What's New: vSphere Virtual Volumes

VMware Storage Business Unit Documentation
v 1.5/August 2015

Contents

Introduction

1.1 Software Defined Storage

VMware's Software-Defined Storage vision and strategy is to drive transformation through the hypervisor, bringing to storage the same operational efficiency that server virtualization brought to compute.

As the abstraction between applications and available resources, the hypervisor can balance all IT resources – compute, memory, storage and networking – needed by an application. With server virtualization as the de-facto platform to run enterprise applications, VMware is uniquely positioned to deliver Software-Defined Storage utilizing the pervasiveness of this software tier.

By transitioning from the legacy storage model to Software-Defined Storage with Virtual Volumes, customers will gain the following benefits:

- **Automation of storage "class-of-service" at scale:** Provision virtual machines quickly across data center using a common control plane (SPBM) for automation.
- **Self-Service capabilities:** Empower application administrators with cloud automation tool integration (vRealize Automation, PowerCLI, OpenStack).
- **Simple change management using policies:** Eliminate change management overhead and use policies to drive infrastructure changes.
- **Finer control of storage class of service:** Match VM storage requirements exactly as needed with class of service delivered per VM.
- **Effective monitoring/troubleshooting with per VM visibility:** Gain visibility on individual VM performance and storage consumption.
- **Non-disruptive transition:** Use existing protocols (Fiber channel, ISCSI, NFS) across heterogeneous storage devices.
- **Safeguard existing investment:** Use existing resources more efficiently with an operational model that eliminates inefficient static and rigid storage constructs.

The goal of Software-Defined Storage is to introduce a new approach that enables a more efficient and flexible operational model for storage in virtual environments. This is accomplished in two ways:

- The abstraction of the Virtual Data Plane enables additional functions that an array may provide to be offered as data services for consumption on a per-VM basis. Current implementation, data services are bound to the array for the most part. Data Services can provide functionality such as compression, replication, caching, snapshots, de-duplication, availability, migration and data mobility, performance capabilities, disaster recovery, and other capabilities. While the data services may be instantiated at any level of the infrastructure, the virtualized data plane allows for these services to be offered via policy on a per-VM basis.

- Implementing an automation layer that enables dynamic control and monitoring of storage services levels to individual virtual machines across heterogeneous devices – VMware refers to this as the Policy-Driven Control Plane

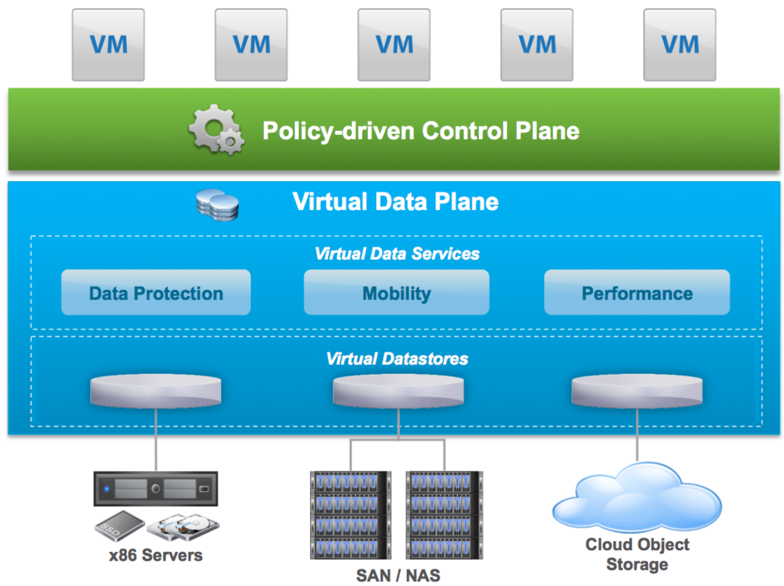

Figure 1: Software-Defined Storage Conceptual Diagram

1.1.1 Virtual Data Plane

The virtual data plane is responsible both for storing data and applying data services (compression, replication, caching, snapshots, de-duplication, availability, etc). While data services may be provided by a physical array or implemented in software, the virtual data plan abstracts the services and will present them to the policy-driven control plane for consumption and applies the resultant policy to the objects in the virtual datastore.

In today's model, the data plane operates on rigid infrastructure-centric constructs (LUNs or storage volumes) that are typically static allocations of storage service levels (capacity, performance and data services), independently defined from applications.

In the VMware Software-Defined Storage model, the data plane is virtualized by abstracting physical hardware resources and aggregating them into logical pools of capacity (virtual datastores) that can be more flexibly consumed and managed.

Additionally, to simplify the delivery of storage service levels for individual applications, the virtual data plane makes the virtual disk the fundamental unit of management around which all storage operations are controlled. As a result, exact combinations of data services can be instantiated and controlled independently for each VM. For each virtual machine that is deployed, the data services offered can be applied individually depending on the vendor implementation: Each application can have its own unique storage service level and capabilities assigned to it at its time of creation.

This allows for per-application storage policies, ensuring both simpler yet individualized management of applications without the requirement of mapping applications to broad infrastructure concepts like a physical datastore.

In the Software-Defined Storage environment, the storage infrastructure expresses the available data services and capabilities (compression, replication, caching, snapshots, de-duplication, availability, etc) to the control plane to enable automated provisioning and dynamic control of storage services levels through programmatic APIs. These storage services may come from many different locations: Directly from a storage array, from a software solution within vSphere itself, or from a third party location via API.

These capabilities are given to the control plane for consumption and expression by policies. The ability to pull in multiple sources of data services and abstract them to a policy engine gives the administrator the ability to create unique policies for each VM in accordance with their business requirements, consuming data services from different providers in each.

VMware's implementation of the virtual data plane is delivered through Virtual Volumes for external SAN/NAS arrays and Virtual SAN for x86 hypervisor-converged storage.

1.1.2 Policy-Driven Control Plane

In the VMware Software-Defined Storage model, the control plane acts as the bridge between applications and storage infrastructure. The control plane provides a standardized management framework for provisioning and consuming storage across all tiers, whether on external arrays, x86 server storage or cloud storage.

The policy-driven control plane is the management layer responsible for controlling and monitoring storage operations. In today's model, the control plane is typically, tied to each storage device – each array is operated in a different way - and implements a "bottom-up" array-centric approach in which storage service levels are aggregated into physical tiers or "classes of services", which are static pre-allocations of resources and data services tied to the infrastructure.

Upon provisioning, an application is rigidly mapped to these pre-configured storage containers. These storage containers are rarely aligned to precise application boundaries, and their capabilities need to be defined broadly to encompass the requirements of a broad set of applications.

This restricts the ability of a storage container to be focused specifically on the business requirements of an individual application. To circumvent this problem, storage administrators may be asked to create numerous purpose-built datastores, increasing the management overhead and complexity associated with storage.

Through Software-Defined Storage, the storage classes of service become logical entities controlled entirely by software and interpreted through policies. Defining and making adjustments to these policies enables automating the provisioning process at scale, while dynamically controlling individual service levels over individual virtual machines at any point in time.

This makes the Software-Defined Storage model able to more flexibly adapt to ongoing changes on specific application requirements. Policies become the control mechanism to automate the monitoring process and to ensure compliance of storage service levels throughout the lifecycle of the application.

The control plane is programmable via public APIs, used to consume and control policies via scripting and cloud automation tools, which in turn enable self-service consumption of storage to application tenants as well as a variety of external management frameworks.

VMware's implementation of the policy-driven control plane is delivered through Storage Policy-Based Management (SPBM).

1.2 vSphere Virtual Volumes

Virtual Volumes is a new virtual machine disk management and integration framework that exposes virtual disks as primary unit of data management for storage arrays. This new framework enables array-based operations at the virtual disk level that can be precisely aligned to application boundaries. Virtual Volumes is composed of two key implementations:

- **Flexible consumption at the logical level**

Virtual Volumes virtualizes SAN and NAS devices by abstracting physical hardware resources into logical pools of capacity (represented as Virtual Datastore in vSphere) that can be more flexibly consumed and configured to span a portion of, one or several storage arrays.

The Virtual Datastore defines capacity boundaries, access logic, and exposes a set of data services accessible to the virtual machines provisioned in the pool. Virtual Datastores are purely logical constructs that can be configured on the fly, when needed, without disruption and don't require formatting with a file system.

- **Finer control at the Virtual Machine level**

Virtual Volumes defines a new virtual disk container (the Virtual Volume) that is independent of the underlying physical storage representation (LUN, file system, object, etc.). In other terms, with Virtual Volumes the virtual disk becomes the primary unit of data management at the array level. This turns the Virtual Datastore into a VM-centric pool of capacity.

It becomes possible to execute storage operations with virtual machine granularity and to provision native array-based data services such as compression, snapshots, de-duplication, encryption, etc. to individual virtual machines. This allows admins to provide the right storage service levels to each individual virtual machine.

To enable efficient storage operations at scale, even when managing thousands of virtual machines, Virtual Volumes uses vSphere Storage Policy-Based Management (SPBM). SPBM is the implementation of the policy-driven control plane in the VMware Software-Defined Storage model.

- **Efficient operations through automation**

SPBM allows capturing storage service levels requirements (capacity, performance, availability, etc.) in the form of logical templates (policies) to which virtual machines are associated. SPBM automates virtual machine placement by identifying available datastores that meet policy requirements and coupled with Virtual Volumes, it dynamically instantiates necessary data services. Through policy enforcement, SPBM also automates service level monitoring and compliance throughout the lifecycle of the virtual machine.

The goal of Virtual Volumes is to provide a simpler operational model for managing virtual machines in external storage while leveraging the rich set of capabilities available in storage arrays. With Virtual Volumes, VMware seeks to transform a "bottom-up" operational model in which hardware is provisioned and then virtual machines are fit as best as possible, to a "top-down" operational model in which virtual machines requirements drive storage provisioning.

Virtual Volume transforms the data plane and control plane of supported SAN/NAS storage systems by aligning storage consumptions and operations with virtual machines. Virtual Volumes makes supported SAN/NAS storage systems VM-aware and unlocks the ability to leverage array based data services and storage array capabilities with a VM-centric approach at the granularity of a single virtual disk.

With Virtual Volumes most of the data operations such as snapshot, cloning, and migrations are offloaded to the storage arrays. New data management and monitoring operations, as well as communication mechanism have been implemented to manage the required communications between vSphere, storage arrays and the Virtual Volumes.

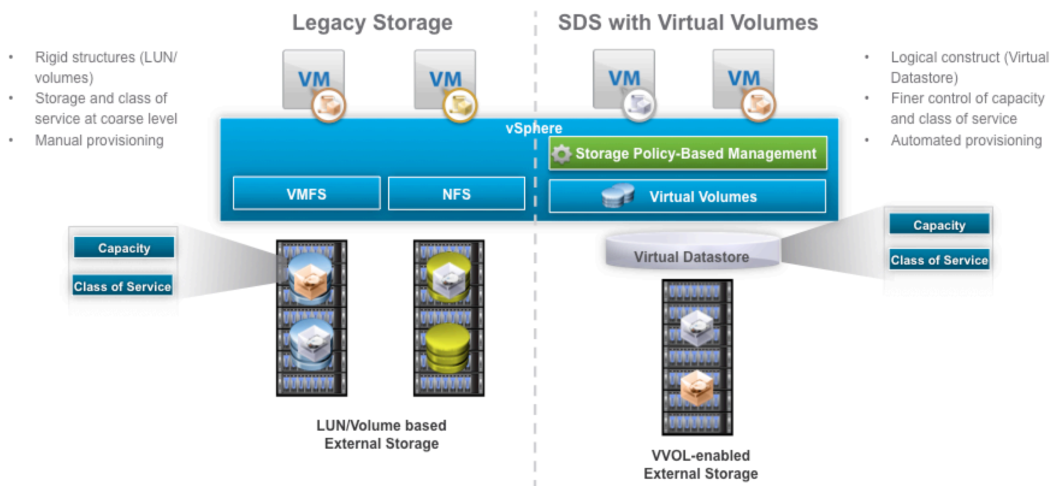

Figure 2: *Differences in Legacy storage and Software-Defined Storage with vSphere Virtual Volumes*

This model eliminates the complexity of managing the storage infrastructure and its services independently from the virtual machine consumers, and introduces a new control plane (SPBM) that centers on the virtual machine. Each virtual machine receives what it needs – no more no less – with a management framework that is common across heterogeneous storage systems.

In contrast, in the legacy model, each virtual machine needs to be mapped to a pre-defined storage tier and pre-configured storage container (LUN/Volume) and any virtual machine placed inside these structures will inherit its class of service. In the legacy model, it is common to manage different tiers of service using LUN/volumes.

For instance a "Gold LUN" would denote the highest level of service and any VMs placed inside this LUN would inherit the class of service of this LUN. In many cases, a LUN/volume with the right combination of capacity and storage class of service may not be available.

So, a virtual machine would be placed in a Gold LUN without requiring all the storage services. This leads to complexity in managing data services that are mandatorily inherited by all VMs in the Gold LUN: Perhaps not all systems need to be replicated but are nonetheless, as the Gold LUN is not only replicated but is also the highest performing, an attribute that is required by those systems. There is an operational cost associated with managing these interdependencies, as well as financial costs associated with the byproducts of overprovisioning, or overconsumption of bandwidth, etc. Thus, transitioning from legacy storage model to Virtual Volumes is similar to moving from a fixed menu consumption model to a la carte consumption model.

Systems often have changing requirements after being provisioned. Perhaps a given workload has end-of-month or seasonal performance spikes, or a less critical system has become mission-critical over time. In a traditional storage model it would require careful analysis of available capacity, LUN management, storage vMotion and possibly even re-provisioning of the application to meet the new requirements.

With Virtual Volumes and SPBM one can simply assign a new policy to the already deployed system, and the array will automatically move the system to a location where its new policy can be met or perhaps change the policy in-place, depending on the implementation. This allows for easy adaptation of even individual applications to meet changing business requirements simply by altering the policy and having the Virtual Volume respond.

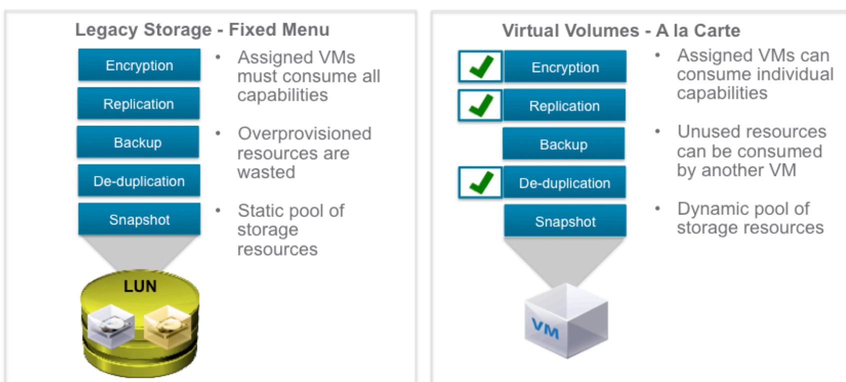

Figure 3: Finer control of class of services

Virtual Volumes enable rapid and dynamic provisioning of storage services to individual virtual machines at scale with a policy based approach that drives high levels of automation and simpler control.

In comparison, in the legacy model, hardware dictated virtual machine class of service and provisioning required manual operations causing slow deployment of virtual machines and a model that was not scalable.

LUNs and volumes, moreover, are pre-allocated storage containers whereas with Virtual Volumes, storage capacity and class of service are delivered dynamically on demand in real-time.

1.3 vSphere Virtual Volumes Architecture

Virtual Volumes creates a significantly different and improved logical storage architecture that allows operations to be conducted at the virtual machine level, while still using native array capabilities.

Additionally, Virtual Volumes eliminates the need for provisioning and managing large numbers of LUNs/Volumes per host to achieve application level services. This increases the manageability of the vSphere infrastructure by reducing operational overhead while enabling scalable data services on a per VM level.

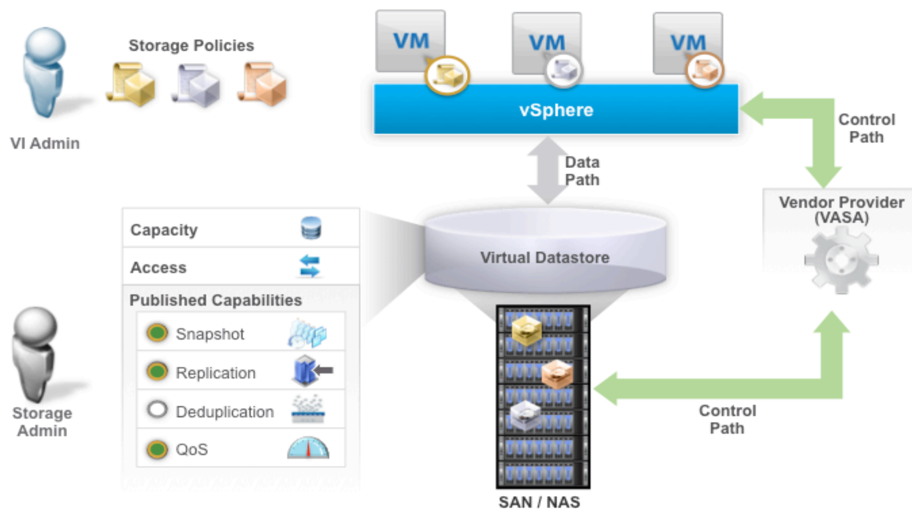

Figure 4: vSphere Virtual Volumes Logical Architecture

Virtual Volumes requires both VMware and storage vendor support. From a VMware perspective, the key elements in delivering Software-Defined Storage are vSphere Virtual Volumes, vSphere Storage-Policy Based Management, and vSphere APIs for Storage Awareness (VASA).

From a storage vendor perspective, arrays must be enabled for Virtual Volumes and vendors must deliver a vSphere API for Storage Awareness (VASA) provider based on the VASA APIs version 2.0. and used as out of band communication.

Once a Storage Admin defines the storage array capabilities that can be consumed, these capabilities are then surfaced to vSphere using the VASA APIs and are made available to the vSphere administrator for consumption. The vSphere administrator would then create storage policies from the set of the published storage array capabilities.

To allow a direct data path from the host to the underlying storage, a new data path construct called the Protocol Endpoint (PE) has also been introduced. This provides a path for the data that is separate for the policy and API communication, and allows for individual vendors to customize data path resilience dependent on their specific capabilities.

In order to manage, monitor and operate the new features and capabilities introduced by Virtual Volumes, a new set of communication mechanism and management constructs have been introduced.

1.3.1 Protocol Endpoints (PE)

Protocol Endpoints (PE) is a transport mechanism or access points that enable and manage access control and communications between ESXi hosts and storage array systems.

Protocol Endpoints are part of the physical storage fabric, and they are capable of establish data paths on demand from virtual machines to their respective virtual volumes. Protocol Endpoints inherit the access functionalities of today's LUNs and thereby can support all industry standard protocols.

This allows the end user to benefit from Virtual Volumes without any disruption to his existing configuration. The goal of Protocol Endpoints is to not only support the existing protocols, but to also help with scale. A single protocol endpoint can now bind to a multitude of Virtual Volumes as opposed to a LUN with a defined upper limit.

Protocol Endpoints are compatible with all SAN/NAS industry standard protocols:

- iSCSI
- NFS v3
- Fiber Channel (FC)
- Fiber Channel over Ethernet (FCoE)

Protocol Endpoints are setup and configured by Storage administrators.

1.3.2 Storage Containers (SC)

Virtual Volumes reside in storage containers. Storage Containers are logical storage constructs that are defined and setup by storage administrators. These containers are used to define:

- Storage capacity allocations and restrictions
- Storage policy settings based on data service capabilities on a per virtual machine basis

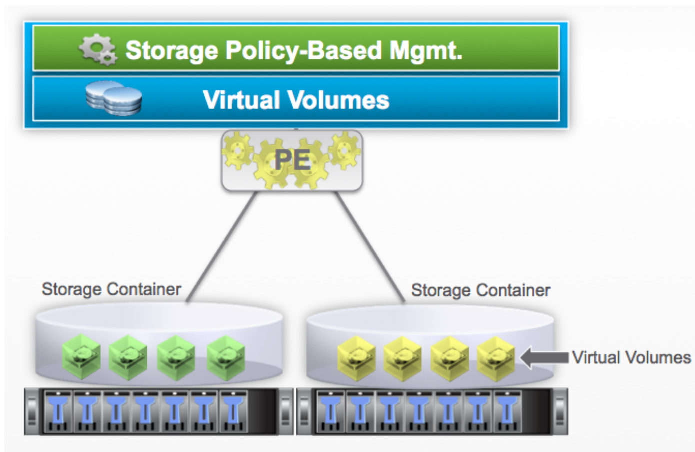

Figure 5: Storage Containers and Virtual Volumes

A Storage Container is a purely logical construct within an array that puts a user-defined limit around a set of storage capacity and capabilities that will be presented to vSphere as a virtual datastore.

This allows the storage administrator the ability to configure capacity and capabilities into logical groupings, and the vSphere administrator the ability to interact with familiar datastores without requiring insight into the physical constructs behind the Storage Container.

A storage Container is made accessible to an ESXi host and have a 1:1 relationship with storage containers created on the storage array. For example if you create three storage containers on the array you will then have three corresponding VVol datastores available to the connecting hosts

Similar to Protocol Endpoints, Storage Containers are setup and configured by the storage administrator and are later automatically discovered in vSphere when adding virtual datastores.

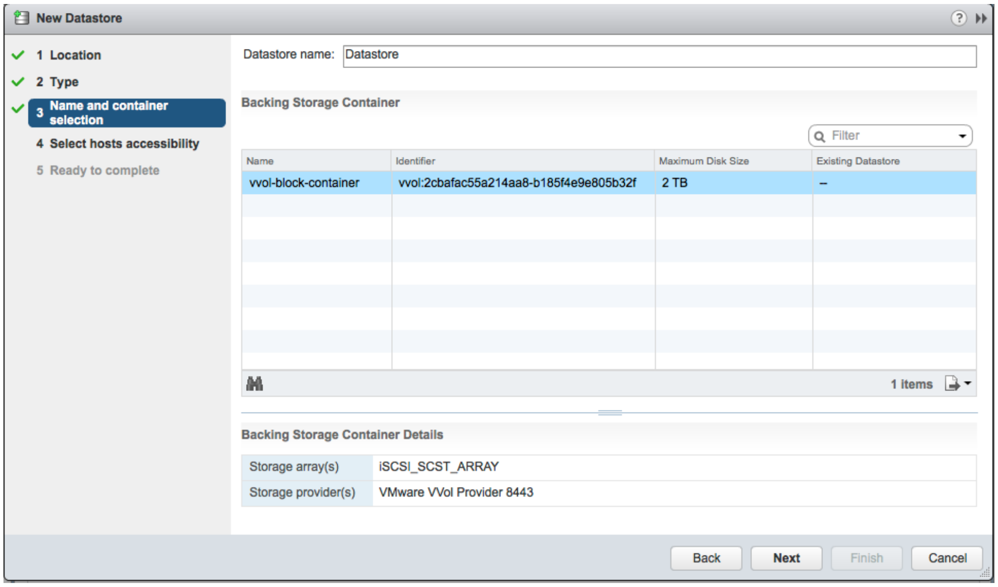

Figure 6: Adding a Virtual Volume Datastore in vSphere

1.3.3 Vendor Provider (VP)

Vendor Providers (VP), also referred to as VASA Provider, is a storage-side software component (Plug-In) that act as a storage awareness service for vSphere. Storage vendors exclusively develop vendor providers.

ESXi hosts and vCenter Server connect to the Vendor Provider and obtain information about available storage topology, capabilities, and status. Subsequently vCenter Server provides this information to vSphere clients, exposing the capabilities around which the administrator might craft storage policies in SPBM.

Vendor Providers can be implemented in the form of a virtual appliance, or directly from the array's management plane. This is the means by which the Vendor Provider communicates the array capabilities is through the vSphere Storage APIs to export storage array features and capabilities and present them to vSphere through the VASA APIs for granular per-VM consumption.

Vendor Providers are typically setup and configured by the vSphere administrator in one of two ways:

- Automatically via the array vendors plug-in
- Manually through the vCenter Server

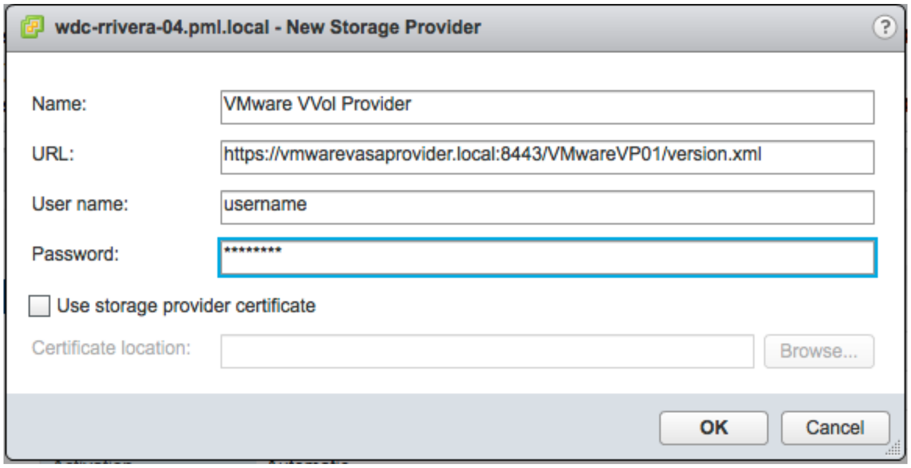

Figure 7: Vendor VASA Provider Manual Registration

1.3.4 Virtual Volumes (VVols)

Virtual Volumes are a new type of virtual machine objects, which are created and stored natively on the storage array. VVols are stored in storage containers and mapped to virtual machine files/objects such as VM swap, VMDKs and their derivatives.

There are five different types of Virtual Volumes object types and each of them map to a different and specific virtual machine file.

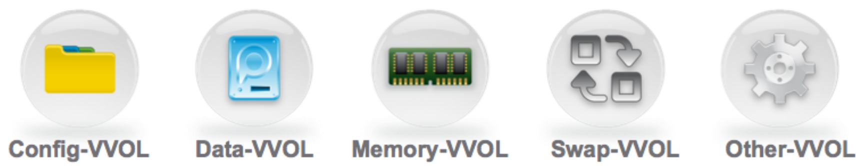

Figure 8: vSphere Virtual Volumes Object Types

- **Config-VVol** – VM Home, configuration related files, logs, etc
- **Data-VVol** – Equivalent to VMDKs
- **Memory-VVol** – Snapshots
- **Swap-VVol** – VM memory swap
- **Other-VVol** – Generic type for solution specific

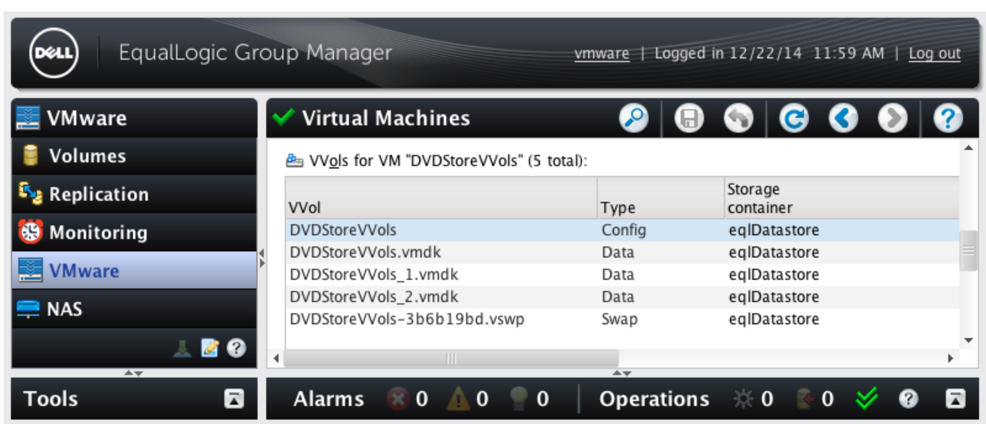

Figure 9: Storage Admin View of VVols on Storage Array

1.4 Architecture Comparison

In traditional vSphere and storage virtualization architectures for virtual machines, the datastore (a LUN formatted as VMFS or an NFS mount point) serves the dual purposes of being an endpoint to receive SCSI or NFS read and write commands, and being a storage container for a number of virtual machines metadata and data files.

In a Virtual Volumes architecture, the Protocol Endpoints act as the network transport mechanism, and the logically defined Storage Containers acts as virtual datastores. A default storage policy can be attached to a storage container and passed to each virtual machine located within.

However, the virtual machines within a storage container do not have to share a set of global properties defined by the membership in that storage container. Rather, each virtual machine can have an individual policy attached that can be persistent across storage containers. Because they are logically defined and not tied to a static LUN or disk group, they can instantiate a virtual machine anywhere with a container that will satisfy the policy attached to it.

A Storage Container may contain storage areas with different RAID configurations, different performance configurations, etc. and will deploy a virtual machine to a location that is appropriate. Multiple data services and different types of capabilities may be present within a logically defined Storage Container.

This very fact allows virtual machines to have their own storage policy, and yet be placed appropriately within the container alongside other virtual machines, each with their own policy that is not mandatorily shared by all VMDKs stored within the container.

By storing a large number of virtual machines onto a single LUN or NFS mount-point provided by a storage container reduces the number of entities on the fabric and their configuration and scalability overhead, but limits the granularity at which data services can be applied between hosts and the underlying storage.

In a standard storage operating model, one may deploy a large number of virtual machines onto a single datastore to reduce the number of LUN or NFS entities on the storage fabric, and minimize the configuration and scalability overhead.

Yet this approach limits the granularity with which data services can be delivered to virtual machines and requires broadly similar profiles across all applications in order to be successful.

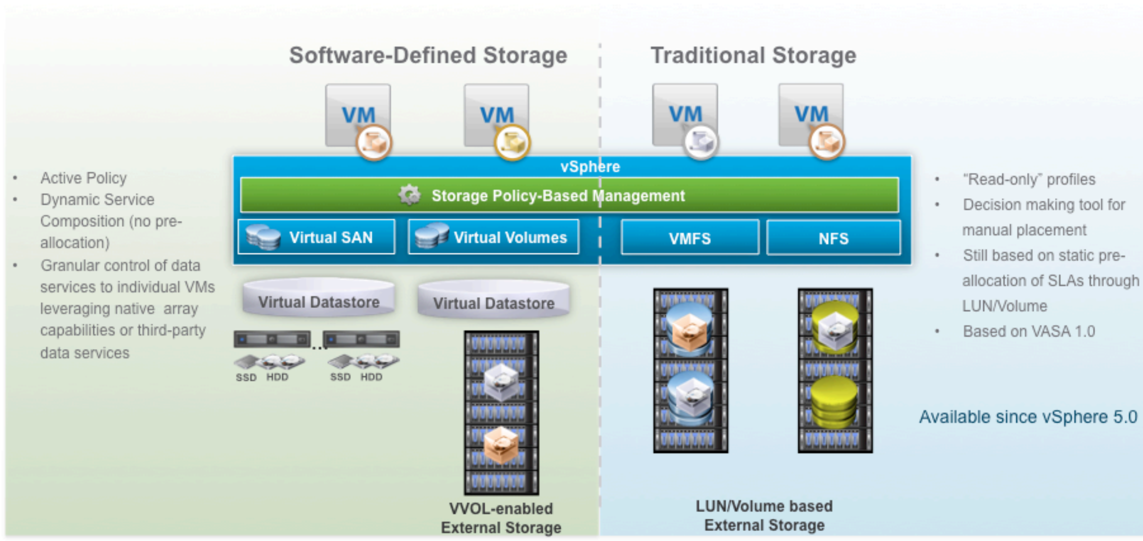

Figure 10: Traditional storage Architectures vs. Virtual Volumes SDS Architectures

In a Software-defined Storage environment with Virtual Volumes, the storage system provides a Protocol Endpoint, which is a discoverable entity on the physical fabric. For example, SAN storage systems could provide a proxy LUN Protocol Endpoint that is discovered using regular LUN discovery commands.

The Protocol Endpoint is accessible by multiple paths. Traffic on these paths will flow as defined by the path selection policy. The VMware path selection policies include:

- Most Recently Used
- Round Robin
- Fixed

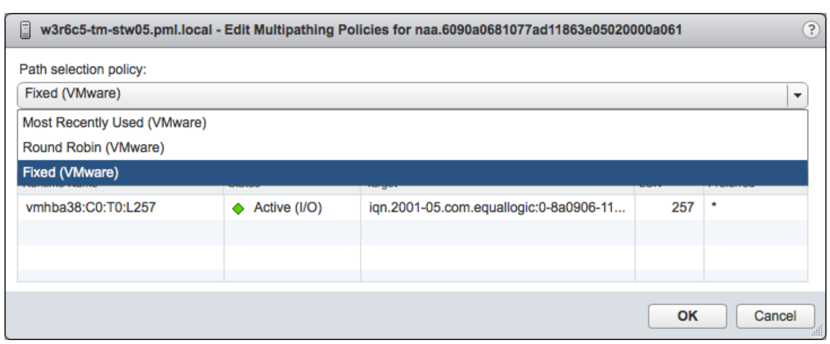

Figure 11: VMware Multipath Policies

An NFS server mount (IP address and a share name) is a Protocol Endpoint for NAS storage systems. Each VMDK deployed to NFS will be associated in a 1-to-1 relationship with a Virtual Volumes and each Virtual Volume will be its own storage container. When running on NFS, each VM becomes not only its own Virtual Volume, but its own storage container resident on the shared mount point.

vSphere feature	Non-VVol SAN/NAS	Virtual Volumes enabled SAN/NAS
VMFS	Clustered file system used to format LUN	Not needed
vSphere API for Storage Awareness (VASA)	Storage capabilities are exposed to the virtualization layer in the form of read-only profiles associated with datastores. This capability is currently available in vSphere through a feature called – Profile-driven storage. In the current implementation profiles are used as tags of virtual disks and datastores for simplified provisioning. They are not used to control the instantiation of individual data services For instance, if replication is set to "on" at the LUN, this capability cannot be modified from the VM policy	Storage Capabilities can be controlled and instantiated through active policy profiles associated with individual VM objects without the need of LUNs or Volumes. For instance, array replication can be can be set to "on" or "off" as required by the virtual machine. VVols uses VASA 2.0
vSphere API for Array Integration (VAAI)	Limited set of Block, NAS, and thin provisioning primitives.	Integrated with VVols.
vSphere Profile-Driven Storage	Allows tagging of virtual disks and based on read-only static profiles.	Evolves into vSphere Storage Policy-Based Management. SPBM will continue to support tagging for traditional storage as well as policy-based automation for Software-Defined Storage.
Storage DRS	Performed at LUN/volume level.	When supported by storage vendors, it will be performed at the Virtual Volumes level. With Storage DRS, an application or VM is moved to meet requirements.
Tiering	Tiers determined by the LUN level (e.g. gold, silver, bronze) and bound by the LUN size limit.	Determined by datastore capabilities provided by vendors. Datastores offer a wide range of capabilities and from this universe of capabilities; many combinations can be derived to create a policy. Capabilities are specific to vendor implementations.
Storage IO Control	Adjusted by administrators to control VM network loads.	When configured, it will allow end users to establish performance thresholds at the Virtual Volume level.

Figure 12 – vSphere Virtual Volumes Comparison Against External Architecture

Virtual Volumes does not replace VMFS or other existing file system and storage access protocols that are supported vSphere. The use of VMFS is still available and supported whenever needed in vSphere. The vSphere platform will be able to simultaneously connect both to Virtual Volumes and to traditional LUN based SAN/NAS arrays.

Depending on vendor implementation, a single array could be capable of support both VVols and LUN based operations. VVols present several advantages over traditional storage solutions from technical limitations perspective as well as operational impact.

1.5 Benefits of vSphere Virtual Volumes

Virtual Volumes delivers value in three main areas: storage operations, service level delivery, and resource utilization.

1.5.1 Simplify Storage Operations

For both the vSphere Administrator and Storage Admin, Virtual Volumes greatly simplifies management over the existing operational model. Virtual Volumes allows the separation of provisioning and consumption of storage for VMs.

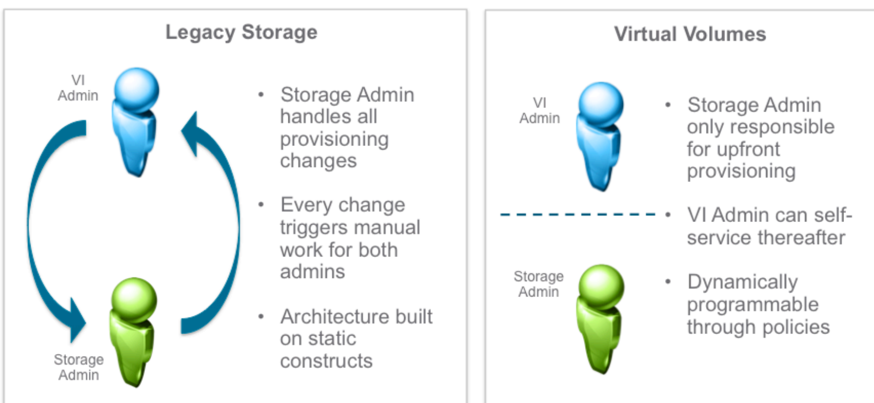

Figure 13: Separation of Consumptions and Provisioning

In the VMware Software-Defined Storage model with Virtual Volumes, the Storage Admin sets up an entity called the storage container. The capacity and data services published from the storage array by the Storage Admin in the storage container become similar to menu items from which the vSphere Administrator can consume on demand.

The Storage Admin retains control of the storage resources, as the vSphere Administrator can only consume the published storage array capabilities. However, the Storage Admin no longer needs to determine which data services should be assigned to an application. Thus, the Storage Admin is responsible for up front setup, but the vSphere Administrator is self-sufficient afterwards.

With Virtual Volumes, the vSphere Administrator gains control and becomes responsible for defining the various storage classes of service for applications. However, the classes of service are no longer physical pre-allocations, but instead they are logical entities controlled and automated entirely by software and interpreted through the mechanism of policies.

By associating one or many virtual machines to the right policy, the provisioning and instantiation of storage service levels are automated for that virtual machine or set of virtual machines. Automated policy enforcement also becomes the mechanism to simplify the monitoring process and to ensure compliance of storage service levels throughout the lifecycle of the application.

Policy-driven automation enables more agile storage consumption for virtual machines, which ultimately delivers faster provisioning for new applications with different requirements and simplifies change management, as the vSphere Administrator no longer depends on the Storage Admin to fulfill infrastructure change requests. The vSphere Administrator can make changes to policies at any time, and the necessary infrastructure changes are configured through automation. This allows for quicker adjustment to business changes.

Storage Admins	**Current**	**With Virtual Volumes**	**Benefit**
Datastore provisioning	Many standard size LUNs based policies – one for each application Capacity mgmt may be tied to LUNs Needs to format VMFS	Small number of datastores; but will need to be shared Capacity planning on a pool	Define capabilities to be consumed without static limitations Vendor mgmt tools many need to be updated
Troubleshooting and Monitoring	Keyed off from datastores	Based on VM/VMDK	VM-awareness on storage systems Vendor mgmt tools many need to be updated
Provisioning	To datastores Using tag based policies	To datastores With granular policy and ability to dynamically change he policies	Storage consumption has been automated. No need to map VMs to LUNs
Troubleshooting and Monitoring	Per VM Datastore driven Manual process to ensure policy compliance	Per-VM Vendor container Automated monitoring of policy compliance	VM-awareness at storage Full visibility of policy compliance VMware tools could be made better with Virtual Volumes

Table 1: Virtual Volumes vs. Current Storage limitations

1.6 Ecosystem Support

vSphere Virtual Volumes is an industry-wide initiative and represents the collaboration towards a new storage operational model for VMs. Several vendors have participated in the design of Virtual Volumes, helping drive the SDS vision for external storage. Although Virtual Volumes is an industry-wide initiative, partners can deliver unique solutions through their implementations.

Since VMworld 2012, partners have been showcasing their unique implementations and in VMworld 2014, three partners announced the availability of storage array models supporting Virtual Volumes.

Through VVols and SPBM, customers can safeguard existing array investments and transition non-disruptively to a simpler and more efficient operational model that works across all storage types.

Figure 14: Partner Ecosystem Supporting vSphere Virtual Volumes

1.7 Use Cases

vSphere Virtual Volumes enables private cloud deployments through integration with tools such as vRealize Automation, PowerShell, and OpenStack. The automation and control capabilities allow organizations to simplify storage management and more quickly deliver value to IT customers.

Through self-service capabilities, application administrators gain control with agility and are no longer constrained to bottlenecks that existed with service fulfillment in the traditional storage operational model. By adopting Virtual Volumes, customers can achieve scalable cloud deployments across all types of workloads.

Environments that have dynamic resource requirements can benefit from Virtual Volumes in conjunction with Storage Policy Based Management. For example during a seasonal workload increase, an automated process might update a given storage policy to deliver higher levels of performance.

The simple change to one policy will be inherited by all virtual machines to which that policy applies, and the implementation of these virtual machines on a Virtual Volume array would allow only those VMs to be adapted to meet the new policy.

Another example is the provisioning multi-tier applications with different storage requirements for each tier. This process no longer requires different storage configurations to satisfy their unique specifications.

Each virtual machine of an application can now simply have an appropriate policy attached to it at time of deployment and will be instantiated as a unique Virtual Volume with the required storage characteristics, even when deployed to the same storage container.

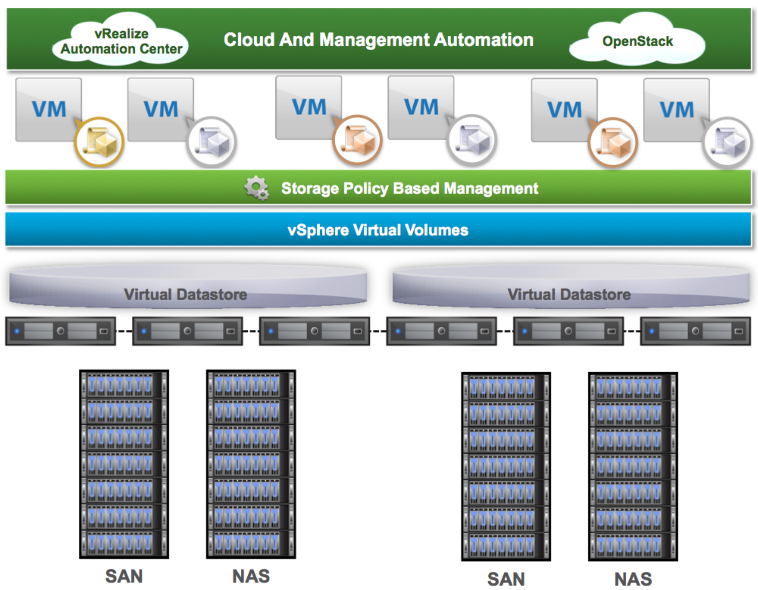

Figure 15: vSphere Virtual Volumes Enables Self-service and Private Cloud Deployments

1.8 vSphere Requirements

1.8.1 vCenter Server & vSphere Hosts

The current or latest release of vSphere 6.0 builds is required to use vSphere Virtual Volumes. The software and all of its related components are available directly from VMware product download site.

1.9 Storage Requirements

1.9.1 Storage Arrays

Virtual Volumes requires a compatible storage array system. In most cases, a software solution such as a virtual storage appliance from one of the supporting vendors is supported for testing management workflows, operations, and functionalities.

Depending on the vendor specific implementation, storage array system may or may not require a firmware upgrade in order to support vSphere Virtual Volumes. Check with your storage vendor for detailed information and configuration procedures.

Conclusion

Virtual Volumes and SPBM bring Software-Defined Storage benefits to SAN/NAS arrays by enabling policy-driven provisioning and control of storage services to individual VMs leveraging native array capabilities.

Virtual Volumes and SPBM expose virtual disks as native storage containers and enables granular array-based operations at the virtual disk level. Through automation and by establishing the VM as the unit of data, storage workflows are streamlined and underlying arrays are utilized more efficiently.

Virtual Volumes is an industry wide initiative, supported by all major vendors and will make provisioning of SAN/NAS arrays in vSphere environments more agile, simpler, and will reduce storage costs.

Acknowledgements

I would like to thank Suzy Visvanathan, Sr. Product Manager, Juan Novella, Product Marketing Manager, Chris McCain, Director of Technical Certification in the Network and Security Business Unit, Ken Werneburg, Sr. Technical Marketing Architect and Chuck Hollis, Chief Strategist of the Storage and Availability Business Unit for their invaluable contributions to this paper.

About the Author

Rawlinson is a Senior Technical Marketing Architect in the Storage and Availability Business Unit at VMware, Inc. focused on Software-Defined Storage technologies such as VMware Virtual SAN, vSphere Virtual Volumes and other storage related products.

Previously, as an architect in VMware's Cloud Infrastructure & Management Professional Services Organization he specialized in VMware vSphere and Cloud enterprise architectures for VMware fortune 100 and 500 customers.

Rawlinson is amongst the few VMware Certified Design Experts in the world and is the author of multiple publications based on VMware and other technologies.

• Follow Rawlinson's blogs:

http://blogs.vmware.com/vsphere/storage
http://www.punchingclouds.com

• Follow Rawlinson on Twitter:

@PunchingClouds

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