FREQUENTLY ASKED QUESTIONS

VMware vSphere Virtual Volumes

Getting Started

Q1.What is Virtual Volumes (VVol)?

A. An integration and management framework delivering a new operational model for external storage (SAN/NAS. It is comprised of a control plane using Storage Policy-Based Management, and a data plane using VASA APIs for external storage and vSphere APIs for IO Filtering for in-hypervisor software data services.

Q2.What are the software requirements for Virtual Volumes based storage?

A. To use VVol storage, you will need vSphere 6.0 or higher and certified array vendor Virtual Volumes software (VASA Provider).

Q3.What are the hardware requirements for Virtual Volumes based storage?

A. Details on supported array models is available at the hardware compatibility guide page:

http://www.vmware.com/resources/compatibility/

Q4.Where can I get the storage array vendor Virtual Volumes software?

A. Storage vendors are providing Virtual Volumes integration in different ways. Please contact your storage vendor for more details or visit your vendor's website for more information on Virtual Volumes integration.

Key Elements of Virtual Volumes Storage Framework

Q5.What is a Protocol Endpoint (PE)?

A. Protocol endpoints are the access points from the hosts to the storage systems, which are created by storage administrators. All path and policies are administered by protocol endpoints. Protocol Endpoints are compliant with both, iSCSI and NFS. They are intended to replace the concept of LUNs and mount points.

Q6.What is a storage container and how does it relate to a VVol Datastore?

A. A storage container is a logical abstraction on to which Virtual Volumes are mapped and stored. Storage containers are setup at the array level and associated with array capabilities. vSphere will map storage containers to VVol Datastores and provide applicable datastore level functionality. The VVol Datastore is a key element and it allows the vSphere Admin to provision VMs without depending on the Storage Admin. Moreover, the VVol Datastore provides logical abstraction for managing a very large number of Virtual Volumes. This abstraction can be used for better managing multi-tenancy, various departments within a single organization, etc.

Q7.How does a PE function?

A. A PE represents the IO access point for a Virtual Volume. When a Virtual Volume is created, it is not immediately accessible for IO. To Access Virtual Volumes, vSphere needs to issue a "Bind" operation to a VASA Provider (VP), which creates IO access point for a Virtual Volume on a PE chosen by a VP. A single PE can be the IO access point for multiple Virtual Volumes. "Unbind" Operation will remove this IO access point for a given Virtual Volume.

Q8.What is the association of the PE to storage array?

A. PEs are associated with arrays. One PE is associated with one array only. An array can be associated with multiple PEs. For block arrays, PEs will be a special LUN. ESX can identify these special LUNs and make sure that visible list of PEs is reported to the VP. For NFS arrays, PEs are regular mount points.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Q9.What is the association of a PE to storage containers?

A. PEs are managed per array. vSphere will assume that all PEs reported for an array are associated with all containers on that array. E.g: If an array has 2 containers and 3 PEs then ESX will assume that Virtual Volumes on both containers can be bound on all 3 PEs. But internally VPs and storage arrays can have specific logic to map Virtual Volumes and storage containers to PE.

Q10.What is the association of a PE to Hosts??

A. PEs are like LUNs or mount points. They can be mounted or discovered by multiple hosts.

VVol Storage Architecture and Technical Details

Q11.Can I have one PE connect to multiple hosts across clusters?

A. Yes. VPs can return the same Virtual Volumes binding information to each host if the PE is visible to multiple hosts.

Q12.I use multi-pathing policies today. How do I continue to use them with Virtual Volumes?

A. All multi-pathing policies today will be applied to a PE LUN. This means if path failover happens it is applicable to all Virtual Volumes bound on that PE. Multi-pathing plugins have been modified not to treat internal Virtual Volume error conditions as path failure. vSphere will make sure that if older MPPs won't claim PE LUNs.

Q13.How many Storage Containers can I have per storage array?

A. It depends on what array is configured, but number of containers is normally small count in (few tens in number). There is a limit of 256 storage containers per host.

Q14.Can a single VVol Datastore span different physical arrays?

A. No. At this point vSphere does not support this.

Q15. We have the multi-write VMDK feature today? How will that be represented in Virtual Volumes?

A. A Virtual Volume can be bound by multiple hosts. vSphere provides multi-writer support for Virtual Volumes.

Q16.Can I use VAAI enabled storage arrays along with Virtual Volumes enabled arrays?

A. Yes. vSphere will use VAAI support whenever possible. In fact VMware mandates ATS support for configuring Virtual Volumes on SCSI.

Q17.Can I use legacy datastores along with Virtual Volumes?

A. Yes.

Q18.Can I replace RDMs with Virtual Volumes?

A. Virtual Volumes do not currently support SCSI-2 reservations. Hence Virtual Volumes cannot be used to substitute RDMs in all environments. Whenever an application requires direct access to the physical device, Virtual Volumes are not a replacement for pass-thru RDM (ptRDM). Virtual Volumes are superior to non-pass-thru RDM (nptRDM) in a majority of virtual disks use cases.

Q19.Can I use SDRS/SIOC to provision Virtual Volumes enabled arrays?

A. No. SDRS is not supported. SIOC will be supported as we support IO scheduling policies for individual Virtual Volumes.

Q20.Is VASA 2.0 a requirements for Virtual Volumes support?

A. Yes. Virtual Volumes does require VASA 2.0. The version 2.0 of the VASA protocol introduces a new set of APIs specifically for Virtual Volumes that are used to manage storage containers and Virtual Volumes. It also provides communication between vCenter, hosts, and the storage arrays.

Q21.How does Virtual Volumes affect backup software vendors?

A. Virtual Volumes are modeled in vSphere exactly as today's virtual disks. The VADP APIs used by backup vendors are fully supported on Virtual Volumes just as they are on vmdk files on a LUN. Snapshots created by backup software using VADP will look the same to both vSphere and the backup software as non-VVol based snapshots, though on-array the snapshots are actually Virtual Volume objects. Backup software using VADP should be unaffected.

Q22.Is Virtual SAN using some of the Virtual Volumes features under the covers?

A. Although VSAN presents some of the same capabilities (representing virtual disks as objects in storage, for instance) and introduces the ability to manage storage SLAs on a per-object level with SPBM, it does not use the same mechanisms as Virtual Volumes. 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 Virtual Volume array's capabilities and present a single, uniform way of managing storage profiles and virtual disk requirements.

Q23.Can you shrink or grow the Storage Container on the fly?

A. Storage Containers are a logical entity only and are entirely managed by the storage array. In theory there's nothing to prevent them from growing and shrinking on the fly. That capability is up to the array vendor to implement.

Q24.Where are the Protocol Endpoints (PE) setup? In the vCenter with vSphere web client?

A. PEs are configured on the array side and vCenter is informed about them automatically through the VASA Provider. Hosts discover SCSI PEs as they discover today's LUNs; NFS mount points are automatically configured.

Q25.Where are the array policies (snap, clone, replicate) applied?

A. Each array will have a certain set of capabilities supported (snapshot, clone, encryption, etc) and these are defined at the storage container level. In vSphere, a policy is a combination of multiple capabilities and when a VM is provisioned, recommended datastores that match a policy are presented

Q26.From the vSphere Admin, is a Virtual Datastore accessed like a LUN? (storage browser, vm logs, vm.vpx config file, etc)

A. You can browse a Virtual Datastore as you browse any other datastore (you will see all config Virtual Volumes, one per VM). It's these config Virtual Volumes that hold the information previously in the VM's directory, like vmx file, VM logs, etc.

Q27. Is there a maximum number of Virtual Volumes or maximum capacity for a SC/PE?

A. Those limits are entirely determined by the array vendor's implementation. The vSphere Virtual Volume implementation does not impose any particular limits.

Q28.Can you shrink the PE as well as increase on the fly?

A. PEs don't really have a size - they're just a conduit for data traffic. There's no file system created ON the LUN, for instance - they're just the destination for data I/O on Virtual Volumes that are bound to it.

Q29.Does the virtual disk VVol contain both the .vmdk file and the -flat.vmdk file as a single object?

A. The .vmdk file (the virtual disk descriptor file) is stored in the config VVol with the other VM description information (vmx file, log files, etc.). The VVol object takes the place of the -flat file you mention - the vmdk file has the ID of the VVol.

Q30.How many Virtual Volumes will be created for a VM? Do we have a different Virtual Volume for flat.vmdk and different VVol for .vmx files etc?

A. For every VM a single VVol is created to replace the VM directory in today's system. That's 1 so-called "config" VVol. Then there's 1 Virtual Volume for every virtual disk, 1 VVol for swap if needed, and 1 Virtual Volume per disk snapshot and 1 per memory snapshot. The minimum is typically 3 Virtual Volumes/VM (1 config, 1 data, 1 swap). Each VM snapshot would add 1 snapshot per virtual disk and 1 memory snapshot (if requested), so a VM with 3 virtual disks would have 1+3+1=5 Virtual Volumes. Snapshotting that VM would add 3+1=4 Virtual Volumes.

Q31.From a vSphere perspective, are snapshots allowed to be created at individual VMDK (Virtual Volume) level?

A. The APIs provide for snapshots of individual Virtual Volumes but note that the UI only provides snapshots on a per-VM basis, which internally translates to requests for (simultaneous) snapshots of all a VM's virtual disks. It is not per LUN, though.

Q32.Will PowerCLI provide support for native Virtual Volumes cmdlets?

A. Yes.

Q33.Is there documentation that shows exactly which files belongs to each type of Virtual Volume object?

A. There are no "files". The Virtual Volumes are directly stored by the storage array and are referenced by the VM as it powers on. There is metadata information linking a VVol to a particular VM so a management UI could be created for the Virtual Volumes of a VM.

Q34. Is there any NFS or SCSI conversions going on under PE or array side?

A. Storage Containers are configured and the PEs set up to use NFS or SCSI for data in advance. The NFS traffic is unaltered NFS using a mount-point and a path. SCSI I/O is directed to a particular Virtual Volume using the secondary LUN ID field of the SCSI command.

Q35. Will it be possible to use storage policy definitions to automate VM's to choose different SC / array like flash, etc?

A. Yes. SPBM will separate all datastores (VMFS, VSAN, VVol, NFS) into "Compatible" and "Incompatible" sets based on the VM's requirement policy.