

85TB Data Warehouse Fast Track Reference Architecture for Microsoft SQL Server 2016 using Dell EMC PowerEdge R730 and SC7020

Configuration and performance results

Dell EMC Storage Applications Engineering
November 2016

Revisions

Date	Description
November 2016	Initial release

Acknowledgements

Author: Mike Matthews, Dell EMC Storage Applications Engineering

Special thanks to Jamie Reding, Sadashivan Krishnamurthy and Jarupat Jisarojito at Microsoft.

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2016 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [11/28/2016] [Reference Architecture] [3229-RA-SQL]

Dell EMC believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of contents

Revisions.....	2
Acknowledgements.....	2
Executive summary.....	5
1 Microsoft SQL Server 2016: data warehousing with improved column-store technology.....	6
1.1 Dell EMC Data Warehouse Fast Track reference architectures for SQL Server 2016.....	6
2 Recommended reference architectures	7
2.1 Single server reference architecture	7
2.2 Highly available reference architecture	9
3 Hardware components	11
3.1 PowerEdge R730 server	11
3.2 QLogic QLE2662 16 Gbps host bus adapter	11
3.3 Brocade 6505 SAN switch.....	11
3.4 SC7020 storage array	11
4 Storage configuration	13
4.1 Cabling.....	13
4.2 I/O ports	14
4.3 Disk folder	14
4.4 Storage type	14
4.5 SC Series storage profiles.....	15
4.6 Read and write cache	15
4.7 SC Series server object.....	15
4.8 SC Series volumes	17
5 Server configuration	19
5.1 System BIOS	19
5.2 Host bus adapters.....	19
6 Windows Server 2016 configuration.....	20
6.1 Power plan.....	20
6.2 Lock pages in memory	20
6.3 Windows volumes.....	20
6.4 MPIO.....	21
7 SQL Server 2016 Enterprise Edition configuration	23
7.1 Startup parameters.....	23

7.2	SQL Server maximum memory	23
7.3	Max degree of parallelism (MAXDOP)	23
7.4	Resource governor	23
7.5	Database configuration.....	23
7.6	Tempdb configuration	23
8	Additional considerations for the Highly Available (HA) reference architecture	24
9	DWFT for SQL Server 2016 certification.....	25
10	Summary	26
A	Technical support and resources	27

Executive summary

Dell EMC and Microsoft®, in cooperation, provide guidelines and 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. Innovative 13th-generation Dell EMC servers, along with robust and cutting edge Dell EMC SC7020 storage arrays, form efficient candidates for a high performing data warehouse solution.

This reference architecture describes the design principles and guidelines used to achieve an optimally balanced 85 TB Data Warehouse Fast Track reference architecture for SQL Server 2016 using Dell EMC PowerEdge™ R730 servers and SC7020 arrays. The configuration used to achieve the performance numbers for the reference configuration is discussed in detail.

The target audience for this paper is database administrators, business intelligence architects, storage administrators, IT directors and data warehousing users seeking sizing and design guidance for their business intelligence solutions with Microsoft SQL Server 2016.

1 Microsoft SQL Server 2016: data warehousing with improved column-store technology

Microsoft SQL Server 2016 has made significant improvements in data warehousing technologies and performance, including column-store features as well as many other improvements.

Column-store indices offer great advantages over traditional row stores for analytics and data warehousing queries. They are ideally suited for the star schemas, and tables with billions of rows which are commonly seen. Among their advantages for analytics are:

- **Up to 10X compression in data size**

Data warehouses are very large by nature and the compression offered by column-store index technologies offers both space and cost savings as well as significantly increased performance. These benefits are possible due to the dramatically reduced I/O requirements given by the compression and coupled by the ability to only scan the specific columns required by each query. This compression also reduces the amount of memory required to hold a given number of rows from the source data warehouse.

- **Additional indices**

SQL Server 2016 adds the capability to add (B-Tree) indices to column store-based tables, which enables efficient single-row lookup.

In addition to these architectural features, we have further optimized the processing of queries in column-store indices in the following ways:

- **Operator pushdown**

Pushdown refers to moving both filter and aggregation query operations closer to the data, so that many of the filters and calculations can be done in the scan operators, dramatically reducing the volume of data which needs to be handled further on in query processing.

- **Batch-mode processing**

SQL Server 2016 includes enhancements in batch-mode processing which processes many rows at a time rather than serially doing calculations on each individual row. These batch operations are further optimized by leveraging Single Instruction Multiple Data (SIMD) vector processing CPU instructions in the Intel® architectures.

1.1 Dell EMC Data Warehouse Fast Track reference architectures for SQL Server 2016

Dell EMC and Microsoft have refreshed the DWFT reference architecture offerings with the latest technology advancements in database, server, and storage technology. Dell EMC PowerEdge 13th-generation servers with Intel Xeon® processors, Dell EMC SC7020 arrays, and Microsoft SQL Server 2016 are the latest additions to the list of reference architecture components.

The Dell EMC DWFT reference architectures for SQL Server 2016 are engineered jointly by Dell EMC and Microsoft. The hardware and software optimizations are tested by Dell EMC and the performance results are crosschecked by Microsoft. This approach presents a faster time-to-value using integrated, balanced, and verified architectures.

2 Recommended reference architectures

The following subsections describe the two different DWFT reference architectures for SQL Server 2016, comprised of PowerEdge R730 servers and SC7020 storage arrays.

2.1 Single server reference architecture

Figure 1 illustrates the single server reference architecture with the major elements and Table 1 lists the component details.




Figure 1 Single server reference architecture

Table 1 Single server reference architecture details

Component	Description	
Server	PowerEdge R730	
	Processors	Two Intel® Xeon® E5-2697v4 (2.3Ghz, 18 cores, 36 threads)
	Total cores	36
	Total logical processors	72 (Hyper-Threading enabled)
	Total memory	1,536 GB
	Network adapters	Minimum of one network adapter (1 Gbps or 10 Gbps based on requirements) Recommended to have more than one network adapter with load balancing configured
	Host bus adapters	Four QLogic® QLE2662 dual port 16 Gbps FC HBA
Software	Operating system	Windows Server® 2016 Standard Edition
	Database software	SQL Server 2016 Enterprise Edition
Storage	Storage array	One SC7020 (SCOS v7.1)
	I/O cards	Two QLogic QLE2694 quad port 16 Gbps FC HBAs (one per controller)
	Disk drives	25 x 1.92 TB read intensive SSDs (2.5" SAS)
	SAN switches	Two Brocade® 6505 with 16 Gbps SFPs

2.2 Highly available reference architecture

For achieving high availability for the database, Windows® failover clustering is recommended. Using Microsoft clustering services, one database server is configured as the primary (active) server and the second server is configured as the secondary (passive) server. The secondary server should have exactly the same configuration as the primary server. Since the database is only 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 illustrates the highly available reference architecture with the major elements and Table 2 lists the component details.




Figure 2 Highly available reference architecture

Table 2 Highly available reference architecture details

Component	Description	
Server	Two PowerEdge R730	
	Processors	Two Intel Xeon E5-2697v4 (2.3Ghz, 18 cores, 36 threads) per server
	Total cores	36 per server
	Total logical processors	72 per server (Hyper-Threading enabled)
	Total memory	1,536 GB per server
	Network adapters	Public network: Minimum of two network adapters (1 Gbps or 10 Gbps based requirements) per server with load balancing configured Private network (cluster): Minimum of one network adapter with 1 Gbps (or more) connectivity per server
	Host bus adapters	Four QLogic QLE2662 dual port 16 Gbps FC HBA per server
Software	Operating system	Windows Server 2016 Standard Edition with Windows Failover Clustering
	Database software	SQL Server 2016 Enterprise Edition configured as a Failover Cluster Instance
Storage	Storage array	One SC7020 (SCOS v7.1)
	I/O cards	Two QLogic QLE2694 quad port 16 Gbps FC HBAs (one per controller)
	Disk drives	25 x 1.92 TB read intensive SSDs (2.5" SAS)
	SAN switches	Two Brocade 6505 with 16 Gbps SFPs

3 Hardware components

3.1 PowerEdge R730 server

The PowerEdge R730 is a highly versatile, two-socket 2U rack server with impressive processor performance, a large memory footprint, extensive I/O options and a choice of dense, high performance storage or low-cost, high-capacity storage. The R730 offers simplified management, purposeful design, and energy efficiency with support for the latest Intel Xeon E5 series Broadwell processors and ECC DDR4 memory with a maximum memory capacity of 1536 GB. PowerEdge R730 provides up to seven Gen3 PCIe slots.

For more information about PowerEdge R730 Servers, visit <http://www.dell.com/us/business/p/poweredge-r730/pd>.

3.2 QLogic QLE2662 16 Gbps host bus adapter

The reference architecture makes use of four QLogic QLE2662 dual port 16 Gbps Fibre Channel (FC) to PCI Express host bus adapters.

For more information about QLogic QLE2662 FC adapters, visit http://www.qlogic.com/OEMPartnerships/Dell/Documents/FAQ_16Gb_Fibre_Channel_2600_Host_Bus_Adapters.pdf.

3.3 Brocade 6505 SAN switch

The Brocade 6505 switch is a 1U, 24 port, rack-mountable Fibre Channel switch providing up to 16 Gbps of bandwidth per port. This switch enables organizations to simplify IT infrastructures, improve system performance, maximize the value of virtual server deployments, and reduce overall storage costs.

For more information about Brocade 6505 SAN switches, visit <http://www.brocade.com/products/all/switches/product-details/6505-switch/index.page>.

3.4 SC7020 storage array

The SC7020 provides more capacity and greater flexibility as you manage a variety of workloads. Realize rack space savings with a 3U, 30-drive bay chassis that can expand up to 500 drives with 12 Gbps SAS expansion enclosures. It is powered by dual Intel Xeon E5 processors with eight cores each for enhanced application performance. The SC7020 multiprotocol connectivity options allow you to expand your front-end network connectivity and increase capacity with expansion enclosures.

Experience real savings with the Dell EMC low-cost-per-gigabyte flash drives (all-flash or a hybrid combination of drive types) to balance both your workload and performance requirements. With automatic policy-based deduplication and compression, true optimization of drive capacity is realized.

Simplify management through a single management console and control cost with minimized rack space and cabling complexities. This single array can reduce required rack space by over 60%¹.

Proven features, like the patented auto-tiering feature Data Progression, leverages diverse drive types to optimize data placement according to actual use, allowing you to achieve your performance and capacity needs at a lower cost. Thin-allocation methods are active by default, and RAID levels are dynamically provisioned, adjusting and rebalancing automatically to help you achieve your most aggressive application objectives with the fewest drives possible.

For more information about SC7020 storage arrays, visit <http://www.dell.com/us/business/p/storage-sc7000/pd>.

¹ Based on Dell EMC internal calculation using the rack space required for two SC9000 and two SC420 configuration compared to a single SC7020.

4 Storage configuration

4.1 Cabling

The hardware components were connected using Dell EMC best practices.

Port 1 on each HBA in the server and ports 1 and 3 on each HBA in the SC7020 were connected to the same Brocade 6505 switch. Port 2 on each HBA in the server and ports 2 and 4 on each HBA in the SC7020 were connected to the other Brocade 6505 switch.




Figure 3 Single server configuration cabling diagram

4.2 I/O ports

The front-end Fibre Channel (FC) ports were configured to use two fault domains in virtual port mode. Ports 1 and 3 from each controller were put into **Fault Domain 1** and ports 2 and 4 from each controller were put into **Fault Domain 2**.

All four SAS ports were configured as back end ports.

4.3 Disk folder

All 25 disks were put into the **Assigned** disk folder. Two disks were defined as hot spares, leaving 23 active disks. All volumes were created in the **Assigned** disk folder.

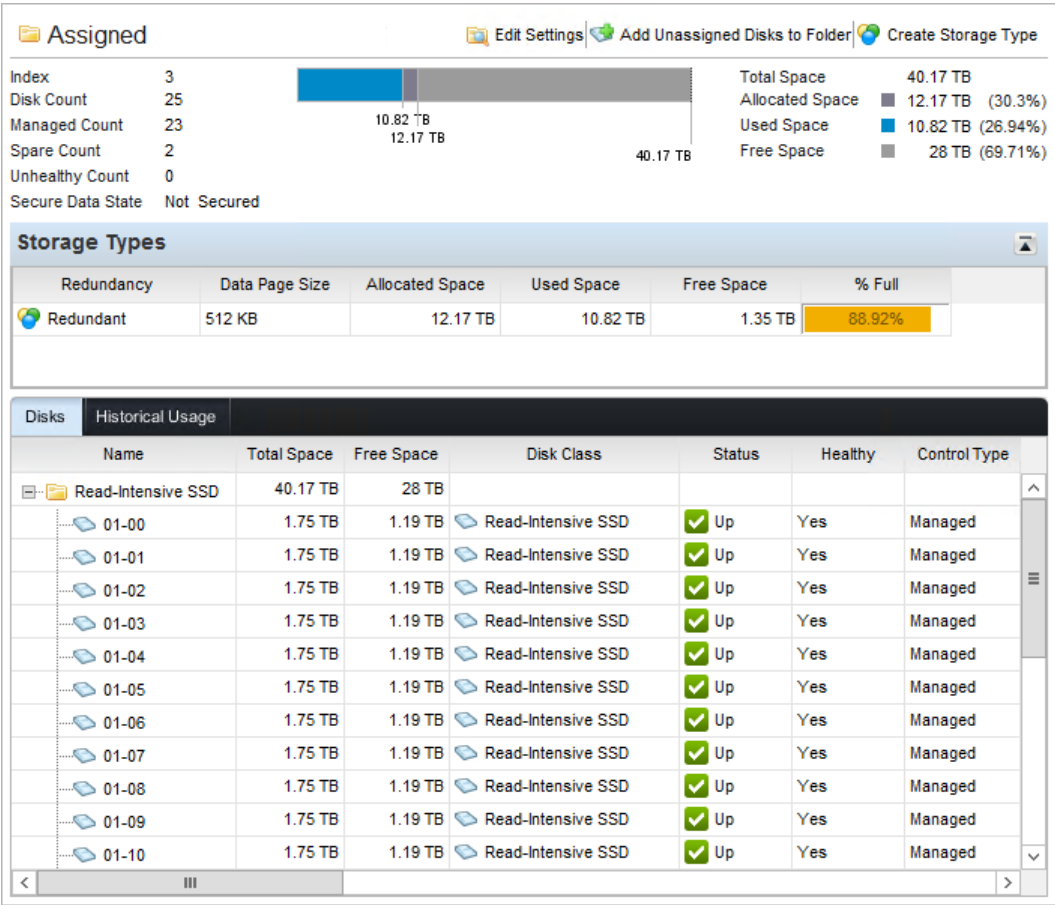


Figure 4 Reference architecture disk configuration

4.4 Storage type

The disk folder storage type defines the page size on the array and whether the disks are configured for redundancy (single or dual). The disk folder can be configured to use a 512KB, 2MB (default), or 4MB page.

The **Assigned** disk folder was configured for single redundancy, using a page size of 512KB.

4.5 SC Series storage profiles

The array uses storage profiles to define the RAID level and tiers where the data is stored for a given volume. By default, newly created volumes use the **Recommended (All Tiers)** storage profile. This storage profile uses RAID 10 for active (writeable) data and RAID 5 for snapshots (replays).

To increase capacity, the reference architecture uses RAID 5 for all data stored on the SQL Server data volumes. A custom storage profile named **RAID 5 All Tiers** was created. SQL Server data volumes on the storage array use this custom storage profile to ensure that RAID 5 is always used for those volumes. All other volumes use the **Recommended (All Tiers)** storage profile.

The RAID 5 stripe width is left at the default setting of **9 Wide (89% Efficient)** to maximize capacity on the array.

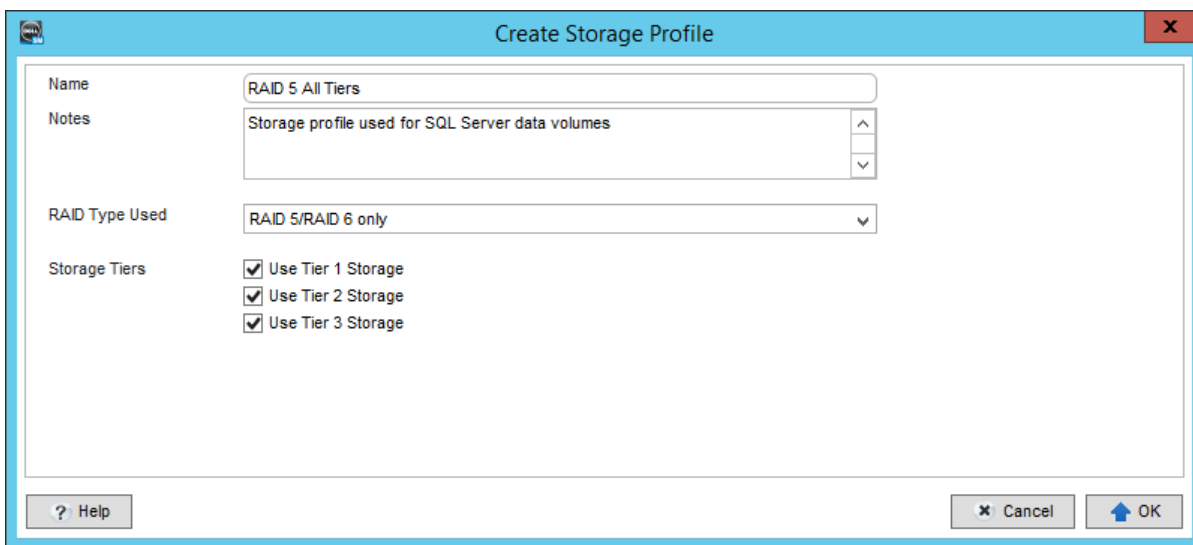


Figure 5 Creating the custom storage profile **RAID 5 All Tiers**

4.6 Read and write cache

The read cache is enabled on the array, which is the default read cache setting. Write cache is disabled on the array, which is a best practice for all flash SC Series arrays. The cache settings on each volume are left at the default values, which enable both read and write cache. Since the write cache is disabled on the array, the default setting enabling write cache on the volume will be overridden.

4.7 SC Series server object

The process of mapping a volume to a server object creates the I/O path (or paths) between a volume and a server. When creating a server object, the operating system of the server is defined. The SC Series array includes both single path and multi-path definitions for each version of Windows.

In this configuration, the server object on the array was created using the **Windows 2012 MPIO** operating system definition. Since the MPIO version of the Windows 2012 definition was chosen, each volume mapping will contain 16 paths.

Create Server [SC 46]

Name

DWFT-2S

Operating System

Windows 2012 MP10

Notes

Dell R730 - DWFT for SQL Server 2016

☒ Alert On Lost Connectivity
 ☒ Alert On Partial Connectivity

Select the HBA that the server uses to send IO requests to the Storage Center.

Host Bus Adapters

	Name	Port Type	Connectivity	ISCSI IPv4 Address
<input checked="" type="checkbox"/>	2001000E1E09D3BB	Fibre Channel	Up	
<input checked="" type="checkbox"/>	2001000E1EC2874F	Fibre Channel	Up	
<input checked="" type="checkbox"/>	2001000E1EC2874E	Fibre Channel	Up	
<input checked="" type="checkbox"/>	2001000E1EC20484	Fibre Channel	Up	
<input checked="" type="checkbox"/>	2001000E1EC20485	Fibre Channel	Up	
<input checked="" type="checkbox"/>	2001000E1E09B932	Fibre Channel	Up	
<input checked="" type="checkbox"/>	2001000E1E09B933	Fibre Channel	Up	
<input checked="" type="checkbox"/>	2001000E1E09D3BA	Fibre Channel	Up	

[Manually Add HBA](#)
☒ Only Show Up Connections

? Help

Cancel

OK

Figure 6 Creating the server object

4.8 SC Series volumes

The **Recommended (All Tiers)** storage profile is assigned to all volumes, with the exception of the SQL Server data volumes. The SQL Server data volumes use the **RAID 5 All Tiers** custom storage profile. This uses RAID 5 for the SQL Server data volumes and RAID 10 for all other volumes.

None of the volumes were configured to use snapshots. The reference architecture does not include hardware resources to support the use of snapshots.

The server is configured to boot from the array. When mapping the boot volume to the server, select **Map volume using LUN 0** in the advanced options. To simplify installation, create the boot volume, map it to the server and configure the HBAs before mapping any other volumes to the server.

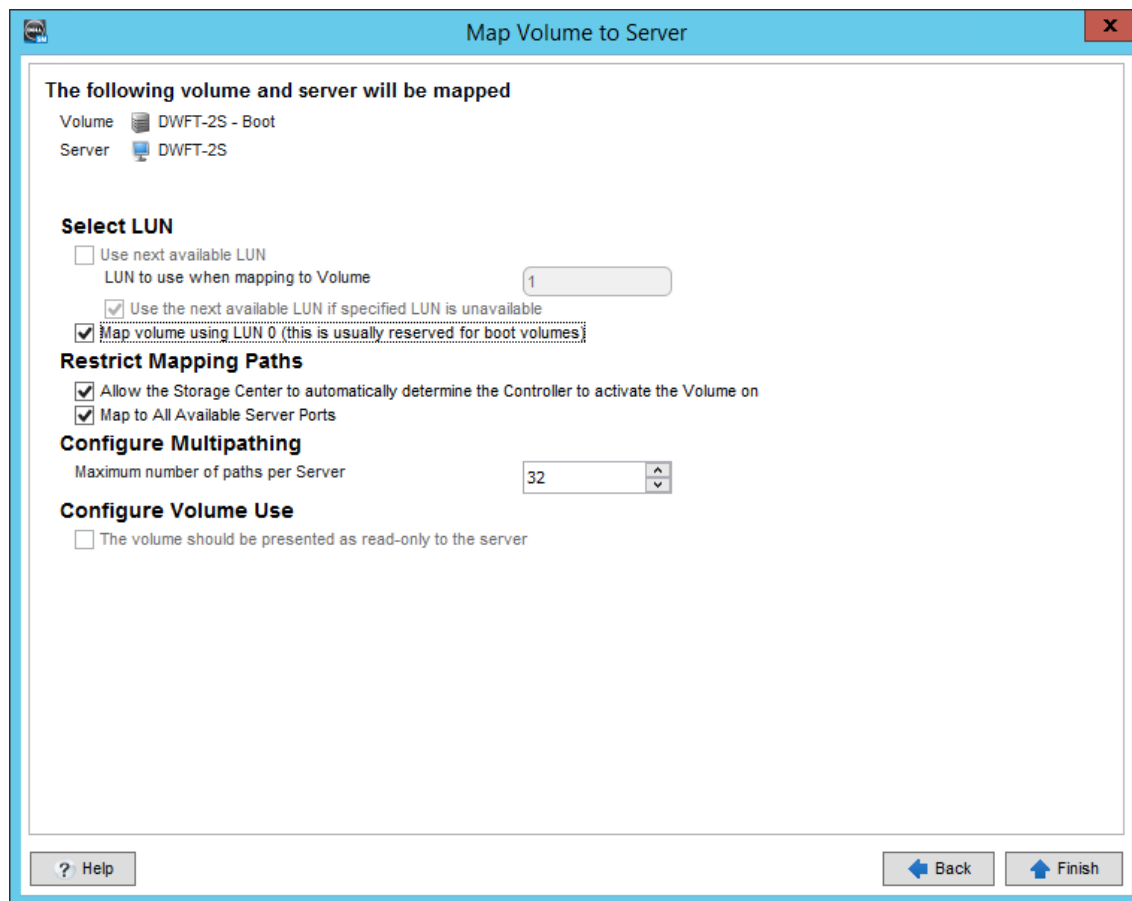


Figure 7 Mapping the boot volume with advanced options

Eight volumes (LUNs) were created to store the SQL Server data files for the data warehouse, with four volumes assigned to each controller. Two volumes were created to store the tempdb data files, with one volume assigned to each controller. If the volumes are created and mapped to the server, one right after the other, the array should automatically distribute the volumes evenly across the controllers. Otherwise, the controller can be manually selected using the advanced options when mapping the volume to the server. To achieve optimal results, the SQL Server data and tempdb volumes must be evenly distributed across the controllers.

Table 3 SC Series volumes created for the reference architecture

Volume name	Storage profile	Volume purpose
DWFT-2S - Boot	Recommended (All Tiers)	Windows boot volume
DWFT-2S - MPHost	Recommended (All Tiers)	Mount point host
DWFT-2S - SQLSystem	Recommended (All Tiers)	SQL Server system databases and files
DWFT-2S - SQLLog	Recommended (All Tiers)	Transaction log files for tempdb and the data warehouse
DWFT-2S - SQLData01	RAID 5 All Tiers	Data files for the data warehouse
DWFT-2S - SQLData02	RAID 5 All Tiers	Data files for the data warehouse
DWFT-2S - SQLData03	RAID 5 All Tiers	Data files for the data warehouse
DWFT-2S - SQLData04	RAID 5 All Tiers	Data files for the data warehouse
DWFT-2S - SQLData05	RAID 5 All Tiers	Data files for the data warehouse
DWFT-2S - SQLData06	RAID 5 All Tiers	Data files for the data warehouse
DWFT-2S - SQLData07	RAID 5 All Tiers	Data files for the data warehouse
DWFT-2S - SQLData08	RAID 5 All Tiers	Data files for the data warehouse
DWFT-2S - SQLTempdb01	Recommended (All Tiers)	Data files for tempdb
DWFT-2S - SQLTempdb02	Recommended (All Tiers)	Data files for tempdb

5 Server configuration

5.1 System BIOS

The system profile is set to **Performance**. All other options, outside of iDRAC configuration, were left at their factory default settings. The logical processor option, under processor settings, is left at its default setting of **Enabled**. This enables hyper-threading, which maximizes the number of logical processors available to SQL Server.

5.2 Host bus adapters

After each HBA port was reset to the factory defaults, several parameters were changed in accordance with Dell EMC best practices for SC Series storage arrays. Those changes are detailed in Table 4.

Table 4 HBA parameter changes

Parameter Menu	Parameter	Setting
Adapter Settings	Host Adapter BIOS	Enabled
	Connection Options	1 (Point to Point only)
Advanced Adapter Settings	Enable LIP Reset	Yes
	Login Retry Count	60
	Port Down Retry Count	60
	Link Down Timeout	30
Selectable Boot Settings (Each HBA port has two paths to the boot volume. The WWN for each path should be selected.)	Selectable Boot	Enabled
	Boot Port Name, LUN	WWN for the 1 st boot volume path
	Boot Port Name, LUN	WWN for the 2 nd boot volume path

6 Windows Server 2016 configuration

6.1 Power plan

To maximize performance, the server was configured to use the **High performance** power plan as shown in Figure 8.

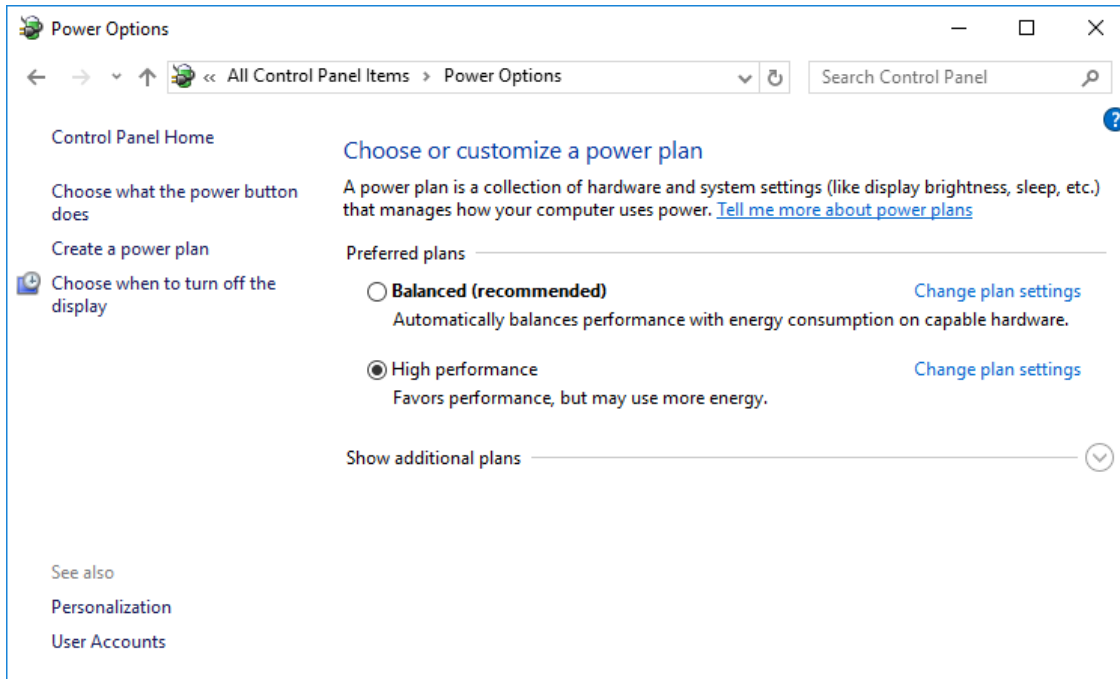


Figure 8 Windows power plan

6.2 Lock pages in memory

To prevent Windows from paging SQL Server memory to disk, the **Lock pages in memory** option was enabled for the SQL Server service account.

For information on enabling this option, visit <https://msdn.microsoft.com/en-IN/library/ms190730.aspx>.

6.3 Windows volumes

A single Windows volume was created on each SC Series volume. All volumes were formatted with the NTFS file system. The boot volume and the mount point host volume used the default allocation unit. All other volumes used an allocation unit of 64KB.

For DWFT reference architectures, Dell EMC recommends using mount points for the volumes instead of drive letters. It is highly recommended to assign appropriate volume and mount point names in order to simplify troubleshooting and performance analysis. Ideally, the mount point names should be assigned in a way that makes it easy to identify the SC Series volume for a given Windows volume.

Table 5 shows the volume labels and access paths used for the reference configuration.

Table 5 Windows volume details

Storage Center volume name	Windows volume label	Access path
DWFT-2S - Boot	Boot	C:\
DWFT-2S - MPHost	MPHost	M:\
DWFT-2S - SQLSystem	SQLSystem	M:\ft\SQLSystem
DWFT-2S - SQLLog	SQLLog	M:\ft\SQLLog
DWFT-2S - SQLData01	SQLData01	M:\ft\SQLData01
DWFT-2S - SQLData02	SQLData02	M:\ft\SQLData02
DWFT-2S - SQLData03	SQLData03	M:\ft\SQLData03
DWFT-2S - SQLData04	SQLData04	M:\ft\SQLData04
DWFT-2S - SQLData05	SQLData05	M:\ft\SQLData05
DWFT-2S - SQLData06	SQLData06	M:\ft\SQLData06
DWFT-2S - SQLData07	SQLData07	M:\ft\SQLData07
DWFT-2S - SQLData08	SQLData08	M:\ft\SQLData08
DWFT-2S - SQLTempdb01	SQLTempdb01	M:\ft\SQLTempdb01
DWFT-2S - SQLTempdb02	SQLTempdb02	M:\ft\SQLTempdb02

6.4 MPIO

MPIO was configured using Dell EMC best practices. MPIO best practices for the SC Series array are documented in *Dell Storage Center - Microsoft Multipath I/O Best Practices* on Dell TechCenter at http://en.community.dell.com/techcenter/extras/m/white_papers/20437917.

To improve throughput, each SQL Server data volume on the server is configured to use a single unique path between the array and the server. To ensure that only one path is used for each SQL Server data volume, the MPIO policy is set to **Fail Over Only** with one active path. The MPIO policy for all other volumes is set to **Round Robin**, allowing those volumes to use all paths.

Figure 9 shows the MPIO policy used for each SQL Server data volume. Figure 10 shows the MPIO policy for all other volumes.

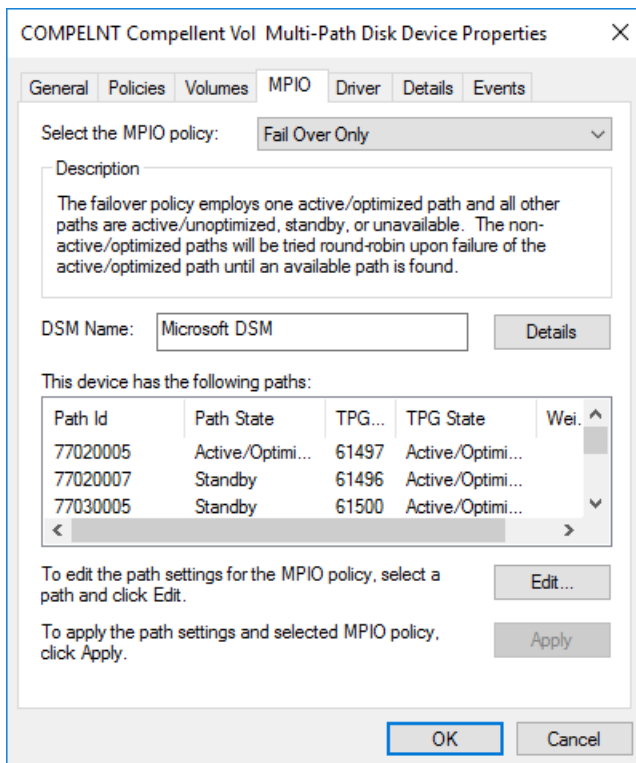


Figure 9 MPIO policy for data volumes

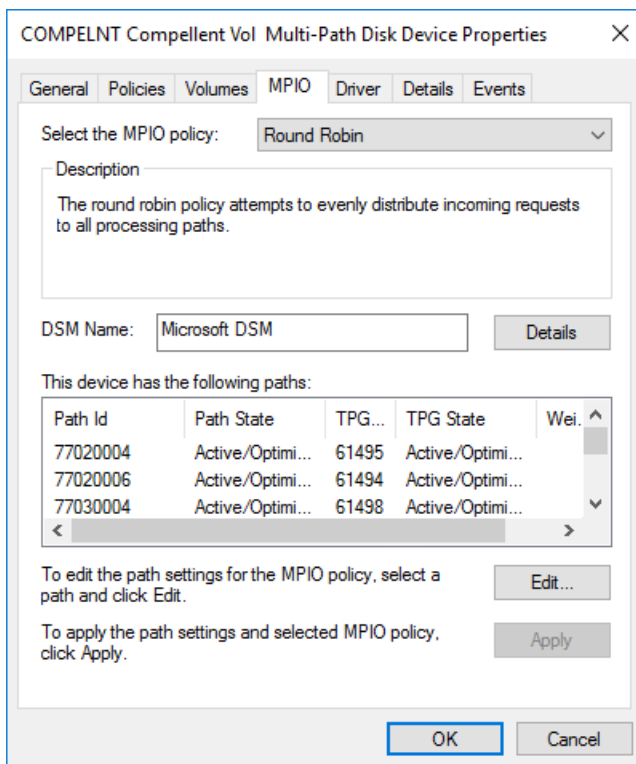


Figure 10 MPIO policy for other volumes

7 SQL Server 2016 Enterprise Edition configuration

7.1 Startup parameters

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

-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.

7.2 SQL Server maximum memory

SQL Server was configured to use 92% of the memory on the server. Maximum server memory for this reference architecture should be set to 1413 GB. If additional applications share the server, adjust the amount of memory left available to the operating system accordingly.

7.3 Max degree of parallelism (MAXDOP)

The max degree of parallelism was set to 36 for the row store tests and 72 for the column store tests.

For information on configuring the max degree of parallelism, visit <https://msdn.microsoft.com/en-us/library/ms189094.aspx>.

7.4 Resource governor

The resource governor was used to limit the maximum memory grant to 5 percent for the row store tests and 25 percent for the column store tests.

For information about the Resource Governor, visit <https://msdn.microsoft.com/en-us/library/bb933866.aspx>.

7.5 Database configuration

The data warehouse database was configured to use multiple file groups, each containing eight files distributed evenly across the eight data volumes. All files were allowed to grow automatically. The file groups were configured with the **AUTOGROW_ALL_FILES** option, to help ensure that all files within a given file group remain the same size.



7.6 Tempdb configuration

The tempdb database was configured to use eight data files of equal size. The data files were evenly distributed across the two tempdb data volumes, with four files stored on each volume. The tempdb transaction log file was placed on the log volume. All files were expanded to the appropriate size and auto grow was enabled.

8 Additional considerations for the Highly Available (HA) reference architecture

The HA reference architecture leverages Windows Failover Clustering to achieve high availability. When configuring a Windows failover cluster, there are additional storage considerations:

- The recommended quorum mode is **Node and Disk Majority**. An additional volume needs to be created and configured as the disk witness. Dell EMC recommends using a 2 GB volume for the disk witness. For more information on quorum modes and voting configurations in a failover cluster, visit <https://msdn.microsoft.com/en-us/library/hh270280.aspx>.
- All volumes need to be mapped to each node of the cluster. It is recommended to use a cluster server object in the SC Series array when mapping volumes to the cluster.
- All volumes need to be configured as a cluster resource and added to the SQL Server cluster resource group.

DWFT Certification #2016-010	Dell EMC R730 with SC7020 - 85TB DWFT Reference Architecture		Report Date: 11/03/2016																			
DWFT Rev. 5.4																						
System Provider	System Name	Processor Type	Memory																			
	Dell EMC PowerEdge R730	Intel Xeon E5-2697v4 2.3 GHz (2S/36C/72T)	1536 GB																			
Operating System		SQL Server Edition																				
Windows Server 2016		SQL Server 2016 Enterprise Edition																				
Storage Provider	Storage Information																					
	Dell EMC SC7020 25 x 1.92TB SSD for OS, data, log and tempdb RAID 10 for OS, log and tempdb, RAID 5-9 for data																					
<table><tr><th colspan="4">Primary Metrics</th></tr><tr><td>Rated User Data Capacity¹ (TB)</td><td>Row Store Relative Throughput²</td><td>Column Store Relative Throughput³</td><td>Maximum User Data Capacity¹ (TB)</td><td></td></tr><tr><td>85</td><td>275</td><td>419</td><td>95</td><td></td></tr></table>					Primary Metrics				Rated User Data Capacity ¹ (TB)	Row Store Relative Throughput ²	Column Store Relative Throughput ³	Maximum User Data Capacity ¹ (TB)		85	275	419	95					
Primary Metrics																						
Rated User Data Capacity ¹ (TB)	Row Store Relative Throughput ²	Column Store Relative Throughput ³	Maximum User Data Capacity ¹ (TB)																			
85	275	419	95																			
<table><tr><th colspan="6">Row Store</th></tr><tr><td>Relative Throughput²</td><td>Measured Throughput (Queries/Hr/TB)</td><td>Measured Scan Rate Physical (MB/Sec)</td><td>Measured Scan Rate Logical (MB/Sec)</td><td>Measured I/O Throughput (MB/Sec)</td><td>Measured CPU (Avg.) (%)</td></tr><tr><td>275</td><td>283</td><td>7,427</td><td>8,790</td><td>8,108</td><td>87</td></tr></table>					Row Store						Relative Throughput ²	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/O Throughput (MB/Sec)	Measured CPU (Avg.) (%)	275	283	7,427	8,790	8,108	87
Row Store																						
Relative Throughput ²	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/O Throughput (MB/Sec)	Measured CPU (Avg.) (%)																	
275	283	7,427	8,790	8,108	87																	
<table><tr><th colspan="6">Column Store</th></tr><tr><td>Relative Throughput²</td><td>Measured Throughput (Queries/Hr/TB)</td><td>Measured Scan Rate Physical (MB/Sec)</td><td>Measured Scan Rate Logical (MB/Sec)</td><td>Measured I/O Throughput (MB/Sec)</td><td>Measured CPU (Avg.) (%)</td></tr><tr><td>419</td><td>2,722</td><td>2,356</td><td>N/A</td><td>N/A</td><td>96</td></tr></table>					Column Store						Relative Throughput ²	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/O Throughput (MB/Sec)	Measured CPU (Avg.) (%)	419	2,722	2,356	N/A	N/A	96
Column Store																						
Relative Throughput ²	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/O Throughput (MB/Sec)	Measured CPU (Avg.) (%)																	
419	2,722	2,356	N/A	N/A	96																	

The reference configuration is a 2 socket system rated for 25TB using SQL Server 2014 and the DWFT V4 methodology

¹ Assumes a data compression ratio of 5:1

² Percent ratio of the throughput to the row store throughput of the reference configuration.

³ Percent ratio of the throughput to the column store throughput of the reference configuration.

* Reported metrics are based on the qualification configuration which specifies database size and SQL Server memory.

10 Summary

Dell EMC, in partnership with Microsoft, enables customers to deploy tested and validated data warehouse solutions using Data Warehouse Fast Track reference architectures for SQL Server 2016. These uniquely designed architectures ensure optimal Business Intelligence (BI) solutions. The end-to-end best practices and recommendations enable the customer to achieve enhanced return-on-investment (ROI) and faster time-to-value with a balanced, better performing data warehouse environment than traditional data warehouse systems.

Dell EMC DWFT reference architectures provide the following benefits to customers:

- Delivers a tested and validated configuration with proven methodology and performance behavior
- Delivers outstanding performance on the Dell EMC 13G server platform with blazing processor speeds and leading edge flash-based Dell EMC storage arrays
- Achieves a balanced and optimized solution at all the levels of the stack by following the best practices for both hardware and software components, achieving faster time-to-value and lower total-cost-of-ownership (TCO)
- Avoids over-provisioning of hardware resources
- Offers high availability at all levels of setup (host, switches, and storage)
- Offers single point of contact/accountability for purchases, services, and support; SQL Server is available to purchase from Dell EMC worldwide
- 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

This paper describes a reference architecture using an R730 server with an SC7020 storage array. By implementing Data Warehouse Fast Track for SQL Server 2016 design principles, this configuration achieved an 85TB rating.

A Technical support and resources

Dell.com/support is focused on meeting customer needs with proven services and support.

[Dell TechCenter](#) is an online technical community where IT professionals have access to numerous resources for Dell software, hardware, and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell EMC Storage platforms.

Additional resources:

- Dell EMC products: <http://www.dell.com>
- Dell SQL Server solutions: <http://www.dell.com/sql>
- Dell EMC Data Warehouse Fast Track for SQL Server 2014 Advisor:
http://www.dell.com/solutions/advisors/us/en/g_5/SQLFastTrack/5/Start?s=biz