

100TB Data Warehouse Fast Track Reference Architecture for Microsoft SQL Server 2017 using Dell EMC PowerEdge R740 and SCv3020

Configuration and performance results

Abstract

This paper describes the design principles and guidelines used to achieve an optimally balanced 100TB Data Warehouse Fast Track reference architecture for SQL Server 2017 using Dell EMC™ PowerEdge™ R740 servers and Dell EMC SCv3020 arrays.

December 2017

Revisions

Date	Description
December 2017	Initial release

Acknowledgements

Author: Doug Bernhardt

Special thanks to Jamie Reding 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.

© 2017 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

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
Table of contents	3
Executive summary.....	5
1 Microsoft SQL Server 2017: Data warehousing with improved column-store technology	6
1.1 Dell EMC Data Warehouse Fast Track reference architectures for SQL Server 2017	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 R740 server	11
3.2 QLogic QLE2662 16 Gbps host bus adapter	11
3.3 Brocade 6505 SAN switch.....	11
3.4 SCv3020 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.....	14
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 2017 Enterprise Edition configuration	22
7.1 Grant perform volume maintenance task privilege.....	22
7.2 SQL Server maximum memory	22

Table of contents

7.3	Max degree of parallelism (MAXDOP)	22
7.4	Resource governor	22
7.5	Database configuration	22
7.6	Tempdb configuration	22
8	Additional considerations for the Highly Available (HA) reference architecture	23
9	DWFT for SQL Server 2017 certification	24
10	Summary	25
A	Technical support and resources	26

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 Dell EMC PowerEdge™ 14th generation (14G) servers, along with robust and cutting-edge Dell EMC SCv3020 storage arrays, form efficient candidates for a high performing data warehouse solution.

This paper describes the design principles and guidelines used to achieve an optimally balanced 100TB Data Warehouse Fast Track (DWFT) reference architecture for SQL Server 2017 using PowerEdge R740 servers and SCv3020 arrays. The configuration used to achieve the performance numbers for the reference configuration is presented in detail.

The target audience for this paper includes database administrators, business intelligence architects, storage administrators, IT directors, and data warehousing users seeking sizing and design guidance for business intelligence solutions with SQL Server 2017.

1 Microsoft SQL Server 2017: Data warehousing with improved column-store technology

Microsoft SQL Server 2017 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 2017 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, Microsoft has 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 that needs to be handled further on in query processing.

Batch-mode processing: SQL Server 2017 includes enhancements in batch-mode processing that handles 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 2017

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 14G servers with Intel® Xeon® processors, Dell EMC SCv3020 arrays, and Microsoft SQL Server 2017 are the latest additions to the list of reference architecture components.

The Dell EMC DWFT reference architectures for SQL Server 2017 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 fast 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 2017, comprised of PowerEdge R740 servers and SCv3020 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.

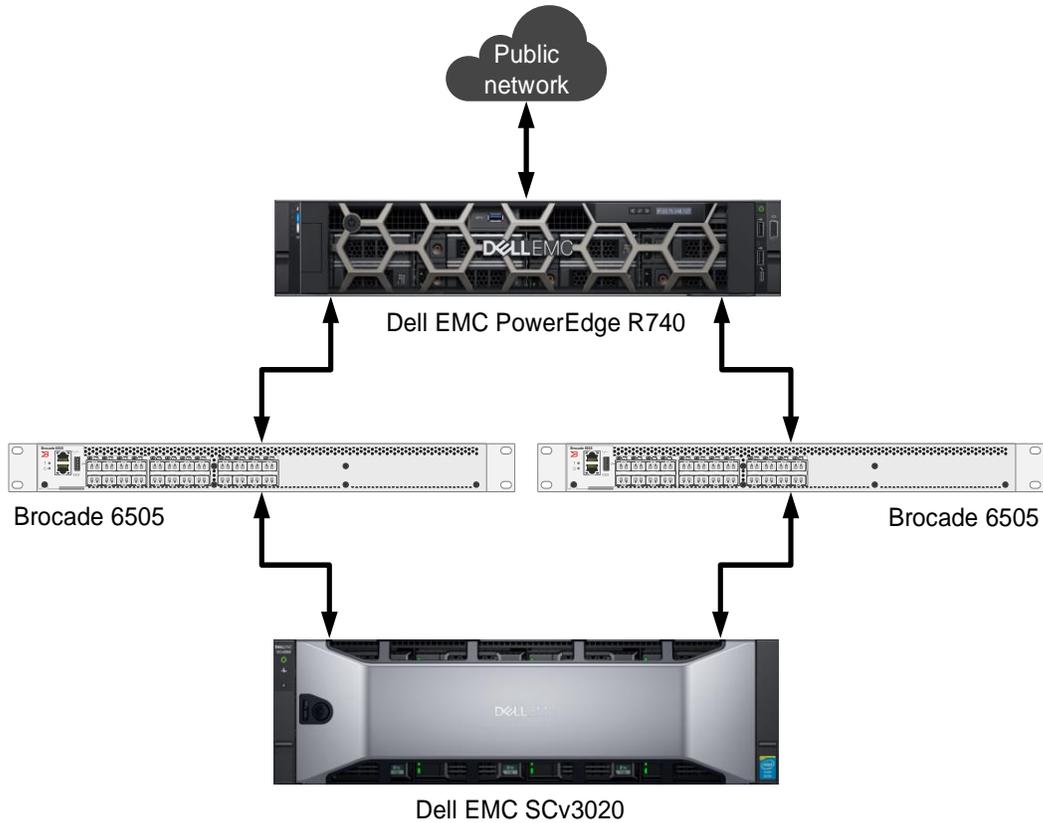


Figure 1 Single server reference architecture

Table 1 Single server reference architecture details

Component	Description	
Server	Dell EMC PowerEdge R740	
	Processors	Two Intel® Xeon® Platinum 8168 Processors (2.7GHz, 24 cores, 48 threads)
	Total cores	48
	Total logical processors	96 (Intel Hyper-Threading Technology enabled)
	Total memory	1152GB
	Network adapters	Minimum of one network adapter (1Gbps or 10Gbps 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 HBAs
Software	Operating system	Windows Server® 2016 Enterprise Edition
	Database software	SQL Server 2017 Enterprise Edition
Storage	Array	Dell EMC SCv3020 (SCOS v7.2)
	I/O cards	Two QLogic QLE2694 quad-port 16 Gbps FC HBAs (one per controller)
	Disk drives	30 x 1.92TB read-intensive SSDs (2.5" SAS)
	SAN switches	Two Brocade® 6505 with 16Gbps SFPs

2.2 Highly available reference architecture

For database high availability, 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.

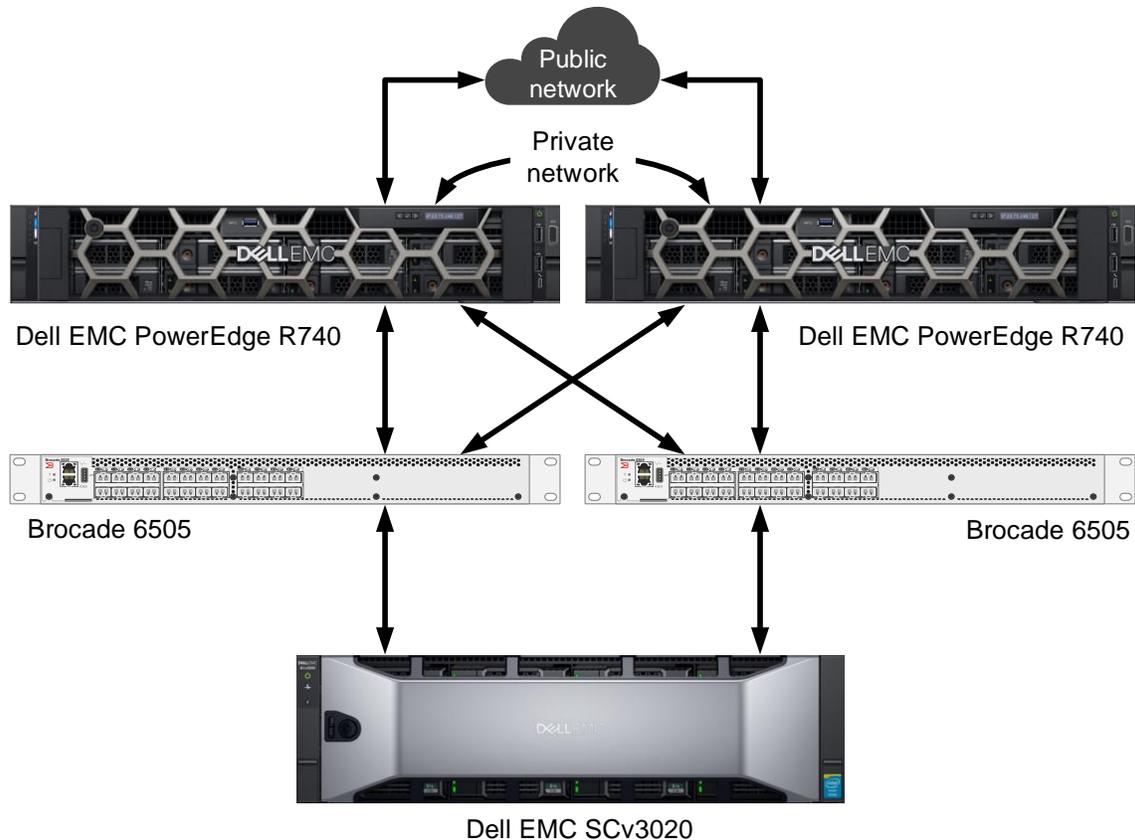


Figure 2 Highly available reference architecture

Table 2 Highly available reference architecture details

Component	Description	
Server	Two Dell EMC PowerEdge R740	
	Processors	Two Intel® Xeon® Platinum 8168 Processors (2.7GHz, 24 cores, 48 threads) per server
	Total cores	48 per server
	Total logical processors	96 per server (Hyper-Threading Technology enabled)
	Total memory	1152GB per server
	Network adapters	Public network: Minimum of two network adapters (1Gbps or 10Gbps based requirements) per server with load balancing configured Private network (cluster): Minimum of one network adapter with 1Gbps (or more) connectivity per server
	Host bus adapters	Four QLogic QLE2662 dual-port 16Gbps FC HBAs per server
Software	Operating system	Windows Server 2016 Enterprise Edition with Windows Failover Clustering
	Database software	SQL Server 2017 Enterprise Edition configured as a Failover Cluster Instance
Storage	Array	Dell EMC SCv3020 (SCOS v7.2)
	I/O cards	Two QLogic QLE2694 quad-port 16 Gbps FC HBAs (one per controller)
	Disk drives	30 x 1.92TB read-intensive SSDs (2.5" SAS)
	SAN switches	Two Brocade 6505 with 16Gbps SFPs

3 Hardware components

3.1 PowerEdge R740 server

The PowerEdge R740 server 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 R740 server offers simplified management, purposeful design, and energy efficiency with support for the latest Intel® Xeon® series processors as well as NVDIMM memory. The R740 server supports up to eight Gen3 PCIe expansion cards.

For more information, see the [PowerEdge R740 server product page](#).

3.2 QLogic QLE2662 16 Gbps host bus adapter

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

For more information, see the QLogic FAQ document, [QLogic 2600 Series 16Gb Gen 5 FC HBAs for Dell PowerEdge Servers](#).

3.3 Brocade 6505 SAN switch

The Brocade 6505 switch is a 1U, 24-port, rack-mountable Fibre Channel switch providing up to 16Gbps 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, see the [Brocade 6505 switch product page](#).

3.4 SCv3020 array

The SCv3020 array makes storage cost savings automatic with a modern architecture that optimizes the data center for economics while delivering transformational SSD, HDD, or hybrid performance.

SC Series storage provides the lowest effective cost per GB for flash and hybrid flash¹, giving organizations of any size the technology advantage needed to compete in the current fast-changing markets. Highlights include:

- Data Progression: Achieve IOPS goals with the least expensive mix of storage media, even as your performance needs evolve.
- Deduplication and compression: Dramatically reduce the raw capacity required to store data.
- RAID tiering: Eliminate manual RAID provisioning, and increase efficiency and utilization.
- Federation: Simplify multi-array environments with quick and seamless data movement, plus proactive load balancing assistance using Live Migrate and Volume Advisor.

¹Net usable capacity of Dell array with 5 years of support, after 4:1 data reduction, vs. major competitors net of data reduction. Street price analysis is based on a variety of sources including analyst data, price sheets when available, and public information as of January 2017.

- Dell ProSupport™ services: Reduce deployment costs with remote installation options that ensure the project is successful the first time.
- Persistent software licensing: Future-proof the investment, and minimize the cost of upgrades and expansions.

Designed as the next-generation successor to the popular Dell Storage SCv2000 Series arrays, the SCv3020 array is a performance powerhouse. With dual 6-core Intel® Xeon® Processors, double the memory, and a 12Gb SAS back end, the SCv3020 delivers:

- Up to 50% more IOPs²
- Up to 3x more bandwidth²
- 50% more capacity
- Double the number of snapshots

The 3U all-in-one chassis includes 30 drive bays plus dual hot-swappable controllers, providing up to 460TB of raw capacity in a single compact unit. A variety of expansion enclosures enables scaling up to 222 drives or 1 petabyte (1PB) per array — with even larger scale-out potential in federated multi-array systems. In addition to fast hardware, the SCv3020 includes all of the SCOS features to be expected from SC Series storage.

For more information see the [SCv3000 Series product page](#).

²Based on June 2017 internal Dell EMC testing, compared to previous-generation SCv2000 Series. Actual performance will vary depending upon application and configuration.

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 2 on each HBA in the SCv3020 were connected to the same Brocade 6505 switch. Port 2 on each HBA in the server and ports 3 and 4 on each HBA in the SCv3020 were connected to the other Brocade 6505 switch.

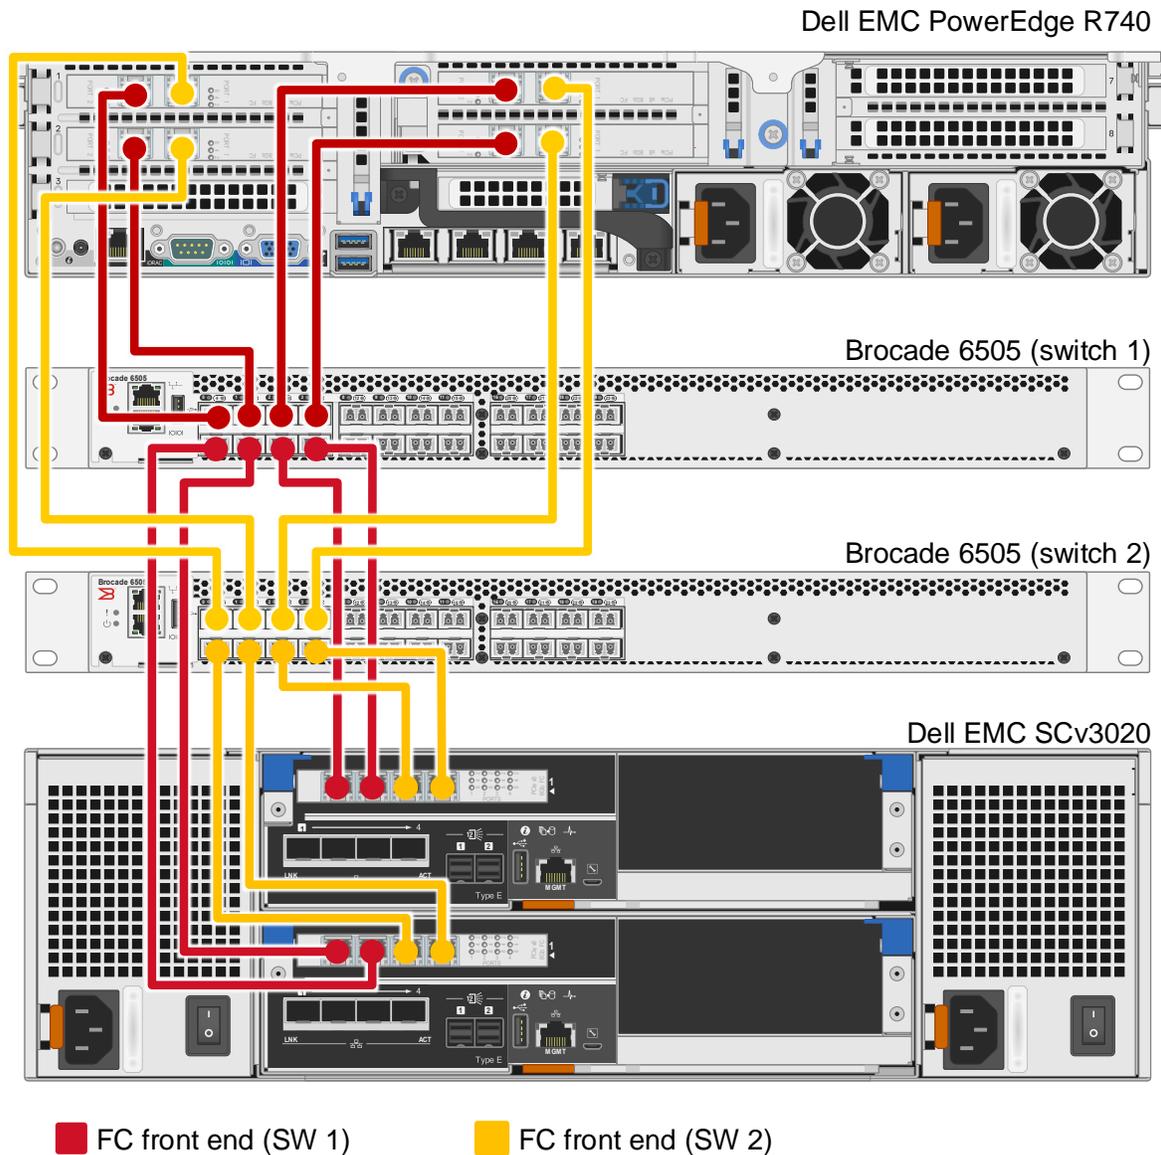


Figure 3 Single server configuration cabling diagram

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 2 from each controller were added to **Fault Domain 1** and ports 3 and 4 from each controller were added to **Fault Domain 2**.

4.3 Disk folder

All seven disks were added to the **Assigned** disk folder. One disk was defined as a hot spare, leaving six active disks. All volumes were created in the **Assigned** disk folder.

4.4 Storage type

The disk folder storage type defines the page size on the array and how 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 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.

Create Storage Profile

General

Name: RAID 5 All Tiers

Notes:

Data Writes

Select the write tier and RAID type to be used for writing data.

Write Tier: Tier 1

Write RAID Type: Parity (RAID 5 / RAID 6)

Snapshot Data

Select the RAID type to use for snapshot data in each tier.

Tier 1 Snapshot RAID Type: Parity (RAID 5 / RAID 6)

Tier 2 Snapshot RAID Type: Parity (RAID 5 / RAID 6)

Tier 3 Snapshot RAID Type: Parity (RAID 5 / RAID 6)

Data will be written to Tier 1 Parity (RAID 5 / RAID 6). Snapshots will remain in Tier 1 Parity (RAID 5 / RAID 6) when they are taken. Data Progression will move data amongst all selected tiers based upon frequency of access.

? Help × Cancel ↑ OK

Figure 4 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 Server 2016 MPIO** operating system definition. Since the MPIO version of the Windows 2016 definition was chosen, each volume mapping will contain eight paths.

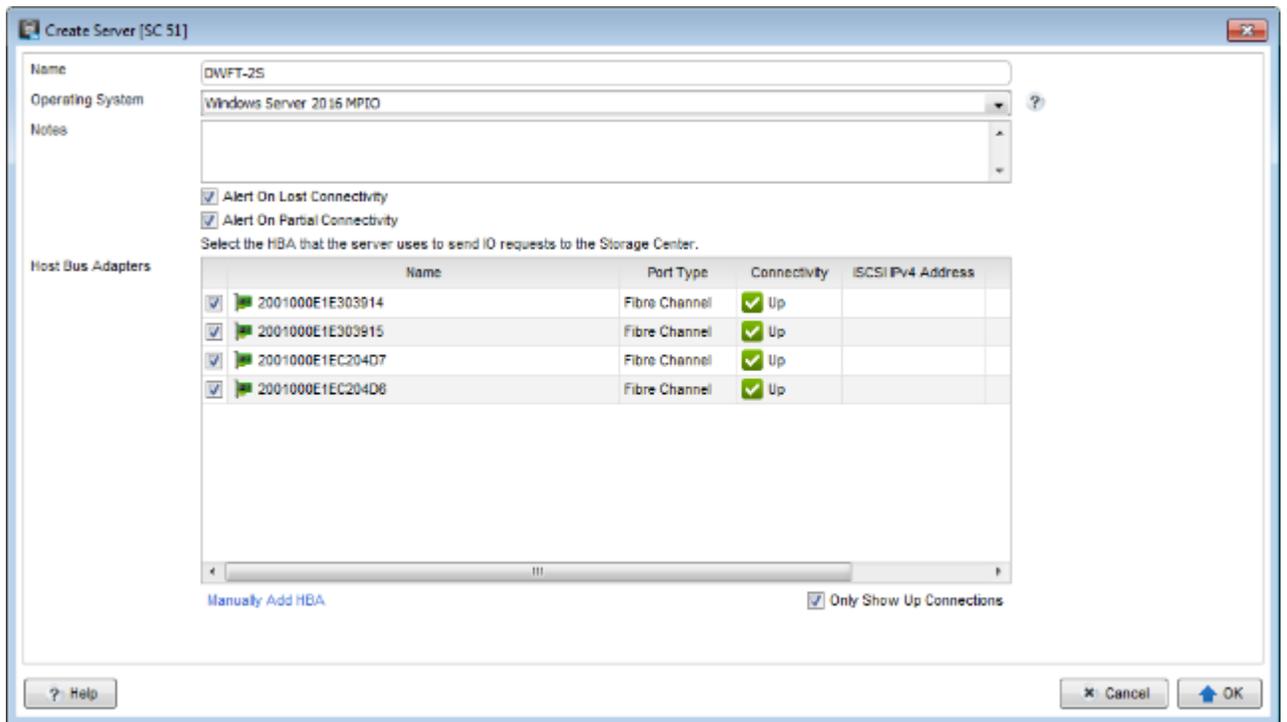


Figure 5 Creating the server object

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.

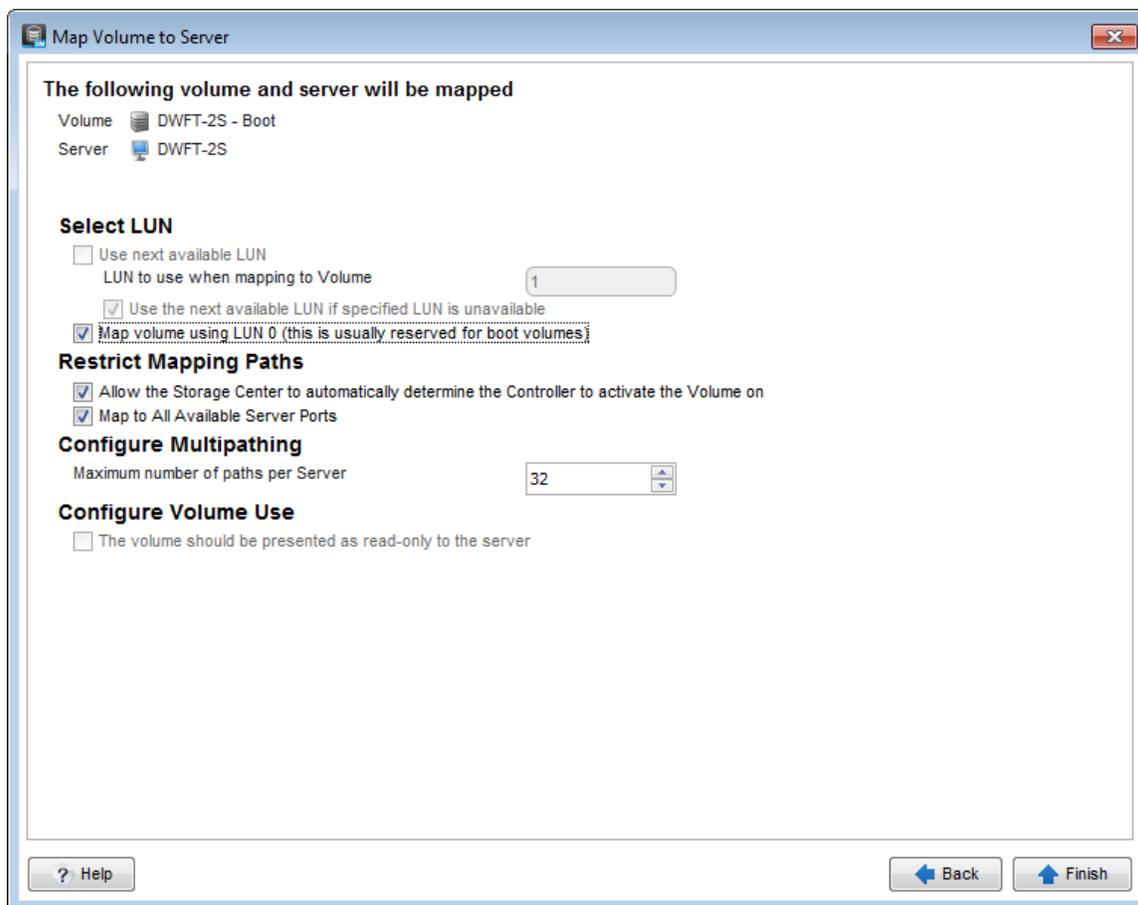


Figure 6 Mapping the boot volume with advanced options

Four volumes (LUNs) were created to store the SQL Server data files for the data warehouse, with two 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 - 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 Intel® Hyper-Threading Technology, 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 first boot volume path
	Boot Port Name, LUN	WWN for the second 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 7.

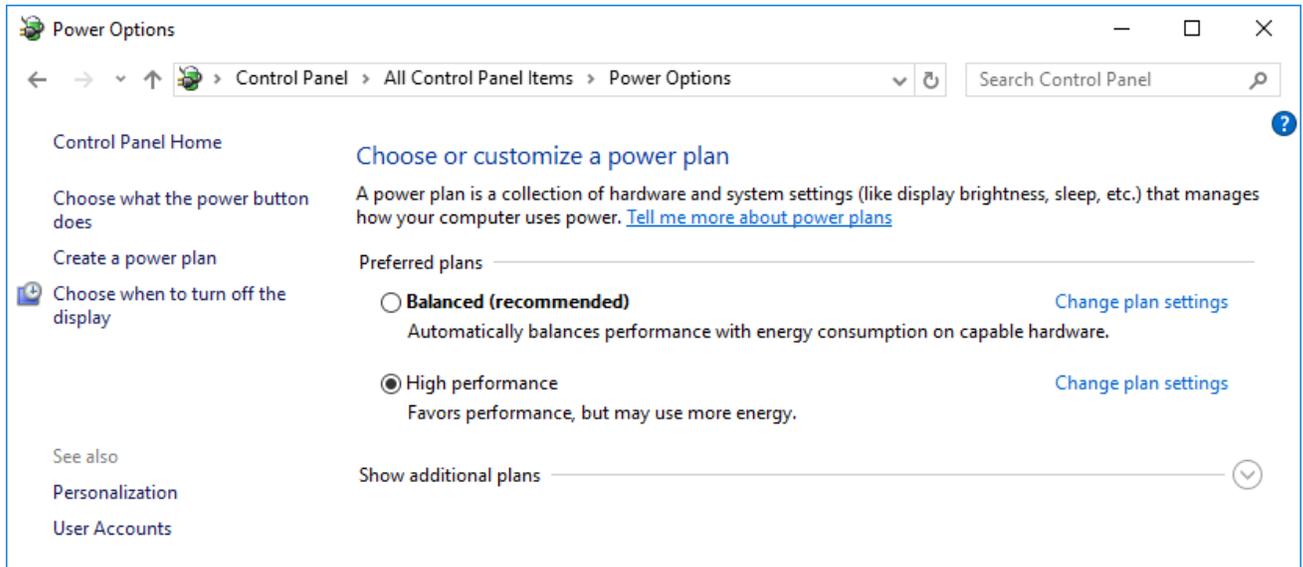


Figure 7 Windows power plan

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 - 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 the paper, [Dell EMC SC Series Storage: Microsoft Multipath I/O Best Practices](#).

The MPIO policy for all volumes is left at the default setting of **Round Robin**, allowing volumes to use all paths between the server and the storage array.

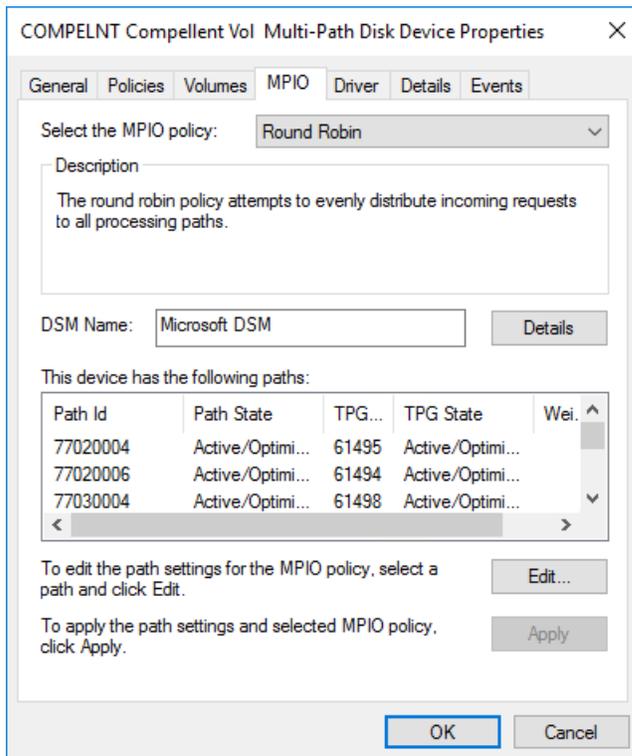


Figure 8 MPIO policy for all volumes

7 SQL Server 2017 Enterprise Edition configuration

7.1 Grant perform volume maintenance task privilege

During installation of SQL Server 2017, the option to grant the **SQL Server Database Engine Service** the **Perform Volume Maintenance Task** privilege was selected.

7.2 SQL Server maximum memory

The maximum server memory for this reference architecture should be set to 1120GB which leaves 32GB for the operating system. 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 **40**.

For more information, see the Microsoft article, [Configure the max degree of parallelism Server Configuration Option](#).

7.4 Resource governor

The resource governor was used to limit the maximum memory grant to 12 percent.

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 four files distributed evenly across the four 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 configuration is to allow all cluster nodes to have quorum votes and use a disk witness. An additional volume needs to be created and configured as the disk witness. Dell EMC recommends using a 2GB volume for the disk witness. For more information on quorum and voting configurations in a failover cluster, see the Microsoft article, [Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster](#).
- 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.

9 DWFT for SQL Server 2017 certification

DWFT Certification #2017-005	Dell PowerEdge R740 DWFT Reference Architecture		Report Date: 11/29/2017		
DWFT Rev. 5.4					
System Provider	System Name	Processor Type		Memory	
	Dell PowerEdge R740	Intel Xeon Platinum 8168 2.7 GHz (2S/48C/48T)		1152 GB	
Operating System			SQL Server Edition		
Windows Server 2016			SQL Server 2017 Enterprise Edition		
Storage Provider	Storage Information				
	Dell EMC SC3020 28 x 1.92TB SSD for OS, data, log and tempdb RAID 10 for OS, log and tempdb, RAID 5-5 for data				
Primary Metrics					
Rated User Data Capacity ¹ (TB)	Row Store Relative Throughput ²	Column Store Relative Throughput ³	Maximum User Data Capacity ¹ (TB)		
100	340	679	174		
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.) (%)
340	429	8,760	11,256	10,008	79
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.) (%)
679	4,628	4,049	N/A	N/A	96
<p>The reference configuration is a 2 socket system rated for 25TB using SQL Server 2014 and the DWFT V4 methodology</p> <p>¹ Assumes a data compression ratio of 5:1</p> <p>² Percent ratio of the throughput to the row store throughput of the reference configuration.</p> <p>³ Percent ratio of the throughput to the column store throughput of the reference configuration.</p> <p>* Reported metrics are based on the qualification configuration which specifies database size and SQL Server memory.</p>					

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 2017. These uniquely designed architectures ensure optimal business intelligence solutions. The end-to-end best practices and recommendations enable the customer to achieve enhanced return on investment and faster time to value with a balanced data warehouse environment that can perform better than traditional data warehouse systems.

Dell EMC DWFT reference architectures provide the following benefits:

- Deliver a tested and validated configuration with proven methodology and performance behavior
- Deliver outstanding performance on the Dell EMC PowerEdge 14G server platform with blazing processor speeds and leading-edge, flash-based Dell EMC storage arrays
- Achieve a balanced and optimized solution at all levels of the stack by following best practices for both hardware and software components, achieving faster time to value, and lower total cost of ownership
- Avoid over provisioning of hardware resources
- Offer high availability at all levels of setup (host, switches, and storage)
- Offer a single point of contact and accountability for purchases, services, and support; SQL Server is available to purchase from Dell EMC worldwide
- Help customers avoid the pitfalls of an improperly designed and configured system
- Reduce future support costs by limiting solution rearchitect efforts due to scalability challenges

This paper describes a reference architecture using a PowerEdge R740 server with an SCv3020 storage array. By implementing Data Warehouse Fast Track for SQL Server 2017 design principles, this configuration achieved a 100TB 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