VMware ESXi on Fault Tolerant Shared PERC8 enabled Dell PowerEdge VRTX

A Dell Technical White Paper

Dell Hypervisor Engineering and Storage Validation

September 2014

September 2014

Table of Contents

Introduction to VRTX

PowerEdge VRTX is a new modular infrastructure targeted for Remote Office/Back Office (ROBO) installations. PowerEdge VRTX combines existing blade servers (M620 and M520) with new Mezzanine bridge cards, I/O and shared Storage infrastructure. It supports up to four blade servers (half-height). With an updated version of PowerEdge VRTX release, we support M820 blade as well. PowerEdge VRTX also includes an IO module slot, a front LCD panel for convenient chassis management, and a KVM port that can be assigned to a server, a DVD drive that can be assigned to a server, redundant power supply units, and fans.

For detailed information on PowerEdge VRTX platform architecture, features, and capabilities, refer to Dell website. VRTX is the first shared IT solution designed specifically for remote and small office environments, with enterprise-class capabilities in a desk-side, space-saving design. It combines servers, shared storage, IO in the form of PCIe slots and networking into a 5U tower (rackable) chassis that is suitable for the small and remote offices it is specifically designed to operate in.

With the latest release, The PowerEdge VRTX system is currently available in two configurations of the Shared PowerEdge RAID Controller (PERC) 8:

1. Single Shared PERC 8 Controller Configuration
2. "Fault Tolerant Shared PERC 8" Controllers Configuration

For more details on the configuration details, refer to the 'Overview' section in PowerEdge VRTX Shared PERC 8 User Guide available at Dell support site. The high level overview of both configurations is as follows.

- Single Shared PERC 8 Controller configuration — In this configuration, the PowerEdge VRTX system is installed with a single Shared PERC 8 card. The default cache policy for virtual disks created in this configuration is write-back.

- Fault tolerant Shared PERC 8 Controllers configuration — In this configuration, the PowerEdge VRTX system is installed with two Shared PERC 8 controllers. Either of the two Shared PERC 8 controllers can access the storage subsystem. If one Shared PERC 8 controller fails, the other Shared PERC 8 controller takes control in a seamless transition. The default virtual disk cache policies created in this configuration is write-through.

NOTE: **Performance Variations in Fault Tolerant Configurations:**

Performance is highly variable across configurations, drive type, applications & IO workloads. The write performance of write-through mode typically is lower than that of write-back mode. Some configurations and workloads exhibit significant performance deltas. Other configurations and workloads exhibit less noticeable deltas.
Users are encouraged to evaluate their needs for maximum I/O performance (a single Shared PERC 8 configuration) versus maximum data protection (a Fault Tolerant Shared PERC 8 configuration).

Note: Based on your needs, you can enable or disable second Shared-PERC8 controller on dual SPERC8 configuration in CMC firmware 1.35 onwards.

Audience and Scope

This white paper is intended to customers who are currently running or planning to run VMware ESXi on Dell PowerEdge VRTX. This document provides an overview of VRTX supported ESXi versions, the newly released fault tolerant shared perc8 feature and the multipathing support of ESXi on VRTX. It also provides the recommended way of configuring VRTX in VMware environment from Dell point of view.

Fault Tolerant Shared PERC 8

New, enhanced capabilities for PowerEdge VRTX provides **Next-level data protection for Shared Storage available in VRTX chassis** by enabling Optional redundant Shared PowerEdge RAID Controller's (SPERC) which provides fault tolerant for Shared PERC8 controllers failover in seamless transition thereby customers doesn't see IO interrupt to Shared Storage. Either of the two Shared PERC 8 controllers can access the storage subsystem. Refer to PowerEdge VRTX Shared PERC 8 user guide to understand more on the features. The CMC window as depicted below indicates the shared PERC controllers.

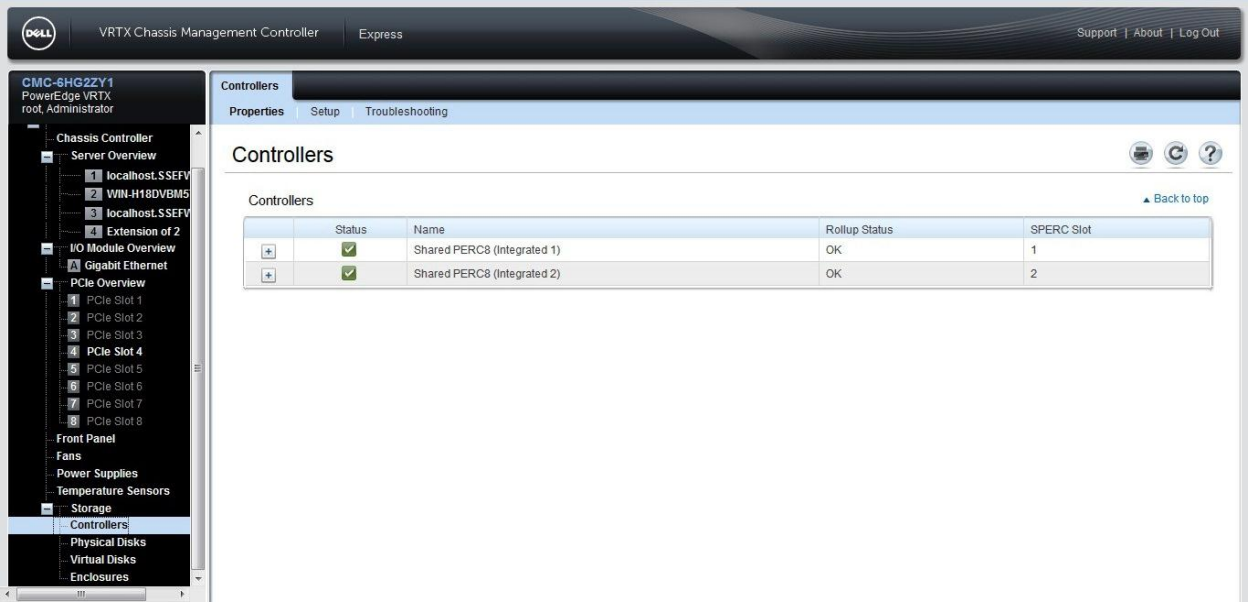

Figure 1 : Dual Shared PERC 8 in VRTX

Supported PowerEdge Blade Models on VRTX

The PowerEdge models supported on VRTX are:-

- PowerEdge M520
- PowerEdge M620
- PowerEdge M820

The modules that are specifically configured for the enclosure, and can be identified by a label marked PCIe on the server module. If you install PowerEdge M520, M620, or M820 server modules that are not configured for the enclosure, an error message is displayed. For more information on configuring a server module for the enclosure, see the Dell PowerEdge VRTX Owner's Manual at dell.com/power edge manuals.

Firmware Pre-requisites

The below are the firmware levels required to run Fault Tolerant Shared PERC8 on PowerEdge VRTX.

- VRTX CMC FW: 1.30 and above
- VRTX Chassis Infrastructure FW : 1.3 and above
- SAS Expander FW-: 1.02 and above
- Shared PERC8 FW: 23.8.10-0059 and above.

NOTE:

1. **These pre-requisites are the minimum required Firmware versions for VRTX to support Dual Shared PERC8 controller configuration in fault tolerance mode.** However, these versions are also supported in VRTX single Shared PERC controller configuration.

2. If VRTX system is updated with CMC FW: 1.30, then update VRTX chassis infrastructure FW to 1.3 and Shared PERC FW to: 23.8.10-0059. If VRTX system is updated with CMC FW: 1.31, then update VRTX chassis infrastructure FW to 1.4 and Shared PERC FW to: 23.8.10-0061.

3. CMC firmware 1.35 provides option to enable or disable second Shared-PERC8 controller on dual SPERC8 configuration (i.e. Change between Non-Fault Tolerant & Fault Tolerant Modes).

4. Recommended to use latest/higher version firmware as available in Dell support website.

Supported VMware ESXi versions

The New, Enhanced VRTX 12G servers supports the following VMware ESXi versions:

- VMware ESXi 5.5 Update 2
- VMware ESXi 5.5 Update 1
- VMWare ESXi 5.5
- VMWare ESXi 5.1 Update 2
- VMware ESXi 5.1 Update 1

The specific details of ESXi version support and the driver version details are outlined below.

ESXi5.1U1

Dell Customized ESXi 5.1U1 A03 Image carries Shared PERC8 Driver Version megaraid-sas-*06.801.52.00* (Single Controller). For Customers who want to upgrade to Dual Controller in VRTX with ESXi5.1U1, need to upgrade Driver version to megaraid-sas- 06.802.71.00. This driver can be downloaded from VMware site. For VRTX to work in Fault Tolerant mode (Dual Controller), FW stack (CMC, Main Board FW and PERC FW) has to upgrade to latest releases available at support.dell.com.

ESXi5.1U2

Dell Customized ESXi 5.1 U2 A03 Image carries Shared PERC8 Driver Version megaraid-sas-*06.802.71.00* (Dual Controllers).

ESXi5.5 U2

Dell Customized ESXi5.5 U2 A00 Image carries Shared PERC8 Driver version megaraid_sas-06.803.52.00(Dual Controllers)

ESXi5.5 U1

Dell Customized ESXi5.5 U1 A01 Image carries Shared PERC8 Driver version megaraid_sas-06.803.52.00(Dual Controllers)

ESXi5.5

Dell Customized ESXi5.5 A04 Image carries Shared PERC8 Driver version megaraid_sas-06.803.52.00(Dual Controllers)

NOTE: Customers who are installing ESXi5.1U2 on VRTX single controller system need to upgrade the FW stack of CMC, Mainboard and PERC8 to latest versions available at Dell support site.

Upgrading from Single Shared PERC to dual controllers

It's strongly recommended to call Dell tech support to get guidance on upgrading from single Shared PERC controller to dual controllers. For VRTX system upgrading from Single Shared PERC8 to Dual Shared PERC8 controller configuration, follow the below steps.

The upgrade steps are documented at PowerEdge VRTX setup guide as well.

CAUTION: The upgrade procedure is complex and can put your data at risk if mistakes are made. It is imperative that critical data is backed up prior to starting the procedure.

For information about version numbers, refer to the "Prerequisites" section. The Firmware/driver modules for VRTX Single Shared PERC8 controller configuration should be update in the following order:

- Chassis Management Controller.
- Chassis Infrastructure Firmware (i.e. Mainboard FW).
- SAS Expanders.
- Shared PERC8 Driver.
- Shared PERC8 Firmware.

NOTE:

- The CMC firmware should be updated prior to updating the server component firmware modules listed here.
- Before updating a single PERC, you must turn off the servers.

VMware Multipath support (Only in Dual Shared PERC8 Controllers)

Multipath support on VMware ESXi running in VRTX is automatically enabled using VMware Native Multi-pathing (NMP). This is applicable only when Dual Shared PERC8 controller's configuration in VRTX systems supporting Dell versions of supported ESXi operating systems.

NOTE: Dell and VMware strongly recommended NMP path policy as "*Most Recently Used (MRU)*" for VRTX shared storage virtual disks.

Setting up Multipath for shared LUNs

The multipathing policy set for the LUNs exposed from Shared PERC8 is MRU (Most Recently Used". This is set by default for the shared LUNs exposed from Shared PERC8 controllers.

As an example, let's take a scenario by creating a LUN from shared PERC 8 and mark it available for nodes running VMware ESXi. From vCenter client/server or by using command line utilities (*esxcfg-mpath* etc) the multiple paths configured for the shared LUN can be listed.

1. Create a LUN via VRTX CMC UI as shown below. As you can see, the LUN has got an active controller and a Redundant/Failover controller through which the fault tolerance is achieved.

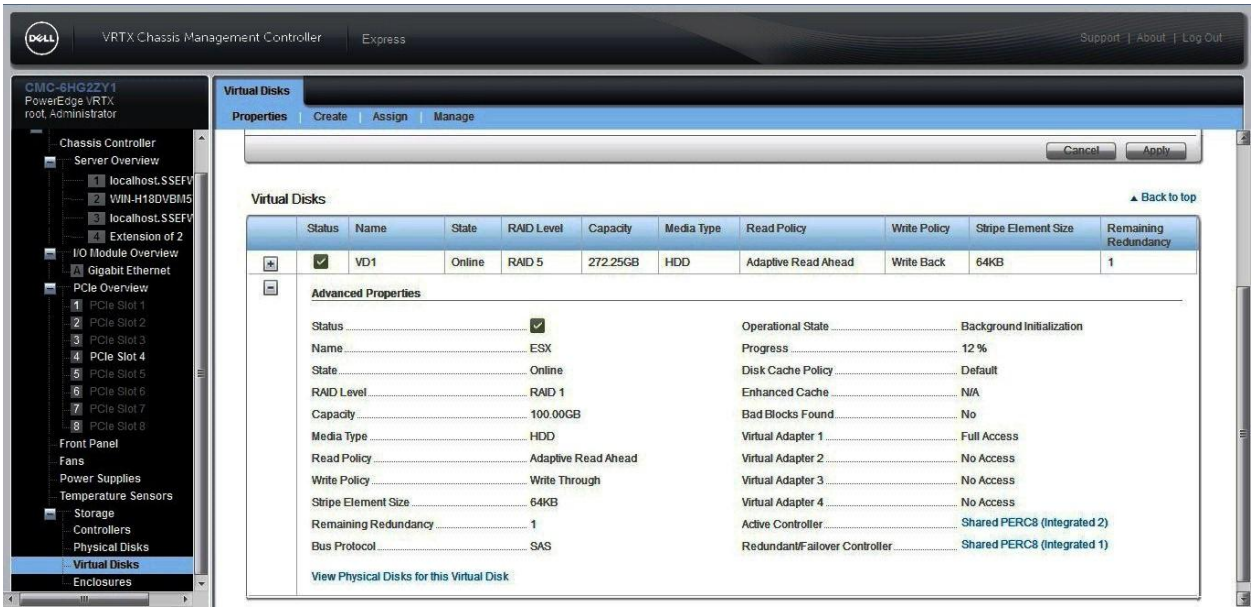

Figure 2 : VRTX CMC Virtual Disks Property page

2. In Dual Shared PERC8 configuration, upon mapping Virtual Function (VF) to ESXi host shows 2 Shared PERC8 controllers as shown in below picture.

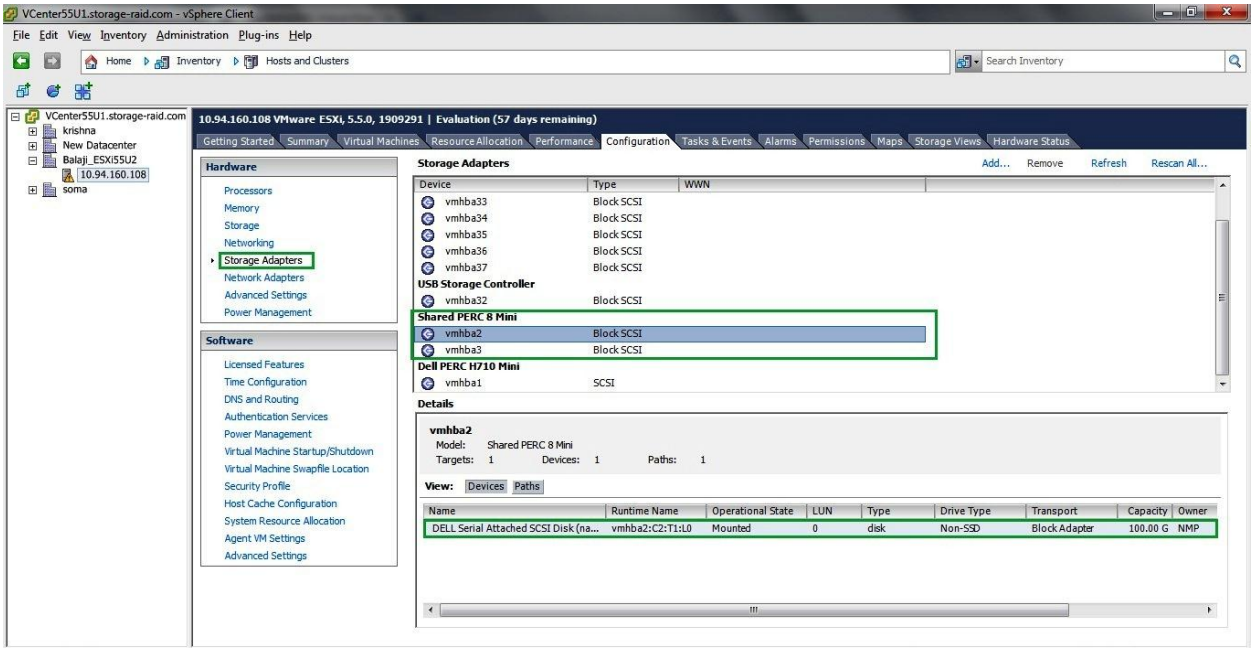

Figure 3 : Shared PERC8 Virtual Functions in ESXi

3. From VMware vCenter client/server, the LUN properties can be viewed as below. As you can see, the LUN has got an *Active (I/O)* path and *Active* path. The multi path policy used is "MRU" as indicated below.

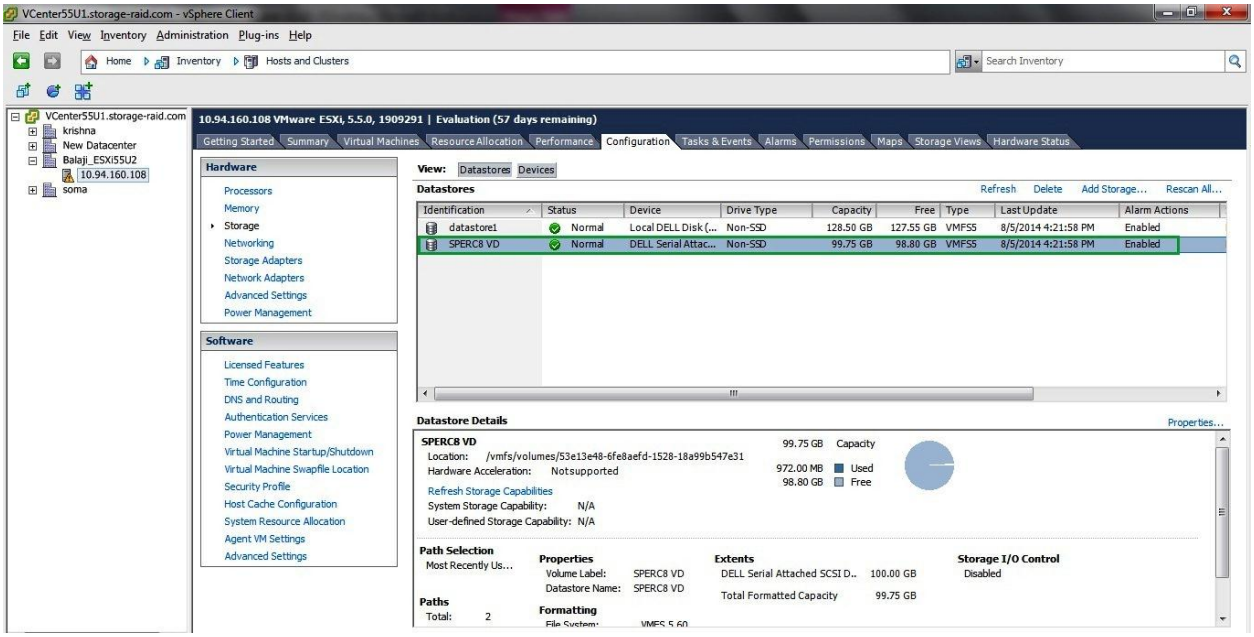

Figure 4 : Shared PERC8 Virtual Disk

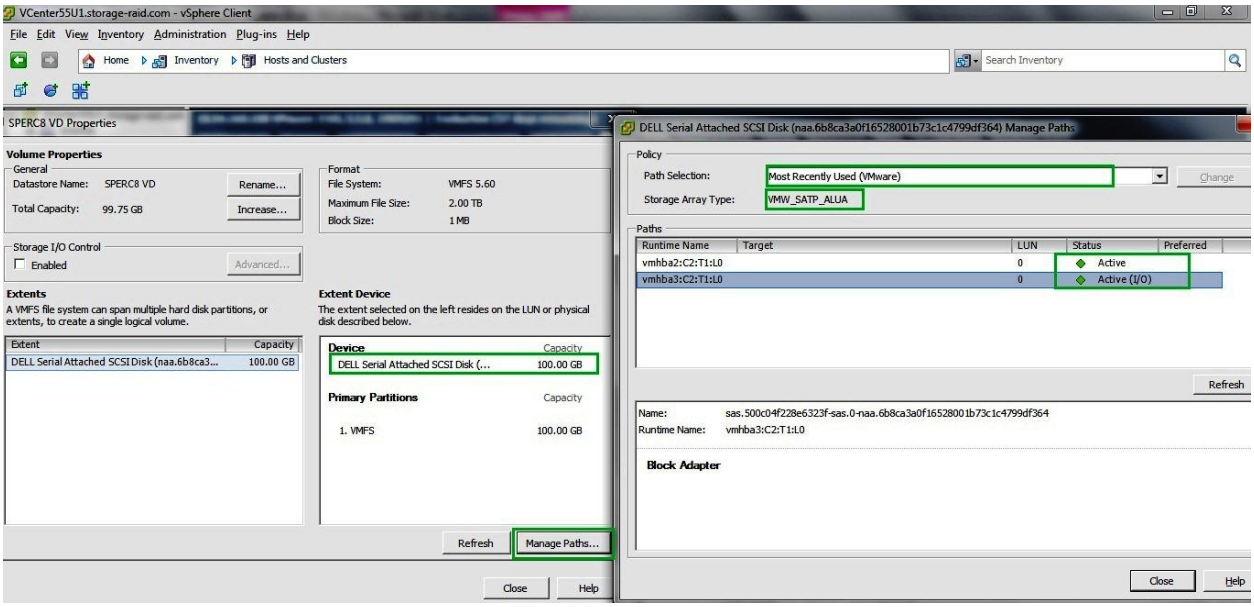

Figure 5 : Multipathing Policy and properties

Known issues

Refer to Dell VMware release notes to understand the known issues w.r.t PowerEdge VRTX and VMware ESXi.

Conclusion

With the Fault Tolerant capability of Shared PERC8 controllers, Dell PowerEdge VRTX provides seamless IO transition between Shared PERC8 controllers accessing the centralized storage array in the event of any one Shared PERC8 controller fails. When configured VRTX in Fault Tolerant mode as mentioned in the paper, VRTX can sustain single shared PERC8 controller failure along with the disk failures, which is being supported by different RAID levels. As described in the paper, the benefits of VRTX can be maximized in VMware ESXi environments with controller configured in Fault Tolerant mode.

References

- Dell tech center blog about VRTX Release
- VRTX Release Notes
- VRTX Setup Guide
- Shared PERC 8 User Guide
- VMware MultiPathing

Authors

Hypervisor Engineering	Storage Validation
Krishnaprasad K Kiran Devarapalli	Vishnu Murthy Kavi Chakkravarthy Neeraj Joshi