

# NetApp Unlocks the Power of VMware VVOLs



**Graham Smith**  
Senior Product Manager



**Nick Howell**  
Virtualization Solutions Architect

While there's no question that virtual server infrastructure has had tremendous benefits for the modern data center, it has introduced new challenges where data storage is concerned.

- It is difficult to directly manage the data storage needs of individual virtualized applications (without third-party add-ons).
- Administrators unfamiliar with storage often have to know a lot about the underlying data storage systems to satisfy specific application performance and availability requirements.

These factors can prevent even the most seasoned administrators from unlocking the potential of their storage infrastructure to support virtualization. VMware is introducing Virtual Volumes (VVOLs, currently part of the vSphere Beta) to address these limitations. This article looks at VVOL architecture in detail, and describes how NetApp integrates policy-based management with VVOLs to deliver the full functionality of the clustered Data ONTAP® storage operating system—including deduplication, Snapshot™ copies, replication and cloning—to VMware users.

The latest version of our VASA Provider is currently in beta along with the next version of vSphere, showcasing our tight integration with VVOLs. Don't delay; the time to learn about the biggest change in virtual storage since VMFS-3 is NOW!

## What Problems Do VVOLs Address?

---

Before virtualization, it was fairly simple to tailor your storage—and your storage policies—for each application. With virtualization came the abstraction of the datastore, which became the de facto level of granularity at which storage is managed.

All application virtual disks (also known as vDisks or VMDKs) in a traditional datastore are treated and managed the same way. Because an ESXi host cannot be certain of the features that the underlying storage can provide to the overall virtualization experience, it often does a lot of storage-related work, when that would be better offloaded to the storage system itself. (This can be accomplished using the storage-offload standard known as vStorage APIs for Array Integration, or VAAI.)

---

## Explore

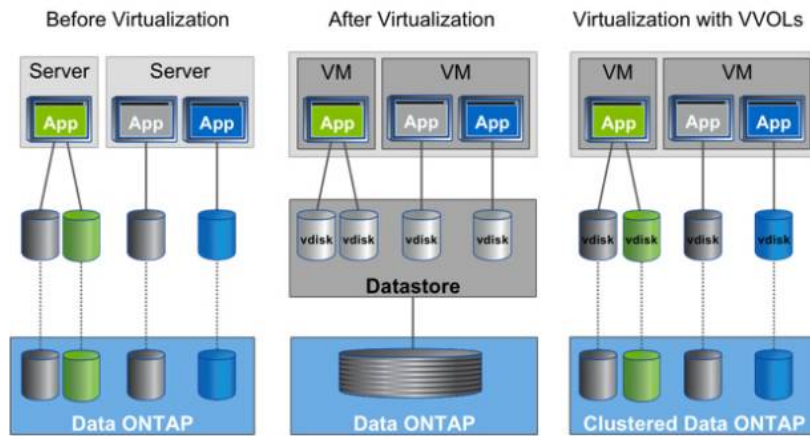
### Check Out VVOLs on NetApp at VMworld 2014

If you're attending VMworld this year, don't miss these opportunities to try out VVOLs running on NetApp storage first hand:

- Hands-On-Labs — HOL-PRT-1471 - Policy-Based Mgmt & VVOLs Tech Preview
- Live Demos in the NetApp booth in the Solutions Exchange
- Technical Staff and Developers in attendance to answer any and all questions

Get full details on all NetApp's VMworld activities [here](#).

---



**Figure 1)** Server virtualization introduced the concept of the datastore, which limited storage flexibility. VVOLs provide granular storage management for greater flexibility plus automated storage placement.

VMware and application administrators had to become increasingly familiar with the details of the underlying storage systems—what protocols were used, what the backup schedule was, whether or not replication was occurring—in order to make sure that application needs were being satisfied.

VVOLs address these issues in two ways: by allowing individual virtual disks to be managed independently of one another, or of the virtual machine entirely, and by introducing much greater automation in the process of matching VM storage requirements to underlying storage system functionality.

VVOLs achieve this through:

- Policy-based provisioning and management of datastores and virtual disks
- Ability to access VVOL datastores through SAN and NAS on an equal footing
- Separation of control and data paths for out-of-band management of storage resources

VVOLs connect VMware and storage administrators with a common platform and, in the process, provide the underpinnings for adopting a software-defined storage (SDS) approach in the data center.

## Understanding VVOLs

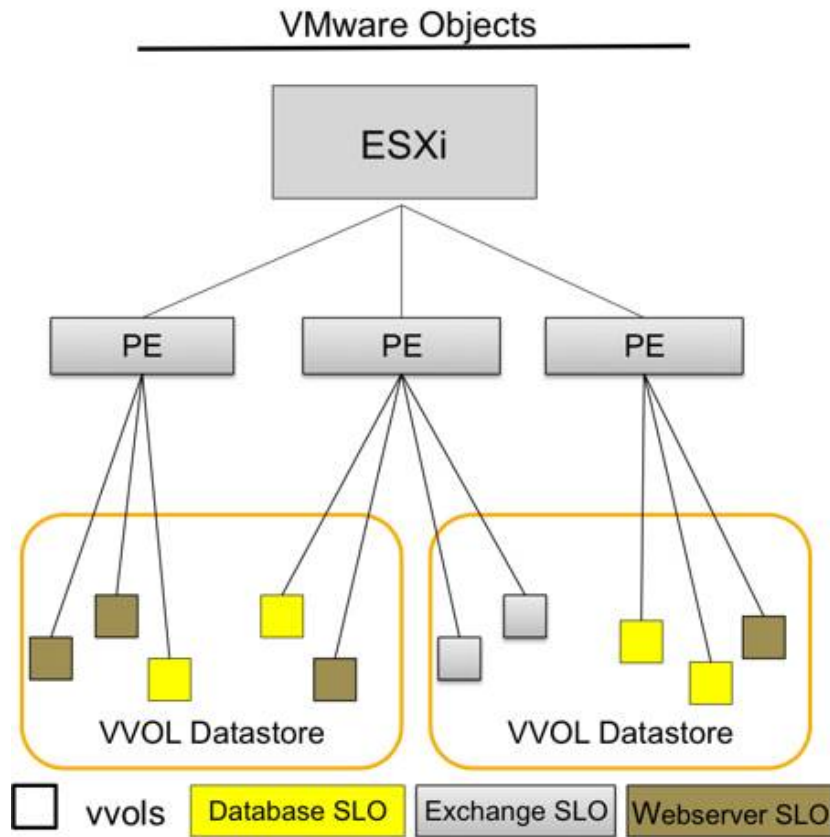
---

The VVOL architecture retains the idea of a datastore, providing a new methodology for managing storage while maintaining familiarity and compatibility with previous implementations for users. However, the virtual disks in a VVOL datastore can have different service level objectives (SLOs) and can be managed independently. The VVOL datastore is decoupled from the underlying storage and exposes all the underlying storage capabilities for use with any virtual disk it contains.

### Protocol Endpoints

ESXi servers connect to the virtual disks in a VVOL datastore through a Protocol Endpoint (PE), whose operation depends on the storage protocol being used. (NFS, iSCSI, Fibre Channel and FCoE are all fully supported.) For NFS, the PE is simply an NFS mount point, and the virtual disks are files beneath that mount point.

For SAN protocols, the PE is a proxy LUN. [vSphere supports only 256 direct LUN connections per ESXi host](#); the proxy LUN acts as a multiplexor, allowing each ESXi host to address thousands of underlying LUNs, each with a unique identifier. Each LUN corresponds to a virtual disk.

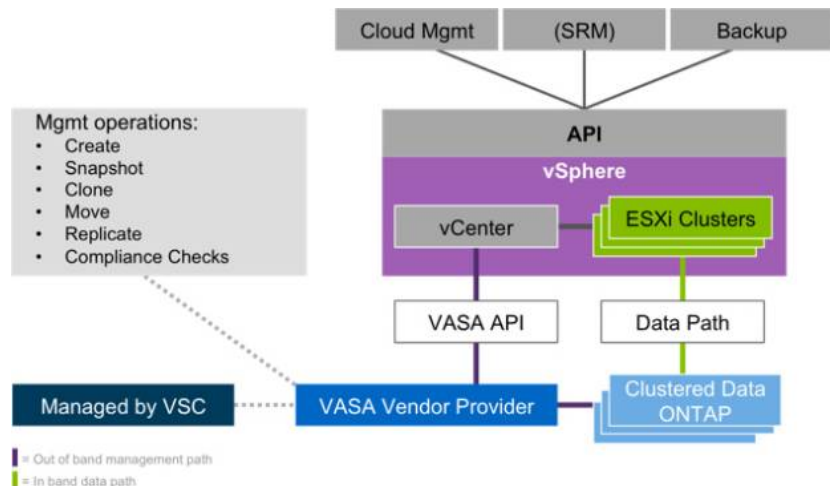


**Figure 2)** The VVOL architecture uses Protocol Endpoints (PEs) to provide uniform access to both SAN and NAS storage.

As a result, virtual disk provisioning is no longer tied explicitly to one protocol or the other. The VVOL datastore is mounted as an abstract entity and traditional datastores are now just volumes inside that logical container.

### Separate Control and Data Paths

On the management side, the control path and data path are logically separated and vSphere APIs for Storage Awareness (VASA) are used to provide policy-based management. Storage systems use a separate VASA Vendor Provider that advertises available storage capabilities to vCenter using the VASA API. Provisioning requests are automatically matched to the best underlying storage to satisfy the stated SLOs.



**Figure 3)** The VVOL architecture has separate data and control paths.

For example, as shown in figure 3, clustered Data ONTAP has a VASA Vendor Provider that is managed by the

NetApp Virtual Storage Console (VSC)—NetApp's proven management plugin for vCenter. Through the Vendor Provider, clustered Data ONTAP can make available all its capabilities, including various types of media, various storage protocols, and data management functions such as create, Snapshot, clone, replicate and so on. You simply create a VVOL datastore through the vCenter interface, add storage volumes, and create rules in VM Storage Profiles that describe the available storage capabilities. The workflows will be familiar to those who have worked with traditional datastores, simplifying VVOL adoption.

When new storage is needed, VMs are provisioned based on VM Storage Profiles, so the requestor only needs to know what capabilities they need. VM Storage Profiles can be created that implement specific policies. For instance, you might have the familiar Bronze, Silver, and Gold policies, or policies with meaningful names such as Exchange, Database, or Webserver. Policies are defined ahead of time by an architecture or application expert.

The VMware administrator just needs to pick the right policy, they don't need to know anything about the underlying physical storage systems. Storage can be provisioned either directly through vCenter or through cloud management tools. vSphere automatically provisions storage resources from the most appropriate VVOL datastore when a request is received.

## NetApp and VVOLs

---

At NetApp, we recognized the need for more granular storage management for virtual machines early on and have been moving in that direction for a number years. [VSC 5.0](#), released in April 2014, was a big step in that direction, and we're excited about the new enhancements that VVOLs make possible.

We've made a big commitment to provide support for VVOLs as quickly as possible. NetApp is a Design and Reference Partner for VVOL technology, which makes our implementation the standard. We are also the only vendor to support both NFS and SAN protocols from a single platform.

We currently have a working solution in beta that implements all of the capabilities shown in Table 1 for both SAN and NFS.

**Table 1)** NetApp VM management operations for both traditional datastores and VVOLs

| Feature                     | Traditional  |         | VVOLS           |                 |
|-----------------------------|--------------|---------|-----------------|-----------------|
|                             | SAN          | NFS     | SAN             | NFS             |
| vDisk Cloning (offloaded)   | Slow         | Fast    | Fast            | Fast            |
| Datastore Snapshots         | Yes          | Yes     | Yes             | Yes             |
| vDisk Snapshots (offloaded) | No           | Limited | Yes             | Yes             |
| Datastore QoS               | Yes          | Yes     | Yes             | Yes             |
| vDisk QoS                   | No           | Yes     | Yes             | Yes             |
| Datastore Live move         | Yes          | Yes     | Yes             | Yes             |
| vDisk Live move             | ONTAP (node) | Host    | ONTAP (cluster) | ONTAP (cluster) |
| Datastore Replication       | Yes          | Yes     | Yes             | Yes             |
|                             |              |         |                 |                 |

|                            |           |           |        |        |
|----------------------------|-----------|-----------|--------|--------|
| vDisk Replication          | Host      | Future    | Future | Future |
| vDisk SLO compliance check | Yes       | Yes       | Yes    | Yes    |
| vDisk SLO compliance solve | Manual    | Manual    | Auto   | Auto   |
| vDisk placement            | Semi-Auto | Semi-Auto | Auto   | Auto   |

## Monitoring and Correcting VVOL Compliance

Note that the NetApp VVOL implementation includes both **SLO compliance check** and **SLO compliance solve**. Compliance check verifies that virtual disks meet the mandated SLO. Compliance solve takes that a step further and automatically corrects compliance problems whenever possible. Remediation steps can include things like changes to QoS settings, increases in replication frequency, or movement of a VVOL to another location.

## Participating in the NetApp Beta Test for VVOLs

NetApp is currently the only storage vendor running an open beta test for VVOLs alongside the current VMware vSphere Beta. If you're a NetApp customer (or partner) you can sign up and begin testing and familiarizing yourself with this new technology right away.

To obtain the required versions of both vSphere and vCenter, you must sign up for the [VMware vSphere Beta program](#).

To utilize VVOLs on NetApp storage systems, you also need to [download two pieces of NetApp software](#):

- VSC 5.0.1 Beta
- VASA Provider 6.0 for clustered Data ONTAP Beta

Here are a few quick tips to help get you started:

- You will need clustered Data ONTAP 8.2.1 or later. [Simulate ONTAP](#), the NetApp simulator, also supports this functionality for test and development purposes.
- You should thoroughly read both the README and the Quick Start Guide to get an introduction to VVOLs on NetApp storage.

## VMworld 2014

---

If you're attending VMworld, you have some unique opportunities to try out VVOLs on NetApp storage first-hand:

- Hands-On-Labs — HOL-PRT-1471 - Policy-Based Mgmt & VVOLs Tech Preview
- Live Demos in the NetApp booth in the Solutions Exchange
- Technical Staff and Developers in attendance to answer any and all questions

We would love to hear your feedback. You can get full details on all NetApp activities at VMworld [here](#).

---

› **Got opinions about VVOLS?**

Ask questions, exchange ideas, and share your thoughts online in NetApp Communities.



**By Graham Smith, Sr. Solutions Product Manager and Nick Howell, Virtual Storage Architect**

Graham focuses on VMware server virtualization solutions for NetApp. While a lot of what he does falls under a product management hat, Graham tries to remain close to the technology and the business problems it can solve. Throughout his career, he has been passionate about innovation and making people aware of the benefits it can bring.

Prior to joining NetApp in 2010, Graham worked at I/O and networking startups. He has also worked in HP's UNIX division in both marketing and R&D positions.

Nick started in IT around Y2K, eventually diving into server administration after receiving his CCNA and MCSE 2000. He discovered virtualization and shared storage in 2007, and immediately began implementing it where he worked. Nick is an outspoken advocate of "virtualizing everything." At NetApp, he helps customers design and deploy advanced virtualization solutions, with a particular focus on management technologies, backup/recovery, and other technologies that make life easier for the everyday administrator.

---

*Quick Links*

- › [Tech OnTap Community](#)
  - › [Archive](#)
  - › [User Groups](#)
  - › [PDF](#)
-