FREQUENTLY ASKED QUESTIONS

Virtual Volumes

Beta Program

Product

Q1. What is Virtual Volumes?

A. Virtual Volumes is a storage management and integration framework with SAN/NAS arrays that exposes virtual disks as native storage containers and enables array-based operations at the virtual disk granular level.

Q2. How does Virtual Volumes fit into VMware's SDS Strategy?

A. VMware defines Software-Defined Storage as the ability to dynamically compose storage services, aligned on application boundaries. There are 2 key architectural elements to SDS:

- **Virtual Data Plane:** This is a new data persistency layer. The data plane is virtualized in that storage containers are aligned with application containers. Just like for compute, this abstraction is independent of the underlying physical array and representation (LUN, file system, object, etc.). In the virtual data plane, characteristics of underlying infrastructure are expressed to the hypervisor to the control plane for service composition. Virtual Data services are composed and applied along precise application container boundaries. Virtual Volumes is an implementation of a Virtual Data plane.
- **Policy-driven control plane:** This is the new management layer that provides common orchestration and automation for service composition, placement, ongoing optimization, and change management. The Policy-driven control plane interprets storage requirements of individual applications specified in policies associated with individual VMs and composes the storage service placing the VM on the right storage tier, allocating capacity and instantiating the necessary data services (snapshots, replication, etc.).

Beta

Q3. How do I sign up for the Beta?

A. The Virtual Volumes Beta is run through the vSphere Beta program. The vSphere Beta is open to everyone to sign up and provide feedback. Sign up in the vSphere Beta community page:

- https://communities.vmware.com/community/vmtn/beta/vsphere-beta/vvols

Q4. What do I need to participate in the Beta?

A. You need the vSphere beta bits and your equivalent array vendor Virtual Volumes beta bits. Please reference the Virtual Volumes Beta community page for the list of Beta participating array vendors:

- https://communities.vmware.com/community/vmtn/beta/vsphere-beta/vvols

Technical

Q5. Where can I get the storage array vendor Virtual Volumes beta bits?

A. The Virtual Volumes Beta community page has a list of participating vendors and a link to their download page and/or contact. Please follow the instructions provided by the partner.

Q6. I need a technical overview of Virtual Volumes. Where can I find it?

A. Please refer to documentation and presentation available in the Virtual Volumes community beta pages.

Q7. What are Protocol Endpoints (PEs)?

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.

Q8. How does a PE function?

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

Q9. What is the association of the PE to storage array? (eg; is it one PE per 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 special LUN. ESX can identify these special LUNs and make sure that visible list of

Virtual Volumes

Beta Program

PEs is reported to VP. For NFS arrays, PEs are regular mountpoints.

Q10. What is the association of a PE to storage containers? (eg; Is it one PE per container or several containers?)

A. Like containers, 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 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 some specific logic to maps Virtual Volumes and storage containers to PE.

Q11. What is the association of a PE to Hosts??

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

Q12. Can I have one PE connect to multiple hosts across clusters?

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

Q13. 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 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.

Q14. What is a storage container?

A. Storage container is logical abstraction to which Virtual Volumes are mapped. vSphere will map storage container to datastore (of type Virtual Volume) and provide applicable datastore level functionality.

Q15. Why do I need storage containers?

A. Storage container provides logic abstraction for managing very large number of Virtual Volumes. This abstraction can be used for managing multi-tenancy, various departments within single organization, etc. Storage containers can also be used to set capacity limits for a given logical grouping of Virtual Volumes Various SPBM capabilities are associated per container.

Q16. Do I still need to setup Datastores and why?

A. Yes. Storage container is equivalent to a datastore in vSphere. The reason for the concept of a Virtual Volume datastore has to do with the vSphere infrastructure built around the VMFS-centric concept of a volume or datastore. For instance, a storage migration of a VM can occur between a VMFS datastore and a Virtual Volume datastore with little conceptual change in the workflow of migration. The reason why there is the dual-concept of a storage container and a datastore though is because a storage container is a combination of a logical pool of storage known by the storage array, an association to protocol endpoints, and a means for providing default QoS levels (default storage policies and storage capabilities). This also takes into consideration the nuances of configuring storage on a per-vendor basis and abstracts it into a uniform means for managing the logical pool from vSphere's perspective. None of these concepts apply to the traditional managed object we know as a datastore itself, hence the duality. Depending on setup, user doesn't want to see all containers on all hosts. Hence by providing datastore setup steps, vSphere allows user to configures containers on hosts of its choice.

Q17. 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.

Q18. Can a single Storage Container span different physical arrays?

A. No. In 2015, vSphere will not support this.

Q19. 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.

Q20. 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.

Virtual Volumes

Beta Program

Q21. Can I use VSAN and Virtual Volumes together?

A. Yes.

Q22. Can I use legacy datastores along with Virtual Volumes?

A. Yes.

Q23. Where can I learn about Storage Profile Based Management (SPBM) and its interaction with Virtual Volumes?

A. Visit the documents section of the Virtual Volumes community and page and download the document titled "vsphere-virtual-volumes-chapter2-b2.pdf" which talks about SPBM in details.

Q24. 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.

Q25. Can I use SDRS/SIOC to provision Virtual Volumes enabled arrays?

A. No. SDRS I not supported. SIOC is support given the support for IO scheduling policies for individual Virtual Volumes).

Q26. Is VASA 2.0 a requirement for Virtual Volumes?

A. Yes. Virtual Volumes requires 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.

Q27. How does Virtual Volume affects backup software vendors?

A. Virtual Volumes are modeled in vSphere exactly as today's virtual disks. The VADP APIs backup vendors use are fully supported on Virtual Volumes just as they are on vmdk files on a LUN. Backup software using VADP should be unaffected.

Q28. Are Virtual Volume features implemented in Virtual SAN?

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 a Virtual Volume array's capabilities and present a single, uniform way of managing storage profiles and virtual disk requirements.

Q29. Can the size of Storage Containers be adjusted?

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.

Q30. 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.

Q31. Where are the array policies (snap, clone, replicate) applied? At the Storage Container or the Protocol Endpoint?

A. The capabilities are associated with the Storage Controller and not the Protocol Endpoints.

Q32. From the vSphere Admin, is a SC accessed like a LUN? (storage browser, vm logs, vm.vpx config file, etc)

A. You can browse a Storage Container as you browse any other datastore (you'll 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.

Q33. 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

Virtual Volumes

Beta Program

implementation does not impose any particular limits.

Q34. From the vSphere Admins point of view, is a PE/SC accessed just like a LUN (e.g. using Datastore browser, accessing VM logs, accessing the .vmx, etc.)?

A. Yes.

Q35. Can the size of a PE can be adjusted?

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.

Q36. Does the virtual disk VIRTUAL VOLUME 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 Virtual Volume with the other VM description information (vmx file, log files, etc.). The Virtual Volume object takes the place of the -flat file you mention - the vmdk file has the ID of the Virtual Volume.

Q37. How many Virtual Volumes will be created for a VM?

A. For every VM a single Virtual Volume is created to replace the VM directory in today's system. That's 1 so-called "config" Virtual Volume. Then there's 1 Virtual Volume for every virtual disk, 1 Virtual Volume 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.

Q38. 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's not per LUN, though.

Q39. Will PowerCLI provide support for native Virtual Volumes cmdlets?

A. Yes.

Q40. Are there files associated for each 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 Virtual Volume to a particular VM, so a management UI could be created for the Virtual Volumes of a VM.

Q41. 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.

Q42. Will it be possible to use storage define policy to automate VM's to choose different SC / array like flash, etc?

A. Yes. SPBM will separate all datastores (VMFS, VSAN, Virtual Volume, NFS) into "Compatible" and "Incompatible" sets based on the VM's requirement policy. Note, of course, that using vendor-specific capabilities will immediately limit that list to that vendor.