Hitachi Storage for VMware vSphere Virtual Volumes (VVol) Frequently Asked Questions (FAQ)

1. What is Virtual Volumes (VVol)?

VMware vSphere Virtual Volumes is based on a new integration and management framework between vSphere and the storage system. Virtual Volumes virtualizes SAN and NAS devices by abstracting physical hardware resources into logical pools of capacity (called Virtual or VVol Datastore) that can be more **flexibly consumed** and configured to span a portion, one or several storage systems. It implements an out-of-band bidirection control path through the vSphere APIs for Storage Awareness (VASA) and leverages unmodified standard data transport protocols for the data path (for example, NFS, iSCSI, Fibre Channel). On the storage system side, two new components are added to the storage environment: "VASA provider" for integration with VASA APIs and "Protocol Endpoint" (PE) that allows the storage controller to direct I/Os to the right Virtual Volumes. On vSphere, there are three dependent features: VASA, Virtual Volumes and SPBM. In order to create policies at the vSphere level, a set of published capabilities must be first defined in the storage system. Once these capabilities are defined, they are surfaced up to vSphere using the storage vendor VASA provider.

The Virtual Datastore (sometimes referred to as VVol Datastore) defines capacity boundaries and access logic, and it exposes a set of data services accessible to the virtual machines (VMs) provisioned in the pool. Virtual Datastores are purely logical constructs that can be configured on the fly, when needed, without disruption: They do not need to be formatted with a file system.

Virtual Volumes defines a new virtual disk container (Virtual Volume) that is independent of the underlying physical storage representation, allowing for **finer control**. In other terms, with Virtual Volumes the virtual disk becomes the primary unit of data management at the storage system level. This turns the Virtual Datastore into a VM--centric pool of capacity. It becomes possible to execute storage operations with VM granularity and to provision native storage-system-based data services to individual VMs. This allows admins to provide the right storage service levels to each individual VM.

2. How has HDS been involved with Virtual Volumes (VVol)?

Virtual Volume (VVol) is a storage management framework devised by VMware. HDS has worked with VMware as a consultative partner to deliver on VVol implementation on both block and file storage over the last 2.5 years. Here is link to VMware press release on the general availability of VVol with vSphere 6 referencing Hitachi Data Systems.

3. Why choose Hitachi for VMware VVol?

Zero worry. Running Virtual Volumes on Hitachi infrastructure ensures a robust, reliable enterprise journey to software-defined, policy-controlled data center.

4. What is SPBM (Storage Policy Based Management)?

To enable efficient storage operations at scale, even when managing thousands of VMs, Virtual Volumes is intertwined with vSphere Storage Policy-Based Management (SPBM). SPBM is the implementation of the policy-driven control plane in the VMware SDS model. The SPBM framework allows both advertising of storage capabilities and capture of storage-service-level requirements (capacity, performance, availability, and so forth) in the form of logical templates (VM storage policies). SPBM automates VM placement by identifying an available VVol datastore that matches the specified policy requirements. When coupled with VVol, it can dynamically instantiate the necessary data services when required. Through policy enforcement, SPBM also automates service-level monitoring and compliance throughout the life cycle of the VM to ensure continued matching of VM required storage policy to the advertised capabilities. See Figure 1.

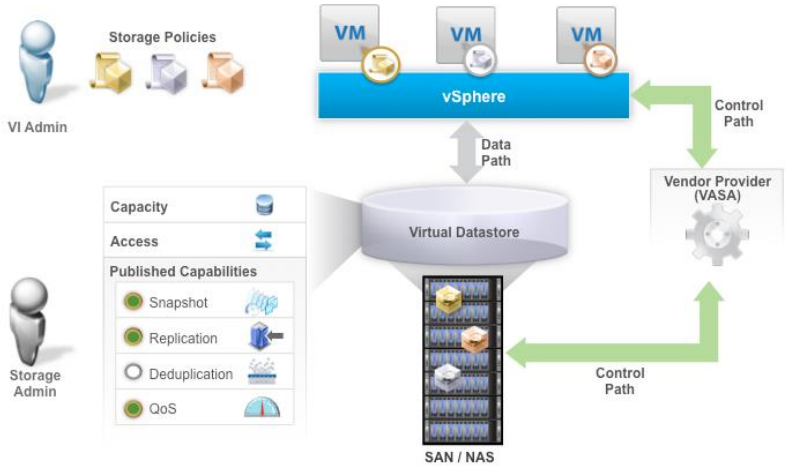

Figure 1. SPBM automates service-level monitoring and compliance.

5. What is HDS value proposition for Virtual Volumes management framework ?

Below is the general value proposition for Virtual Volumes.

a. Simplifies storage operations.
   i. Separates consumption and provisioning.
   ii. Enables end-to-end visibility.
b. Simplifies delivery of storage service levels.
   i. Allows finer control, dynamic adjustment.
   ii. Improves resource utilization.
   iii. Maximizes storage offload potential.

The HDS approach is to provide an enterprise-level, trusted, reliable, zero-worry implementation (see Figures 2 and 3). We embolden the storage administration team with rich SPBM storage capabilities control and leverage the VVol and SPBM framework to further elevate the rich data services that we bring to a VMware environment (global-active device, Hitachi Virtual Infrastructure Integrator, efficient data movement technologies, efficient storage-system-based cloning).

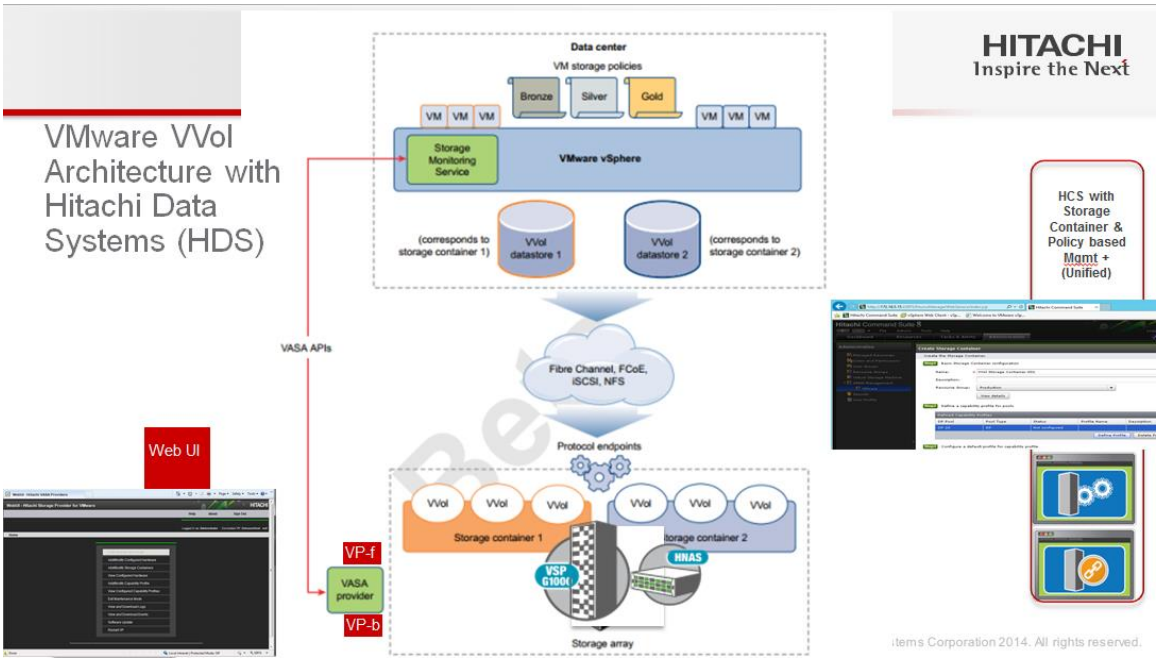

HCS = Hitachi Command Suite, VSP G1000 = Hitachi Virtual Storage Platform G1000, HNAS = Hitachi NAS Platform

Figure 2. VMware VVOL Architecture, Based on HDS Storage

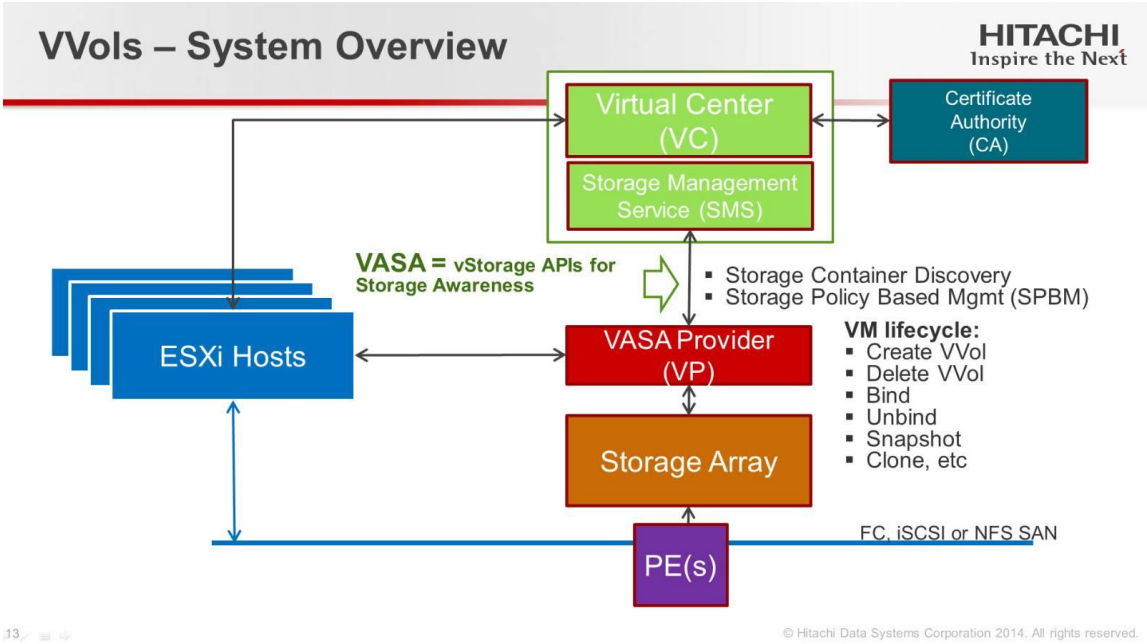

Figure 3. Running Virtual Volumes on Hitachi storage will bring a robust reliable enterprise journey to software-defined, policy controlled data center.

6. **What is the Hitachi Data Systems road map for supporting VVol in its storage systems?**

   HDS would support VVol in Hitachi NAS Platform (HNAS) systems by April 2015. Hitachi Virtual Storage Platform G1000 (VSP G1000) will support VVol integration by August 2015. For other Hitachi storage platforms [VSP, Hitachi Unified Storage (HUS) 100 series and HUS VM)], the options are virtualizing with VSP G1000, or virtualize with HNAS 4000 cluster gateway to surface VVol support.

7. **What are the key components of VVol enablement?**

   VASA Provider: Hitachi Storage Provider for VMware vCenter sets up a communication management path between vCenter and Hitachi storage platform(s). It operates as a virtual appliance(s) in the environment. It translates those VC management operations such as Create VVol, 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.

   Protocol Endpoint (PE): Protocol Endpoint provides I/O data path connectivity between ESXi hosts and Hitachi storage systems. Protocol Endpoints are compliant with both Fibre Channel (Virtual LUN) and NFS (mount point). Multipathing occurs against the PEs. In the initial release, HDS VVol solutions will support Fibre Channel and NFS PEs with support for iSCSI and FCoE following, subsequently.

   Web UI: Tailored interfaces for VM admin to manage the VASA Provider virtual appliances and entry-level SPBM.

   Hitachi Command Suite (HCS): Enterprise-level interface extension that allows storage admins to manage storage containers and storage capabilities.

8. How does a Protocol Endpoint (PE) function?

PE represents the data path or I/O access point for a Virtual Volume. All I/O flows through one or more PEs. Hosts discover SCSI PEs as they discover today's LUNs; NFS mount points are automatically configured. When a Virtual Volume is created, it is not immediately accessible for I/O. To access Virtual Volume, vSphere needs to issue a "Bind" operation to a VASA Provider (VP), which creates an I/O access point for a Virtual Volume on a PE chosen by a VP. A single PE can be an I/O access point for multiple Virtual Volumes and PEs can be visible to all ESXi hosts. An "Unbind" operation will remove this I/O access point for a given Virtual Volume. vCenter is informed about PEs automatically through the VASA Provider.

9. Do I require separate a VASA Provider for each Hitachi storage system?

We will be providing a Unified 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 and SPBM framework to manage where the VMs get placed to best match the policy requirements.

10. I have Hitachi NAS Platform serving VMs over NFS today in production. What do I need to enable VVol?

Table 1 outlines requirements to enable VVol.

Table 1. Hardware and Software Requirements

	Item	Requirements
Hardware	HDS NAS Gateway	Hitachi NAS Platform: HNAS 4060, HNAS 4080, HNAS 4100
	Back-End Storage Platforms	Hitachi Virtual Storage Platform (VSP), VSP G Series, Hitachi Unified Storage (HUS) VM, HUS 100 series
Software	HDS	HNAS OS 12.2.3753.08 or later HNAS SMU 12.2.3753.10 or later HNAS File Clone license
	VMware	vSphere 2015 (also known as vSphere 6.0)

11. What is the association of a Protocol Endpoint to hosts?

PEs are like LUNs or NFS mount points. They can be mounted or discovered by multiple hosts or multipathing applied. PEs are a pass-through device. In a Fibre Channel implementation, it's the sub-LUN behind the PE that is actually storing the data.

12. What are the key points of a VVol implementation by HDS?

a. Reliable, robust implementation of VASA Provider (VP).
b. Unified VP package for file and block (Q2 2015 planned).
c. Flexible storage capability catalogue to deliver application-required data services.
d. Tailored interfaces for both VM admins (Web UI) and storage admins (HCS).
e. Multitenant (resource group for block and IP multitenancy for file).
f. VVol scalability:
    1. Up to 100 m snapshots and clones (HNAS).
    2. 400,000 VVols in first release, architectural support and 10 million VVols (HNAS).
    3. 64,000 VVols, 1 M snapshots or clones (VSP G1000).

13. Can one PE connect to multiple hosts across clusters?

Yes. PEs can be visible to all hosts. VPs can return the same Virtual Volume binding information to each host if the PE is visible to multiple hosts.

14. Why are storage containers needed?

Storage containers are a collection of one or more storage resources and you can define multiple storage containers. Storage containers provide a logic abstraction for managing very large numbers of Virtual Volumes. In the Hitachi implementation, they consist of one or more pools of storage (think of one or more Hitachi Dynamic Provisioning and/or Hitachi Dynamic Tiering pools) or one or more file systems (for example, regular FS#1 and Tiered Flash Filesystem#7 group as storage container) that the storage admin can define. In the first release, they don't span outside a storage system, but that will change in subsequent releases. This abstraction can be used for managing multitenant environments, various departments within single organization or dedicated resources for certain apps. There is a 1-1 mapping between a storage container and a VVol datastore. Here is link to our architecture slide as a refresher

https://community.hds.com/servlet/JiveServlet/showImage/38-1707-5232/VVol+arch+%232.png.

15. How many storage containers can be mapped to one host?

256 storage containers per host is the limit.

16. **Describe HDS storage capabilities that will be advertised.**

In the initial releases, HDS is supporting three classes of capabilities.

- Auto-generated capabilities (for example, RAID type, encryption).
- Managed storage capabilities (for example, performance class, availability class, cost class).

- Custom capabilities (for example, availability zone).

17. Explain more about the managed storage capabilities. Are they the same across vendors? How I can use cost as a parameter?

Table 2 shows the managed storage capabilities that we are initially focused on, which map to our target market and customer needs. Other vendors may or will have alternatives but we are sharing below to encourage the community to gravitate towards a common set. With the Hitachi implementation, storage admins can now discriminate the resources within the storage container and between storage containers based on the all-important performance, availability, cost and operational recovery class attributes. For example, they can now have storage resources that will only entertain VMs being provisioned on them and that specifically request Tier 1 performance class and Tier 1 availability class [hence, avoiding noisy or resource-hogging non-mission-critical VM resources mistakenly being placed in the vicinity]. Or, they can separately have storage resources that can support a minimum of Tier 2 performance class for other business essential apps. The cost class is a control valve that infrastructure teams can use to avoid the issue of business groups or tenants requesting Tier 1 performance for all their VMs when indeed lower-cost class resources would indeed meet their business app performance needs. The snapshot backup class will be an interesting attribute capability for those familiar with Hitachi Virtual Infrastructure Integrator.

Table 2. Auto-Generated and Managed Storage Capabilities: HDS Focus

Category	Category	Capability	Value Range	Note
Auto-Generated Capabilities			Informational Auto Generated Capabilities	
		Array Type / Serial #	Model / serial	
		RAID Level	[RAID10 / RAID5 / RAID6] values are according to storage specification.	Familiar to storage admin. (Some capabilities used in VMworld2015 Demo)
		Pool Type	[HDT / HDP]	
		Drive Type/Drive Speed	[SAS / SSD / Flash(HAF) / External]	
		Encryption	[Yes / No]	
		Space Efficiency	Thin or N/A	
		Snapshot	[Yes / No]	

Category	Category	Capability	Value Range	Note
Managed Storage Capabilities	Performance	**Performance IOPS – Class**	[Tier1_IOPS / Tier2_IOPS / Tier3_IOPS]	Label to differentiate storage pools/FS based on IOPS
		Performance Latency – Class	[Tier1_Latency / Tier2_Latency / Tier3_Latency]	Label to differentiate storage pools/FS based on Latency
	Availability	**Availability - Class**	[Tier1 / Tier2 / Tier3]	Label to differentiate storage pools/FS based on availability
		Cost - Class	[10 - 1000]	Maximum cost that VM admin/tenant willing to pay for storage service. Used for showback and chargeback services..Storage/Cloud Infra team determines their currency/multiple to use
	Operational Recovery	**Snapshot Backup Importance – Class**	[High / Medium / Low]	what Class of snapshot services required. Tie into Hitachi Virtual Infrastructure Integrator (V2I) to manage policy
Custom Tag Attribute Capabilities	Custom	**tag name**	User Definable [xx,x,x]	Some examples that user may use for illustration. DC Location=[,] Availability Zone=[,] Tenant=[,] MaintenanceWindow=[,] PolicyName = [,]

HDP = Hitachi Dynamic Provisioning, HDT = Hitachi Dynamic Tiering, HAF = Hitachi Accelerated Flash

18. Why do I need to differentiate storage services?

Consider the following analogy (and video): Imagine your CIO tasked you to find a fast, energetic soccer player for his team with the words "don't disappoint me." From the pool of soccer "storage" resources at your disposal, would't it be helpful if their skills were advertised before choosing that player?
https://youtu.be/EI5sE7y4PVQ

19. Can I use legacy or traditional datastores along with Virtual Volumes datastores?

Yes. You can provision legacy or traditional datastores along with Virtual Volumes datastores in same storage system.

20. Where are the storage system policies (snap, clone, replicate) applied? At the storage container or the Protocol Endpoint?

The capabilities are advertised as part of the storage container and not the Protocol Endpoints. PEs are just the communication path.

21. VSAN also promises policy-based management, so how is VSAN different from VVol?

VSAN is storage management framework for server-attached storage (hyper-converged) whereas VVol framework is meant for external NAS or SAN storage systems. Different customers 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 a storage system's VASA Provider to manage Virtual Volumes on that storage system 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 storage system's capabilities and present a single, uniform way of managing storage profiles and virtual disk requirements.

22. Are there any capacity limitations on Virtual Volumes?

Virtual Volumes can expand as big as the capacity of storage container.

23. How is storage-policy-based-management implemented by VVol?

During provisioning, a VM storage policy is selected, which ensures that the storage-selected maps are delivered to the VM storage policy. During the course of the life of that VM, compliance checks are performed to ensure the VM is still being served from storage that maps to that storage capability.

24. How many Virtual Volumes will be created for a VM?

At least three and potentially more depending on VM configuration will be created. There is a Virtual Volume for every virtual disk, one Virtual Volume for swap, if needed, one Virtual Volume per disk snapshot and one Virtual Volume per memory snapshot. The minimum is typically three Virtual Volumes perVMs (one config, one data, one swap). Each VM snapshot would add one virtual disk and one memory snapshot (if requested).

25. Would VVol framework be compatible with older versions of vSphere?

VVol will only work with vSphere 6.0 edition and upward. Customers running vSphere 5.5 and older versions will have to upgrade to vSphere 6.0 to be able to use VVol and deploy a supporting storage system with corresponding storage VASA Provider implementing the VASA version 2 APIs.

26. You mentioned the storage capabilities that Hitachi will support. How do I as a storage admin define those?

Below is a screenshot of the interface within Hitachi Command Suite (HCS) dedicated to storage container and SPBM aspects. You (as part of storage or infrastructure team) can leverage this interface to define the capabilities on the individual storage resources within the storage container(s). In the example in Figure 4, the storage admin has identified that his storage resource can advertise that it will support maximum of Tier 2 performance class and Tier 2 availability class, with minimum cost of $500, and so forth. The system auto-determines the other capabilities, such as encryption and RAID level if VM admin wants to leverage in their VM storage policies.

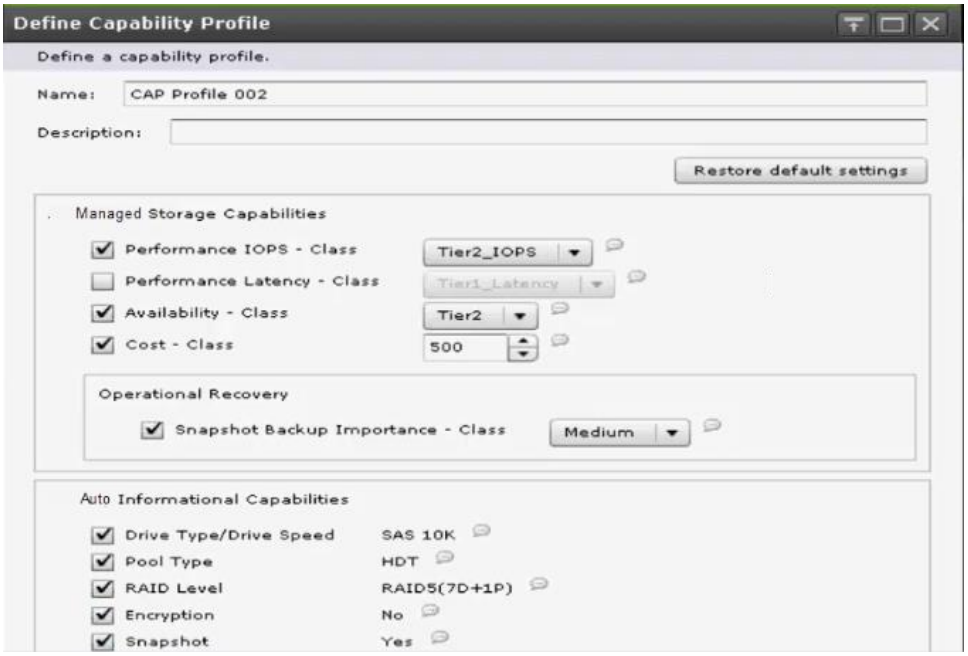

Figure 4. Hitachi Command Suite Interface

27. As a VM admin, how do I access these capabilities?

Figure 5 presents a screenshot of the vSphere Web client, where VM admins can create VM storage policies. In the example, the VM admin has decided to create a standard VM storage policy that all production Oracle databases being deployed must leverage. The policy states that VMs get Tier 1 performance, and either Tier 1 or Tier 2 availability with no restriction in the cost of the storage resource being allocated.

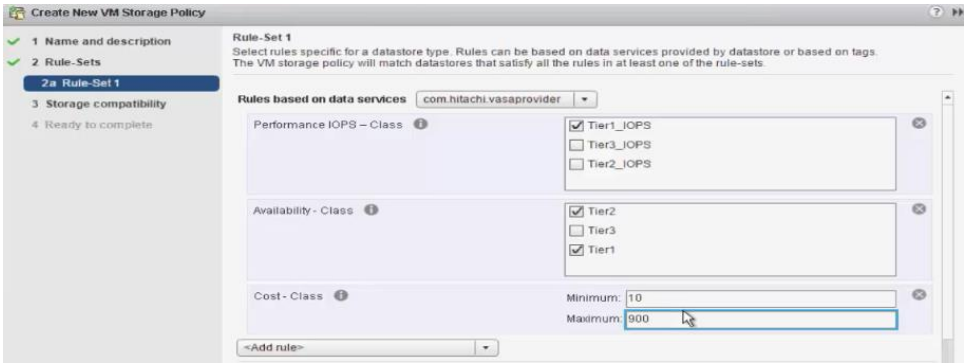

Figure 5. VSphere Web Client

28. What licensing is required in getting VVol support in storage systems?

VVol management framework will be offered as no charge item. (That is, there is no licensing fee for Hitachi Storage Provider for VMware vSphere or VASA Provider.) Customer needs to ensure they have the appropriate Hitachi platform that supports VVol (that is, VSP G1000 class systems and/or HNAS gateway) For other or older Hitachi storage platforms (HUS VM, VSP, HUS 100

series) or indeed third-party storage, we provide option to virtualize with VSP G1000 or virtualize with HNAS 4000 cluster gateway to surface VVol support.

29. **Will PE manage features such as primary dedupe?**
    No, Protocol Endpoint (PE) has no relevance here, as it only deals with in-band related activities, such as all important I/O data path and has no interaction with storage capabilities. Storage capabilities such as dedupe become an advertised capability of a resource within the storage container. Dedupe is normally on by default on our storage resource but you could have situation where you create or advertise a resource as nondedupe enabled.

30. **Is dynamic tiering used with storage containers ? Will HNAS VVol support Hitachi Dynamic Tiering?**
    Hitachi Dynamic Tiering (HDT) and Hitachi Data Provisioning (HDP) will be surfaced as part of "Auto-generated Informational Capabilities." Storage containers (for block storage) should be able to expose these capabilities by default. Other "Informational Capabilities" include RAID Level, Pool Type, Drive Type and Speed, Encryption Y/N, and Snapshot Y/N (for HNAS, you also get "Space Efficiency," that is, thin or thick). As far as the question related to HDT on HNAS implementation over NFS, it is unrelated to this discussion. However, it is not generally recommended as HNAS converts VM random I/O into sequential behavior on back-end storage device for optimum performance.

31. **Will there be a best practice guide published for VVol deployment ?**
    Yes, there will be several VVol related papers, in the following order of availability, which will also cover deployment related practices:
    a. VVol on HNAS Tech Note, available in April 2015.
    b. Storage-Policy-Based Management on VVol Tech Note (block and file VASA Providers and their capabilities, showcasing various use cases), available in May or June 2015.

32. **Would it be possible to migrate current customer's data stores with VMs to storage containers with VVols?**
    That was in fact part of the original design specs and core requirements. As part of the firmware upgrade for our storage systems, our VASA Providers and VVol implementations will allow for hardware-accelerated nondisruptive migration between source VMFS and NFS datastores and destination VVol storage container, via storage vMotion.

33. **Which Hitachi NAS Platform models will have VVol support?**
    The recommended Hitachi NAS Platforms for VVol in production are HNAS 4060, HNAS 4080 and HNAS 4100. There is a provision to support HNAS 3080 and HNAS 3090 for proof-of-concept (POC) or evaluation use cases.

34. **What could be considerations for choosing right storage protocols for VVol?**
    For vast majority of use cases, we expect protocol agnostic deployments. However, for use cases such as test and development, fast provisioning cloud related VM workloads NFS may be

preferred over Fibre Channel. For latency variability sensitive applications, Fibre Channel may be preferred over NFS.

35. Will snapshot and replication be offered as capabilities for storage containers?

Yes. We offer a snapshot backup class and other replication-based capabilities are coming in subsequent releases.

36. Will VVol have different way of handling VM snapshot/backup/restore history during VMotion event?

Yes, VVol will preserve and move associated snapshots even with Storage vMotion events. This is another benefit of having VMDK-based storage.

37. Will VMFS cease to exist after introduction on VVol?

For VVol, there is no concept of VMFS as you know it today as a hypervisor-based shared filesystem. With VVol, vCenter is communicating with VMDK storage and previous operations that used to occur in the hypervisor/VMFS (such as snapshot), are now handled by the storage array. Obviously, VMFS will continue to co-exist with VVol to continue to support traditional VMFS datastores.

38. What will be the management interface for VVol?

In the first release Web UI will used as management interface. For VM admins, the Web UI will be the management interface. For storage admins, HCS will be the management platform.

39. Will your backup and recovery solution Hitachi Virtual Infrastructure Integrator support VVol objects?

Virtual Infrastructure Integrator in addition to Hitachi Data Instance Director (HDID) will be key differentiator for HDS. These offerings allow customers to benefit from specified data protection as part of the provisioning process while benefiting from enhanced automated backup and recovery of VMs in granular manner. They eliminate the challenges that VMware environments have experienced with hypervisor-based snapshots and subsequent disk consolidation of delta snapshots which can cause havoc. VMware vCenter has initiated the native offloading of snapshot and clone operations to Hitachi storage in the initial release. The actual Virtual Infrastructure Integrator managed integration will be in subsequent release after the initial release.

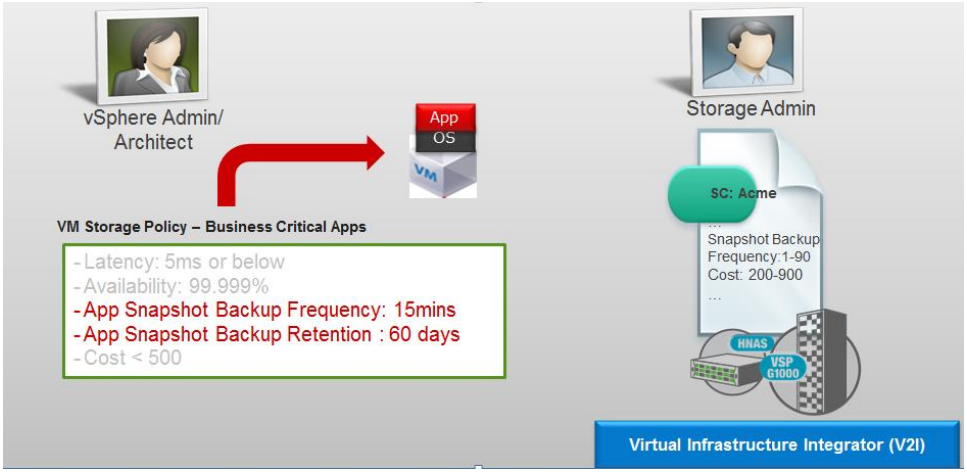

HNAS = Hitachi NAS Platform, VSP = Hitachi Virtual Storage Platform G1000

Figure 6. Virtual Infrastructure Integrator helps to eliminate snapshot challenges in VMware environments.

40. What are the performance and scalability limits of the Protocol Endpoint? How will this affect VM density ?

Protocol Endpoints are used to access I/O. VVol architecture implementation predisposes that PE doesn't become 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, we are free to choose any number of PEs to bind their VVols to (that is, they have full control over the number of PEs deployed, which could be very many). Second, VMware is allowing for greater queue depth for PE LUNs to accommodate a possibly greater I/O density. However, considering that we already provide 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 single digit number of PEs deployed.

41. Do you know if we plan to have a mechanism to limit the usage of a container for a VMware admin? If there is no control, my customers are afraid that the VMware guys will quickly use all the capacity available. How can they control this?

The storage admin controls the creation of the storage container(s) and can decide what capacity to impose when creating those storage containers. The philosophy is that storage admin provides this storage resource, assign capabilities and then the VM admins are free to consume to meet their VM provisioning needs. There are alarms and events to notify admins as capacity gets consumed.

42. Aren't there concerns about number of LUs to a port given the likely explosion in number of LUs presented to ESX?

No, ESXi hosts only see ALU (Protocol Endpoints) presented. PEs are mapped to ports and never the back-end LUs for VVols. The backed VVols are secondary logical units (SLU). VVol over file or NFS is similarly unaffected as it's dealing with PEs as mount points and VVols as files.

43. Are the VVols a replacement for the virtual machine disk (VMDK) or separate VVols inside a VMDK? And, since the LUN is specific to the guest, does this create more work for customers?

From VI admin perspective, VVols can be considered a 1-to-1 relationship to VMDK (see Figure 7). VI admins still manage VMs, and storage admins manage storage containers, with increased visibility to the VVols consuming the storage container capacity. Internally, we use file or DPVol association to a VVol-VMDK depending on the Hitachi storage platform. It actually creates less work, allowing one-time multipathing for PE and efficient provisioning with the SPBM framework to take away the guesswork.

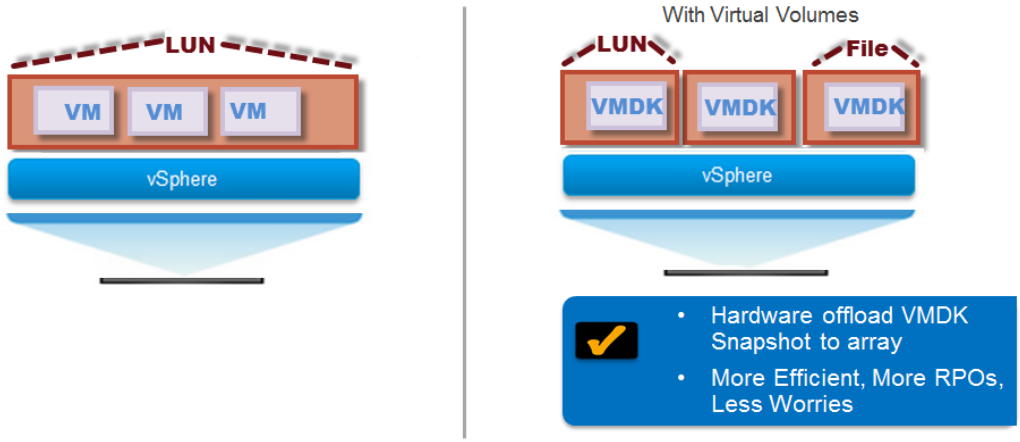

Figure 7. VVols can be considered a 1-to-1 relationship to VMDK.

44. When will HDS support VSP G1000? Also, do we support Hitachi Universal Replicator (HUR) and the global-active device feature?

VSP G1000 will have VVol integration in Q2 2015. Managing integrated capabilities, such as enabling global active device, Virtual Infrastructure Integrator and HUR will be supported in phased manner.

45. Are VVols going to be supported on FCOE for others, such as Cisco? Or, this could be an advantage for vendors who do not "push" FCOE by default?

Yes. VVol will support existing protocols that are supported across VSP G1000 and HNAS. As VSP G1000 supports FCoE, VVol will support FCoE.

46. If a VM needs 50GB is that what it gets? Or, is the container actually bigger? If the container is small, doesn't that use more addressing from us? And how would that possibly affect host storage domains (HSDs) on the ports?

The container is actually bigger (thin provisioned).

47. Will HNAS 4040 support VVol?

Recommended Hitachi NAS Platforms are HNAS 4060, HNAS 4080 and HNAS 4100. We will support VVol on HNAS 3080, HNAS 3090 and HNAS 4040 for test/dev or POC purposes only.

48. **Will a customer still use Hitachi Dynamic Link Manager (HDLM) if they are using PE?**
PE will be used for multipathing so customer will continue to use either native multipathing or HDLM.

49. How does PE affect replication?
None. PE is used for I/O communication. Replication services flow through different ports so there is no association between protocol endpoints and replication

50. If the VP dies, do we just lose the ability to do new provisioning, and so forth, or do we lose access to existing LUNs?
If VP is unavailable, only storage management operations will be impacted (clone, snapshot, power off). HDS is deploying VP in N+1 model for higher availability. VP down situation does not impact VM I/O as that flows through PEs.

51. Do we have a road map of storage system replication support? When do we support local snaps or clones, Hitachi Universal Replicator or Hitachi TrueCopy, remote snap or clone, VMware (Site Recovery Manager), and so forth?
Snapshot, clones and HUR services will be supported at general availability (GA). SRM orchestration will be supported in next release when VMware begins supporting SRM on VVol.

52. How does a VVol coexist and uses HDT if they are defined by containers?
A storage container could contain 2 HDP and 1 HDT pools for example. A VM storage policy may end up placing that VVol on the HDT pool part of the storage container because it matched the advertised capabilities.

53. Is there a correlation between the number of LUNs available on a storage system and the number of PEs available?
There is no direct relation between storage containers and PEs. One PE can be used to manage multiple storage containers. PEs are used to access I/Os and for multipathing and zoning purposes. In the VVol paradigm, LUNs will be replaced by storage containers.

54. Is there a mechanism to limit the usage of a container for a VMware admin? If there is no control, my customers are afraid that the VMware guys will quickly use all the capacity available. How can they control this?
The storage admin controls the creation of the storage containers and can decide what capacity to impose when creating those storage containers. The philosophy is that storage admin provides this storage resource and assigns capabilities, and then VM admins are free to consume from this container to meet their various VM provisioning needs. There are alarms and events to notify them as capacity gets consumed.

55. How does VVol work and co-exist with Hitachi Unified Compute Platform Director?

UCP Director has a road map item to add support for vSphere 6 and VVol and VASA.

56. If a customer has an SQL VM where he might want the drives on different storage types (for example, logs on SSD, backup on SATA) is that possible with VVOL?

Yes. The customer can apply different storage policies on different VVol. Log files and the SQL database can reside on different mediums from the operating-system-based on storage policies defined by the VI admin.

57. Can a storage system that only supports 2048 LUNs support many more VVols?

Yes and No. The key part is how may LUNs can be assigned. For example, VSP G1000 supports 64,000 addressable VVols but has the ability to support another 1 million Hitachi Thin Image (HTI) snapshot VVols. When one of those snapshot VVols gets presented, it takes one of the available address spaces from the 64,000.

58. When will VVol get SRM support?

VMware has stated that SRM support is not in initial release with a promise that it will be delivered in subsequent release. An important point: This does not mean that replication is not supported. For example, existing Hitachi replication technologies such as HUR and TrueCopy (file or object replication) still provide disaster recovery and replication capabilities for VVol between Hitachi platforms. Customers will have to wait for VMware to update SRM to add the orchestration layer for those who use that offering for orchestration.

59. Will customers need to create policies for every VM?

No. Higher-level VM administrators or architects will create storage policies based on their environment needs, SQL production, Tier 1 or Tier 2 services. These policies then become available to be selected as part of the VM provisioning process.

60. Will sub-LUN replication be supported with VVol on HNAS?

In an HNAS storage array, each VVol-VMDK becomes a file based operation. We continue to provide same file-based replication services that we do today. On block, each VVol-VMDK is a DPVol, so we can support existing data services, such as replication.

61. Will we have FCoE support for VVol?

Yes.

62. Is Hitachi NAS File Clone new license outside of HNAS license pack?

The NAS File Clone license is available as part of a value package or on standalone SKU. If you order HNAS VM bundles, this license is taken care of as part of package.

63. How many VVols are included per VMware cluster?

VMware conceptually supports 19 million VVols.

64. As VVols are created within storage, what is impact on monitoring?

We will be using DPVol concept for VVols, we will continue to use DPVol monitoring tools, such as vRealize Operations Manager and native DPVol monitoring tool Hitachi Tuning Manager.

65. Will a command device become a VVol and need to be applied to each VM?

Yes a command device could or should become a VVol, but there is no need for command device to become VVol immediately, as customers in vSphere 6.0 can continue to use Resource Device Mapping (RDM).

66. Could VMware policy change effect a change in VVol placement?

Yes. If storage policy required certain storage capabilities and those capabilities changed over time, the effect of that change would be an action to bring that VM back into compliance and potentially change the location of VVol (potentially, to move from one HDP pool to a HDT pool within a storage container).

67. Do we have VVol technical demo available?

Yes. The Demo can be found at:
https://www.youtube.com/watch?v=3OBS2DgyKuw

68. Where can I get additional information regarding VVol?

HDS.com - http://www.hds.com/solutions/virtualization/vmware-vsphere/storage-solutions-for-vmware-virtualization/virtual-volumes.html
VMware.com - http://www.vmware.com/products/virtual-volumes
HDS Server and Desktop Community site: https://community.hds.com/community/products-and-solutions/server
Notable Blogs by VMware SMEs: http://cormachogan.com/2015/02/17/vsphere-6-0-storage-features-part-5-virtual-volumes/

69. Who do I contact for further questions on VVol?

Submit questions to VMwareQuestions@hds.com or post questions to HDS Server and Desktop Community site: https://community.hds.com/community/products-and-solutions/server

BR-12459 August 2015