



VMware vSphere Virtual Volumes on Dell PS Series

A technical introduction to VMware vSphere Virtual Volumes as implemented on Dell PS Series storage

Dell Engineering
March 2015

Revisions

Date	Description
March 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.

© 2015 Dell Inc. All rights reserved. Dell, the DELL logo, and the DELL badge are trademarks of Dell Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others.



Table of contents

Revisions.....	2
Executive summary	4
1 Introduction to Virtual Volumes.....	5
2 Virtual Volumes or Virtual Machine File System.....	6
2.1 VMFS limitations.....	6
3 Why Virtual Volumes?	8
3.1 Virtual Volumes terminology	8
3.2 Types of Virtual Volumes.....	11
3.3 Coexisting with VMFS and migrating to Virtual Volumes	12
3.4 Changes in storage policy-based management	12
3.5 Licensing.....	13



Executive summary

The VMware vSphere® 6.0 Virtual Volumes (VVols) feature is a significant change from a volume-centric approach of providing services in a virtual environment to one where SANs are VM-aware. This alters how shared storage is utilized in virtualized environments, and therefore how virtualized environments are designed. VVol SAN based services are enabled for a more granular application at a per-virtual-machine basis, and in some cases at a per-virtual-disk basis. The result is a faster, more efficient, better-protected virtual environment, and ultimately a more scalable environment.

Virtual Volumes enable data protection for each virtual machine using SAN-based snapshots, more granular end-to-end monitoring of virtual machines and their virtual disks, more granular storage profiles, and faster deployment of virtual machines.

With support for Virtual Volumes, a SAN providing block storage to vSphere becomes aware of the individual virtual machine and the virtual disks associated with that virtual machine, enabling the storage to:

- Provide data protection in the form of hardware-based SAN snapshots
- Track detailed per-virtual-disk storage metrics from the perspective of the SAN
- Leverage the SAN existing volume clone capabilities to rapidly deploy new virtual machines



1 Introduction to Virtual Volumes

Enabled by the second generation VMware® APIs for Storage Awareness (VASA) provider included with Virtual Storage Manager 4.5, Virtual Volumes enables storage to be virtual machine aware, and for virtual machines to be first class objects in the SAN.

With Virtual Volumes, little changes in the day-to-day activities of a vSphere administrator; a virtual machine is still a virtual machine, and the workflows within VMware vCenter™ do not change. However, on the storage side, a virtual machine now consists of a collection of volumes on the SAN, rather than a collection of virtual machines sharing a single volume. This enables SAN volume manipulation capabilities to perform storage-related virtual machine tasks such as virtual machine cloning and virtual machine deployment.



2 Virtual Volumes or Virtual Machine File System

Virtual Machine File System (VMFS) and SAN volumes have been doing a satisfactory job for over a decade; why the need for change? Rather than accepting the satisfactory job that VMFS has been doing, we need to consider the limitation that VMFS has placed on the virtual environment, and the work-around or best practices that have come into play to smooth over these limitations.

2.1 VMFS limitations

- Cumbersome management of virtual machine to volume relationship

Queue depth, SCSI-2 conflicts, capacity, data-protection needs, and other variables all play a role in deciding which traditional datastore contains a virtual machine. However, obtaining this information is time consuming and complex. In an attempt to respond to an issue before a significant problem arises, rules of thumb are invoked and metrics are observed. Virtual Volumes move away from VMFS, enabling per-virtual-machine data protection, and leveraging singular large storage containers and always thin virtual disks to simplify virtual machine placement.

- No insight to a virtual machines true space and performance consumption

The performance and consumption statistics that the vSphere administrator and storage administrator see are different. While the vSphere administrator can see individual Virtual Machine Disk (VMDK) I/O metrics, the storage administrator can only see the I/O metrics for the entire volume where that VMDK resides (which will include the metrics of several other VMDKs). This difference in information can result in delayed diagnosing and resolution of a performance issue. With Virtual Volumes, both administrators have visibility to the same level and detail of I/O metrics in their preferred tool.

- All data services at a volume level

SAN storage array provides a number of data protection capabilities, such as snapshots. However, a SAN only preforms these on a per-volume basis. With traditional datastores, the vSphere administrator has to choose between short-lived but per-virtual-machine, VMware-based snapshots, or long life SAN-based snapshots that included all virtual machines on the volume. With Virtual Volumes, vSphere administrators gain the granularity of per-virtual-machine SAN-based snapshots.

- Need for over provisioning

When sizing datastores for capacity, space must be set aside to accommodate growth of thin provisioned VMDKs, vswap files, and VMware snapshot delta files. Even though this space may infrequently or never used, it must be set aside on each datastore. With the shift to Virtual Volumes, multiple traditional VMFS datastores are replaced with one large storage-container-backed datastore, resulting in this slack space being sharable across more virtual machines and less slack space needing to be maintained.



- Frequent data migrations

In previous versions of vSphere, VMware provided a datastore cluster feature. This feature groups similar datastores under one datastore cluster object. The advantage of this is that a vSphere administrator could deploy a virtual machine to the datastore cluster, and let vSphere place the virtual machine on a datastore that has sufficient capacity. However, as thin provisioned VMDKs or I/O latencies grow, virtual machines need to be migrated using vMotion to other datastores within the cluster. While this does provide a hands off approach to datastore capacity management, a storage penalty is gained on the SAN because the blocks on the source datastore are not automatically unmapped by vSphere. The blocks are marked as used by the SAN until a manual unmap operation is run.

- Wasted resources, wasted time, high costs

The above limitations require manual intervention from vSphere administrators, and in some instances direct cooperation with storage administrators. When these problems occur, it takes a long time to resolve them and consumes valuable administrator time. The flexibility and granularity of Virtual Volumes enables vSphere and storage administrators to deploy more efficient and dynamic environments, empowering them to meet the ever-changing demands of business with ease.



3 Why Virtual Volumes?

Without VVols, cloning a virtual machine or deploying a virtual machine from a template is a large-file copy operation. While the VMware vSphere Storage APIs – Array Integration (VAAI) Full Copy primitive provides acceleration of these tasks, these operations are even faster with VVols because they become a SAN volume clone operation. A volume clone operation (the manipulation of some block pointers and reserving of space on the SAN) is completed within a matter of seconds. The result for the vSphere administrator is the ability to have a new virtual machine in seconds rather than minutes.

Without VVols, VMware warns that virtual machine snapshots may decrease performance and recommends limiting their use to no more than 24 to 72 hours, with 2-3 delta files in a chain. With VVols the workflow remains unchanged, but the old delta file snapshots now become efficient pointer-based snapshots on the SAN. This results in rapid creation of snapshots that can be kept for an indefinite period, plus restoring a virtual machine from a snapshot becomes a rapid operation.

Note: While the Dell PS Series SAN firmware permits a volume to have 512 snapshots, the current vSphere VVol implementation is limited to 32 snapshots. Even with the limitation, this enables vSphere administrators to complement a current backup strategy with more frequent and rapidly restorable snapshots.

With VVols, a virtual machine is a collection of volumes on the SAN. This enables the existing EqualLogic SAN Headquarters (SAN HQ) performance monitoring tool to provide detailed I/O analysis on a per-virtual-machine and per-virtual-disk level. It is worth noting that the individual VM and VMDK performance metrics provided in vCenter are generated from the host side, and cannot normally show the impact of I/Os, latency, and block size on the underlying physical disks. However, SAN HQ is a powerful tool provided with each PS Series SAN (SAN HQ coupled with the EqualLogic vRealize Operation (formerly vCenter Operation) Manager Adapter) makes this detailed information available within vCenter Operation. This enables both the vSphere administrator and PS Series storage administrator to see the same information from their respective preferred interfaces.

Another benefit of VVols is that Storage Policy-Based Management (formerly Profile-Driven Storage) has become more granular in this implementation. Rather than creating a policy based on all the capabilities of an underlying volume, vSphere administrators can create storage policies based on the individual capabilities advertised by the pool of the storage container.

3.1 Virtual Volumes terminology

An understanding of the following VVols-relevant terminology is important while reading this paper.

- VASA Provider:
 - The VASA Provider plays an important role in enabling a VVol environment. The VASA Provider offers out-of-band management access to the SAN from vCenter. It enables vCenter to communicate with the SAN in ways that the current SCSI protocol does not. Through this



communication channel, vCenter sends operational requests for interacting with the virtual volumes that back virtual volume based virtual machines.

- The Dell EqualLogic VASA Provider shipped as part of the Virtual Storage Manager plugin for vCenter, which also provides enhanced storage management functionality to vCenter.

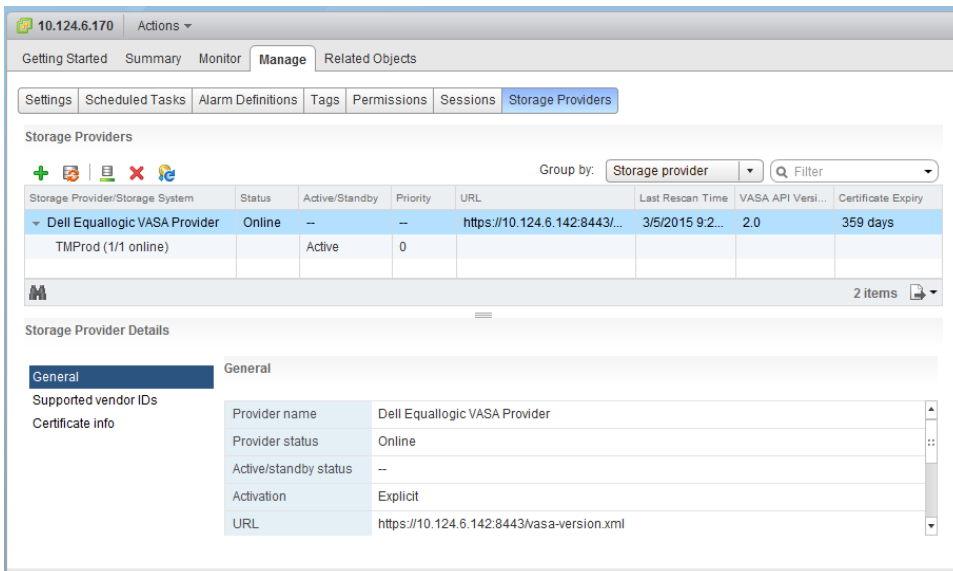


Figure 1 VASA Provider status as shown in vSphere Web Client

- Protocol Endpoint:

- The Protocol Endpoint is a unique volume on the SAN, it has a size of zero megabytes and a LUN ID of 256.
- It is the SAN endpoint of the communication between the ESXi host and the Virtual Volumes on the SAN. The transport protocol (iSCSI in the case of PS Series) endpoint is where the communication is turned over to internal SAN protocols. This unique volume can be thought of as a multiplexer LUN that acts as both the target and the initiator, and enables ESXi hosts to see a single volume while multiple independent volumes fan out behind it in the SAN. These multiple independent volumes are included in a virtual machine.
- The protocol endpoint is also where access controls are placed and initiators are queried to insure that they are permitted access to the storage container and virtual volumes. VSM manages these access controls directly from vCenter.



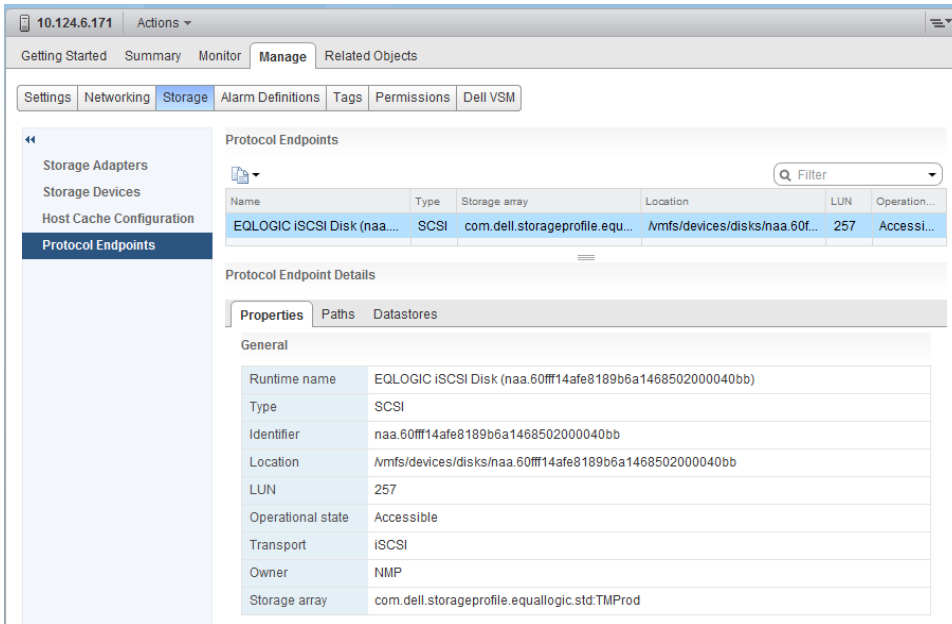


Figure 2 Protocol Endpoint as shown in vSphere Web Client

- Storage Container:
 - A Storage Container is reserved space on the SAN that can be increased and decreased as needs change, PS Series storage requires thick (or 100 percent) space reservation). It can also be conceptualized as a type of folder object on the SAN for organizing multiple volumes together.
 - Storage Containers are seen and treated as regular datastores by vSphere, and are referred to as a VVol-type datastore. They can be browsed as normal as may be done when seeking a virtual machine log files. The enable many vSphere workflows to remain unchanged even though significant changes have occurred.
 - Multiple storage containers can exist within a PS Series group, up to 32, but they cannot span pool or be migrated from one pool to another pool.




Figure 3 Datastores of type VMFS-5 and of type VVOL as shown in VSM plugin for vSphere Web Client

- Virtual Volumes:
 - At a high level, an individual virtual volume is a regular volume to the SAN and it can be manipulated by various SAN functions. However, from the vSphere perspective, an individual virtual volume is part of a virtual machine; a complete virtual machine consists of several virtual volumes of different types. The SAN is aware of the virtual volumes that belong to other



virtual volumes because of information communicated from vSphere through the VASA Provider.

- See section 3.2, "Types of Virtual Volumes" for details on the different types of virtual volumes that are in a virtual machine.

3.2 Types of Virtual Volumes

Traditionally, a virtual machine consisted of a VMX (configuration) file, one or more VMDK files, a VSWP (memory swap) file, log files, and other miscellaneous files. With Virtual Volumes, a virtual machine consists of a collection of virtual volumes on the SAN that consume space from the Storage Container space reservation. A VVol based virtual machine consists of the following types of virtual volumes:

- **Config:** This small VMFS-formatted thick provisioned 4 GB volume hosts the VMX file, log files, and other miscellaneous files.
- **Data:** The equivalent of a VMDK, one exists for each virtual disk attached to the virtual machine. Data virtual volumes, sometimes referred to as VMFS virtual volumes, are always thin provisioned.
- **Swap:** The equivalent of the VSWP file, it is created when the virtual machine is powered on. It is thick provisioned and is the same size as the memory assigned to the virtual machine, less any memory reservation.

When a VMware snapshot is taken of the virtual machine, vSphere offloads this operation to the SAN that creates a hardware-based snapshot. Two more virtual volume types that can exist in this case are:

- **Snapshot:** Created for each Data virtual volume associated with the virtual machine snapshot, and stores the delta of changes since the previous snapshot was taken. Snapshot virtual volumes are hidden in the PS Series interface, but the number of snapshots associated with a data virtual volume is indicated by the number in the snapshot column.
- **Memory:** Created if the option to include a memory dump with the snapshot is selected, it is equal to the size of the memory assigned to the virtual machine.



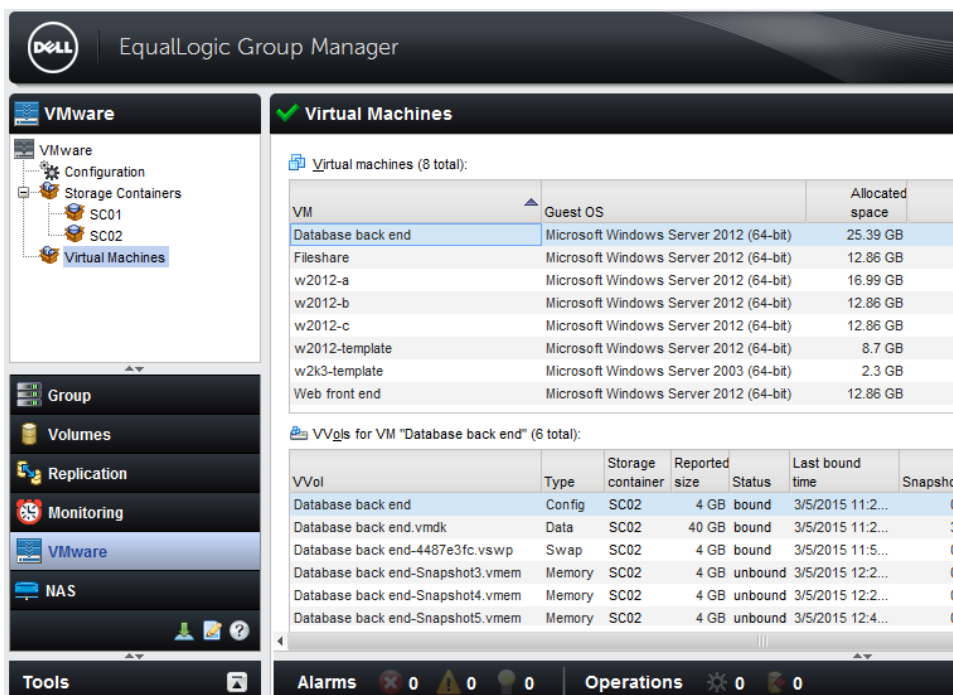


Figure 4 Virtual machines and virtual volumes as seen in EqualLogic Group Manager

3.3 Coexisting with VMFS and migrating to Virtual Volumes

Dell PS Series SANs can simultaneously provide both traditional volumes for VMFS datastores, and storage containers for virtual volume datastores. However, snapshots cannot be taken of a virtual machine that has virtual disks on both a VMFS datastore and a virtual volume datastore.

A Storage vMotion operation is used to migrate an existing virtual machine from a traditional VMFS datastore to a virtual volume datastore. This operation can be performed with the virtual machine in either running or powered off state. Virtual machines can also be migrated from virtual volume datastore to VMFS datastores. In both cases if the virtual machine has a snapshot, the snapshot is also be migrated.

Note: Traditional volumes and virtual volumes count equally against the PS Series group volume limit of 1024 volumes.

3.4 Changes in storage policy-based management

With VASA 1.0, a concatenated string of information about a volume was that the only thing available to storage vendors to advertise the capabilities of a volume to the vSphere administrator. This limited the number of profiles that could be created. With VASA 2.0, which is compatible with Virtual Volumes and Virtual SAN, the ability to advertise multiple individual capabilities about the storage exists. This enables the vSphere administrator to create granular storage policies that contain only the capabilities needed to meet their business needs.



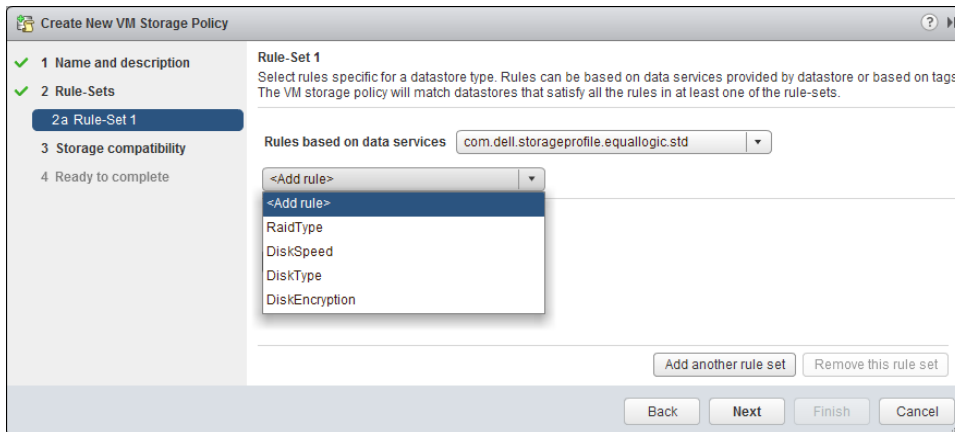


Figure 5 Storage Policy options advertised by PS Series

3.5 Licensing

The PS Series licensing policy is all-inclusive. As long as a customer is under a current support contract, they are entitled to features that are added to the array firmware or the host-side tools. Virtual Volumes requires Virtual Storage Manager 4.5 and PS Series firmware version 8.0, which is supported on PS Series hardware PS4000 and PS60x0 and above.

VMware has several tiers of licensing, and Virtual Volumes and Storage Policy-Based Management are available at the Standard, Enterprise, and Enterprise Plus levels. This change in licensing from VMware makes Virtual Volumes available to a much broader range of customers and environments; it significantly helps drive adoption of virtual volumes in the industry.