



# Dell PS Series Arrays: Advanced Storage Features in VMware vSphere

An introduction to VMware vStorage APIs for Array Integration, vStorage APIs for Storage Awareness, Storage APIs for Multipathing, and Storage I/O Control

Dell Storage Engineering  
June 2015

## Revisions

Version	Date	Description
1.0	March 2011	Initial release
1.1	April 2012	Updated to include new features in VMware vSphere 5.0 and Dell PS Series firmware 5.1
1.2	September 2012	Updated to include new features in PS Series firmware 6.0
1.3	October 2013	Updated to include changes to thin provisioning space reclaim command in vSphere 5.5
1.4	June 2015	Updated to include vSphere 6.0 changes

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.

© 2011–2015 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
1 Introduction.....	4
1.1 Audience.....	4
2 vStorage APIs for Array Integration .....	5
2.1 Full copy .....	5
2.2 Block zeroing.....	5
2.3 Hardware-assisted locking.....	6
2.4 Thin provision stun .....	6
2.5 Thin provisioning space reclamation .....	7
2.5.1 Space reclamation process on vSphere 5.0u1 and 5.1.....	7
2.5.2 Space reclamation process on vSphere 5.5.....	8
3 vStorage APIs for Multipathing .....	9
4 Storage I/O Control.....	10
5 vStorage APIs for Storage Awareness for vSphere 5.x .....	11
5.1 Storage-Based Policy Management.....	11
5.2 Storage Distributed Resource Scheduler.....	11
6 Introduction to Virtual Volumes for vSphere 6.0.....	13
6.1 Benefits of Virtual Volumes .....	13
6.2 Understanding the Virtual Volume storage changes.....	14
7 Summary .....	15
A Configuration details.....	16
A.1 VMware licensing .....	16
A.2 Software and firmware revisions.....	16
B Additional resources.....	17
B.1 Technical support and customer service .....	17
B.2 Dell online services .....	17
B.3 Dell PS Series storage solutions .....	17
B.4 Related documentation .....	17
B.5 Feedback .....	18



# 1 Introduction

With VMware® vSphere® virtualization, even though the hypervisor isolates the virtual machine operating system from the hardware, this does not mean that the hardware no longer matters. All components of the hardware environment are a shared resource in the virtualized data center, so having the right hardware is critical.

Thanks to cooperation between engineering teams at VMware and Dell, VMware vSphere is able to leverage the intelligence within Dell PS Series arrays to maximize the benefit to the virtualized environment. The release of VMware vSphere 4.1 enabled an intelligent awareness of storage, a capability that has been expanded even further with subsequent releases of vSphere. The advanced storage features included in vSphere enable faster deployment of virtual machines, increased densities of virtual machines per datastore, and increased bandwidth to the array.

This technical white paper provides an introduction to the storage-related features in VMware vSphere releases and discusses features facilitated by PS Series arrays to help you understand the benefits and advantages of running VMware vSphere on Dell PS Series storage systems.

## 1.1 Audience

The information in this guide is intended for administrators of VMware vSphere environment utilizing Dell PS Series iSCSI arrays.



## 2 vStorage APIs for Array Integration

Historically, VMware has treated all storage as just storage—whether local disk, DAS, SAN, or NAS. VMware did not utilize any of the advanced capabilities that exist in the more-intelligent storage products, but this changed with the release of vSphere 4.1. VMware now recognizes and uses intelligent storage arrays, such as the Dell PS Series, and can leverage the array capabilities to the benefit of the virtualized environment through a set of SCSI primitives. These SCSI primitives, called vStorage APIs for Array Integration (VAAI), are supported on all PS Series arrays running firmware 5.0 or above.

The current SCSI primitives included in the VAAI integration stack include:

- Full copy
- Block zeroing
- Hardware-assisted locking
- Thin provision stun (requires PS Series firmware 5.1 or above)
- Thin provisioning space reclamation (requires PS Series firmware 5.2 or above)

### 2.1 Full copy

A common day-to-day IT task involves deploying servers to support new business applications. Virtualization changed this from a labor-intensive task of racking a server and installing the OS to a simple task that required only a couple of mouse clicks to deploy a virtual machine from a preconfigured template. While this change has resulted in substantial time savings, there was still a significant amount of time spent watching the progress bar as the virtual machine deployed.

Traditionally, the process of deploying a virtual machine involved all of its data being read from the array, across the network to the VMware ESXi™ host, and then written back across the network to the array. This placed a non-production workload on both the network and the ESXi host, in addition to the production workload of the running environment. Now, with the full copy primitive, ESXi can offload this task to the PS Series array where it can be completed much more efficiently, with a significant workload reduction for the ESXi host and the network.

The benefits of full copy do not end with deploying virtual machines from templates. They also extend to virtual-machine tasks such as Storage vMotion, and virtual machine cloning. For more information, see the document, [Virtual Machine Protection with Dell EqualLogic Virtual Storage Manager v4.0](#).

### 2.2 Block zeroing

Fault-tolerant virtual machines require virtual machine disks (VMDKs) that are eager-zeroed thick. These differ from standard thick or thin VMDKs in that the blocks are zeroed out at the time that the VMDK is created. For large disks, this can take a significant amount of time as each zero is written from the server to the array, and an acknowledgment of each write is sent back from the array to the server.

With the block zeroing primitive, the ESXi host offloads to the PS Series array the task of zeroing out the blocks, to which the array immediately issues an acknowledgment, and permits the host to continue



creating the fault-tolerant virtual machine while the array completes the zeroing task in the background. By offloading the block zeroing to the PS Series array, fault-tolerant virtual machines can be created much faster.

## 2.3 Hardware-assisted locking

To protect Virtual Machine File System (VMFS) metadata, the hardware-assisted locking primitive provides a more granular method than SCSI reservations. Previously, whenever a virtual machine powered on, powered off, grew a thin provisioned virtual disk, or was moved with vMotion to another host, a SCSI reservation lock would be issued by the ESXi host to the underlying volume of the datastore. This prevented other hosts from also issuing a SCSI reservation to service a similar request. While SCSI reservations are short-lived, the impact can be noticed when powering on a large number of virtual machines simultaneously, as typically observed in a virtual desktop infrastructure (VDI) environment.

The hardware-assisted locking primitive resolves this by working with the PS Series SAN to lock only the necessary blocks rather than the entire volume. This enables other hosts to perform similar operations against that same volume at the same time.

## 2.4 Thin provision stun

The thin provision stun primitive is relevant for environments leveraging thin provisioning both on the virtual machine VMDK as well as on the datastore volume on the array. When both forms of thin provisioning are used together, the likelihood of an out-of-space condition, in which the thin provisioned volume has insufficient free space to support the growth required by a thin provisioned VMDK, is greatly increased.

When this occurs without the thin provision stun primitive, the volume is taken offline to preserve existing data and prevent further write attempts. This is disruptive and can be destructive to the virtual machines using the datastore. Because the underlying storage is unexpectedly inaccessible, eventually these virtual machines will crash.

When this shortage of space occurs on a volume configured with the thin provision stun primitive, the virtual machine that is attempting to grow its VMDK is *stunned* (but not taken offline), and the VMware vCenter™ administrator is notified. By enabling the array to communicate through the SCSI storage stack with the initiator, the array can inform vSphere when the utilization of volume space passes an in-use warning limit threshold, set in the PS Series interface. This generates a datastore usage on disk alert within vCenter in addition to the traditional warning on the array. Should the volume reach the second warning mark within the PS Series interface (the maximum in-use space limit), or if there is insufficient free space to permit the volume to grow, the array informs the initiator of the out-of-space condition rather than take the volume offline.

vCenter then stuns the individual virtual machine that is attempting to grow its thin provisioned VMDK, and display a warning message to the vCenter administrator. Only the virtual machines that are attempting to grow their thin provisioned VMDKs are stunned; the other virtual machines on the datastore are unaffected and continue to operate. This process requires the administrator to allocate additional space to



the volume, migrate virtual machines off the volume, or delete virtual machines from the volume. Once space on the datastore has been made available, the stunned virtual machine can be un-stunned, and it will continue operating again under normal conditions.

**Note:** VAAI thin provisioned stun requires Dell PS Series firmware 5.1 or above.

## 2.5 Thin provisioning space reclamation

The thin provisioning space reclamation primitive enables thin provisioned datastores to be manually re-thinned to only consume the actual space they are using on a PS Series array—allowing thin provisioned volumes to remain thin. Traditionally, the size of a thin provisioned volume, as shown at the storage layer, reflects the maximum space consumption that occurred at some point since it was created. This is because the operating system could not inform the array that particular blocks of data were no longer required by the filesystem and were no longer needed to be stored by the array.

When the T10 SCSI committee added the unmap specification standard to the SCSI Block Commands 3 (SBC3), this allowed operating systems to communicate to the array, through the SCSI storage stack, that blocks of data no longer needed to be stored. This unmap ability is referred to as thin provisioning space reclamation by VMware.

With vSphere 5.0 update 1 and above, thin provisioning space reclamation has become available to vSphere administrators as a command-line tool, enabling them to re-thin their thin provisioned datastores so that they are only consuming the space that they require. This can improve efficiency of storage capacity and lower the total cost of ownership of PS Series arrays.

**Note:** VMware recommends limiting thin provisioning space reclamation operations to maintenance windows, as this can be a storage-intensive task on some arrays. Dell PS Series firmware manages unmap operations as additional I/O, so unmap operations should be avoided during periods of peak I/O.

### 2.5.1 Space reclamation process on vSphere 5.0u1 and 5.1

To re-thin a thin provisioned datastore and reduce the amount of disk space consumed on an array, a vSphere administrator must run the following command from the ESXi CLI from within the particular datastore:

```
vmkfstools -y <percentage of deleted blocks to be reclaimed>
```

The percentage of deleted blocks to be reclaimed is a percentage of the free space of the datastore, and not a percentage of the total size of the datastore. VMware suggests, but does not enforce, that a maximum value of 60 percent be used. During the space reclamation operation, a temporary balloon file is created using available storage on the datastore. If a large value such as 100 percent is specified, the balloon file would consume all available space. This would prevent virtual machines located on that datastore from powering on because the datastore would be unable to create its vswap file and a virtual machine would not be able to grow a thin provisioned VMDK.



To calculate the percentage of deleted blocks to be reclaimed, it is necessary to know the amount of free space on the datastore as seen from vCenter, and the amount of free space in the volume as seen on the array interface.

$$\frac{vCenter \text{ datastore free space} - \text{array volume free space}}{vCenter \text{ datastore free space}} * 100$$

In an example with a 1000GB thinly provisioned volume, the array shows 200GB free and vCenter shows the datastore as having 300GB free.

$$\frac{300GB - 200GB}{300GB} * 100 = \frac{100GB}{300GB} * 100 = 33.33\%$$

## 2.5.2 Space reclamation process on vSphere 5.5

With the release of vSphere 5.5, VMware changed the command for performing the unmap operation and the behavior of the unmap command. The command is now a part of the existing esxcli command set, which enables the command to be accessed from many scripting tools and to be called remotely with vSphere vMA or vCLI. In addition to this new capability, vSphere administrators can specify which datastore to unmap and can use the operation to unmap the entire datastore, making it easier to keep thinly provisioned datastores thin.

The full syntax of the command is:

```
esxcli -server=server_name storage vmfs unmap --volume-label=volume_label | --  
volume-uuid=volume_uuid [--reclaim-unit=number]
```

To specify the datastore to be unmapped, use either `--volume-label` or `--volume-uuid`. The optional parameter, `--reclaim-unit`, has a default value of 200 units and specifies the size of the temporary balloon file that is created per iteration with a default of 200 1MB blocks. VMware will continue to create temporary balloon files of the specified size until it has reached the end of the datastore. Dell recommends using the default value for the `--reclaim-unit` command.

**Note:** VAAI thin provisioning space reclamation requires Dell PS Series firmware 6.0 or above.





## 3 vStorage APIs for Multipathing

VMware vSphere offers many new and advanced enhancements to the software iSCSI initiator beyond basic iSCSI SAN connectivity. The most significant of these enhancements is the API support for third-party multipathing plugins. This provides a framework that enables the PS Series Multipath Extension Module (MEM) to build upon VMware round-robin multipathing, and to more intelligently route and efficiently load balance iSCSI traffic across multiple NICs.

The PS Series MEM offers:

- Ease of install
- Increased bandwidth
- Reduced network latency
- Automatic load balancing across multiple active paths
- Automatic connection management
- Automatic failure detection and failover
- Multiple connections to a single iSCSI target

The PS Series MEM will automatically create iSCSI sessions to all members that a volume spans. As the storage environment changes, the MEM will respond by automatically adding or removing iSCSI sessions as needed.

As storage I/O requests are generated on the ESXi hosts, the MEM will intelligently route these request to the array member best suited to handle the request, resulting in efficient, non-statistical load balancing of the iSCSI storage traffic.

For more information about the PS Series MEM, read [Configuring and Installing the EqualLogic Multipathing Extension Module for VMware vSphere and PS Series SANs](#).



## 4 Storage I/O Control

Storage I/O Control (SIOC) ensures that the excessive storage I/O demands of a particular VMDK do not negatively impact the storage I/O needs of other VMDKs residing on the same datastore. Previously, this has been resolved through administrative tasks such as careful VM placement, reactive monitoring of VMDK I/O, and over-sizing of the environment to handle occasional I/O spikes.

With SIOC, the reactive monitoring task is conducted by vSphere across all ESX/ESXi hosts and the reactive action is performed automatically and instantaneously by vSphere, enabling administrators to more efficiently utilize their storage environments.

The advantages of using SIOC include:

**Performance protection:** SIOC ensures that all VMDKs receive a fair share or assigned share of I/O needs regardless of the I/O they demand during period of congestion.

**Better utilization of storage assets:** The storage environment no longer needs to be oversized to cover occasional I/O peaks. Rather, these peaks are leveled out by SIOC.

Offering two deployment options, SIOC can provide equal treatment to all VMDKs sharing the datastore, or it can favor or disfavor certain VMDKs through the use of disk shares. Regardless of the granularity of implementation, SIOC operates in the same way, but the impact to VMDKs with excessive I/O consumption may differ.

SIOC works by monitoring the I/O latency of a datastore. When that latency exceeds the threshold that has been set (30ms by default), SIOC will engage and enforce the assigned disk shares. SIOC achieves this by restricting the number of queue slots available to the VMDKs that are consuming more than their assigned share, and provides the previously deprived VMDKs with improved storage performance.

While SIOC does not eliminate the need for SAN monitoring, it means that the SAN does not need to be actively monitored, freeing up the storage administrator to deal with more important tasks. If SIOC is engaging for significant periods of time, the administrator may need to add additional I/O capacity or relocate I/O-intensive VMDKs.



## 5 vStorage APIs for Storage Awareness for vSphere 5.x

With the release of vSphere 5.0, VMware added vStorage APIs for Storage Awareness (VASA) to its set of APIs. VASA enables vCenter to communicate with the PS Series array, through the PS Series VASA Provider, to discern the capabilities of each volume presented to the vSphere environment. This enables the vCenter client to display pertinent information about a datastore, such as the ability of the underlying volume to perform snapshot or replication data protection, the RAID type such as heterogeneous or homogeneous, or the type of drives such as SSD. This enables vSphere administrators to be more aware of the storage capabilities in the virtual environment and allows administrators to better match virtual machine storage needs with the correct datastore. Without this feature, administrators must track this information manually. This storage capability information is displayed in the Datastores and Datastore Clusters view, and also can be viewed from the PS Series Datastore Manager plugin.

While the basic capability of VASA does not require a particular vSphere license, Enterprise Plus licensing is required by two vSphere features that leverage this information from VASA: Profile-Driven Storage and Storage Distributed Resource Scheduler. See Table 1 for more information on vSphere licensing requirements.

### 5.1 Storage-Based Policy Management

Storage-Based Policy Management (formerly Profile-Driven Storage) uses VASA-provided information to ensure that virtual machines reside on the appropriate datastores. An administrator can create various storage profiles that reflect particular data-protection needs. When a virtual machine is created, the administrator can select the storage profile that meets the requirements of the virtual machine, and place the virtual machine on datastore that provides these capabilities. If the virtual machine is later migrated to a datastore that does not meet the storage profile, or if the capabilities of the datastore change and no longer meet the storage profile, the virtual machine will fail its storage profile compliance check. The compliance status can be seen on the summary page for the individual virtual machine and on the VM storage profile for all virtual machines assigned a storage profile. Storage-Based Policy Management through the PS Series VASA Provider helps administrators quickly ensure that virtual machines are placed on the right datastore and continue to reside on datastore that meet their needs.

### 5.2 Storage Distributed Resource Scheduler

Storage Distributed Resource Scheduler (Storage DRS) takes the VMware CPU and memory resource management concepts and applies them to datastores. Like the traditional DRS capability of VMware, Storage DRS groups datastores with similar performance characteristics into a datastore cluster. When a virtual machine is deployed, it is not deployed to a particular datastore, but rather to a datastore *cluster*. Storage DRS determines on which datastore to place the virtual machine, based on space utilization and I/O load. Like traditional DRS, Storage DRS continuously monitors the cluster space utilization and also monitors the I/O load using Storage I/O Control. If thresholds for space utilization or I/O response time are exceeded, or if there is a significant difference in space utilization between the datastores within the datastore cluster, Storage DRS attempts to relocate a virtual machine using Storage vMotion.



While Storage DRS is aware of the datastores involved (and also the volumes) it does not identify where these volumes exist in the PS Series storage system. Prior to initiating a Storage vMotion action on a virtual machine, Storage DRS consults with the VASA Provider and queries to find out whether the migration of the virtual machine and its workload would benefit the overall I/O workload distribution of the PS Series array. If the migration will not result in an improvement in the distribution of the I/O workload (for example, if the volumes involved reside on the same PS Series group members), the VASA Provider will inform Storage DRS not to perform the migration. Conversely, if the VASA Provider agrees that the migration will result in an improvement in the distribution of I/O (for example, if the volumes involved reside on different PS Series group members), the provider will approve of the migration request. In this case, Storage DRS leverages Storage vMotion to move the virtual machine and its I/O workload to the selected datastore.

In another parallel to the classic VMware DRS feature, Storage DRS has the concept of *maintenance mode*. When a datastore in a datastore cluster is placed in maintenance mode, the virtual machines and VMDKs residing on the datastore are moved to other datastores within the datastore cluster through Storage vMotion. Storage DRS will ensure that the I/O workload and space utilization remains balanced across the remaining datastores not in maintenance mode.

Storage DRS also has a placement-constraint rule that is enforced during migrations. The default option is the *intra-VM VMDK affinity rule*, which keeps specific VMDKs of a virtual machine together on the same datastore. The inverse of that rule, the *VMDK anti-affinity rule*, keeps specific VMDKs on separate datastores within the datastore cluster. Finally, there is the *VM anti-affinity rule* which prevents certain virtual machines from sharing the same datastore.



## 6 Introduction to Virtual Volumes for vSphere 6.0

VMware vSphere 6.0 introduces Virtual Volumes, a significant change in how storage is utilized in a virtualized environment. Enabled by the second-generation VASA Provider included with Virtual Storage Manager 4.5, this feature enables storage to be virtual-machine aware and for virtual machines to be first-class citizens in the storage array.

With Virtual Volumes, the day-to-day activities of a vSphere administrator change very little. A virtual machine is still a virtual machine, and the work flows within vCenter do not change. However, there are significant changes on the storage side. As described in the following sections, a virtual machine now consists of a group of volumes on the array. This moves a number of storage-centric tasks into the domain of the array.

### 6.1 Benefits of Virtual Volumes

Without Virtual Volumes, cloning a virtual machine or deploying a virtual machine from a template involves a large file copy operation, which can be accelerated with the VAAI primitive full copy. With Virtual Volumes, this becomes a volume clone operation, which is completed within a matter of seconds because it only involves manipulating some block pointers and reserving space. For the vSphere administrator, this means that creating a new virtual machine takes only seconds rather than minutes.

VMware recommends using traditional virtual machine snapshots for no more than 72 hours with 2-3 delta files in a chain, and warns that the snapshots may decrease performance. The workflow remains unchanged, but the old delta file snapshots now become efficient, pointer-based snapshots on the array. This results in rapid creation of snapshots that enable quick restores and can be kept for an indefinite period of time. While the array firmware permits a volume to have 512 snapshots, the current vSphere Virtual Volume implementation is limited to 32 snapshots. Even with this limitation, this enables vSphere administrators to complement their current backup strategy with more-frequent and rapidly restorable snapshots.

With Virtual Volumes, a virtual machine is now a group of volumes on the array. This enables the existing PS Series SAN Headquarters (SAN HQ) array-performance monitoring tool to provide detailed I/O analysis on a per-virtual-machine and per-virtual-disk level. While similar performance metrics can be seen in vCenter, these are generated from the host side and cannot show the same level and detail that can be seen on the array side, including the impact of I/Os, latency, and block side on the underlying physical disks. When coupled with the PS Series vRealize Operations (formerly vCenter Operations) Manager Adaptor, SAN HQ makes this detailed information available within vCenter Operation. This allows both the vSphere administrator and PS Series array administrator to see the same information from their preferred interfaces.

Storage-Based Policy Management, which was introduced in vSphere 5.0, has been updated to also provide additional granularity. With VASA 1.0, storage providers were limited to providing a single string of information about volume capabilities, around which a vSphere administrator created a storage profile. With VASA 2.0, this enabled vSphere administrators to select from the advertised capabilities of the storage



container, and its underlying physical storage, to be used in creating a particular profile. This allowed the vSphere administrator to create storage profiles that meet their particular business needs.

## 6.2 Understanding the Virtual Volume storage changes

Traditionally, when deploying storage to a vSphere environment, a volume is created on the array, and this becomes a datastore within vCenter upon which virtual machines are placed. With Virtual Volumes, some of this remains the same: Virtual machines are placed within datastores, and all work flows that depend upon this remain unchanged. What has changed is what is backing the datastore. Backing the datastore is a new storage object known as a storage container, and while it does contain virtual machines and backs a datastore, it should not be thought of as a volume. The storage container is more like a reservation of space on the array. It is from this reservation that virtual machines will consume their storage needs, with the different Virtual Volumes potentially consuming different performance capabilities and different data-protection options. With Virtual Volume storage containers, it is possible to have just a single storage container for the entire environment because the size of a storage container is limited only by the size of the pool on the PS Series group. However, non-technical issues, such as a preference for keeping virtual machines separate from different departments or different projects, may drive a preference for multiple storage containers from a single pool. In an environment with a PS Series group that has multiple pools, multiple storage containers will be needed because storage containers do not span pools.

Typically, a virtual machine has consisted of a VMX file (configuration file), one or more VMDK files (virtual disk), a VSWP file (memory swap file), and other miscellaneous files including log files. This changes with Virtual Volumes, since a virtual machine now consists of a group of volumes on the array consuming space from the storage container space reservation.

A Virtual-Volume-based virtual machine consists of the following types of virtual volumes:

- Config: A small VMFS-formatted 4Gb volume that hosts the VMX and other miscellaneous files including log files
- Data: The equivalent of a VMDK, with one existing for each virtual disk attached to the virtual machine
- Swap: The equivalent of the VSWP file, existing only when the virtual machine is powered on

If a VMware snapshot of the virtual machine is taken, which is now an array-based snapshot, two more virtual volume types may exist:

- Snapshot: A hidden virtual volume that exists for each data virtual volume included in each virtual machine snapshot that stores the delta of changes since the previous snapshot was taken
- Memory: A virtual volume that is created if selecting the option to include a memory dump with the snapshot



## 7 Summary

Combined with Dell PS Series storage, VMware vSphere can significantly accelerate and automate the virtualized data center. Advanced vSphere storage features leverage the intelligence that exists in PS Series arrays, enabling more-efficient use of the server and networking infrastructure and more-efficient use of IT staff administering these environments. A strong partnership between Dell and VMware assures that future products will continue to build upon these features and bring further efficiencies to the data center.



## A Configuration details

### A.1 VMware licensing

The storage features discussed in this document are dependent upon the vSphere licensing requirements listed in Table 1.

Table 1 vSphere licensing requirements

vSphere storage feature	vSphere license required
vStorage APIs for Array Integration (VAAI)	Enterprise and above
vStorage APIs for Multipathing	Enterprise and above
Storage I/O Control (SIOC)	Enterprise Plus
vStorage APIs for Storage Awareness (VASA)	None
VASA-dependent features in vSphere 5.x: Storage DRS and Profile-Driven Storage	Enterprise Plus
VASA-dependent features in vSphere 6.0: Virtual Volumes and Storage-Based Policy Management	Standard and above

**Note:** For vSphere 6.0, vStorage APIs for Array Integration (VAAI) and vStorage APIs for Multipathing are available at Standard and above licensing.

For additional information on VMware licensing and vSphere editions, see:

<http://www.vmware.com/products/vsphere/compare>

### A.2 Software and firmware revisions

Table 2 lists the software and firmware revisions used for the preparation of this document.

Table 2 Software and firmware revisions

Vendor	Model	Revision
Dell	PS Series SAN	5.0.2, 5.1.x, 5.2.x, 6.x, 7.x, and 8.0.x
Dell	PS Series Multipathing Extension Module for VMware vSphere	1.2 and 1.3
VMware	vSphere ESXi and vSphere vCenter Server	4.1, 5.x, and 6.0





## B Additional resources

### B.1 Technical support and customer service

Offering online and telephone-based support and service options, Dell support service can answer your questions about PS Series arrays, groups, volumes, array software, and host software. Availability varies by country and product, and some services might not be available in your area.

Visit [Dell.com/support](http://Dell.com/support) or call 800-945-3355 (United States and Canada).

For international support of Dell PS Series products, visit <http://www.dell.com/support/contents/us/en/555/article/Product-Support/Dell-Subsidiaries/equallogic>

**Note:** If you do not have access to an Internet connection, contact information is printed on your invoice, packing slip, bill, or Dell product catalog.

### B.2 Dell online services

Learn more about Dell products and services using this procedure:

1. Visit [Dell.com](http://Dell.com) or the URL specified in any Dell product information.
2. Use the locale menu or click on the link that specifies your country or region.

### B.3 Dell PS Series storage solutions

To learn more about current and upcoming Dell PS Series solutions, visit the Dell TechCenter site: <http://delltechcenter.com/page/EqualLogic>. Here you can find articles, demos, online discussions, technical documentation, and more details about the PS Series product family.

### B.4 Related documentation

Table 3 lists the referenced or recommended documents related to this white paper.

Table 3 Referenced or recommended documents

Vendor	Document
Dell	<a href="#">Configuring and Installing the EqualLogic Multipathing Extension Module for VMware vSphere and PS Series SANs</a>
Dell	<a href="#">Dell EqualLogic Configuration Guide</a>
Dell	<a href="#">Dell PS Series compatibility list</a> <b>Note:</b> Requires login to eqsupport.dell.com account.



VMware	<a href="#">Virtual Machine Protection with Dell EqualLogic Virtual Storage Manager v4.0</a>
VMware	<a href="#">Storage I/O Control Technical Overview and Considerations for Deployment</a>
VMware	<a href="#">What's New in VMware vSphere 5.1 – Storage</a>

## B.5 Feedback

We encourage readers of this publication to provide feedback on the quality and usefulness of this information by sending an email to [StorageSolutionsFeedback@Dell.com](mailto:StorageSolutionsFeedback@Dell.com).

