancy. In addition, this configuration provides the flexibility to add clustered nodes with Microsoft Cluster Service (MSCS) to provide server redundancy for data centers that require highly available configurations.

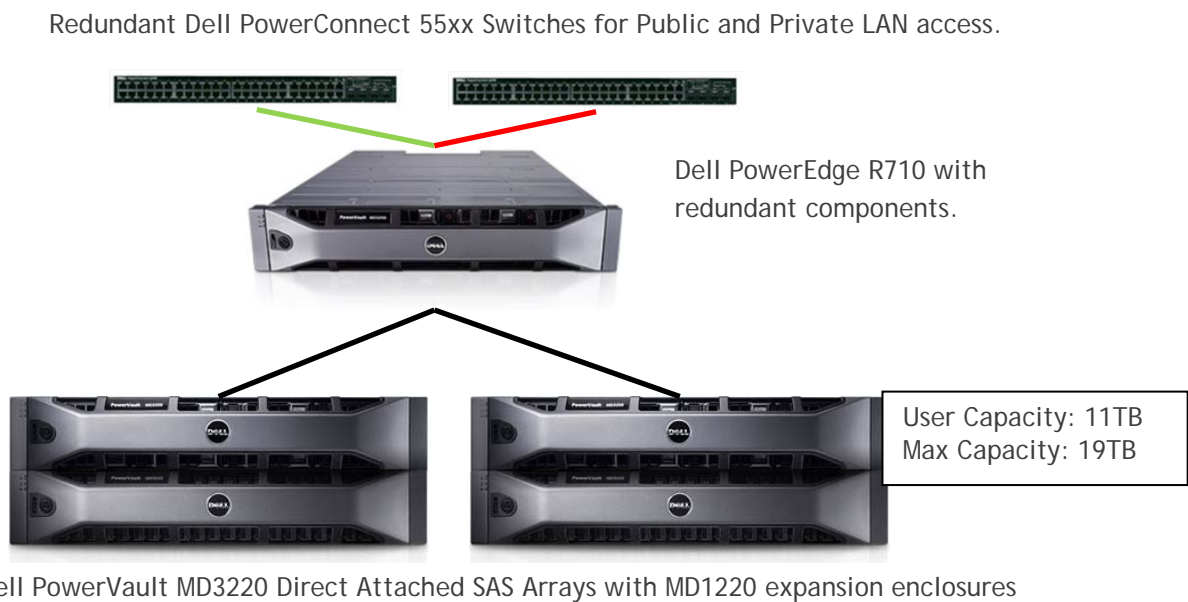
The complete Dell SQL Server Fast Track Reference Architecture may be represented as in Figure 2 for single-node configurations.

Figure 2. Reference Architecture



For Data Warehouse environments that require larger capacities, two options can be implemented. One option is to maintain higher rotational speed drives of 15K and add an expansion enclosures to create additional disk groups depicted in Figure 3 below. Another option is to implement slower rotational speed drives of 10K with larger capacities of 300GB and 600GB. Careful evaluation and testing is required to ensure bandwidth is maintained at consistent levels.

Figure 3. High Capacity Configuration



The Fast Track architecture with high availability Microsoft Cluster Service (MSCS) is depicted in Figure 4 below.

Figure 4. Reference Architecture with Clustering




Figure 4 above shows an active-passive SQL Server configuration using Microsoft clustering technology. Microsoft clustering enables the passive database node to host the database service if the primary database server (the active node) fails or is taken offline for maintenance. This configuration provides a highly available configuration at the database, network, and storage layers.

Server

Server selection for Fast Track implementations is based on the principals for building a balanced system from the storage to the server. Given the multicore capabilities of today's servers, a starting point to build a balanced configuration is based on the maximum consumption rate (MCR) a system can process.

The PowerEdge R710 server is a dense 2U server with dual-socket Intel® Xeon® 5500 and 5600 series quad- and six-core processors, Intel 5520 I/O Hub (IOH) with QuickPath architecture, up to 18 DIMMS or 288GB of DDR3 memory capacity, DIMM thermal sensors, PCI Express Generation 2, and two dual-port embedded Gigabit Ethernet controllers for uplink connectivity.

CPU

This configuration describes balanced performance for the PowerEdge R710 server with the Intel® Xeon® 5600 Series X5680 3.33GHz six-core processors.

This configuration was tested with Logical Processors enabled. Logical Processors provide an additional hyper-threaded core for each processor core. This configuration is running with 24 total logical processors.


Using a 300MB per second per core of page-compressed data to determine the MCR of the 12-core system, the PowerEdge R710 server yields an MCR of 3.6GB/s. The MCR provides a starting point to determine the bandwidth from the storage to maximize the CPU capacity. This methodology helps

define a watermark of performance when designing a balanced configuration free of bottlenecks throughout the data path; for example, storage spindles, storage controllers, Storage HBAs, CPUs and memory. Refer to the *Microsoft Fast Track Data Warehouse 3.0 Reference Guide*

Memory

Memory sizing for a Fast Track Data Warehouse configuration depends on the workload characteristics of data access. Microsoft recommends the minimum amount of memory for this configuration of 4GB of RAM per logical core. For a 24-core system, 96GB RAM can be configured with twelve 8GB DIMMs running optimized at 1333MHz. Figure 5 below show the memory population per socket.

Figure 5. Minimum Memory Configuration with 8GB DIMMs




The above figure shows the DIMM slots used to populate highlighted. For 8GB DIMMS, two memory modules per channel ie; [A1,A2,A3] and [A4,A5,A6] for CPU 1 support up to 1333MHz speeds.

Internal Storage

Customers have a choice to select from two types of chassis for the PowerEdge R710 server, with up to eight 2.5" 6Gbps SAS drives.

Figure 6 shows both chassis configured with the first two drives in a RAID 1 configuration for operating system installation. The remaining drives are used for staging or for backup space and can be ordered with larger capacities. Depending on the space required, customers can configure the drives in the following two configurations.

Figure 6. Internal Storage Configuration



*Space configured with RAID5 or other RAID types can be configured depending on space or performance requirements.

It is recommended to utilize the PERC/i or PERC H700 internal storage controller. This controller provides internal hardware RAID capabilities (0, 1, 5, 6, 10, 50, and 60) and mixed RAID configurations; for example, RAID 1 for operating system drives and RAID5 or RAID10 for staging/backup.

For configurations in which RAID 1 is sufficient for the staging/backup area, the SAS6/iR or H200 offers RAID 0 and RAID 1 capabilities. Note that the SAS6/i and H200 only offers a maximum of two RAID 1 sets.

Internal PCIe I/O Channel

Fast Track configuration utilizes two 6Gbps SAS HBAs. The PowerEdge R710 server planar provides two x16 PCI Express connector slots. Riser 1 and Riser 2 provide x4 and x8 PCI Express connectors. The HBAs should be placed on slots 3 and 4 to take advantage of PCIe Gen 2 x8 speeds.

Planar slot specifications:

- Two x4 PCIe Gen2 slot full height with x4 physical connector (Slot 1 & 2)
- Two x8 PCIe Gen2 slots full height with x8 physical connector (Slots 3 & 4)
- One x4 PCIe Gen1 slot for dedicated storage controller card (PERC 6/i or SAS6/iR) —connected to the IOH

Figure 7 below shows the back of the chassis to demonstrate the physical HBA placement. Each slot is identified by a number labeled on the chassis

Figure 7. HBA Slot Placement



Database Settings

Microsoft provides configuration guidelines and best practices for deploying a data warehouse using the Fast Track approach. This configuration includes minimum indexing, table partitioning, and table compression considerations. The major database settings that would enable the SQL Server database to be optimized for a sequential database workload may include the following database startup parameters:

- -E : Increases the number of extents allocated to a database table
- -T1117: Ensures the even growth of all files in the file group

For the complete configuration guidelines, see the *Microsoft Fast Track Data Warehouse 3.0 Reference Guide* listed in the References section of this white paper.

Storage Settings

The storage configuration plays a major role in the database performance. This section discusses storage Disk Groups, Virtual Disks and data file layout, HBA cabling, and Multipath I/O configuration.

The *Fast Track Reference Configuration Guide 3.0* provides guidelines for creating RAID configurations and storage virtual disks on an external storage array. In addition, the distribution of the storage virtual disks between the storage processors of the array plays a significant role in delivering the optimum output out of the array.

Figure 7 shows the front and back of the PowerVault MD3220 chassis.

Figure 8. Dell PowerVault MD3220 Chassis



Figure 8 above shows the front and rear of the Dell PowerVault MD3220 storage array. The MD3220 is configured with redundant active/active controllers and power supplies.

Figure 9 below shows the physical disk group layout across both storage arrays. Each disk group is composed of groups of five hard drives configured as (4+1) RAID 5 groups. Microsoft specifies utilizing RAID1/0 mirrored sets for database data and RAID 1 for log files.

To attain a better balance across storage controllers, we have tested this configuration with (4+1) RAID 5 Disk Groups. This configuration results in four data disk groups per array where we can assign two disk groups per controller. Configuring the disk groups with RAID 1/0 results in 3 data disk groups on the first controller and 2 data disk groups in the second controller. Configuring data disk groups in RAID 1/0 is optimizes data loading with reduced capacities.

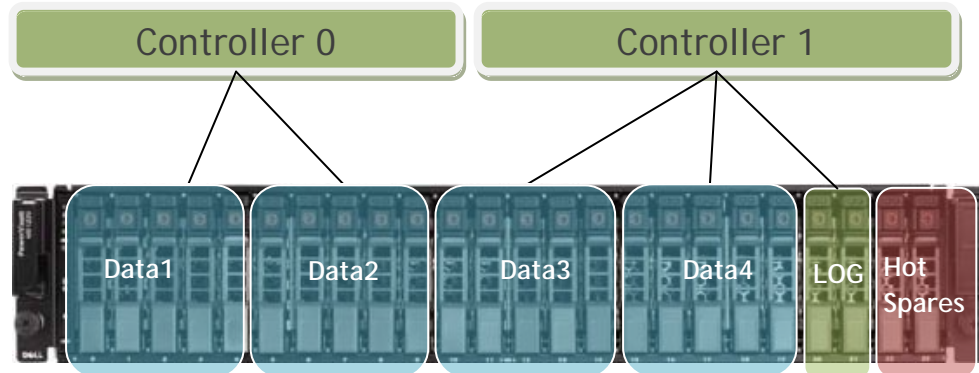
Note: RAID 5 may introduce a write performance impact for environments that require maximum load times. Depending on the amount of data change and load windows, careful evaluation is required to ensure proper service level agreements (SLA's) are met.

Figure 9. Disk Group Layout



A Virtual Disk (VD) is configured for each RAID disk group and utilizes the entire disk group capacity. This configuration with RAID 5 yields four VDs, assigned two per storage controller as shown by lines connecting each VD to the storage controller representing controller ownership of the VD.

Figure 10. Storage Controller Disk Group Assignment



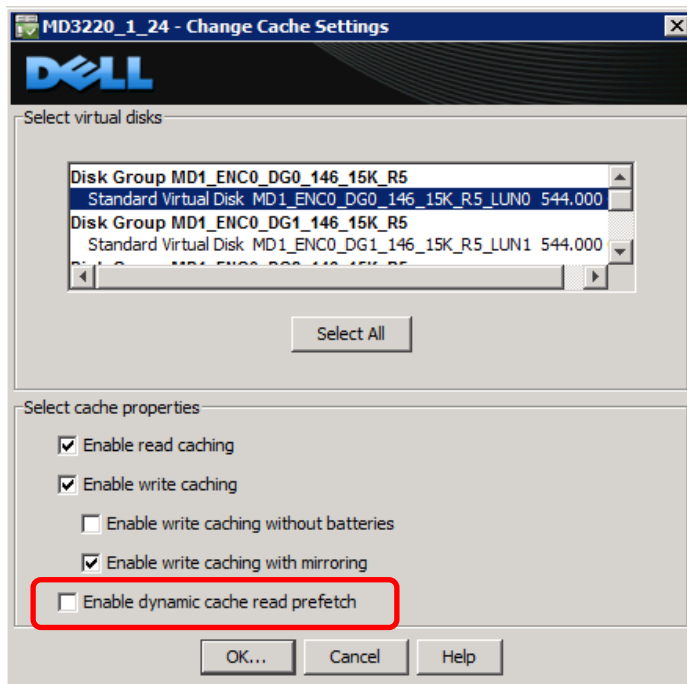
The figure above depicts a logical representation of disk group assignment to each storage controller. In environments where log file is heavily used, it would be recommended to use RAID 1/0 disk groups for data, and assign more hard drives to the log file disk group either with internal storage or via an expansion enclosure as shown by the high capacity Fast Track configuration.

Storage Cache Configuration

The Dell PowerVault storage array has 2GB of cache per controller, totaling 4GB of cache. Some performance gains may be achieved by configuring 100% of cache to writes during data load windows. Temporary files may also benefit from storage cache configuration. However, since **tempdb** files share the same spindles with database files, it is recommended to increase the available memory to allow sorts to happen in memory as opposed to disk which may affect the overall performance.

Figure 11 below shows the Cache Settings properties window accessed through the Dell Modular Disk Storage Management (MDSM) Client. Depending on the workload, it is recommended to test a combination of Cache settings to obtain maximum performance. As a first step, during Fast Track evaluation, reading small amount of data from cache to obtain the maximum consumption rate or (MCR), validates against bottlenecks that may have been introduced during deployment and configuration. Refer to the Microsoft Fast Track Data Warehouse Reference Guide 3.0 for test procedure specifics listed in the Reference Section of this white paper.

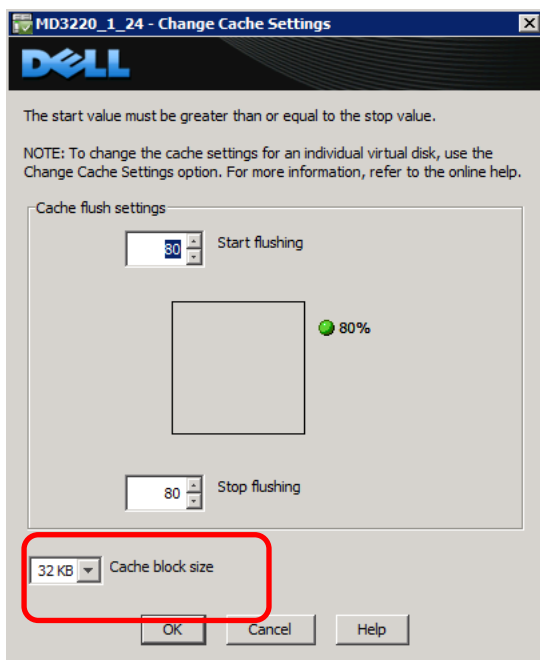
Figure 11. Storage Cache Settings



For large sequential scans it is recommended to disable dynamic cache read prefetch setting on each of the data disks. This setting might impose an overhead in some cases. Workload testing is advised.

For large sequential workloads, Dell recommends increasing the Controller Cache Block size to 32KB for optimal performance. Figure 12 below shows the Storage Array Cache Settings.

Figure 12. Large Cache Block Size Setting



High Performance Tier (Turbo Mode) Configuration

The Dell PowerVault MD3220 is enabled with a premium feature called High Performance Tier also known as Turbo Mode. This licensed feature is enabled by default for Microsoft SQL Fast Track configurations to maximize performance. This feature enables increased bandwidth capabilities for high data transfer workloads. Other premium features can be included depending on customer needs such as Snapshot Virtual Disk and Virtual Disk Copy which are optional.

Figure 13. Cabling Diagram Configuration




Figure 13 above shows the cabling diagram from a single server. Each server HBA port connects to each Port 0 on each controller. In configurations with Microsoft Cluster Server, the second cluster node HBA ports attach to Port 1 in a similar fashion.

Expansion units can easily be cascaded behind the MD3220 utilizing the Dell PowerVault MD1220 storage as shown in Figure 14 below.

Figure 14. Cascading MD3220 with Expansion



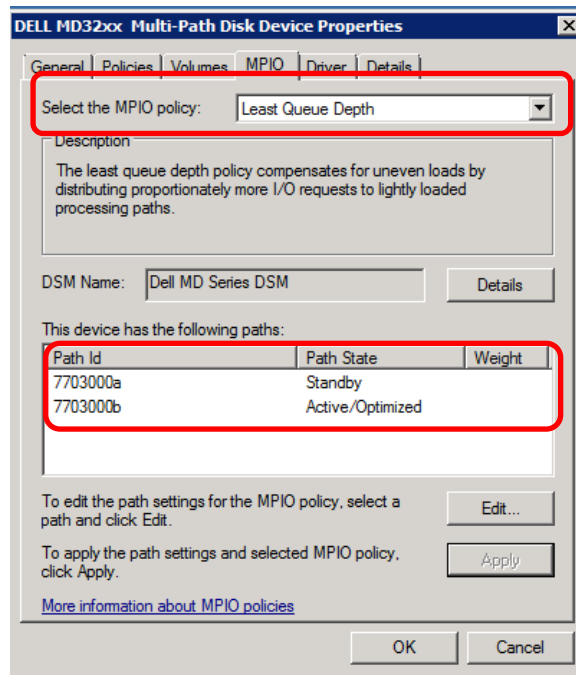
Note: This configuration only expands the capacity of the balanced Fast Track configuration. Depending on the workload and processor utilization, no additional bandwidth may be realized by adding an expansion enclosure.

Multipathing Options and Settings

Microsoft MPIO

Microsoft references MPIO as the multipathing solution for Fast Track implementations as a fully-integrated solution for external storage connectivity. In addition, Dell Multipath Device Specific Module (DSM) easily installs through the Dell Resource DVD included with the PowerVault MD3220 Storage.

Figure 15. Microsoft MPIO Settings



Microsoft recommends the use of the default Least Queue Depth function as the preferred load balancing policy. The Figure 15 above shows the disk properties and selections to enable the setting. Ensure that there a single Active/Optimized path to the disk and the second as Standby. This ensures that large I/O's are not divided into smaller segments to send on each active path.

Performance Characteristics

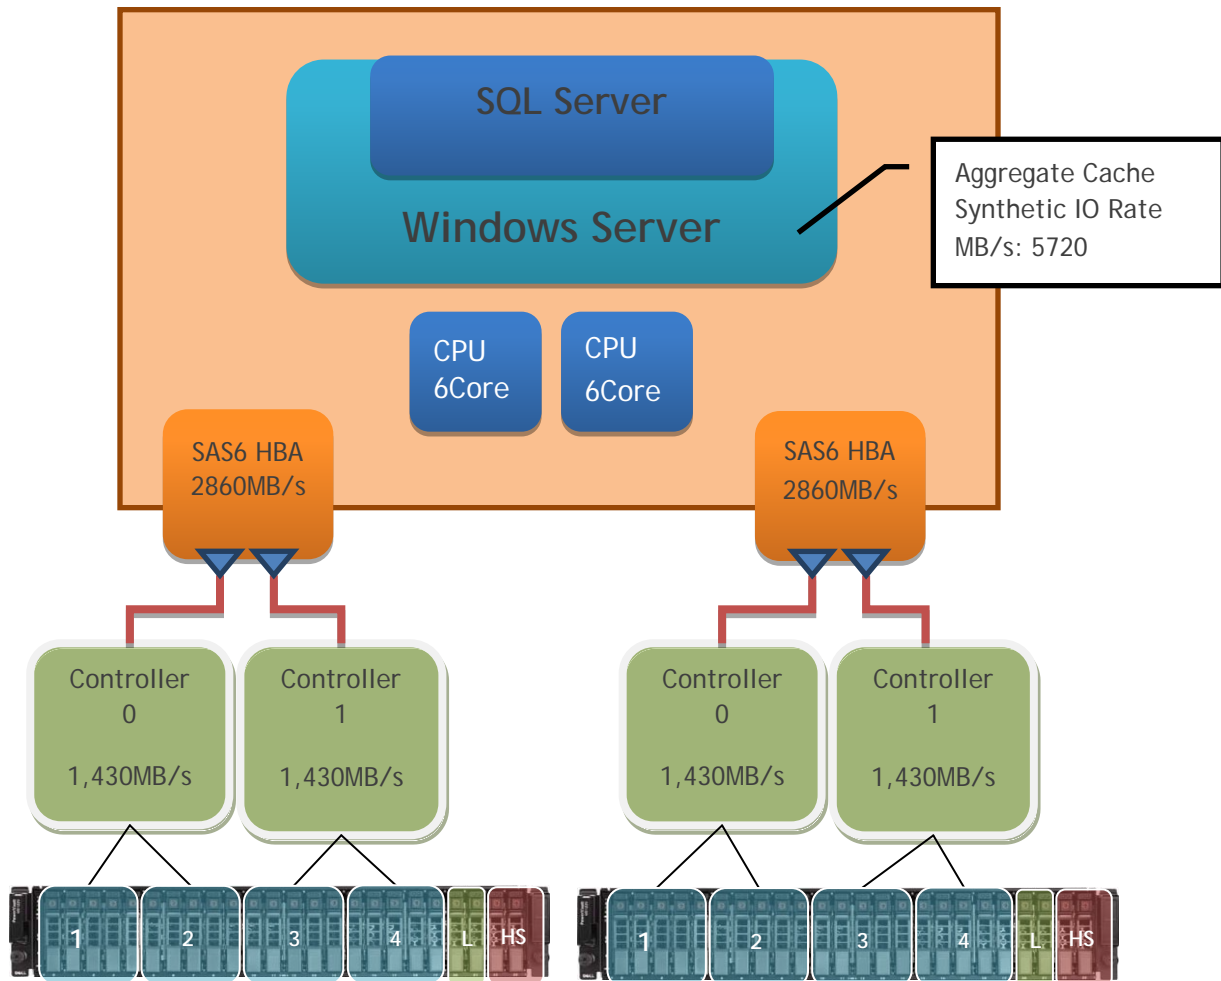
This section details the performance characteristics of this configuration following Microsoft Fast Track guidelines. The first step is to validate the hardware with a synthetic block level IO workload to ensure the hardware has been properly configured and the components are performing to their maximum listed capacities. The second step is to validate the system by loading a defined workload on SQL Server to ensure that access to data is performed as efficiently as possible for the given hardware configuration and capacity.

Baseline Hardware Validation

This step validates aggregate bandwidth and establishes a baseline of the system and validates the components ratings to their listed capabilities. It is a synthetic this test and is performed with a block level IO tool such as SQLIO. The first hardware validation test is reading 512KB sequentially from cache.

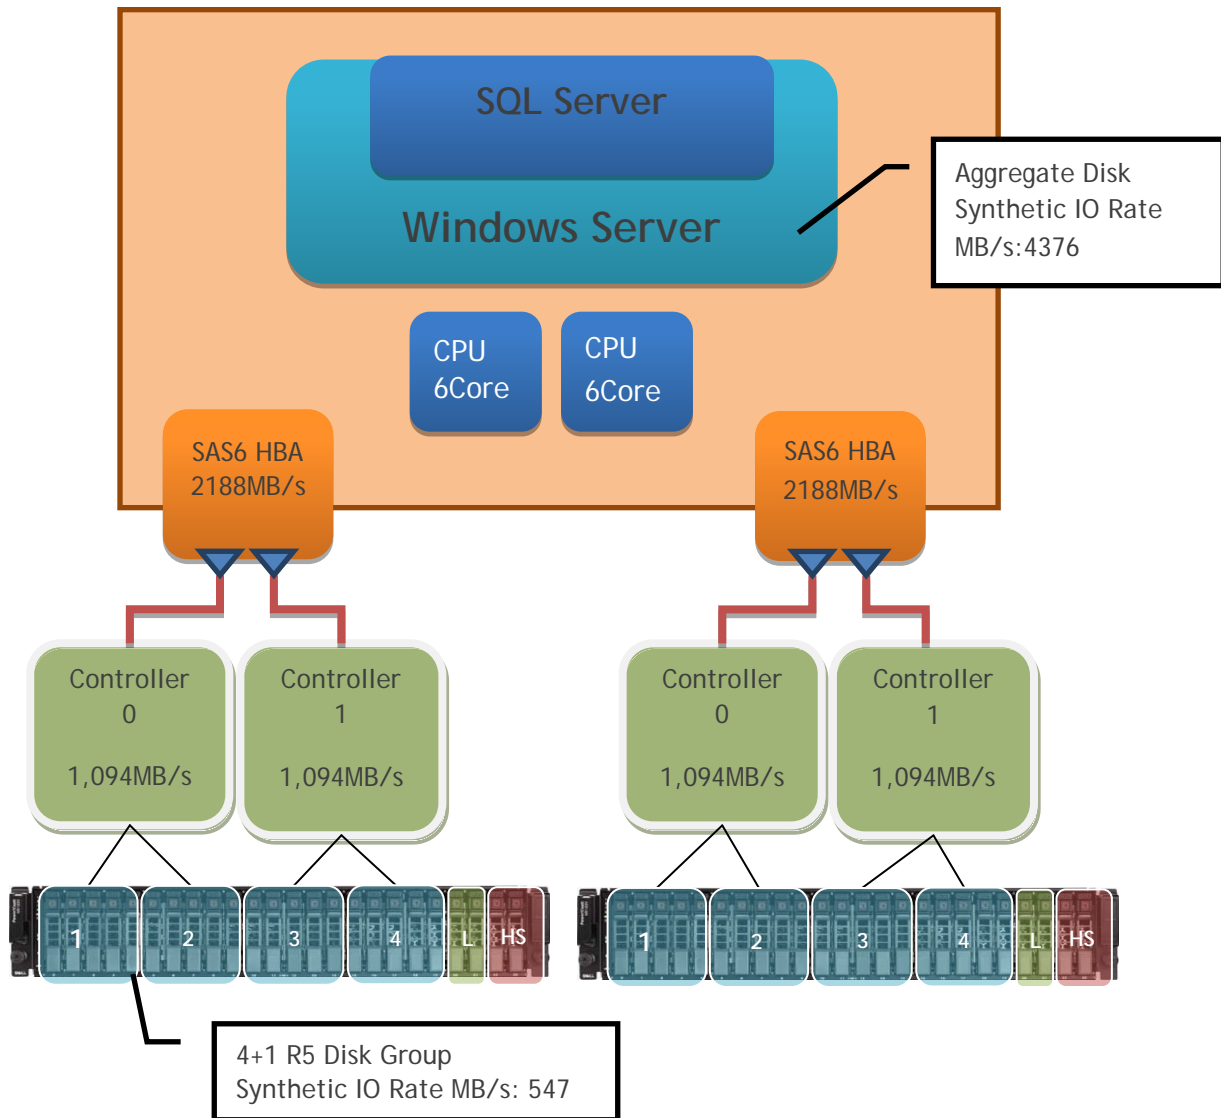
Refer to the Microsoft Fast Track Data Warehouse Reference Guide 3.0 for test procedure specifics listed in the Reference Section of this white paper.

Figure 16. Synthetic I/O Rate Reads from Cache



The second validation reads from hard disk utilizing a large file of 25GB in size to avoid cache access to determine the maximum baseline of reading from disk.

Figure 17. Synthetic I/O Rate Reads from Hard Disks



The figure above shows a sustained rate achieved for this configuration of 547MB/s per RAID 5 disk group. Note: For workloads that require large amount of load new data in the least amount of time, it is recommended to configured the storage with (2+2) RAID 1/0 disk groups.

Conclusion

Microsoft SQL Server Fast Track Architecture presents a balanced configuration of processor core and disk performance capabilities, optimized for a sequential workload. Dell, as an efficient hardware partner, adds value to the configuration by providing the best practices and recommendations at all the hardware layers. Following these recommendations helps to ensure a completely balanced configuration with optimized performance specifically designed for sequential workloads. In addition, Dell provides specialized services to help customers implement Data Warehouses and Business Intelligence Solutions.

To summarize, Dell/Microsoft Fast Track Reference Architecture provides the following value additions and advantages, along with a powerful, high-performing configuration:

- Best practices at all hardware and software layers
- Tested and validated configuration with proven methodology
- High availability at every level of the Fast Track configuration
- Complete Planning and Deployment Services

Dell and Microsoft together provide expertise at the hardware and software layers to design and build balanced reference architectures for Fast Track Data Warehouse deployments, ensuring better out-of-box performance.

References

Dell SQL Server Solutions
www.dell.com/sql

Dell Services
www.dell.com/services

Dell Support
www.dell.com/support

Dell Business Intelligence and Data Warehousing
www.dell.com/BI

Microsoft Fast Track Data Warehouse and Configuration Guide Information
www.microsoft.com/fasttrack

An Introduction to Fast Track Data Warehouse Architectures
<http://msdn.microsoft.com/en-us/library/dd459146.aspx>

Introduction to New Data Warehouse Scalability Features in SQL Server 2008
[http://msdn2.microsoft.com/en-us/library/cc278097\(SQL.100\).aspx](http://msdn2.microsoft.com/en-us/library/cc278097(SQL.100).aspx)

Best Practices for Data Warehousing with SQL Server 2008
<http://msdn.microsoft.com/library/cc719165.aspx>

Dell PowerVault MD3220 Storage Array Documentation
<http://support.us.dell.com/support/edocs/systems/md3200/en/index.htm>