

Evaluation of Dell™ PowerEdge™ VRTX Shared PERC8 in Failover Scenario

Evaluation report prepared under contract with Dell

Introduction

Dell introduced its PowerEdge VRTX integrated IT solution for remote-office and small-office environments. This solution integrates servers, storage, networking and management in a single chassis designed with dimensions, acoustics and security for the office setting. The PowerEdge VRTX is an ideal platform for clustered servers because it contains the network and storage infrastructure needed for clustered servers within a single chassis. In order to provide enterprise-class fault tolerance for clustered servers, the PowerEdge VRTX can be configured with an optional, redundant PERC that eliminates any single point of failure in access to internal shared storage, and provides next-level data protection inside the chassis.

With the PowerEdge VRTX solution, Dell brings enterprise class high availability shared storage in a low-cost platform that does not require an external storage network, which can be cost prohibitive for these environments. Important applications such as databases, email applications and others can be deployed onto the PowerEdge VRTX platform with the redundant RAID controllers in a high availability configuration for relatively low cost compared to larger platforms that require additional, often expensive, equipment.

In order to validate the fault-tolerance of the Dell SPERC8 storage adapters, Dell commissioned Demartek to build a Windows Server 2012 R2 failover cluster with a PowerEdge VRTX system using two SPERC8 adapters, and run a database workload, testing the failover functions. The redundant PERC adapters were configured in an active/standby configuration where the second PERC was hot and ready to take over the workload in the event of failure of the first PERC adapter.

Key Findings

We found that the PowerEdge VRTX system with the dual SPERC8 adapters performed the failover process very smoothly. The SQL Server application observed a delay of 39 seconds for the first failover and only three seconds for the failback, and the application kept running without manual intervention.

This type of high availability solution is frequently available in more expensive equipment found in large datacenters. With the PowerEdge VRTX solution, high availability is available in lower cost solutions for the remote office, branch office and small office environments.

Dell PowerEdge VRTX

The PowerEdge VRTX system is an integrated IT solution that combines the best of rack-server technologies with some of the denseness of blade-server technologies in a form factor that is suitable for remote-office and small-office environments. The PowerEdge VRTX chassis combines up to four servers, 1 GbE or 10 GbE networking technology and internal storage that can be shared by the servers. The entire chassis can be managed using the Chassis Management Controller (CMC) via a command-line or a graphical user interface (GUI).

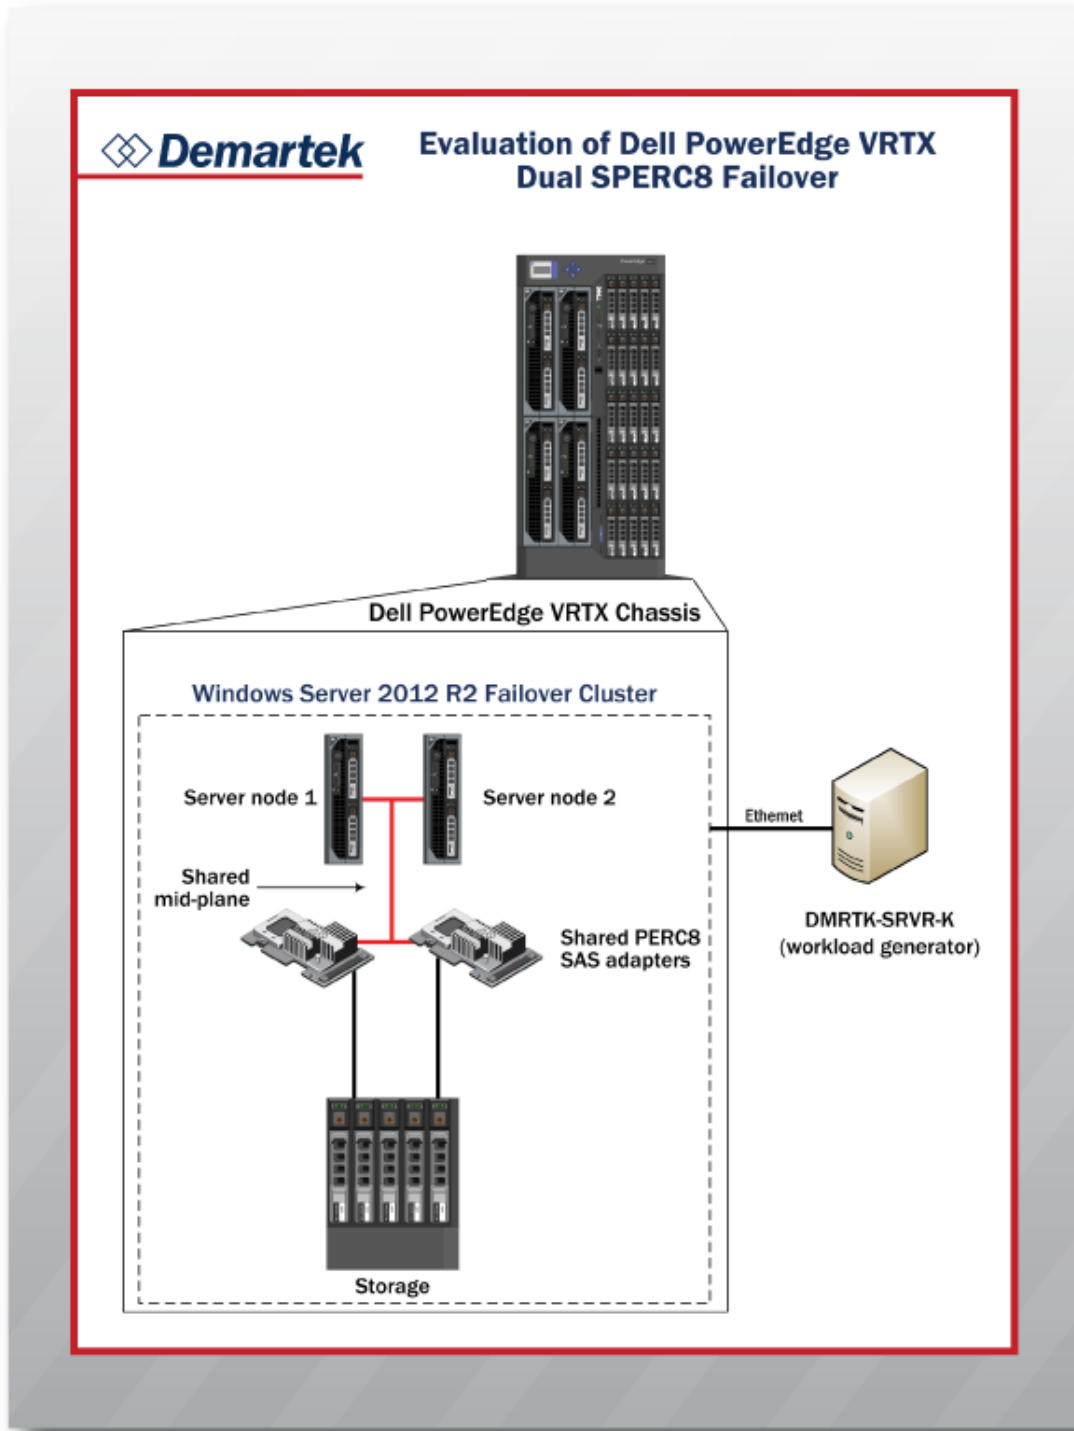
The PowerEdge VRTX chassis is designed in such a way that the eight PCIe I/O slots can be assigned to any of the server nodes within the chassis. This makes it possible to provide the benefit of one expensive I/O adapter to multiple servers.

The PowerEdge VRTX unit that we tested included the following:

- ◆ Qty. 2 – Dell PowerEdge M620 servers, each server having:
 - 2x Intel Xeon E5-2620, 2.0 GHz, 12 total cores, 24 total threads
 - 32GB RAM
- ◆ Qty. 2 – Dell Shared PERC (SPERC8) 6 Gb/s SAS/SATA RAID controllers, configured in fault-tolerant mode
- ◆ Qty. 5 – Seagate 300GB 10K RPM 6 Gb/s SAS-interface HDDs
- ◆ Qty. 1 – Dell R1-2401 PowerEdge VRTX 1Gb switch module
 - Qty. 8 – external 1Gb Ethernet ports
 - Qty. 16 – internal 1Gb Ethernet ports
- ◆ Qty. 4 – 1100 watt power supplies

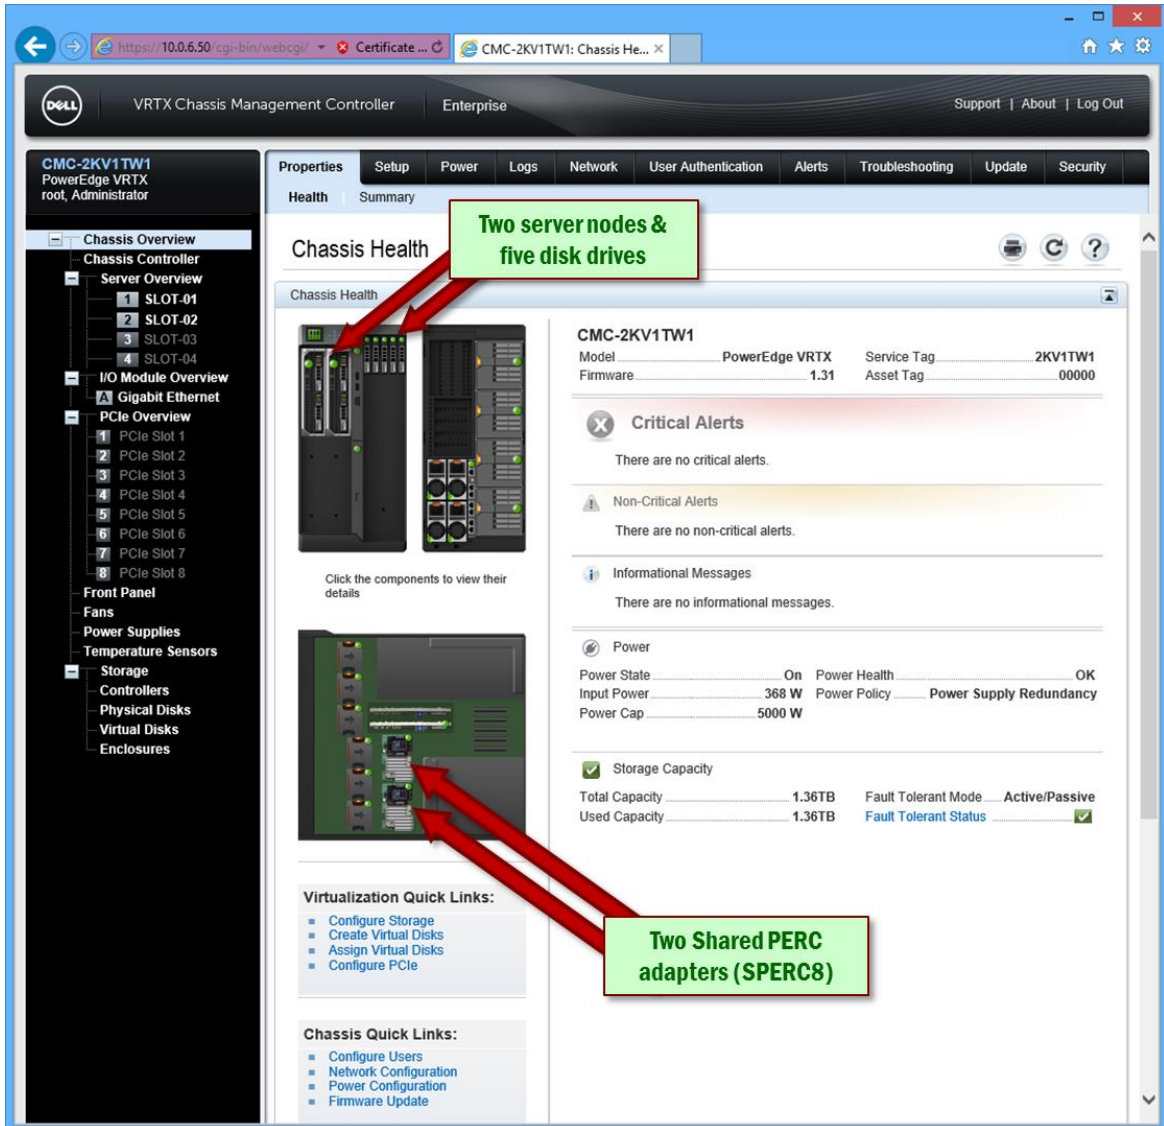
Test Configuration

The test configuration included the PowerEdge VRTX chassis and components, along with one external server that acted as a workload generator.



Management Interface

The CMC GUI provides management functions for the PowerEdge VRTX chassis and all of its components, either directly or indirectly. The chassis overview provides a clean picture of the components installed in the system, and their operational status.



The screenshot displays the Dell VRTX Chassis Management Controller (CMC) GUI. The interface includes a navigation menu on the left with categories like Chassis Overview, Server Overview, I/O Module Overview, PCIe Overview, Front Panel, Fans, Power Supplies, Temperature Sensors, Storage, Controllers, Physical Disks, Virtual Disks, and Enclosures. The main content area shows the 'Chassis Health' page, which includes a 'Chassis Health' summary, a 'Critical Alerts' section, a 'Non-Critical Alerts' section, 'Informational Messages', and 'Power' and 'Storage Capacity' sections. A red box highlights the text 'Two server nodes & five disk drives' with arrows pointing to the server rack image. Another red box highlights the text 'Two Shared PERC adapters (SPERC8)' with arrows pointing to the storage controller image. The 'Power' section shows: Power State: On, Power Health: OK, Input Power: 368 W, Power Policy: Power Supply Redundancy, Power Cap: 5000 W. The 'Storage Capacity' section shows: Total Capacity: 1.36TB, Used Capacity: 1.36TB, Fault Tolerant Mode: Active/Passive, Fault Tolerant Status: .

From the CMC, additional detail for any of the components can be obtained by clicking on the component. In the examples below, the details for one of the server and one of the SPERC8 adapters is shown.

Server Status

Jump to: [Properties](#) | [I/O Fabric Interfaces](#) | [iDRAC System Event Log](#) | [Common iDRAC Network Settings](#) | [IPv4 iDRAC Network Settings](#) | [IPv6 iDRAC Network Settings](#) | [WWN/MAC Addresses](#)

[Launch iDRAC GUI](#) [Launch Remote Console](#)

Properties

Attribute	Value
Slot	1
Slot Name	SLOT-01
Present	Yes
Health	<input checked="" type="checkbox"/>
Server Model	PowerEdge M620
Service Tag	2KK1TW1
iDRAC Firmware	1.40.40 (Build 17)
CPLD Version	1.0.6
BIOS Version	1.7.6
Host Name	D
Operating System	W
CPU Information	2 x Intel(R) Xeon(R) CPU E5-2620 0 @ 2.00GHz
Total System Memory	32.0 GB

I/O Fabric Interfaces

Location	Installed	Type	Model
Fabric A	Yes	10 GbE KR	BRCM 10GbE 2P 57810s bNDC
Fabric B	Yes	PCIe Bypass Generation 2	DELL PCIe Mezzanine
Fabric C	Yes	PCIe Bypass Generation 2	DELL PCIe Mezzanine

iDRAC System Event Log

Severity	Date/Time	Description
<input checked="" type="checkbox"/>	Fri Apr 11 2014 17:56:22	OEM software event.
<input checked="" type="checkbox"/>	Fri Apr 11 2014 17:56:22	C: boot completed.
<input checked="" type="checkbox"/>	Fri Apr 11 2014 17:52:21	OEM software event.

Controllers

Controllers

Status	Name	Rollup Status	SPERC Slot
<input checked="" type="checkbox"/>	Shared PERC8 (Integrated 1)	OK	1
<input checked="" type="checkbox"/>	Shared PERC8 (Integrated 2)	OK	2

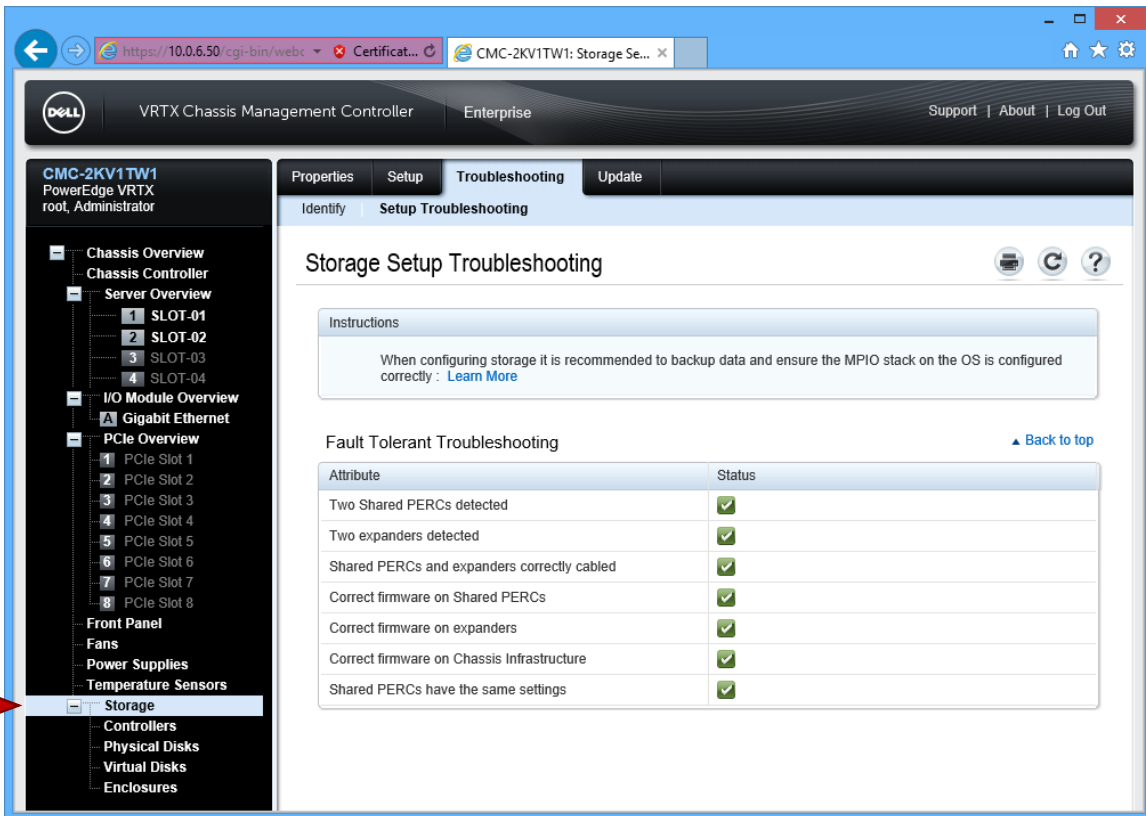
Advanced Properties

Status	<input checked="" type="checkbox"/>	Patrol Read Mode	Automatic
Name	Shared PERC8 (Integrated 1)	Patrol Read State	Stopped
Rollup Status	OK	Check Consistency Mode	Normal
SPERC Slot	1	Check Consistency Rate	30%
Firmware Version	23.8.10-0059	Copy Back Mode	On
Fault Tolerant Mode	Active/Passive	BGI Rate	30%
Fault Tolerant Status	<input checked="" type="checkbox"/>	Rebuild Rate	30%
Peer Controller	Shared PERC8 (Integrated 2)	Preserved Cache	Not Present
Cache Memory Size	1024MB	Battery Status	<input checked="" type="checkbox"/>
SAS Address	5848F690ED42EA00	Battery State	Ready
Capable Speeds	6.0Gb/s & 3.0Gb/s		

[View Physical Disks for this Controller.](#)
[View Virtual Disks for this Controller.](#)

Fault Tolerance

The PowerEdge VRTX system supports fault tolerance for the storage by the use of the two SPERC8 adapters connected to the same storage. The CMC can determine if the two SPERC8 adapters have been correctly configured, as shown below.



The screenshot shows the Dell VRTX Chassis Management Controller (CMC) web interface. The browser address bar shows the URL `https://10.0.6.50/cgi-bin/webc`. The page title is "CMC-2KV1TW1: Storage Setup Troubleshooting". The left sidebar contains a navigation tree with the following items: Chassis Overview, Chassis Controller, Server Overview (SLOT-01 to SLOT-04), I/O Module Overview (Gigabit Ethernet), PCIe Overview (PCIe Slot 1 to 8), Front Panel, Fans, Power Supplies, Temperature Sensors, **Storage** (highlighted with a red arrow), Controllers, Physical Disks, Virtual Disks, and Enclosures. The main content area is titled "Storage Setup Troubleshooting" and includes an "Instructions" box with the text: "When configuring storage it is recommended to backup data and ensure the MPIO stack on the OS is configured correctly : [Learn More](#)". Below this is a "Fault Tolerant Troubleshooting" section with a "Back to top" link and a table of results.

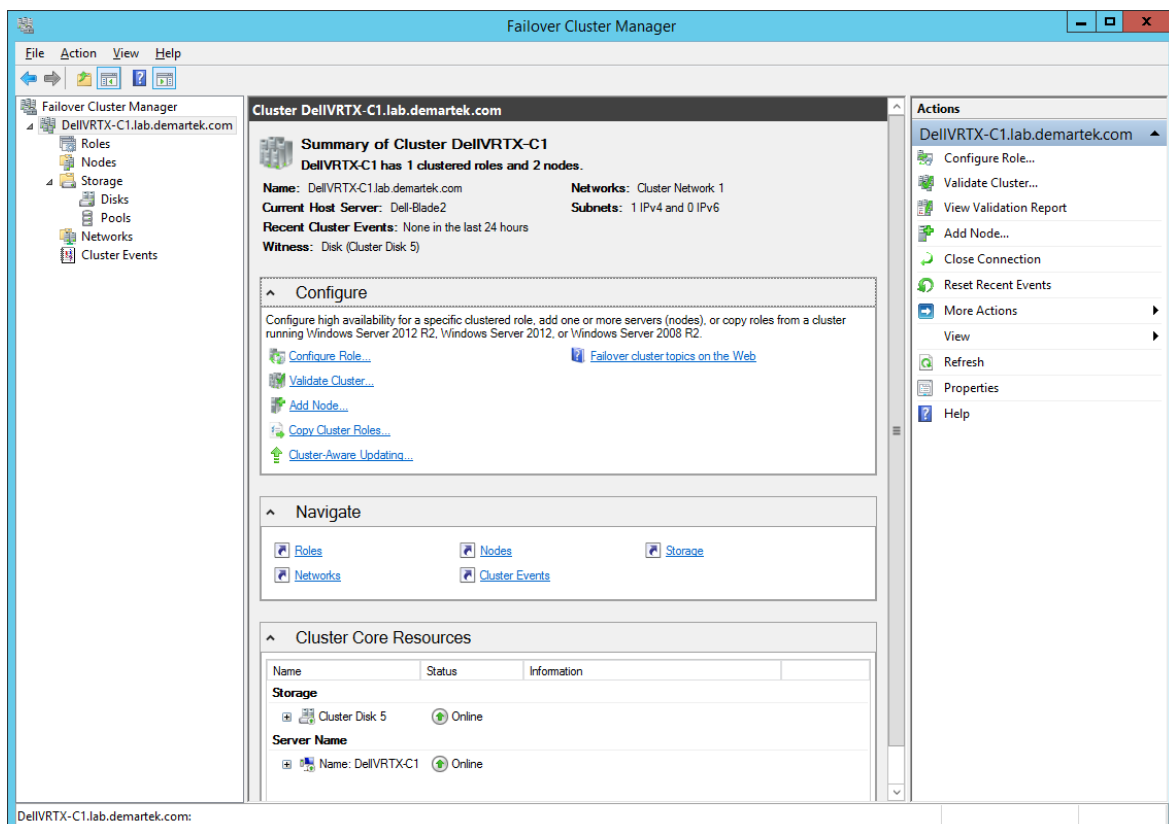
Attribute	Status
Two Shared PERCs detected	✓
Two expanders detected	✓
Shared PERCs and expanders correctly cabled	✓
Correct firmware on Shared PERCs	✓
Correct firmware on expanders	✓
Correct firmware on Chassis Infrastructure	✓
Shared PERCs have the same settings	✓

High-Availability Direct-Attached Storage for Windows Server

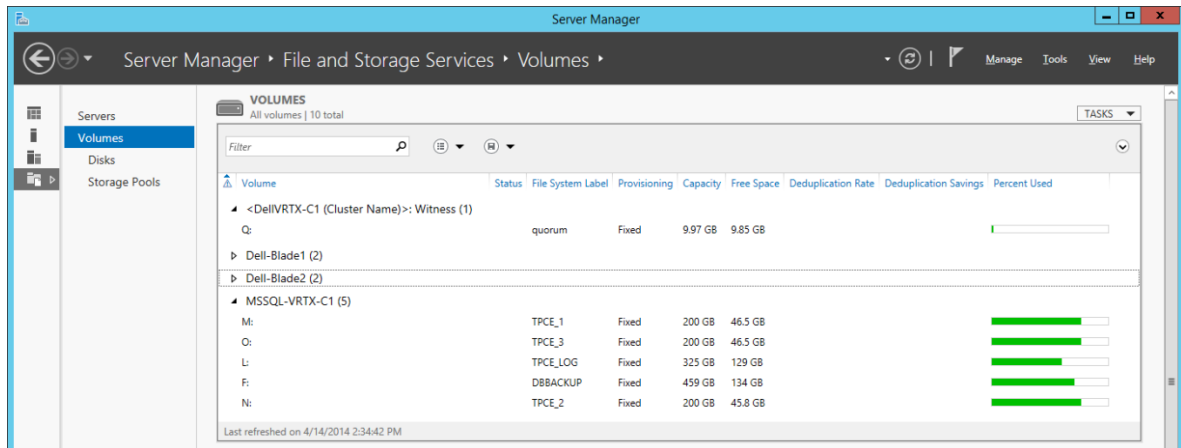
One of the technical capabilities provided by the PowerEdge VRTX solution is the ability to have high availability direct-attached storage (DAS) in a Windows Server environment. Small and medium-sized businesses and cloud datacenters that desire continuous application uptime can have access to their data in the event that a RAID controller was to fail.

For this test, Demartek constructed a Windows Server 2012 R2 Failover Cluster using two server nodes with shared storage controlled by the pair of SPERC8 adapters, entirely within the PowerEdge VRTX chassis. Microsoft SQL Server 2012 was installed on each of the server nodes and an online brokerage application workload was deployed using shared storage within the PowerEdge VRTX chassis. An external server was used as a workload generator communicating with the cluster. Because of the architecture of the PowerEdge VRTX chassis, both server nodes can take advantage of the SPERC8 adapters.

The Windows Server 2012 R2 Failover Cluster was deployed in the Demartek lab domain, as shown below. The cluster quorum disk and all the application storage was contained in the shared storage within the PowerEdge VRTX chassis.



The database application data was configured the same for each server node, with the database spread across three data volumes, one log volume, and a backup volume. The database consisted of 40,000 customers. The database consumes approximately 460GB and the log consumes approximately 196GB for a combined total of approximately 656GB. The drive letters used for the data volumes are shown below.

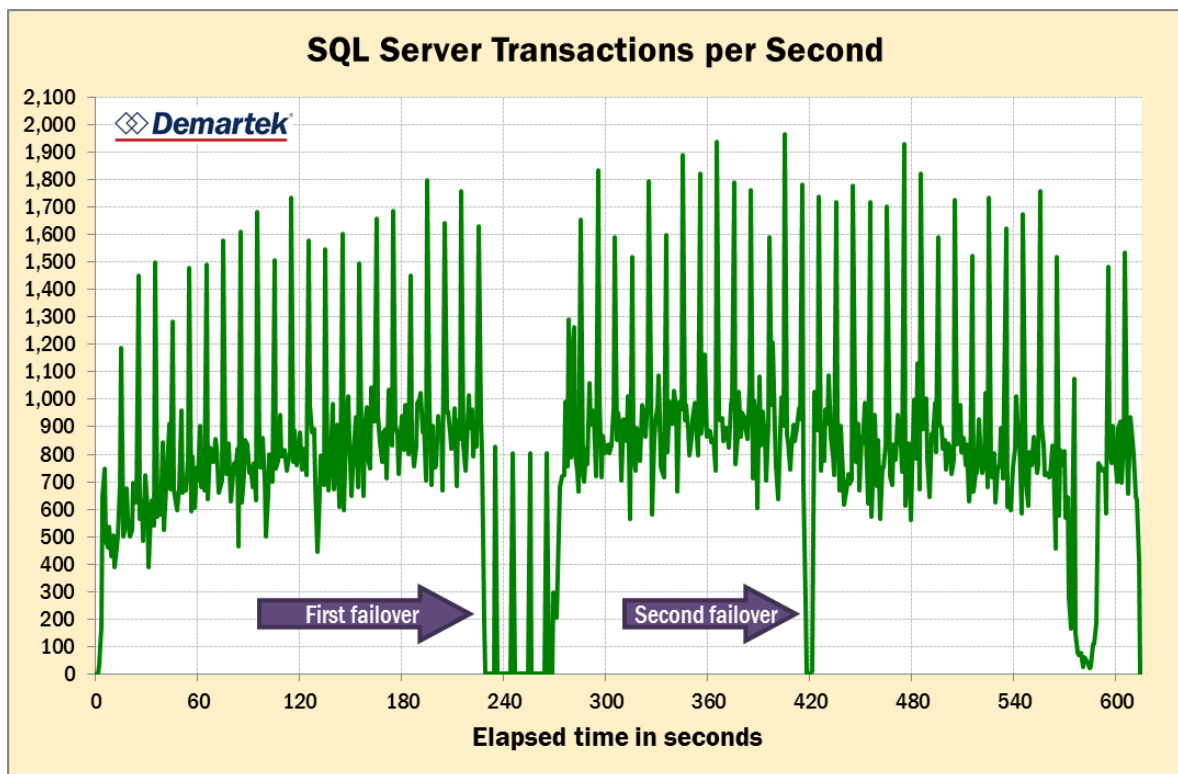


Failover Tests

To test the high availability of the SPERC8 adapters, we started the SQL Server application workload, and while it was running, caused a failover to occur with the SPERC8 cards by issuing the appropriate commands. After a few minutes, we caused the failover to occur on the other SPERC8 adapter. During this time, we captured the performance statistics from the SQL Server application and the Windows Server node.

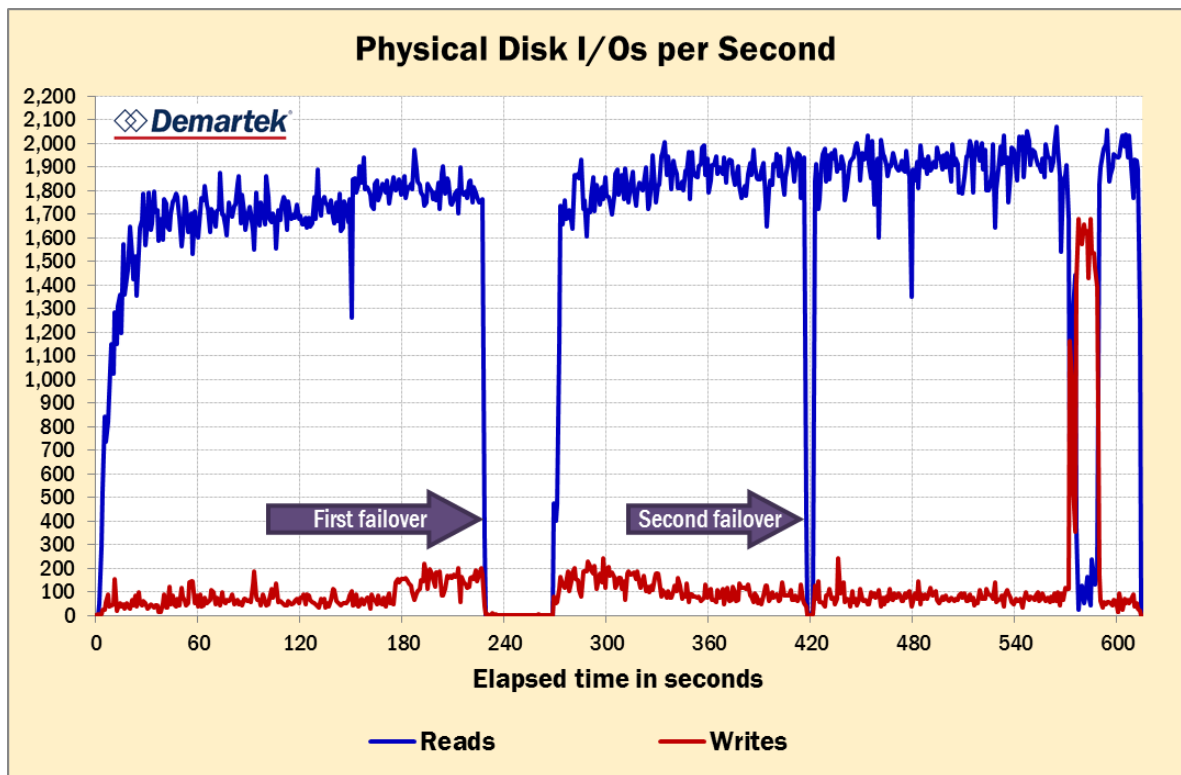
SQL Server Transactions per Second

The SQL server transactions per second momentarily dipped during both occurrences of the failovers. The first failover took a bit longer than the second failover, but the application kept running.



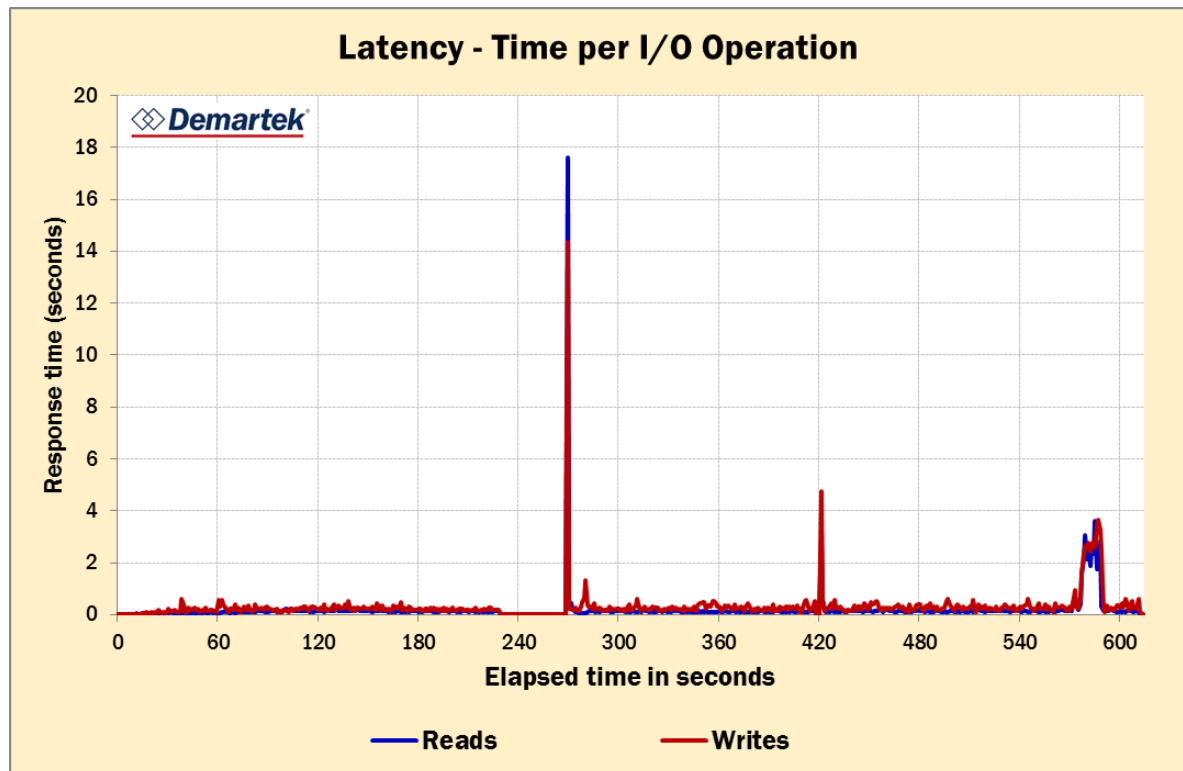
Physical Disk I/Os per Second

The physical disk performance statistics are provided by the standard Windows Performance Monitor (Perfmon). The chart below shows the aggregate physical disk statistics for all of the database and log volumes. As expected, during each failover no disk activity was reported, but the application kept running.



Latency – Response Time

One of the measures of application performance is the time taken to complete I/O requests. In the case of a high availability solution, during the failover, the application should pause but not error or receive a “timeout” condition. In this case, SQL Server running on a Windows Server Failover Cluster simply paused, and did not error out while the failovers occurred. The default settings for timeout values in the Window Server Failover Cluster and SQL Server were used for this test. Even during a failover that took several seconds to complete, the application did not fail or receive I/O errors.



Summary and Conclusion

The Dell PowerEdge VRTX is an excellent solution for the remote office or small office that needs fault tolerance and high availability for its storage. The dual SPERC8 adapters provide this high availability, as we were able to demonstrate. When system uptime and data availability are crucial criteria, customers can rely on PowerEdge VRTX in the remote office and small and medium business environments.

Some environments such as email, web servers and others are frequently run in a 24-by-7 mode regardless of the size of organization using these applications. The PowerEdge VRTX solution brings the ability to run these types of applications in a lower-cost solution, without the need for expensive data center equipment.



The most current version of this report is available at http://www.demartek.com/Demartek_Dell_PowerEdge_VRTX_2014-05.html on the Demartek website.

Dell[™], the Dell logo and PowerEdge[™] are trademarks of Dell Inc.

Demartek is a registered trademark of Demartek, LLC.

All other trademarks are the property of their respective owners.