

---

# Microsoft SQL Server 2012 Fast Track Reference Configuration using PowerEdge R720 and PowerVault MD3620F

---

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

Narasimha Reddy Gopu

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.

© 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, PowerVault, 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 SQL 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

Executive summary .....	4
Introduction to Microsoft SQL Server Fast Track Reference Architectures .....	4
What's New with Dell Microsoft Fast Track Reference Architectures? .....	4
Dell Microsoft SQL Server 2012 Fast Track Reference Architectures using PowerEdge R720 and PowerVault MD3620F .....	5
Single Server Reference Architecture .....	6
Highly Available Reference Architecture .....	7
Reference Architecture Configuration Hardware Component Details .....	8
Reference Architecture Configuration Tuning Details .....	10
Reference Architecture Performance Details .....	15
Summary .....	15

## Tables

Table 1. Recommended Reference Architectures with Solution IDs .....	5
Table 2. Single Server Reference Architecture Details .....	6
Table 3. Highly Available Reference Architecture Details .....	8
Table 4. Switch Zoning Settings Corresponding to each Storage array .....	12
Table 5. Mount Point Naming and Storage Enclosure Mapping .....	14
Table 6. Performance Metrics .....	15

## Figures

Figure 1. Single Server Reference Architecture .....	6
Figure 2. Highly Available Reference Architecture .....	7
Figure 3. Cabling Diagram for Single Server Configuration .....	9
Figure 4. Enabling Storage High Performance Tier(Turbo Mode) .....	10
Figure 5. Modifying Storage Cache Block Size .....	11
Figure 6. Modifying the Virtual Disk Segment Size .....	11

## Executive summary

Dell™ and Microsoft®, in cooperation, provided 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 entire database stack to provide optimal throughput to match the database capabilities of the specific setup.

This white paper describes the architectural design principles and guidelines to achieve an optimally balanced SQL Server Data Warehouse solution according to the Microsoft Fast Track Data Warehouse 4.0 guidelines.

## 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 aligning to these recommended configuration best practices and guidelines:

- Accelerated data warehouse projects with pre-tested hardware configurations
- Reduced hardware and maintenance costs as a result of 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 ensure that the capability and throughput for the entire system is maximally utilized. These 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](http://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™ 12<sup>th</sup> 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 the 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 with the database, which is very beneficial for the data warehouses because of the huge amount of data

handled. Column store indexes also benefit the common data warehousing queries such as filtering, aggregating, grouping, and star-join queries.

Visit [msdn.microsoft.com/en-us/library/gg492088.aspx](http://msdn.microsoft.com/en-us/library/gg492088.aspx) for more information on Column Store Indexes.

The latest Dell Microsoft Fast Track reference architectures incorporate the benefit of Column Store Indexes for improved query performance.

Dell's 12th generation server platforms comprised of 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 MD3620F

We are proposing two different fast track reference architectures comprising of SQL Server 2012, PowerEdge R720, and PowerVault™ MD3620F. Table 1 lists the Proposed Reference architectures along with the assigned Solution IDs.

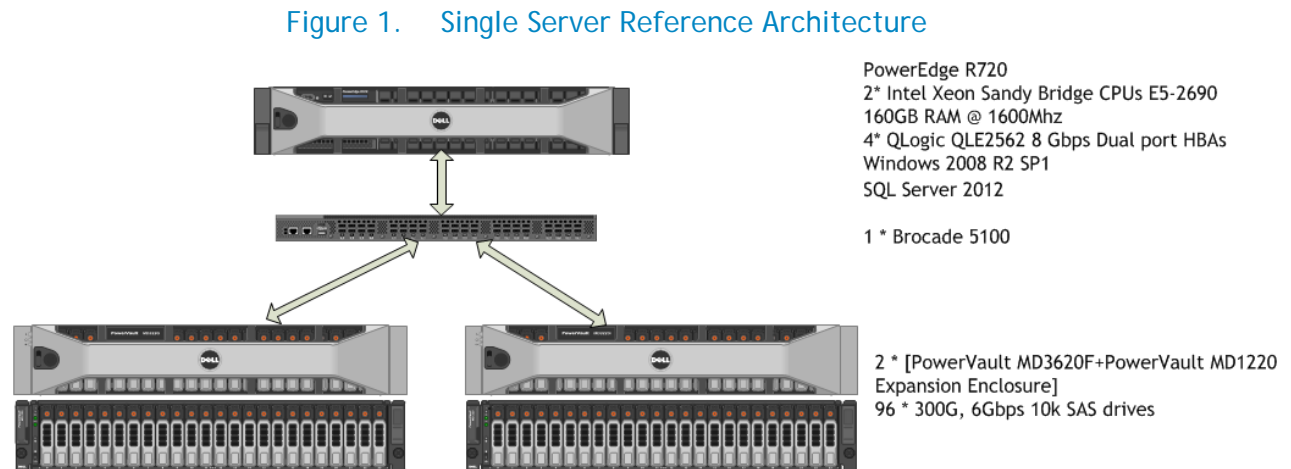
**Table 1. Recommended Reference Architectures with Solution IDs**

Reference Architecture	Solution ID
Single Server Reference Architecture	2319046
Highly Available Reference Architecture	2316699

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

## Single Server Reference Architecture

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




**Table 2. Single Server Reference Architecture Details**

Hardware Components	Details
Server	PowerEdge R720
Processors	2* Intel Xeon Sandy Bridge CPUs E5-2690 @ 2.9Ghz
Total No of Cores per Socket	8
Total No of Logical Processors (HT Enabled)	32
Total Installed Memory	160GB @ 1600Mhz
HBAs	4* QLogic QLE2562 8 Gbps Dual port HBAs
Multipathing Software	Dell MD Series DSM
Multipathing Policy	Least Queue Depth(Default)
Fiber Channel Storage	2 * MD3620F, each with a MD1220 Expansion Enclosure
Disks	96* 300G, 6Gbps 10k SAS drives
Fiber Channel Switch	1 * Brocade 5100
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 would be configured as the primary (active) server and the second server would be 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) would be 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.



**Table 3. Highly Available Reference Architecture Details**

Hardware Components	Details
Server	2 * PowerEdge R720
Processors	[2* Intel Xeon Sandy Bridge CPUs E5-2690 @ 2.9Ghz ] per Server
Total No of Cores per Socket	8
Total No of Logical Processors (HT Enabled)	[2 * 32 ] per Server
Total Installed Memory	160GB @ 1600Mhz per Server
HBAs	[4* QLogic QLE2562 8 Gbps Dual port HBAs] per Server
Multipathing Software	Dell MD Series DSM
Multipathing Policy	Least Queue Depth(Default)
Fiber Channel Storage	2 * MD3620F, each with a MD1220 Expansion Enclosure
Disks	96* 300G, 6Gbps 10k SAS drives
Fiber Channel Switch	2 * Brocade 5100
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 12<sup>th</sup> 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](http://dell.com/us/enterprise/p/poweredge-r720/pd) for more information on PowerEdge R720 Servers.



### QLogic QLE2562 8Gbps Host Bus Adapter

The proposed reference architectures make use of four dual port QLogic QLE2562 8Gbps HBA cards with the default HBA settings. QLogic Sansurfer FC HBA manager tool may be used to verify the HBA settings.

Visit [qlogic.com](http://qlogic.com) for more information on QLogic QLE2562 FC adapters.

### Brocade 5100 SAN Switch

Brocade 5100 switch is a rack mountable 8Gbps FC switch with 40x8-Gbps ports. It consists of redundant and hot-pluggable components and supports non-disruptive software upgrades.

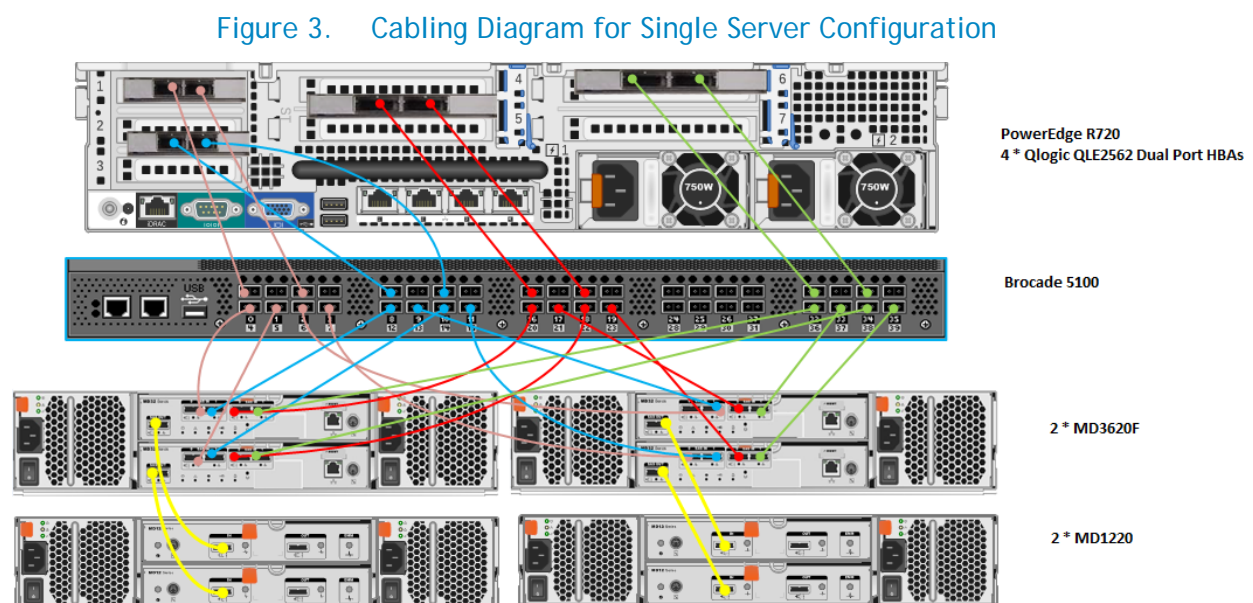
Visit [brocade.com/products/all/switches/product-details/5100-switch/index.page](http://brocade.com/products/all/switches/product-details/5100-switch/index.page) for more information on Brocade 5100 SAN switches.

### PowerVault MD3620F FC Storage Arrays

PowerVault MD3620F storage arrays support 2.5 inches 6Gbps SAS disks, which enable better backend performance. MD3620F supports up to two controllers. Each controller has four 8Gbps FC host side ports and one SAS out port, which is used to connect to any additional storage enclosures. Each RAID controller module contains 2GB of cache that is mirrored with the other controller's cache for high availability and is protected by a battery-powered cache offload mechanism.

Dell PowerVault MD3620F storage comes with its own Device Specific Module (DSM) software, included in the MD3620F's Resource DVD, which is fully integrated with Microsoft MPIO and helps you configure the multipath solutions.

Figure 3 depicts how the front and rear cables are connected in the configuration.



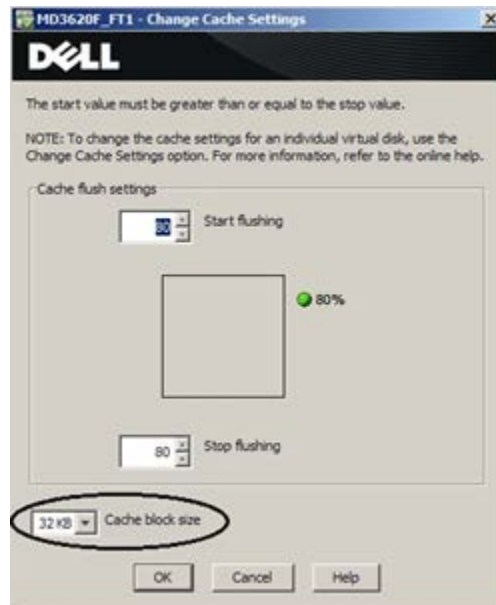
Visit [dell.com/us/enterprise/p/PowerVault-FC-SAN](http://dell.com/us/enterprise/p/PowerVault-FC-SAN) for more information on PowerVault MD3620F storage arrays.



### Set the Storage Cache Block Size to 32KB

In the 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 5).

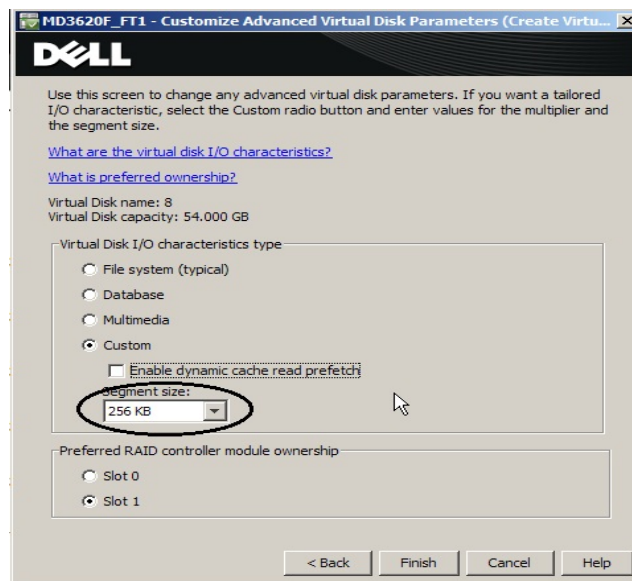
Figure 5. Modifying Storage Cache Block Size



### Set the Virtual disk segment Size to 256KB

For the recommended Fast Tract 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 6).

Figure 6. Modifying the Virtual Disk Segment Size



### Brocade 5100 Switch Tuning (Zoning)

For the reference configuration, Dell recommends using the WWN zoning for the switch. Zones must be created on the switch such that each zone comprises of a single initiator and single target. For each connected storage array, eight zones (one for each host HBA port) should be created and named following proper naming conventions, shown in Table 4. Similarly, eight more zones should also be created for the second storage array. Therefore, 16 total zones are created on the switch for the configuration.

Table 4. Switch Zoning Settings Corresponding to each Storage array

Zone Name	Zone Members
Zone_HBA_0-0	Server_HBA_0-0,Array_Ctrl-0-0
Zone_HBA_0-1	Server_HBA_0-1,Array_Ctrl-1-0
Zone_HBA_1-0	Server_HBA_1-0,Array_Ctrl-0-1
Zone_HBA_1-1	Server_HBA_1-1,Array_Ctrl-1-1
Zone_HBA_2-0	Server_HBA_2-0,Array_Ctrl-0-2
Zone_HBA_2-1	Server_HBA_2-1,Array_Ctrl-1-2
Zone_HBA_3-0	Server_HBA_3-0,Array_Ctrl-0-3
Zone_HBA_3-1	Server_HBA_3-1,Array_Ctrl-1-3

### 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 extents 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 147.2GB.
- **Resource Governor:** By default, SQL Server 2012 Resource Governor provides a maximum of 25% of SQL Server memory resources to each session. Resource Governor Settings can be used to modify the maximum memory consumed per session. 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 15%. Visit [msdn.microsoft.com/en-us/library/ee151608.aspx](https://msdn.microsoft.com/en-us/library/ee151608.aspx) for more information on Resource Governor.

- **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. A proper analysis of the database system configuration is required to appropriately configure this option. For the test configuration, the 'max degree of parallelism' setting of 16 was concluded to be delivering the best performance. Visit [support.microsoft.com/kb/2023536](http://support.microsoft.com/kb/2023536) for more information on 'max degree of parallelism' option.

## 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:

- Six RAID 5 disk groups were created, each consisting of seven disks (three disk groups each from the base and the expansion enclosures)
- One virtual disk (LUNs) was created out of each disk group. These LUNs were evenly distributed across the storage controllers.

Therefore, in each storage array, there were six LUNs (virtual disks) for the primary user data. Each storage controller was assigned three of the LUNs. Therefore, in the entire configuration, we have 12 RAID 5(6+1) virtual disks (six from each storage array) for primary data and temp database files.

On each storage array, for the primary user and system temp database transaction logs, one RAID10 disk group was created of four disks. This RAID group was dedicated to host the database transaction log files.

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

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 make the troubleshooting and performance analysis simpler. The mount point names should be assigned in such a way that the logical file system reflects the underlying physical storage enclosure mapping. Table 5 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 5. 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	SE1 (Base)	1	3	3	SA1-SE1-SP1-DG3-v3	PRI	C:\FT\PRI\SA1-SE1-SP1-DG3-v3
SA1	SE2 (Expansion)	2	4	4	SA1-SE2-SP2-DG4-v4	PRI	C:\FT\PRI\SA1-SE2-SP2-DG4-v4
SA1	SE2 (Expansion)	1	5	5	SA1-SE2-SP1-DG5-v5	PRI	C:\FT\PRI\SA1-SE2-SP1-DG5-v5
SA1	SE2 (Expansion)	2	6	6	SA1-SE2-SP2-DG6-v6	PRI	C:\FT\PRI\SA1-SE2-SP2-DG6-v6
SA2	SE1 (Base)	1	1	1	SA2-SE1-SP1-DG1-v1	PRI	C:\FT\PRI\SA2-SE1-SP1-DG1-v1
SA2	SE1 (Base)	2	2	2	SA2-SE1-SP2-DG2-v2	PRI	C:\FT\PRI\SA2-SE1-SP2-DG2-v2
SA2	SE1 (Base)	1	3	3	SA2-SE1-SP1-DG3-v3	PRI	C:\FT\PRI\SA2-SE1-SP1-DG3-v3
SA2	SE2 (Expansion)	2	4	4	SA2-SE2-SP2-DG4-v4	PRI	C:\FT\PRI\SA2-SE2-SP2-DG4-v4
SA2	SE2 (Expansion)	1	5	5	SA2-SE2-SP1-DG5-v5	PRI	C:\FT\PRI\SA2-SE2-SP1-DG5-v5
SA2	SE2 (Expansion)	2	6	6	SA2-SE2-SP2-DG6-v6	PRI	C:\FT\PRI\SA2-SE2-SP2-DG6-v6
SA1	-	1	7	7	SA1-SE2-SP1-DG7-v7	LOG	C:\FT\SA1_LOG\SA1-SP1-DG7-v7
SA2	-	1	7	7	SA2-SE2-SP1-DG7-v7	LOG	C:\FT\SA2_LOG\SA2-SP1-DG7-v7

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

## Reference Architecture Performance Details

Table 6 shows the performance numbers reported for the recommended reference configuration.

Table 6. Performance Metrics

Metric	Value	Description
FTDW Rated Capacity (TB)	20	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)	2930	Core performance metric for validation. Is the midpoint of Physical and Logical I/O.
Benchmark Scan Rate Logical (MB/s)	3540	Reflects actual user query throughput which includes reads from RAM/Buffer cache.
Benchmark Scan Rate Physical (MB/s)	2310	Reflects physical IO read from disk during benchmark.
FTDW Peak I/O (MB/s)	4660	Maximum observed IO rate.
FTDW Rated CSI (MB/s)	5020	Represents potential throughput using Columnstore Index.

## Summary

Dell, in partnership with Microsoft, enables the customers to enhance the ROI (Return of Investment) on the 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 benefits than the 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 the 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.