
Dell SMB Reference Configuration for Microsoft® SQL Server® 2012 Fast Track Data Warehouse on PowerEdge R720xd


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

Database Solutions Engineering

Anthony Fernandez

Mayura Deshmukh

May 2012



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.

© 2012 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell, the Dell logo, and PowerEdge are trademarks of Dell Inc. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Microsoft, Windows, and Windows Server are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

May 2012 | Rev 1.0

Contents

Introduction	4
Microsoft SQL Server Fast Track Data Warehouse	4
Dell Fast Track Data Warehouse Reference Architecture	4
Hardware Component Architecture	5
Internal Storage Controller (PERC H710P Mini) Settings	7
Application Configuration	10
Performance Benchmarking.....	11
Baseline Hardware Characterization using Synthetic I/O.....	11
Fast Track Database Validation	13
Fast Track Database Validation with Column Store Index (CSI)	14
Reference Architecture Performance Details.....	14
Conclusion	15
References	16

Tables

Table 1. Dell Fast Track Reference Architecture Solution Details	4
Table 2. Tested Dell Fast Track Reference Architecture Component Details	5
Table 3. Mount Point Naming and the Storage Enclosure Mapping	9
Table 4. Performance Metrics.....	14

Figures

Figure 1. Proposed Dell Fast Track Reference Architecture.....	5
Figure 2. Memory Slot Locations	6
Figure 3. Virtual Disk Settings	7
Figure 4. Internal Storage Controller Settings	8
Figure 5. RAID Configuration.....	8
Figure 6. Storage System Components	10
Figure 7. SQLIO Line Rate Test from Cache (Small File)	12
Figure 8. SQLIO Real Rate Test from Disk (Large File)	13

Introduction

A data warehouse is a large store of data accumulated from a wide range of sources. The stored data is analyzed for trend analysis, business intelligence reporting, and various types of predictive analysis. With today's never ending data growth and complexity, it is becoming a tedious job for customers to balance capacity and performance within the data warehouse system. Growing data volumes and loading challenges, OLAP query complexity, and number of users are causing response times to increase. IT executives are looking for solutions that offer lower cost, easier management, and better performance.

There are many challenges in designing a database configuration for OLAP workloads. One is ensuring an optimal balance of I/O, storage, memory and processing power.

Dell™ and Microsoft® jointly developed guidelines and design principles to assist customers in designing and implementing a balanced configuration specifically for Microsoft SQL Server® data warehouse workloads to achieve “out-of-box” scalable performance.

This white paper describes the architecture design principles needed to achieve a balanced configuration for the Dell PowerEdge™ R720xd server using the Microsoft Fast Track Data Warehouse 4.0 guidelines.

Microsoft SQL Server Fast Track Data Warehouse

In order to overcome the limitations of traditional data warehouse systems, Microsoft has developed a cost-effective solution that optimally balances the hardware and software capabilities of the system. It provides an easy to deploy data warehouse infrastructure by mainly focusing on storage tuning and database layout. Fast Track Data Warehouse (FTDW) implements data warehouse solutions differently. As most data warehouse queries scan large volumes of data, FTDW designs are optimized for sequential scans and reads. These methodologies yield performance much better than that of traditional data warehousing systems. Based on this, Dell has developed a reference guide that helps customers implement FTDW on Dell hardware.

Dell Fast Track Data Warehouse Reference Architecture

In order to optimize data warehouse stack component performance, each layer must be properly tuned. Table 1 lists the Proposed Reference architecture along with the assigned Solution IDs. The following sections explain the tuning of selected hardware and software.

Table 1. Dell Fast Track Reference Architecture Solution Details

Solution Description	Configuration ID
Dell Fast Track 4.0 Configuration PowerEdge R720xd	2209618

Hardware Component Architecture

Redundant and robust tests have been conducted on Dell's PowerEdge servers to determine best practices and guidelines for building a balanced FTDW system.

Figure 1. Proposed Dell Fast Track Reference Architecture



Table 2. Tested Dell Fast Track Reference Architecture Component Details

Component	Details
Server	PowerEdge R720xd
CPU	(1) Intel® Xeon® E5-2643 CPU @3.3GHz (HT Enabled)
Number of sockets used	1
Total Number of CPU Cores	4
Memory	96GB RAM (4 X 16GB DDR3 DIMMs @1333MHz 4 X 8GB DDR3 DIMMs @1333MHz)
Internal Hard Drives	22x 300GB 10K 2.5" SAS (18 Data, 4 Logs) 2x 600GB 10K 2.5" SAS (2 Hot Spares) 2x 600GB 10K 2.5" SAS (2 drives OS) Rear Bay
Operating System	Microsoft Windows® 2008 R2 SP1 Enterprise Edition
Database Software	Microsoft SQL Server 2012 Enterprise Edition

Dell PowerEdge R720xd Server

The Dell PowerEdge R720xd server is a 2-socket, 2U high-capacity, multi-purpose rack server offering an excellent balance of internal storage, redundancy, and value in a compact chassis.

The PowerEdge R720xd server was developed with a purposeful design, energy-optimized options and enterprise-class manageability. For more technical specifications of the R720xd Server, refer to the *Power Edge R720xd Technical Guide*, a link to which is provided in the References section of this document.

Processors

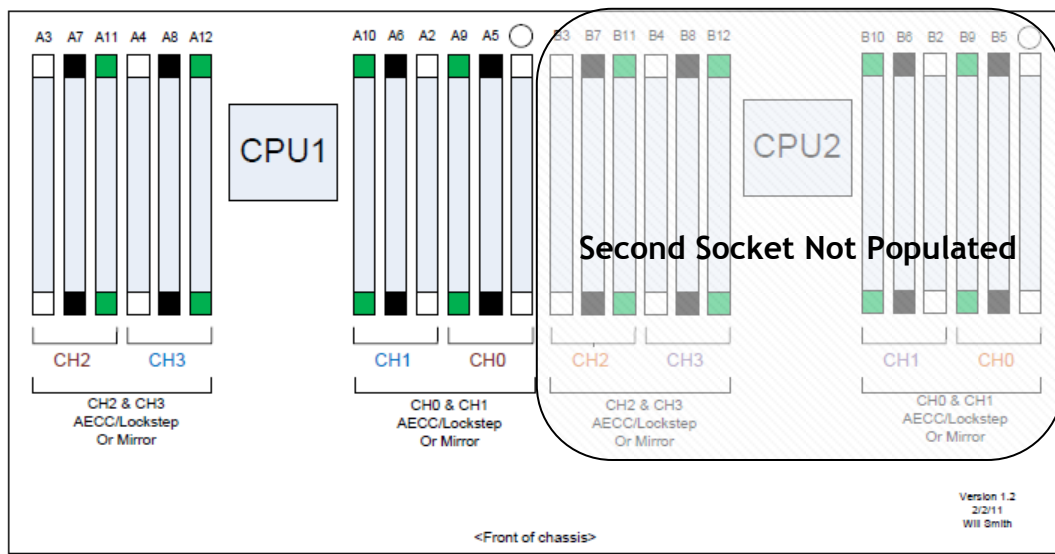
The *Microsoft Fast Track 4.0 Reference Guide* describes how to achieve a balance between components such as storage, memory, and processors. In order to balance the available internal storage and memory for the Dell PowerEdge R720xd, a single Intel Xeon E5-2643 four core processor operating at 3.3GHz speed was used.

Memory

For SQL 2012 reference architectures, Microsoft recommends using 64GB to 128GB of memory for one socket configuration. With enough memory installed on the system, the large-scale queries involving Hash joins and sorting operations will benefit from SQL Server offloading operations from the Tempdb to Memory. Selection of memory DIMMS will also play a critical role in the performance of the entire stack. In our test configuration, we have configured the database server with 96GB of RAM running at 1333 MHz speed.

To achieve 96GB of RAM on the PowerEdge R720xd server, four 8GB RDIMMS are placed on slots A1-A4 (white connectors) and four 16GB RDIMMS are placed on slots A5-A8 (black connectors). See the figure below for memory slot locations.

Figure 2. Memory Slot Locations



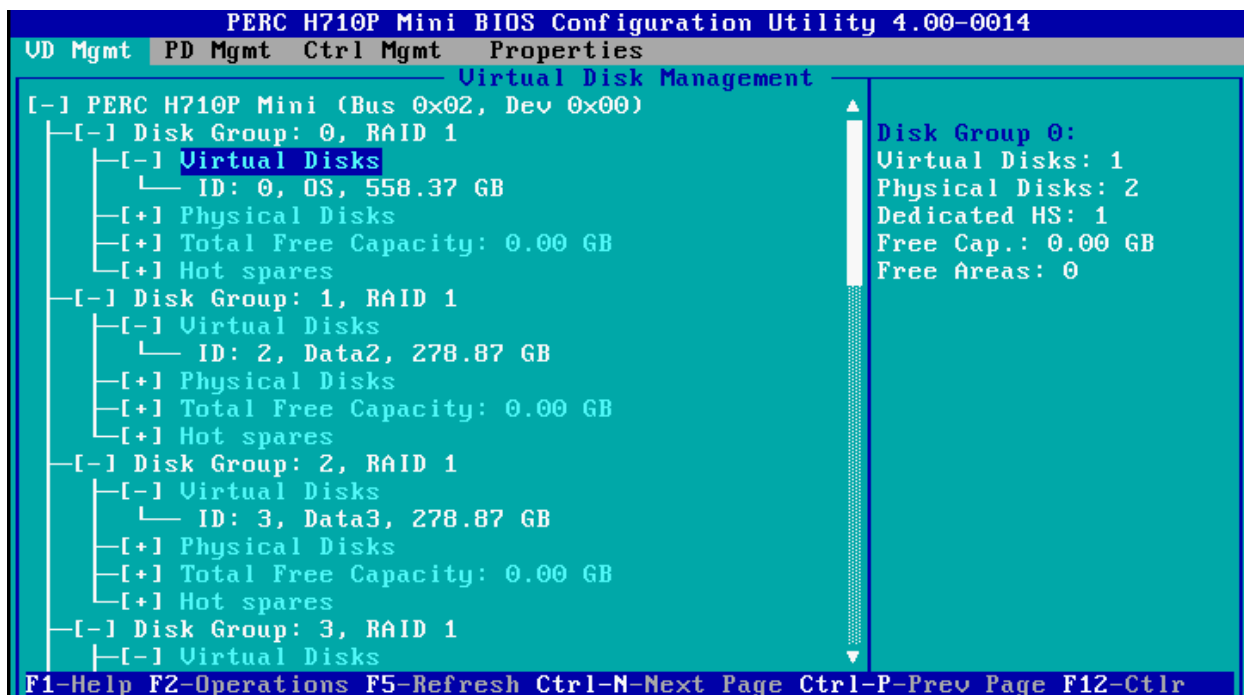
*White Connector = 1st DIMM population
Black Connector = 2nd DIMM population
Green Connector = 3rd DIMM population

Refer to the *Microsoft Fast track 4.0 Reference Guide* for detailed recommendations on system memory configuration.

Internal Storage Controller (PERC H710P Mini) Settings

The Dell PERC H710P Mini is an enterprise-level RAID controller that provides disk management capabilities, high-availability, and security features in addition to improved performance of up to 6GB/s throughput. Figure 3 shows the management console accessible through the BIOS utility.

Figure 3. Virtual Disk Settings



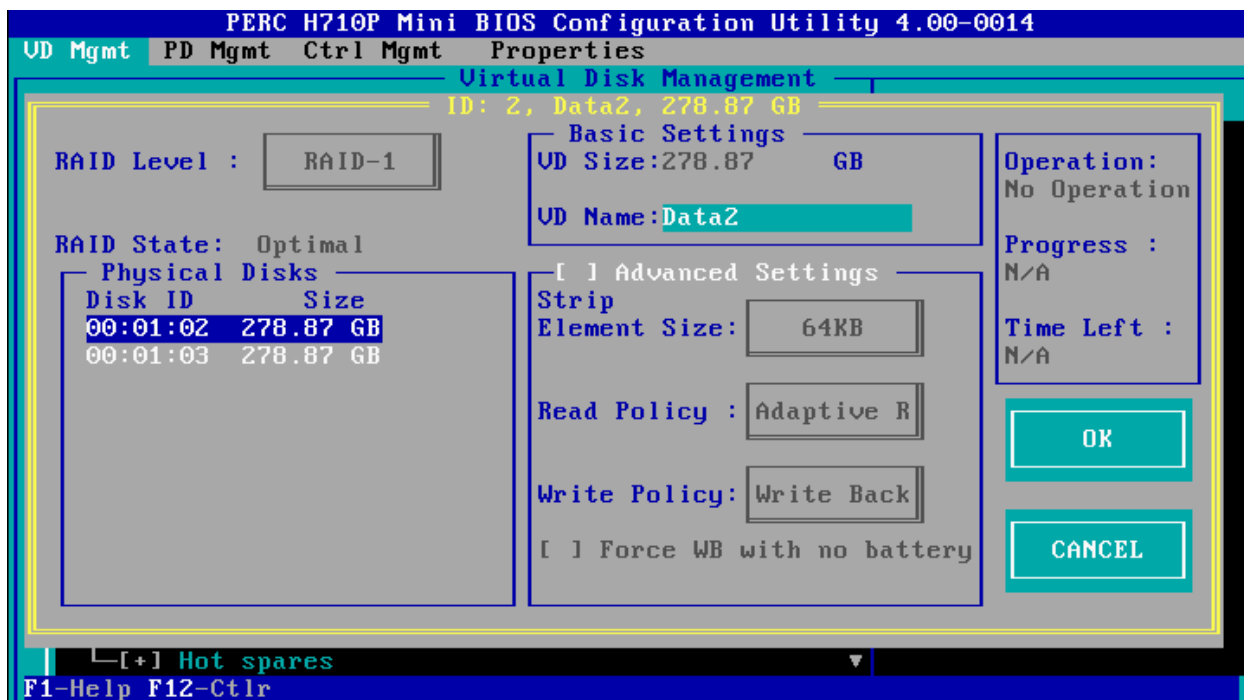
Stripe element size

By default, the PERC H710P Mini creates virtual disks with a segment size of 64KB. During the Fast Track validation testing, Dell tested stripe element sizes of 64KB and 256KB to compare and contrast any performance improvements. For most workloads, the 64KB default size will provide an adequate stripe element size. Dell recommends testing various stripe sizes depending on the workload characteristics of your configuration.

Read policy

The default setting for the read policy on the PERC H710P Mini is "adaptive read ahead." Configuration was tested with "adaptive read ahead," "No read ahead," and "Read Ahead" settings. During testing, we observed that the default setting ("adaptive read ahead") gave the best performance.

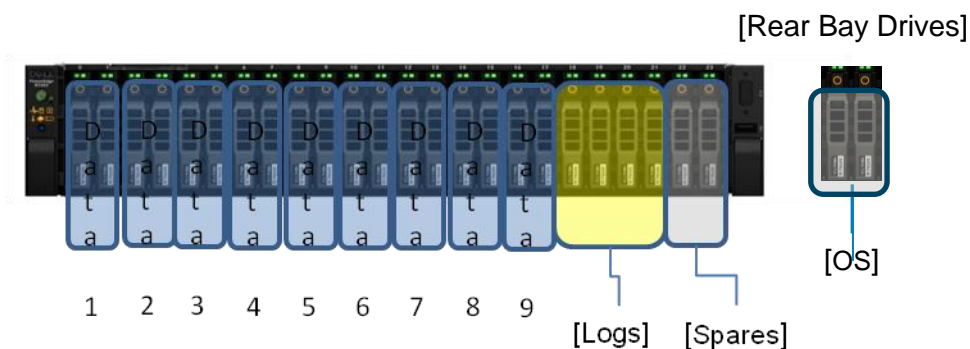
Figure 4. Internal Storage Controller Settings



RAID configuration

One of the most critical decisions to make when deploying a new storage solution is which RAID type(s) to use, because that choice heavily impacts the performance of the application. We have configured the proposed Fast Track configuration using RAID 1 disk groups for database data files and RAID 10 for database log files. Nine RAID 1 data disk groups and one RAID 10 log disk group were created each with a single Virtual Disk. Figure 5 shows the proposed RAID configuration.

Figure 5. RAID Configuration



Drive Slot configuration:

- Slot 0-17: Nine RAID 1 disk groups were created, each configured with a single Virtual Disk dedicated for the primary user data.
- Slot 18-21: One RAID 10 disk group was created of four disks and a single Virtual Disk dedicated to host the database log files.
- Slot 22-23: The remaining two disks were assigned as global hot spares.
- Slot 24-25: Rear Bay Drives, One RAID 1 disk group for operating system and stage space.

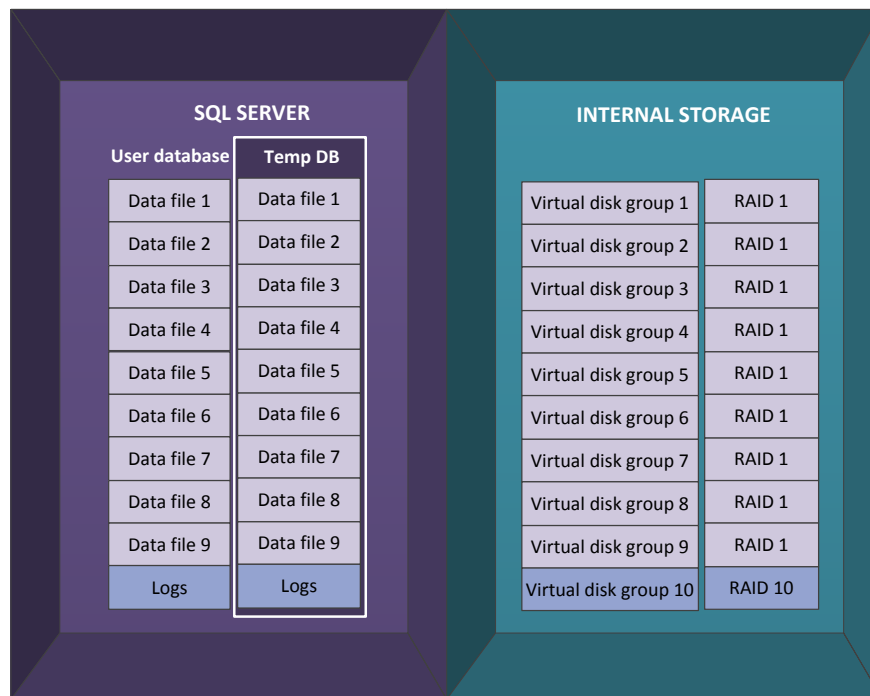
For Fast Track architectures, Dell recommends that you use mount-point rather than drive letters for storage access. It is also very important to assign the appropriate virtual disk and mount-point names to the configuration in order to simplify troubleshooting and performance analysis. Mount-point names should be assigned in such a way that the logical file system reflects the underlying physical storage enclosure mapping. Table 3 shows the virtual disk 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 3. Mount Point Naming and the Storage Enclosure Mapping

Disk Group	Virtual Disk	Virtual Disk Label	Logical Label	Full Volume Path
1	1	Cage1-Card1-vData1	Data1	C:\FT\PRI\Cage1-Card1-vData1
2	2	Cage1-Card1-vData2	Data2	C:\FT\PRI\Cage1-Card1-vData2
3	3	Cage1-Card1-vData3	Data3	C:\FT\PRI\Cage1-Card1-vData3
4	4	Cage1-Card1-vData4	Data4	C:\FT\PRI\Cage1-Card1-vData4
5	5	Cage1-Card1-vData5	Data5	C:\FT\PRI\Cage1-Card1-vData5
6	6	Cage1-Card1-vData6	Data6	C:\FT\PRI\Cage1-Card1-vData6
7	7	Cage1-Card1-vData7	Data7	C:\FT\PRI\Cage1-Card1-vData7
8	8	Cage1-Card1-vData8	Data8	C:\FT\PRI\Cage1-Card1-vData8
9	9	Cage1-Card1-vData9	Data9	C:\FT\PRI\Cage1-Card1-vData9
10	10	Cage1-Card1-vLog	Log	C:\FT\LOG\Cage1-Card1-vLog

Figure 6 represents the storage system configuration for the proposed Fast Track reference.

Figure 6. Storage System Components



The production, staging, and system temp databases were deployed per the recommendations provided in the *Microsoft Fast Track Data Warehouse 4.0 Reference Guide*.

Application Configuration

The sections below explain the settings applied for the operating system and database layers.

Windows Server 2008 R2 SP1

Settings used for Windows 2008 R2 SP1 operating system:

- **Enable Lock Pages In Memory:** This Policy setting is enabled to prevent the system from paging memory to disk.
For more information, refer to [How to: Enable the Lock Pages in Memory Option](#) in the Reference Section.

SQL Server Configuration

- The following startup options were added to the SQL Server Startup options:
 - **-E:** This parameter increases the number of contiguous extents that are allocated to a database table in each file as it grows. This improves sequential access.
 - **-T1117:** This trace flag ensures the even growth of all files in a file group when auto growth is enabled. It should be noted that the Fast Track reference guidelines recommend that you pre-allocate the data file space rather than allow auto grow.
 - **-T834:** This should be evaluated on a case-by-case basis. This can improve throughput rates for many DW workloads. This flag enables large page allocations in memory for the SQL Server buffer pool. However, at this time SQL Server 2012 does not support the use of -T834 in the case that Column Store Indexes (CSI) are in use on the database. If CSI use is planned, do not use this trace flag. For more information about this and other trace flags, refer to [SQL Server Performance Tuning & Trace Flags](#) in the Reference Section.
- **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 operating system should be adjusted accordingly. For this reference architecture, the maximum server memory was set at 90112 MB, i.e., 88GB.
- **Resource Governor:** For SQL Server 2012, Resource Governor provides a maximum of 25% of SQL Server memory resources to each session. The Resource Governor setting can be used to reduce the maximum memory consumed per query. While it can be beneficial for many data warehouse workloads to limit the amount of system resources available to an individual session, this is best measured through analysis of concurrent query workloads. For the test configuration, Resource Governor memory grant was set at 19%. For more information, refer to [Using the Resource Governor](#) in the Reference Section.
- **Max Degree of Parallelism:** The SQL Server configuration option 'max degree of parallelism' controls the number of processors used for the parallel execution of a query. For the test configuration, the 'max degree of parallelism' was set at 12. For more information, refer to [Maximum degree of parallelism configuration option](#) in the Reference Section.

Performance Benchmarking

Microsoft Fast Track guidelines help to achieve optimized database architecture with balanced CPU and storage bandwidth. The following sections describe the performance characterization activities carried out for the validated Dell Microsoft Fast Track reference architecture.

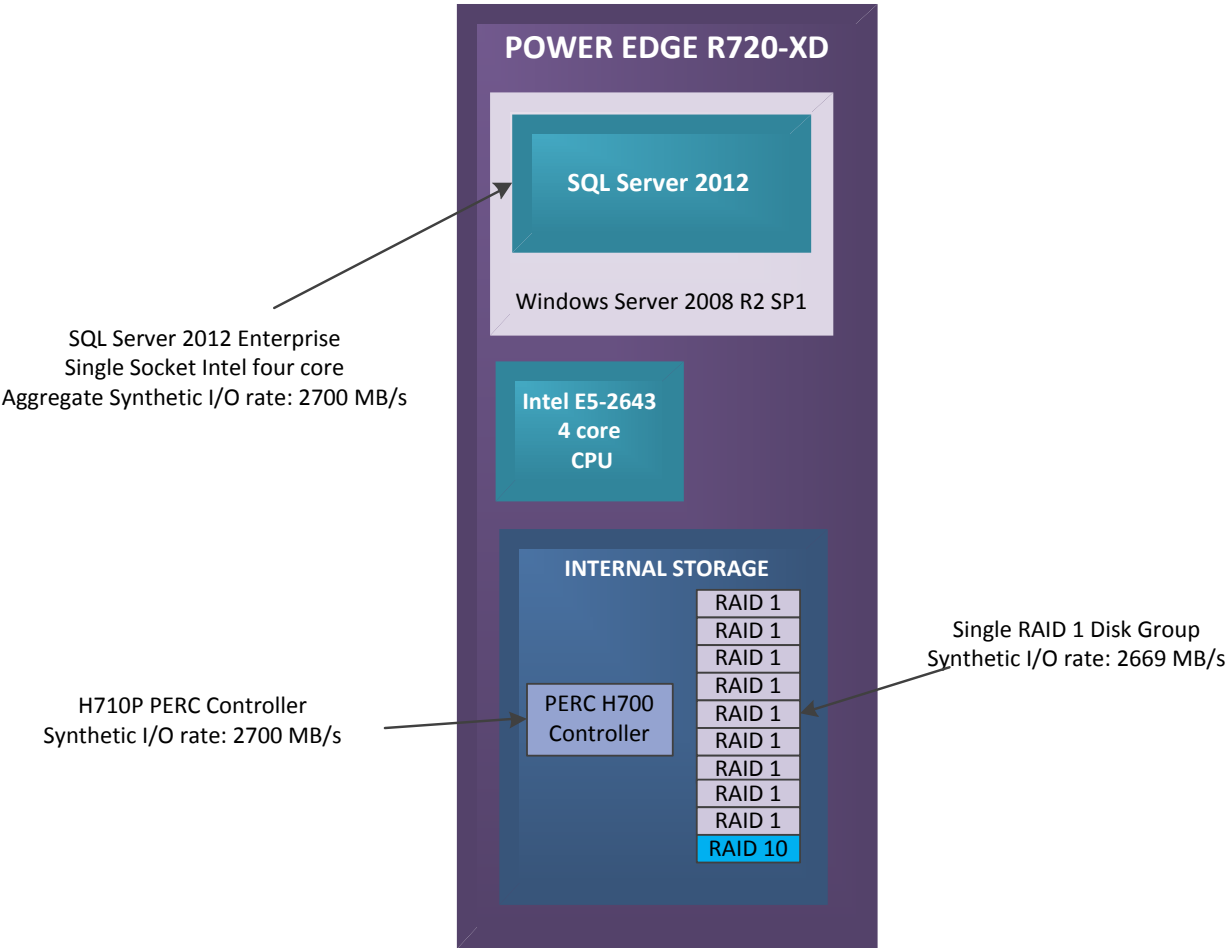
Baseline Hardware Characterization using Synthetic I/O

The goal of hardware validation is to determine actual baseline performance characteristics of key hardware components in the database stack. You must thoroughly analyze the storage hardware to make sure that the backend storage is capable of delivering the maximum possible throughput. This will ensure that the performance of the system is not bottlenecked in any of the intermediate layers.

The disk characterization tool, SQLIO, was used to validate the configuration. Please refer to the *Fast Track Reference Guide* (link provided in the reference section) for detailed guidelines. Figure 7 and Figure 8 show the baseline performance numbers achieved for the validated reference architecture. The results in Figure 7 show the maximum baseline that the system can achieve from a cache called Line Rate. A small file is placed on the storage, and large sequential reads are issued against it with

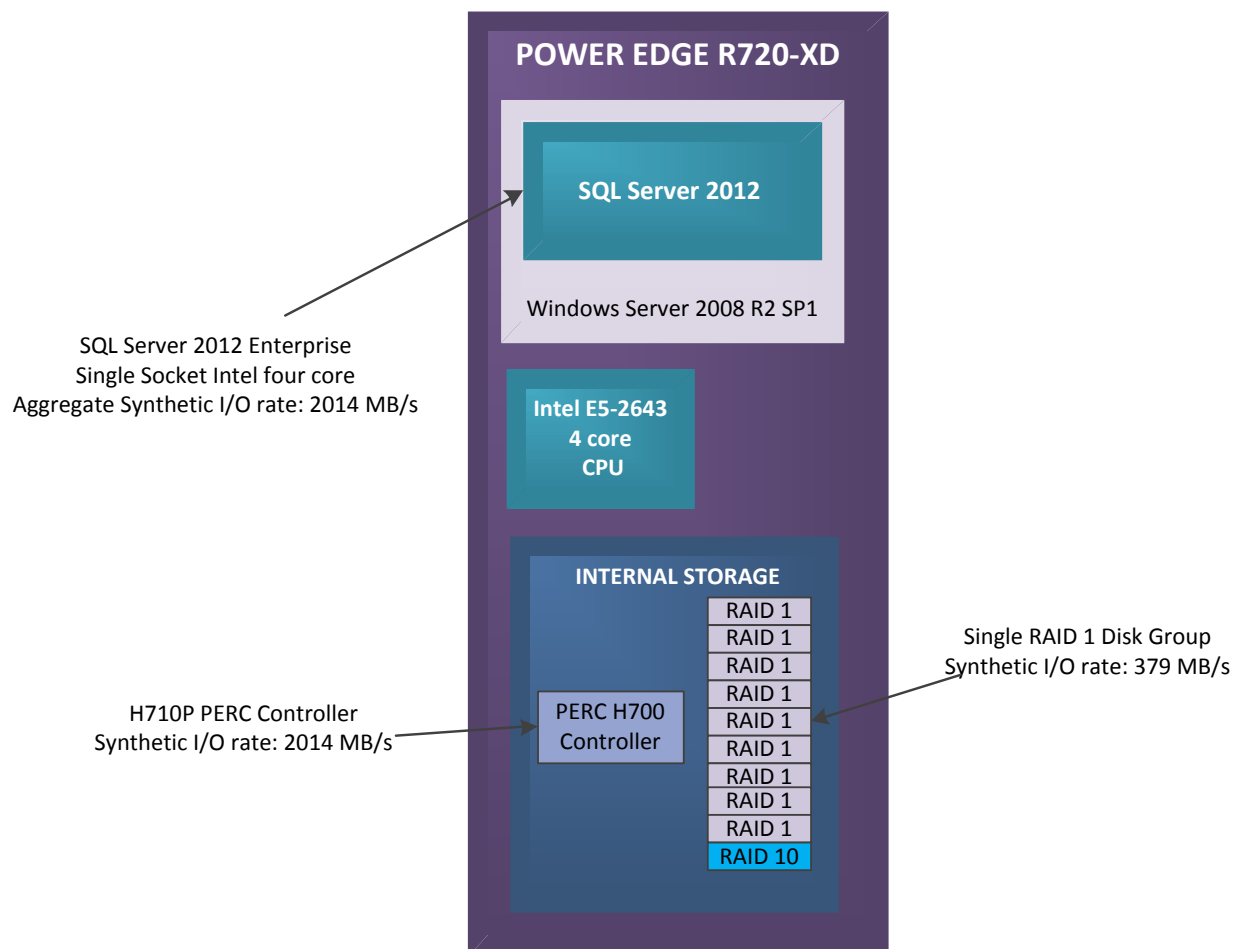
SQLIO. This test verifies the maximum bandwidth available in the system to ensure no bottlenecks are within the data path.

Figure 7. SQLIO Line Rate Test from Cache (Small File)



The second synthetic I/O test with SQLIO was performed with a large file to ensure reads are serviced from the storage system hard drives instead of from cache. Figure 8 shows the maximum real rate that the system is able to provide with sequential reads.

Figure 8. SQLIO Real Rate Test from Disk (Large File)



Fast Track Database Validation

The performance of a Fast Track database configuration is measured using two core metrics: Maximum CPU Consumption Rate (MCR) and Benchmark Consumption Rate (BCR).

Calculating MCR

MCR indicates the per core I/O throughput in MB or GB per second. This is measured by executing a pre-defined query against the data in the buffer cache, and measuring the time taken to execute the query against the amount of data processed in MB or GB. MCR value provides a baseline peak rate for performance comparison and design purposes.

For the validated configuration with one Intel E5-2643 four core processors, the system aggregate MCR was 1244 MB/s. The realized MCR value per core was 311 MB/s.

Calculating BCR

BCR is calculated in terms of total read bandwidth from the storage hard drives and not from the buffered cache as in the MCR calculation. This is measured by running a set of standard queries specific to the data warehouse workload. The queries range from I/O intensive to CPU and memory intensive, and provide a reference to compare various configurations.

For the validated Fast Track configuration, the aggregate BCR was 950.75 MB/s.

During the evaluation cycle, the system configuration was analyzed for multiple query variants (simple, average, and complex) with multiple sessions and different degrees of parallelism (MAXDOP) options to arrive at the optimal configuration.

Fast Track Database Validation with Column Store Index (CSI)

SQL Server 2012 implements CSI technology as a nonclustered indexing option for pre-existing tables. FTDW for SQL Server 2012 system design and validation is based on non-CSI benchmarks. FTDW systems are designed to run effectively in the case that no columnar optimization is achieved for any given period of time. Significant performance gains are often achieved when CSI query plans are active and this performance can be viewed as incremental to the basic system design.

After the test configuration was validated, CSI was added and the I/O and the CPU saturation packages were run. The configuration was analyzed for multiple query variants (simple, average, and complex) with multiple sessions, degrees of parallelism (MAXDOP) set at 12 and Resource Governor memory grant set at 19%.

Reference Architecture Performance Details

Table 4 shows the performance numbers reported for the recommended reference configuration.

Table 4. Performance Metrics

Metric	Value	Description
FTDW Rated Data Warehouse Capacity (TB)	5	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)	950.75	Core performance metric for validation. Is the average of Physical and Logical I/O.
Benchmark Scan Rate Logical (MB/s)	1156.5	Reflects actual user query throughput which includes reads from RAM/Buffer cache.
Benchmark Scan Rate Physical (MB/s)	745	Reflects physical IO read from disk during benchmark.
FTDW Peak I/O (MB/s)	1512	Maximum observed IO rate.
FTDW Rated CSI (MB/s)	1662.5	Represents potential throughput using Columnstore Index.

Conclusion

The Dell Microsoft Fast Track Data Warehouse architecture provides a uniquely well-balanced data warehouse implementation solution. By following the best practices at all the layers of the stack, a balanced data warehouse environment can be achieved with a greater performance benefit than the traditional data warehouse systems.

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

- Tested and validated configuration with proven methodology and performance behavior.
- A balanced and optimized system at all levels of the stack by following the best practices of hardware and software components.
- Avoidance of over-provisioning of hardware resources.
- High-availability at all the levels of setup (host, switches, and storage).
- Help to avoid the pitfalls of improperly-designed and configured systems.
- Reduced future support costs by limiting solution re-architect efforts because of scalability challenges.

References

Dell SQL Server Solutions

www.dell.com/sql

Dell Services

www.dell.com/services

Dell Support

www.dell.com/support

OLTP and OLAP

<http://datawarehouse4u.info/OLTP-vs-OLAP.html>

Microsoft Fast Track Data Warehouse and Configuration Guide Information

www.microsoft.com/fasttrack

http://download.microsoft.com/download/B/E/1/BE1AABB3-6ED8-4C3C-AF91-448AB733B1AF/Fast_Track_Configuration_Guide.docx

An Introduction to Fast Track Data Warehouse Architectures

<http://msdn.microsoft.com/en-us/library/dd459146.aspx>

How to: Enable the Lock Pages in Memory Option

<http://go.microsoft.com/fwlink/?LinkId=141863>

SQL Server Performance Tuning & Trace Flags

<http://support.microsoft.com/kb/920093>

Using the Resource Governor

<http://msdn.microsoft.com/en-us/library/ee151608.aspx>

Maximum degree of parallelism configuration option

support.microsoft.com/kb/2023536

Power Edge R720xd Technical Guide

<http://www.support.dell.com/support/edocs/systems/per720/en/index.htm>