

# Delivering Simplified IT with VMware vSphere Virtual Volumes and Hitachi Storage

## WebTech Q&A Session, April 15, 2015

**1. Can I use legacy/traditional datastores along with Virtual Volumes (VVols) datastores?**

Yes, you can provision legacy/traditional datastores (Virtual Machine File System [VMFS], Network File System [NFS]) along with Virtual Volumes datastores in the same array with VMware vSphere 6.

**2. Would it be possible to migrate current customers' VMFS datastores with virtual machines (VMs) to storage containers with VVols?**

As part of the firmware upgrade for Hitachi Data Systems arrays, our vStorage APIs for Storage Awareness (VASA) providers and VVol implementations will allow for hardware-accelerated nondisruptive migration between source VMFS/NFS datastores and a destination VVol storage container, via the VMware vSphere Storage vMotion feature.

**3. Will customers need to create VM storage policies for every VM?**

No, higher-level VM administrators/architects will typically create best-practice VM storage policies based on their environment or tenant needs. For example, you might define a VM storage policy for all potential MySQL production VMs or pre-production VMs. These policies would then become available to be selected as part of the VM/vApp provisioning process.

**4. I have Hitachi NAS Platform (HNAS) serving VMs over NFS today in production. What do I need to enable VVols?**

Deploy HNAS OS 12.2.3753.08 or later and download and install the Hitachi VASA provider virtual appliance, available as an OVA.

VASA provider (VP) virtual appliance: Hitachi VASA provider sets up a communication management path between vCenter and storage platform(s). It operates as a virtual appliance(s) in the environment. It translates those VC management operations (such as Create VVol or Snapshot VVol) into HDS-specific calls or offload operations. It also provides the interface to share storage capabilities for storage containers between storage platforms and VC.

WebUI: Available after OVA install, tailored interfaces for VM administrators to manage the VP virtual appliances and entry-level Storage Policy Based Management (SPBM).

Hitachi Command Suite (HCS): 8.1.4+: Enterprise-level interface extension for storage administrators to manage storage containers and storage capabilities.

**5. What are the performance and scalability limits of the protocol endpoint?**

Protocol endpoints are used to access I/O. VVol architecture implementation predisposes that PE doesn't become a bottleneck. There has been some concern raised regarding queue depths and VVols. Traditional LUNs and volumes typically do not have very large queue depths, so if there are a lot of VVols bound to a PE, doesn't this impact performance? This is addressed in a number of ways. First, you are free to choose any number of PEs to bind VVols to (i.e., you have full control over the number of PEs deployed, which could be very many). Secondly, VMware is allowing for greater queue depth for PE LUNs to accommodate a possibly greater I/O density. However, considering that HDS already provides a choice regarding the number of PEs per storage container, and storage container size, this increased queue depth may not be relevant in many situations. We don't expect more than a single-digit number of PEs to be deployed.

**6. Since VASA provider appears to be central to data operations, how are you ensuring continuous availability of it, more so since many arrays will be supported by one VP?**

If VP is unavailable, only storage management operations will be impacted (create, clone, snapshot, power off). HDS will be deploying VP in N+1 model for higher availability. A VP down situation does not impact VM, as data I/O flows through PE(s).

**7. Any support planned for older disk arrays, such as Hitachi Unified Storage VM (HUS VM) and Hitachi Virtual Storage Platform (other than VSP G1000)?**

For Hitachi storage platforms (HUS VM, VSP, HUS 100 series), VVols support can be surfaced by virtualizing with either a VSP G1000 or a Hitachi NAS Platform (HNAS) 4000 series cluster gateway.

**8. What is the difference between vSAN and VVol?**

vSAN is a storage management framework for server-attached storage (hyper-converged), whereas a VVol framework is meant for external NAS/SAN arrays. Different customer segments will want to use one or the other (or both). vSAN has VSA datastores and VVol has VVol datastores. They are quite similar with respect to SPBM. Virtual Volumes uses VASA 2.0 to communicate with an array's VASA provider to manage Virtual Volumes on that array, but Virtual SAN uses its own APIs to manage virtual disks. SPBM is used by both, and SPBM's ability to present and interpret storage-specific capabilities lets it span vSAN's capabilities and a Virtual Volume array's capabilities and present a single, uniform way of managing storage profiles and virtual disk requirements.

- 9. How does queue depth work for VVols? We are a FC shop, and when planning a large, 30,000-plus IOP workload we have to build datastores of multiple LUNs to get the queue depth we need to support that IOP and bandwidth requirement?**

The queuing is handled by the protocol endpoints. (Queuing not relevant to NFS PE).

VMware is allowing for greater queue depth for PE LUNs (ALU) to accommodate greater I/O density. Think of the PE as a pass-through device mapped to many ports which will have x number of VVols behind it (i.e., one VVol datastore could equal the many multiple LUN/datastores you create today). You are also free to choose any number of PEs to bind VVols to (i.e., you have full control over the number of PEs deployed, which could be very many). HDS provides a choice regarding the number of PEs per storage container, and storage container size, so this increased queue depth may not be relevant in many situations. We don't expect more than a single-digit number of PEs to be deployed.

- 10. It sounds like a storage container is what HDS has historically called a pool. Is that correct?**

Storage containers are an overlay on top of the storage pools. You could have multiple pools as part of one storage container. For example, a storage container could contain two HDP pools and one HDT pool.

- 11. Can you use Storage vMotion to migrate from a LUN-based datastore to a VVol?**

Yes you can. You can do that seamlessly.

- 12. Does vCenter Site Recovery Manager support per-VM replication now?**

vCenter Site Recovery Manager is not supported in first release with VVols, but HDS does support VVol-level replication such as remote backup of VVols to a secondary datacenter leveraging HDS array and file-based remote replication.

- 13. Tell me about replication and how organizations may transition from current SRM/SRA capabilities based on LUNs and CGs?**

SRM support for VVol is a road-map item from a VMware perspective. The challenge here is the SRAs available today do not communicate with the VASA provider in order to interact correctly with the storage and vSphere platforms. We (HDS) are working closely with VMware in this area to plan this deliverable.

- 14. Will there be multiple pools from multiple arrays?**

You can create multiple storage containers (logical data stores) as per your business operations requirements. However, in current release, storage containers cannot span array boundaries. This may change in future releases.

- 15. When will this be available on the HUS VM platform via FCP block storage?**

For Hitachi Storage platforms (HUS VM, VSP, HUS 100 series) or 3rd-party storage, we provide an excellent option to virtualize an FC platform with VSP G1000 or to virtualize with a HNAS 4000 cluster gateway to surface VVol support. Important to note, we expect customers will go through a transition where existing VMFS/NFS datastores on HUS VM (for example) will gradually transition over existing/new VMs to VVol-based datastores using supported Hitachi storage arrays.

- 16. What are the replication capabilities with VVols?**

Snapshots (or internal replication, as some call it) is supported. Remote backup of VVols to a secondary datacenter leveraging HDS array-based remote replication is also supported with the Hitachi solution on

Day 0. Active-active VVol configuration within campus/metro datacenters leveraging Hitachi global-active device (bidirectional mirror) is being qualified. Disaster recovery implementation for VVol leveraging HDS array-based remote replication is currently in architecture review.

**17. Can HDP/HDT be shared as a storage capability in a storage container and between vdevs attached to physical servers?**

Yes, HDT and HDP can coexist within the same storage container, and it will be surfaced as a capability as part of "Auto-generated Informational Capabilities."

**18. What is the purpose of HNAS if NFS and CIFS are going to VVols?**

VVols support access via FC, FCoE, NFS or iSCSi protocols. We leverage HNAS (our enterprise NFS offering) to provide that NFS access to our VVol implementation. SMB/CIFS is not really relevant in VVol data access context.

**19. Does the VVol look like one large datastore? Can you browse it and copy/modify VM files?. If a snapshot is created, is a delta file created or is it some kind of native array snapshot?**

Two concepts, a VVol = VMDK to keep it simple, and a VVol datastore (sometimes called virtual datastore). You can browse a VVol datastore to see the VMs within it. Snapshot question: No more VMFS, no more delta files, and VM snapshots are handled natively within the array.

**20. Is SRM support with VVols expected in a specific quarter for 2015, and if so, which quarter?**

SRM support for VVol is a road-map item from a VMware perspective. The challenge here is the SRAs available today do not communicate with the VASA provider in order to interact correctly with the storage and vSphere platforms.

**21. Does VAAI still play a role in VVol?**

The default operation with VVol always defaults to the new VVol primitives. When an array is VAAI-compliant, those capabilities will be leveraged when possible and needed.

**22. In a multitenant environment, does each tenant require their own storage container to ensure security/privacy?**

That depends on the requirements. Is logical isolation from a VM perspective important? Or is it about having different administrators not being able to browse a storage container?

**23. On what platforms do you have primary dedupe available? Is this only a feature of HNAS??**

Primary dedupe is available as an NAS feature.

**24. If encryption is chosen as a characteristic of a VVol, how are the encryption keys managed in vSphere?**

Encryption is handled within the Hitachi storage array. There are no encryption keys for vSphere to manage. It simply says, because you have chosen a VM storage policy with encryption, that VM provisioning operation will get assigned a storage resource that will encrypt the VM/VVol data at rest.

**25. We spread the workload across multiple back-end-array ports. Can we continue to do this with VVols?**

You can use the same concept with VVols. You can have all back-end ports allocated to a particular protocol endpoint (PE) or selectively choose which back-end ports are allocated to which PEs. In this case, you would use different protocol endpoints when this is supported by the vendor's implementation.

**26. Does this architecture introduce new dependencies or points of failure? For example, what happens if the VASA appliance fails?**

If the VP is unavailable, only storage management operations will be impacted (create, clone, snapshot, power off). HDS will be deploying VP in N+1 model for higher availability. A VP down situation does not impact VM, as data I/O flows through PE(s).

**27. I'm more interested in the technical details than the policy capabilities. Can you direct me to online training / documentation?**

General details, plus additional FAQs answered in a recent technical blog that address your technical specifics, are available here <https://community.hds.com/community/products-and-solutions/server/blog/2015/03/16/hitachi-and-vmware-virtual-volumes-part-1> [Post a question on that blog if not specifically addressed, and I will get it answered].

Technical Demo - <https://www.youtube.com/watch?v=3OBS2DgyKuw>

**28. When a VM is deleted from vCenter, will the V2I snapshots automatically delete?**

No, we will preserve them for a period of time. This is to allow for the situation that a VM is deleted accidentally and can use the snapshots to recover.

**29. Prior to VVols, VMFS datastores were presented to hosts and represented as LUN IDs from the storage arrays. With VVols, are LUN IDs associated with storage containers through the PEs or individual VVols through the PEs?**

VVol Datastore = Storage Container, which is a logical construct and does not have a LUN ID associated with it. Recall that VVols are bound to a particular PE and that data access is through the PE. ESXi knows how to send I/O for a VVol to a particular PE (ALU or NFS export) and the storage array ensures I/O gets to the relevant VVol object within the array (SLU or File).

**Prior, I'm for some reason just not getting it. I'm the vSphere administrator. We have a Hitachi Virtual Storage Platform G1000 in house and have been using that storage for a long time. I'm not so concerned about how the array side works but how the vSphere side does.** Review this technical demo to see if it addresses your query: <https://www.youtube.com/watch?v=3OBS2DgyKuw>

**30. Do I require a separate VASA provider for each Hitachi storage array?**

**No, our implementation allows one VP to manage multiple HDS storage arrays/NAS nodes.** HDS will be providing a Unified VASA provider (VP) OVA package, which bundles both VPs together to maximize the breadth of storage capabilities, but an administrator will likely start up at least one instance of VP-f (file) and one instance of (VP-b). Each VP instance can manage multiple storage targets. I believe

the majority will use both and leave the storage capabilities/SPBM framework to manage where the VMs get placed to best match the policy requirements.

© Hitachi Data Systems Corporation 2015. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. All other trademarks, service marks, and company names are properties of their respective owners.