



80TB Data Warehouse Fast Track Reference Architecture for Microsoft SQL Server 2014 using PowerEdge R920 and Dell Storage SC9000

Dell configuration and performance results

Dell Storage Engineering
December 2015

Revisions

Date	Description
December 2015	Initial release

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Copyright © 2015 Dell Inc. All rights reserved. Dell and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.



Table of Contents

1	Introduction to Data Warehouse Fast Track reference architectures for SQL Server 2014	6
1.1	Dell Data Warehouse Fast Track reference architectures for SQL Server 2014	6
2	Recommended reference architectures	8
2.1	Single server reference architecture.....	8
2.2	HA reference architecture	10
3	Hardware components.....	12
3.1	Dell PowerEdge R920 server	12
3.2	QLogic QLE2662 16 Gbps host bus adapter.....	12
3.3	Brocade 6505 SAN switch	12
3.4	Dell Storage SC9000 controller	13
3.5	Dell Storage SC420 disk enclosure.....	13
4	Storage configuration	14
4.1	Back-end cabling	14
4.2	Front-end cabling	15
4.3	Local ports.....	16
4.4	Disk folder	17
4.5	Storage type	18
4.6	Storage profile	18
4.7	Storage Center server object	19
4.8	Storage Center volumes	20
5	Server configuration.....	23
5.1	System BIOS.....	23
5.2	Host bus adapters	23
6	Windows Server 2012 R2 configuration	24
6.1	Power plan	24
6.2	Lock pages in memory.....	24
6.3	Windows volumes	24
6.4	MPIO.....	25
7	SQL Server 2014 Enterprise Edition configuration	27
7.1	Startup parameters	27
7.2	SQL Server maximum memory	27



7.3	Max degree of parallelism (MAXDOP)	27
7.4	Resource Governor	27
7.5	Tempdb configuration	27
8	Additional considerations for the HA reference architecture.....	28
9	DWFT for SQL Server 2014 certification.....	29
10	Summary	30
A	Additional Resources.....	31



Acknowledgements

Dell Storage Applications Engineering: Mike Matthews

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

Executive Summary

Dell™ 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. High performance 4-socket Dell servers, along with robust and cutting edge Dell Storage SC9000 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 80 TB Data Warehouse Fast Track reference architecture for SQL Server 2014 using the Dell PowerEdge R920 server and Dell Storage SC9000 storage array. 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 enterprise, business intelligence solutions with Microsoft SQL Server 2014.



1 Introduction to Data Warehouse Fast Track reference architectures for SQL Server 2014

Data Warehouse Fast Track (DWFT) reference architectures provide tested and validated configurations and associated resources to help customers identify the right environment for their data warehouse solutions. Tangible benefits of following these recommended configuration best practices and guidelines are:

- Accelerated data warehouse projects with pretested hardware and SQL Server configurations
- Reduced hardware and maintenance costs by purchasing a balanced hardware solution and optimizing it for a data warehouse workload
- Reduced planning and setup costs by leveraging the certified reference architecture configurations
- Ensured predictable performance by configuring the system correctly and taking advantage of the tuning directions

The DWFT reference architectures provide a balanced server, memory, network, and storage hardware configuration and avoid the risk of improperly designed and configured hardware systems. These guidelines ensure that the capability and throughput for the entire system is maximally utilized. Configurations are designed specifically for the data warehouse or Business Intelligence (BI) systems.

1.1 Dell Data Warehouse Fast Track reference architectures for SQL Server 2014

Dell and Microsoft have refreshed the DWFT reference architecture offerings with the latest technology advancements in database, server, and storage technology. Dell PowerEdge 4-socket servers, Dell Storage SC9000 arrays, and Microsoft SQL Server 2014 are the latest additions to the list of reference architecture components.

The Dell 4-socket server platforms, featuring enhanced onboard memory, storage, and processor speeds, have advanced features that boost data warehouse performance. The latest Intel® Xeon® E7 series processors, larger memory capacities, higher memory speed, and third generation PCI Express slots on the newer PowerEdge platforms ensure faster database throughput.

The SC9000 storage array builds on the success of the Dell SC8000 platform, offering higher levels of performance. The SC9000 delivers a rock-solid enterprise hardware architecture providing enterprise-class performance and reliability, intelligent automation, and simplified storage management. Dell Storage Center offers exceptional performance for both sequential and transactional applications with the ability to meet stringent requirements.

Microsoft SQL Server 2014 Enterprise Edition comes with several exciting features that directly benefit data warehouse environments. One of these features is columnstore indexes. Introduced in SQL Server 2012, columnstore indexes enable storing data in columnar fashion, providing better query performance and better compression rates. This is very beneficial for data warehouse environments as they typically handle



large amounts of data. In SQL Server 2014, columnstore indexes were enhanced to allow clustering and updating, making them even easier to incorporate. For more information on columnstore indexes, visit msdn.microsoft.com/en-us/library/gg492088.aspx. Microsoft DWFT for SQL Server 2014 guidelines incorporate the benefit of columnstore indexes for improved query performance.

The Dell DWFT reference architectures for SQL Server 2014 are engineered jointly by Dell and Microsoft. The hardware and software optimizations are tested by Dell 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 2014, comprised of PowerEdge R920 servers and Dell Storage SC9000 storage arrays.

2.1 Single server reference architecture

Figure 1 illustrates the single server reference architecture with the major elements and 0 lists the component details.




Figure 1 Single server reference architecture

Table 1 Single server reference architecture details

Component	Description	
Server	PowerEdge R920	
	Processors	4 x Intel Xeon E7-8891v2 (3.2Ghz, 10 cores, 20 threads)
	Total cores	40
	Total logical processors	80 (Hyper-Threading enabled)
	Total memory	1536 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	4 x QLogic® QLE2662 dual port 16 Gbps FC HBA
Software	Operating system	Windows Server® 2012 R2 Standard Edition
	Database software	SQL Server 2014 Enterprise Edition
Storage	Controllers	2 x Dell Storage SC9000 (SCOS v6.7)
	I/O cards per controller	2 x QLogic QLE2662 dual port 16 Gbps FC HBA (front-end)
		2 x LSI SAS9300-8E dual port 12 Gbps SAS (back-end)
	Disk enclosures	1 x Dell Storage SC420
	Disk drives	24 x 1.92 TB read intensive SSDs (2.5" 12 Gbps SAS)
	SAN switches	2 x Brocade® 6505 with 16 Gbps SFPs



2.2 HA reference architecture

To achieve high availability for the database, Microsoft 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 in Section 2.1).

Figure 2 illustrates the HA reference architecture with the major elements and Table 2 lists the component details.

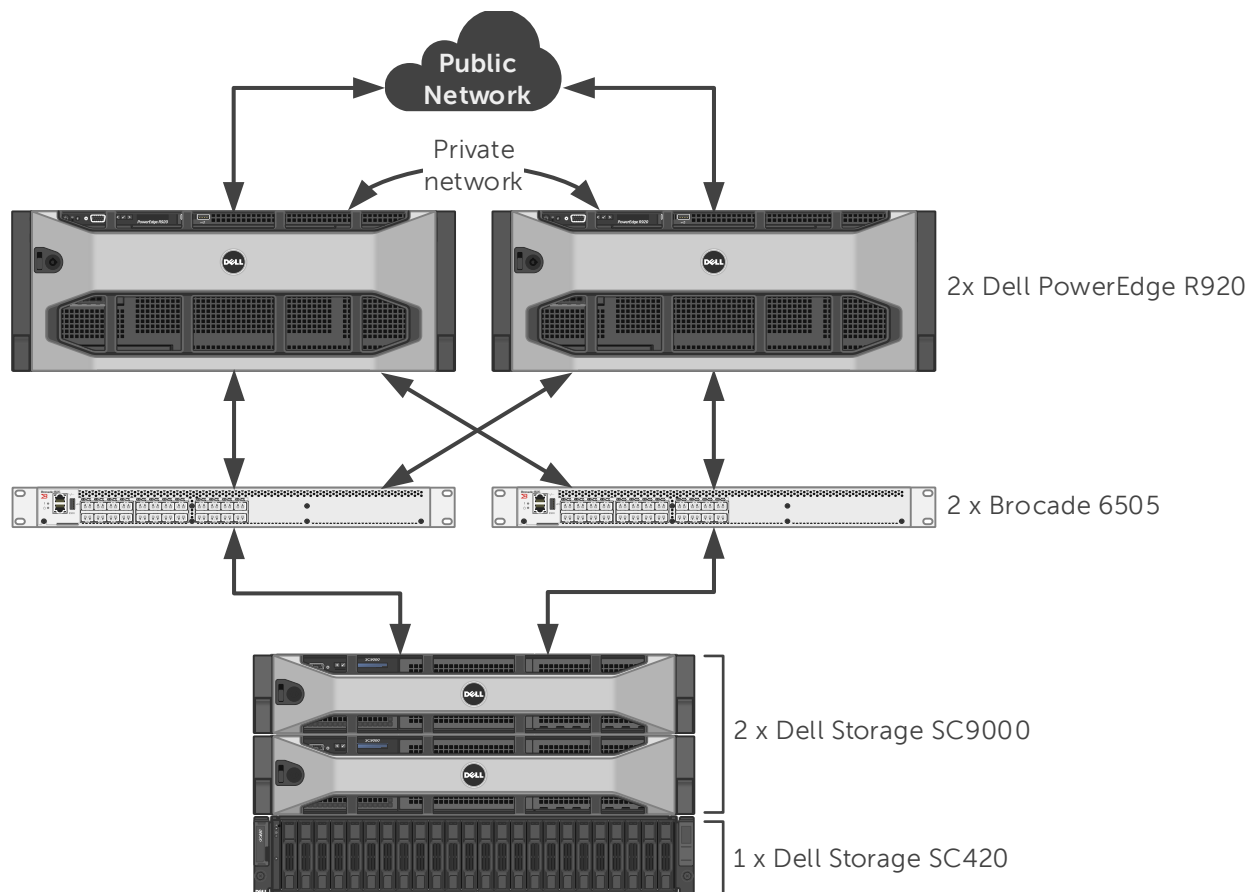


Figure 2 HA reference architecture

Table 2 HA reference architecture details

Component	Description	
Server	2 x PowerEdge R920	
	Processors	4 x Intel Xeon E7-8891v2 (3.2Ghz, 10 cores, 20 threads) per server
	Total cores	40 per server
	Total logical processors	80 per server (Hyper-Threading enabled)
	Total memory	1536 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	4 x QLogic QLE2662 dual port 16 Gbps FC HBA per server
Software	Operating system	Windows Server 2012 R2 Standard Edition with Windows Failover Clustering
	Database software	SQL Server 2014 Enterprise Edition configured as a Failover Cluster Instance
Storage	Controllers	2 x Dell Storage SC9000 (SCOS v6.7)
	I/O cards per controller	2 x QLogic QLE2662 dual port 16 Gbps FC HBA (front-end)
		2 x LSI SAS9300-8E dual port 12 Gbps SAS (back-end)
	Disk enclosures	1 x Dell Storage SC420
	Disk drives	24 x 1.92 TB read intensive SSDs (2.5" 12 Gbps SAS)
	SAN switches	2 x Brocade 6505 with 16 Gbps SFPs



3 Hardware components

3.1 Dell PowerEdge R920 server

The PowerEdge R920 enables you to boost business-application effectiveness, deliver exceptional database performance and efficiently consolidate workloads with large-scale virtualization. The 4-socket, 4U server, with built-in expandability, is designed to provide expanding scalable performance for medium and large enterprises. R920 supports up to:

- 96 DIMMS of memory
- 10 PCIe Gen3/Gen2 slots
- 24 internal 12Gbps SAS drives
- 8 PCIe flash drives (with the optional PCIe backplane)

The Dell PowerEdge R920 is powered by the Intel Xeon processor E7 v2 product family. It is ideal for mission-critical applications, and high performing workloads. PowerEdge R920 supports up to 30 physical cores in a half-loaded configuration and up to 60 physical cores in a full-loaded configuration. With up to 60 physical cores, the PowerEdge R920 server provides abundant computing power for compute-intensive data warehouse solutions and deployments.

For more information on Dell PowerEdge R920, visit
<http://www.dell.com/us/business/p/poweredge-r920/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 on 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 FC 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 on Brocade 6505 SAN switches, visit:
<http://www.brocade.com/products/all/switches/product-details/6505-switch/index.page>.



3.4 Dell Storage SC9000 controller

The Dell flagship storage controller, based on its 13G server architecture, provides the ideal solution for large-scale systems, high-end performance workloads and distributed enterprise environments. Whether you are building a private or hybrid cloud, a data warehouse or a secure OLTP platform, SC9000 gives you quick, strong return on investment by optimizing, accelerating and protecting the data that is most important to your business goals. The SC9000 builds on the success of the popular SC8000 Series arrays with a fully upgraded, more scalable hardware platform powered by the new SCOS 6.7 operating system.

With non-disruptive, modular expansion to over 3PB of raw capacity, the SC9000 can also be linked to other SC Series arrays in larger federated systems under unified management. All-flash and hybrid SSD/HDD configurations revolutionize the speed and impact of your mission-critical applications. You can easily evolve from one configuration to another without having to replace drives. New 12 Gbps SAS back-end and enterprise-class front-end network connections provide plenty of throughput headroom for growth, making the advantages of SSD storage practical at any scale. Redundant, hot-swappable components make the SC9000 a resilient, flexible and easy-to-deploy solution. The 2U controller unit functions in an active/active pair for maximum performance and availability. Adding capacity is simple, with five supported expansion enclosure options — including two new models with 12 Gbps SAS interconnect.

For more information on Dell Storage SC9000 array controllers, visit <http://www.dell.com/us/business/p/storage-sc9000/pd>.

3.5 Dell Storage SC420 disk enclosure

The SC420 is a 12Gbps SAS, 2U enclosure that supports up to 24 - 2.5" drives to enable a broad mix of storage options in one chassis. A single SAS chain can support up to seven fully populated SC420 enclosures.

For more information on Dell Storage SC420 disk enclosures, visit http://www.dell.com/learn/us/en/04/shared-content~data-sheets~en/documents~fy16q3_387_expansion_enclosure_spec_sheet.pdf.



4 Storage configuration

The configuration described in this section refers to the test environment used during validation of the reference architecture.

4.1 Back-end cabling

Each module in the SC420 disk enclosure was connected to the SC9000 controllers using 12 Gbps SAS. On the first module, port 1 was connected to the first controller and port 2 was connected to the second controller. On the second module, port 1 was connected to the second controller and port 2 was connected to the first controller. Since there were two SAS cards in each controller, only one port was used on each SAS card. The cabling diagram in Figure 3 shows how the SC420 was connected to the SC9000 controllers.




Figure 3 Single server configuration back-end cabling diagram

4.2 Front-end cabling

The R920 server was connected to the SC9000 storage array using 16 Gbps FC. Port 1 on each FC HBA in the server and controllers were connected to the same Brocade 6505 switch. Port 2 on each FC HBA in the server and controllers were connected to the other Brocade 6505 switch. The cabling diagram in Figure 4 shows how the R920 was connected to the SC9000 storage array.

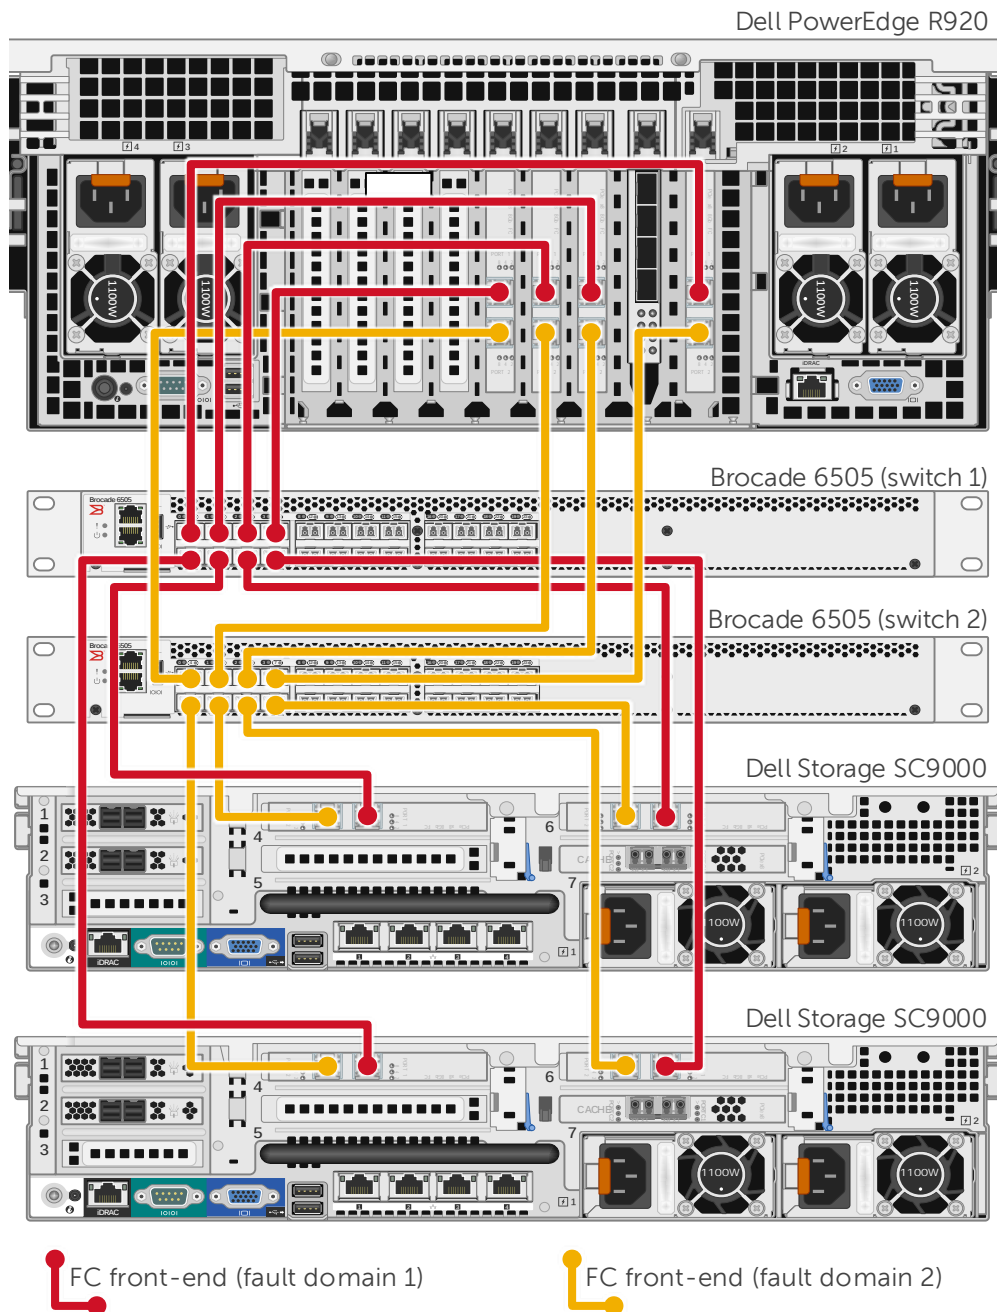


Figure 4 Single server configuration front-end cabling diagram

4.3 Local ports

The front-end FC ports and back-end SAS ports were configured using the Configure Local Ports wizard.

The FC ports were configured to use two fault domains in virtual port mode. Port 1 from each controller was put into fault domain 0 and port 2 from each controller was put into fault domain 1. In addition, all eight ports were configured as front-end ports by selecting **Front End** for the purpose.


All four SAS ports were configured as back end ports by selecting **Back End** for the purpose.



4.4 Disk folder
























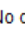
All 24 disks were put into the Assigned disk folder. One disk was defined as a hot spare, leaving 23 active disks. All volumes were created in the Assigned disk folder.

Figure 5 shows the disk configuration used in the reference architecture.

 **Assigned**

Storage Type: Redundant 512 KB

Tier 1 Storage
23 managed disks (40.17 TB) | 1 spare disk (1.75 TB) | Redundancy: Single Redundant

Position or Name	Capacity	Free Space	Classification	Enclosure	Status	Health	Control Type
 02-00	1.75 TB	1.75 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Spare
 02-01	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-02	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-03	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-04	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-05	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-06	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-07	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-08	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-09	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-10	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-11	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-12	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-13	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-14	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-15	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-16	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-17	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-18	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-19	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-20	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-21	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-22	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed
 02-23	1.75 TB	1.42 TB	Read-Intensive SSD	Enclosure - 2	Up	Healthy	Managed

Tier 2 Storage
No disks in this tier

Tier 3 Storage
No disks in this tier

Figure 5 Disk configuration

4.5 Storage type

The disk folder storage type defines the page size on the SC Series array and the level of redundancy (single or dual). A disk folder can be configured to use a 512KB, 2MB (default) or 4MB page.

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

4.6 Storage profile

Storage Center software uses storage profiles to define the RAID level and tiers on which 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 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 used this custom storage profile to ensure that RAID 5 was always used for those volumes. All other volumes used the **Recommended (All Tiers)** storage profile.

Figure 6 shows the creation of the custom storage profile.

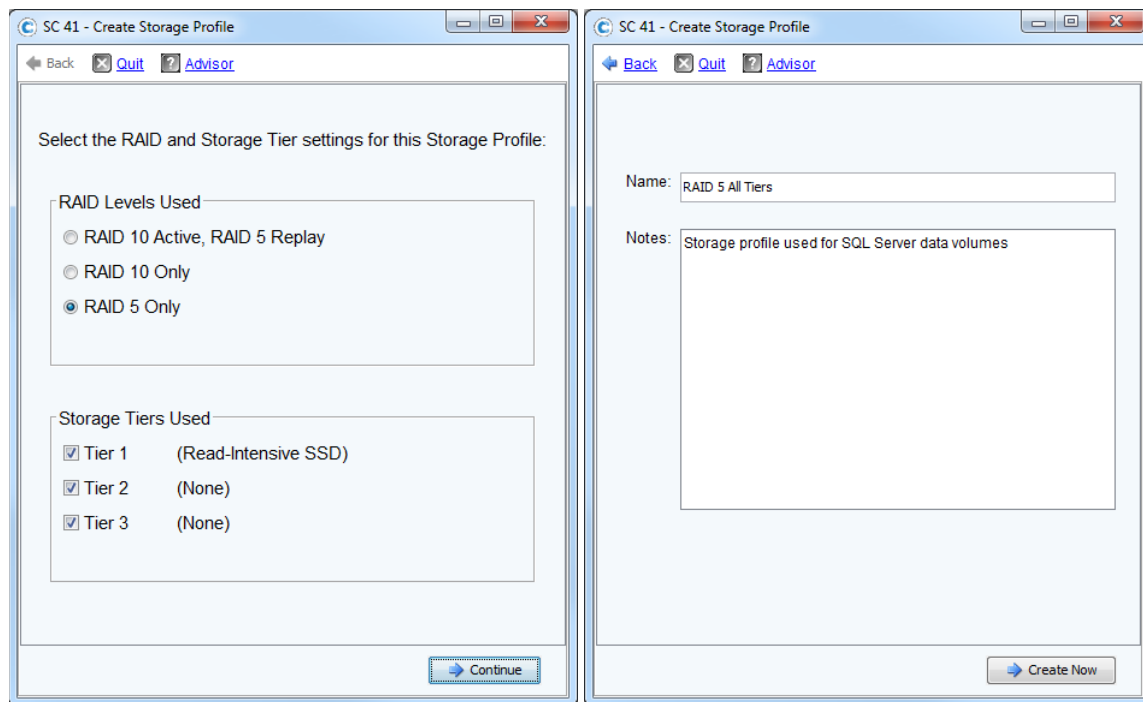


Figure 6 Creating the custom storage profile **RAID 5 All Tiers**

4.7 Storage Center 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. Storage Center includes both single path and multi-path definitions for each version of Windows.

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

Figure 7 and Figure 8 show the creation of the server object.




Figure 7 Creating a server object, part 1: Select the HBAs

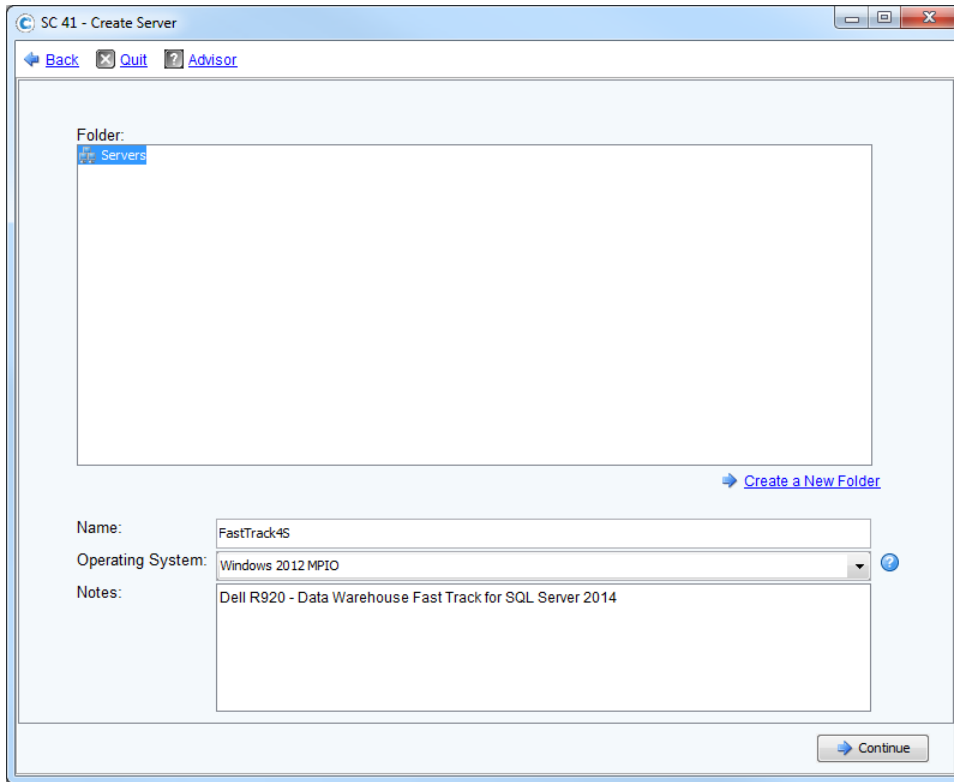


Figure 8 Creating a server object, part 2: Enter the name and select the operating system

4.8 Storage Center volumes

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

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

The server was configured to boot from the Storage Center. When mapping the boot volume to the server, the **Map volume using LUN 0** option was selected in the advanced dialog box. To simplify installation when booting from SAN, map the boot volume to the server and configure the HBAs to boot from SAN before mapping any other volumes to the server.

Figure 9 shows the advanced dialog box that is used when mapping the boot volume.

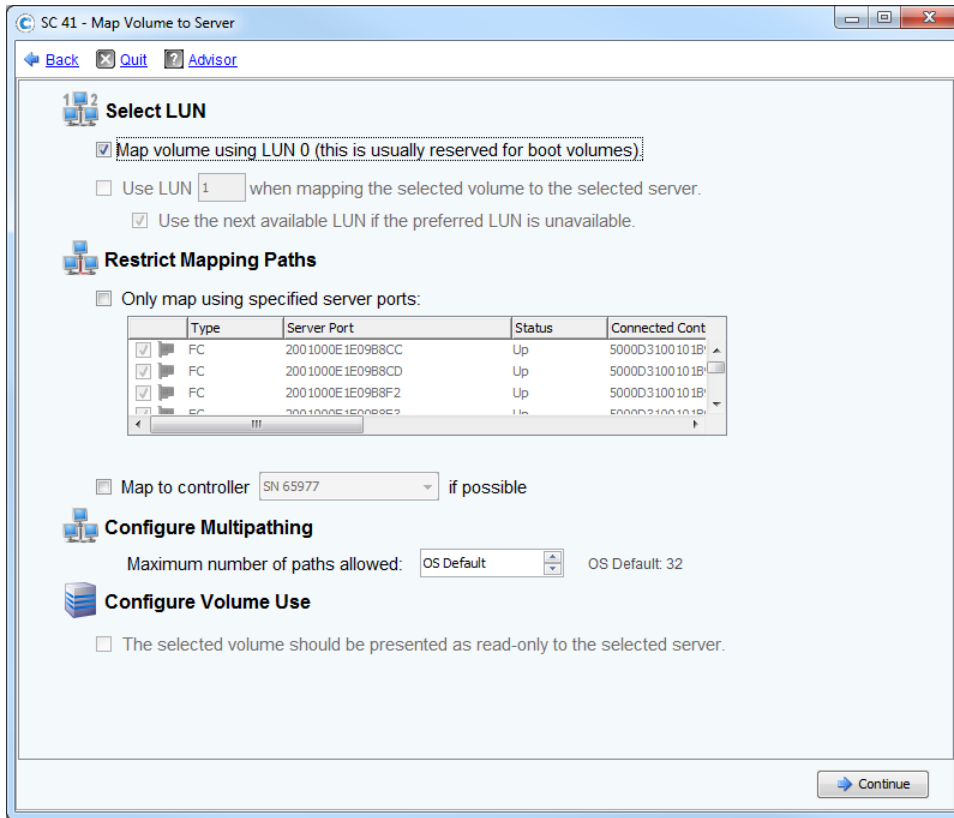


Figure 9 Advanced options for mapping the boot volume

In the test environment, 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 at the same time, one after the other, the Storage Center automatically distributes the volumes evenly across the controllers. Otherwise, the controller can be manually selected in the advanced dialog box when mapping the volume to the server. To achieve optimal results, the SQL Server data volumes must be evenly distributed across the controllers.

Table 3 Storage Center volumes created for the reference architecture

Volume name	Storage profile	Volume purpose
FastTrack4S - Boot	Recommended (All Tiers)	Windows boot volume
FastTrack4S - MPHost	Recommended (All Tiers)	Mount point host
FastTrack4S - SQLSystem	Recommended (All Tiers)	SQL Server system databases and files
FastTrack4S - SQLLog	Recommended (All Tiers)	Transaction log files for tempdb and the data warehouse
FastTrack4S - SQLData01	RAID 5 All Tiers	Data files for the data warehouse

Volume name	Storage profile	Volume purpose
FastTrack4S - SQLData02	RAID 5 All Tiers	Data files for the data warehouse
FastTrack4S - SQLData03	RAID 5 All Tiers	Data files for the data warehouse
FastTrack4S - SQLData04	RAID 5 All Tiers	Data files for the data warehouse
FastTrack4S - SQLData05	RAID 5 All Tiers	Data files for the data warehouse
FastTrack4S - SQLData06	RAID 5 All Tiers	Data files for the data warehouse
FastTrack4S - SQLData07	RAID 5 All Tiers	Data files for the data warehouse
FastTrack4S - SQLData08	RAID 5 All Tiers	Data files for the data warehouse
FastTrack4S - TempdbData01	Recommended (All Tiers)	Data files for tempdb
FastTrack4S - TempdbData02	Recommended (All Tiers)	Data files for tempdb



5 Server configuration

5.1 System BIOS

The system profile was set to **Performance**. All other options, outside of iDRAC configuration, were left at their factory default settings. The logical processor option, under processor settings, was left at its default setting of **Enabled**. This enabled 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 Storage best practices. 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 2012 R2 configuration

6.1 Power plan

To maximize performance, the server was configured to use the **High performance** power plan as shown in Figure 10.




Figure 10 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 Storage Center volume. Disks were initialized as GPT disks and then volumes were formatted on those disks 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 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 Storage Center 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
FastTrack4S - Boot	Boot	C:\
FastTrack4S - MPHost	MPHost	M:\
FastTrack4S - SQLSystem	SQLSystem	M:\ft\SQLSystem
FastTrack4S - SQLLog	SQLLog	M:\ft\SQLLog
FastTrack4S - SQLData01	SQLData01	M:\ft\SQLData01
FastTrack4S - SQLData02	SQLData02	M:\ft\SQLData02
FastTrack4S - SQLData03	SQLData03	M:\ft\SQLData03
FastTrack4S - SQLData04	SQLData04	M:\ft\SQLData04
FastTrack4S - SQLData05	SQLData05	M:\ft\SQLData05
FastTrack4S - SQLData06	SQLData06	M:\ft\SQLData06
FastTrack4S - SQLData07	SQLData07	M:\ft\SQLData07
FastTrack4S - SQLData08	SQLData08	M:\ft\SQLData08
FastTrack4S - TempdbData01	TempdbData01	M:\ft\TempdbData01
FastTrack4S - TempdbData02	TempdbData02	M:\ft\TempdbData02

6.4 MPIO

MPIO was configured using Dell Storage best practices. MPIO best practices for the Storage Center 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/download.

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

Figure 11 shows the MPIO policy used for each SQL Server data volume. Figure 12 shows the MPIO policy for all other volumes.



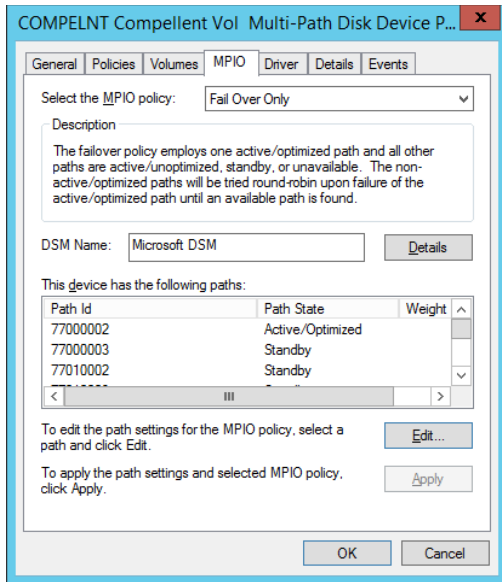


Figure 11 MPIO policy for data volumes

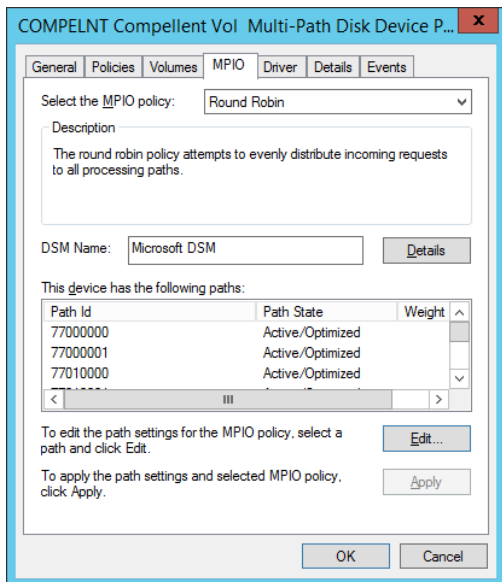


Figure 12 MPIO policy for other volumes

7 SQL Server 2014 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.

-T1117

This trace flag ensures the even growth of all files in a file group when auto growth is enabled. Note that it is recommend to expand the data files to their appropriate size rather than to depend on auto grow.

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 was 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 20 for the row store tests and 80 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 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 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 model is **Node majority with witness (disk or file share)** using a disk witness. An additional volume needs to be created and configured as the disk witness. Dell recommends using a 2 GB volume for the disk witness. For more information on quorum configurations in a failover cluster, visit <https://technet.microsoft.com/en-us/library/jj612870.aspx>.
- All volumes need to be mapped to each node of the cluster. It is recommended to use a cluster server object in the Storage Center 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 2014 certification

DWFT Certification #2014-038	Dell R920 with SC9000 - 80 TB			Report Date: 11/03/2015																			
DWFT Rev. 5.4	DWFT Reference Architecture																						
System Provider	System Name	Processor Type		Memory																			
	Dell PowerEdge R920	Intel Xeon E7-8891v2 3.2 GHz (4S/40C/80T)		1536 GB																			
Operating System		SQL Server Edition																					
Windows Server 2012 R2		SQL Server 2014 Enterprise Edition																					
Storage Provider	Storage Information																						
	Dell Storage SC9000 24 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 colspan="2"></td></tr><tr><td>80</td><td>295</td><td>351</td><td>98</td><td colspan="2"></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)			80	295	351	98				
Primary Metrics																							
Rated User Data Capacity ¹ (TB)	Row Store Relative Throughput ²	Column Store Relative Throughput ³	Maximum User Data Capacity ¹ (TB)																				
80	295	351	98																				
<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>295</td><td>285</td><td>7,950</td><td>9,459</td><td>8,705</td><td>93</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.) (%)	295	285	7,950	9,459	8,705	93
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.) (%)																		
295	285	7,950	9,459	8,705	93																		
<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>351</td><td>2,283</td><td>2,007</td><td>N/A</td><td>N/A</td><td>99</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.) (%)	351	2,283	2,007	N/A	N/A	99
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.) (%)																		
351	2,283	2,007	N/A	N/A	99																		
The reference configuration is a 2 socket system rated for 25TB using 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, in partnership with Microsoft, enables customers to deploy tested and validated data warehouse solutions using Data Warehouse Fast Track reference architectures for SQL Server 2014. 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.

The Dell Microsoft DWFT reference architecture provides the following benefits to customers:

- Delivers a tested and validated configuration with proven methodology and performance behavior
- Delivers outstanding performance on Dell 4-socket servers with blazing processor speeds and leading edge flash-based Dell storage arrays
- Achieves a balanced and optimized solution at all 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 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 R920 server with an SC9000 storage array. By implementing Data Warehouse Fast Track for SQL Server 2014 design principles, this configuration achieved an 80TB rating.



A Additional Resources

- Dell Products:
<http://www.dell.com>
- Dell Services:
<http://www.dell.com/services>
- Dell Support:
<http://www.dell.com/support>
- Dell SQL Server Solutions:
<http://www.dell.com/sql>
- Dell Data Warehouse Fast Track for SQL Server 2014 Advisor:
http://www.dell.com/solutions/advisors/us/en/g_5/SQLFastTrack/5/Start?s=biz
- Dell Storage technical content on Dell TechCenter:
<http://en.community.dell.com/techcenter/storage/w/wiki/2631.storage-infrastructure-and-solutions-team>

