VSPHERE VIRTUAL VOLUMES (VVOLS)

VM-Centric Storage

MAY 2015

THE NET NET

In this Market Landscape Report, Taneja Group presents a survey of eleven companies that are supplying vSphere Virtual Volumes (VVol) arrays. The objective of this survey was to determine the capabilities that these storage arrays have with respect to their VVol implementation. In the end, we found that each vendor has surfaced up different capabilities to their arrays in different ways. It is our hope that this survey will enable senior business and technology leaders to decide which VVol-enabled array is best suited for their datacenter needs.

We based the information in this survey on a questionnaire made up of 32 questions, investigating public information about the vendors' VVol implementation, and, if necessary, one-on-one communication (either verbal or written) with representatives of these vendors. In some cases, we were able to get hands-on experience with a vendor's VVol-enabled array to get a deeper understanding of its implementation.

To assess each vendor's VVol offering, we evaluated factors such as the capabilities that they are able to surface up to vSphere, how those features are surfaced up, the technical aspects of their implementation like storage containers (SC) and protocol endpoints (PEs) as well as other unique features of their implementation of VVol. To summarize our findings, we categorized these offerings into three types of implementations using VM-level storage capabilities as a way to distinguish between types. A key takeaway is that VVol does not enable any pre-existing array limitations to magically disappear.

As a note of caution: VVol is a new capability, and many of the companies that we surveyed had yet to actually release a product. We took their answers in good faith and asked for clarification on their responses if they did not seem consistent with other information that we were able to obtain. The upshot of this methodology is that some of the results we obtained and conclusions we reached are based on "roadmap" features; as these products come to market, you should verify our information by checking out the released product.

INTRODUCTION TO VSPHERE VIRTUAL VOLUMES (VVOLS)

In February of 2015, VMware announced vSphere 6.0. For the most part, this was an evolutionary release for VMware, with one major exception: it included vSphere Virtual Volumes (VVol), a revolutionary new storage feature. VVol represents one of the biggest changes in vSphere's storage stack since its inception, and revolutionizes the way virtual machines (VMs) interact with storage.

VVol is revolutionary in the way it allows virtual machines (VMs) to consume storage. With VVol, the storage is managed on a per-VM basis instead of on a per-datastore basis. Prior to VVol storage in a virtual infrastructure was managed at the level of datastores, which were mapped to LUNs or volumes. Datastores had attributes such as performance and protection levels assigned to them, and the VMs assigned to the datastore inherited these attributes. VVol allows any of the attributes (capabilities) of the attached storage arrays to be surfaced up to vSphere and the VMs can use any of the capabilities of the storage rather than having these capabilities be based indirectly on the LUN or share. This is called VM-centric storage.

A key point to keep in mind when talking about VVol is that it is only a framework for abstracting the capabilities of an array and vSphere. Each storage vendor has the ability to present their pool of storage and its underlying capabilities as they see fit. vSphere is only concerned with the end result, not the underlying construct. In our research, we were surprised at how the different storage vendors went about implementing the VVol framework. In order to best summarize our research results, we created three categories of VVol implementation in which VVol-enabled arrays are placed. Before we discuss these categories, however, let's look at the components and architecture of VVol.

VVOL ARCHITECTURE AND COMPONENTS

VMware users have been struggling with storage since the inception of ESX. The basic problem is that all of our current storage schemes were designed in an age when we had a single server associated with a single storage array. VMware turned this paradigm on its head as we now have hundreds of virtual machines accessing a single array. vSphere worked at the VM level of granularity and the storage worked at the LUN level. We were forced to group VMs unnaturally into LUNs that had the capabilities we needed. What we needed, and what we lacked was VM-centric storage. VVol solves this problem by allowing each VM to have the storage capabilities that is needed to meet its SLAs.

Architecture

VVol allows a layer of abstraction and control between the hypervisor and storage, and separates storage's data plane from its control plane. "Protocol Endpoints" (PEs) are used by the data plane to transfer data between the hypervisor and storage. VASA 2.0 (vSphere APIs for Storage Awareness) is used by the control plane. A PE allows data to flow from the hypervisor to the storage array in a protocol agnostic manner.

The first generation of VASA (VASA 1) only allowed unidirectional communication from vSphere to the storage array, whereas VASA 2.0 allows bidirectional communication to/from vSphere and its attached storage. By allowing bidirectional communication, the storage can surface capabilities from the storage back to vSphere.

On the array, the storage administrator carves out pools of storage into storage containers (SCs) that have capabilities such as performance, encryption, protection, etc. The SCs are used to create a VVol datastore and a storage policy can then be created based on the SCs capabilities. An SC can have more than one policy attached to it. Each policy can have one or more rule-sets created from the SC's attributes that will be used to meet the requirements of a virtual machine (VM). When a VM is created, a storage policy is selected and all the datastores that can meet the requirements are displayed, one of which can be selected for use. Throughout the life of a VM, the storage requirements of a VM will often change; in order to accommodate these changes, the VM can be reassigned a storage policy.

vSphere Storage Policy Based Management (SPBM) is the basis of the control plane of VVol. SPBM allows the Virtual Infrastructure (VI) administrator to define the characteristics of the storage on which a VM will exist, based on a policy built on "rule-sets," in order to allow them to meet the SLAs of a VM. Without the benefit of SPBM it would require the VI administrator to communicate with the storage administrator about the needs of a VM and have them provision storage with the desired characteristics and then manually communicate back to the VI administrator the location of the storage and its characteristics back to the VI administrator. With SPBM the capabilities of the underlying storage is surfaced up to vSphere and the VI administrator. When a VM is created, it can use a policy that has the capabilities required and, through SPBM, the storage that has those capabilities will be presented to the VI administrator for use with that VM.

When the VM is created, the attributes that the VM needs are passed to the array and it is up to the array to meet these requirements. This is a new layer of abstraction for vSphere and the underling storage. Below is a diagram (Figure 1) that VMware has put together that explains the relationship of the various VVol components.

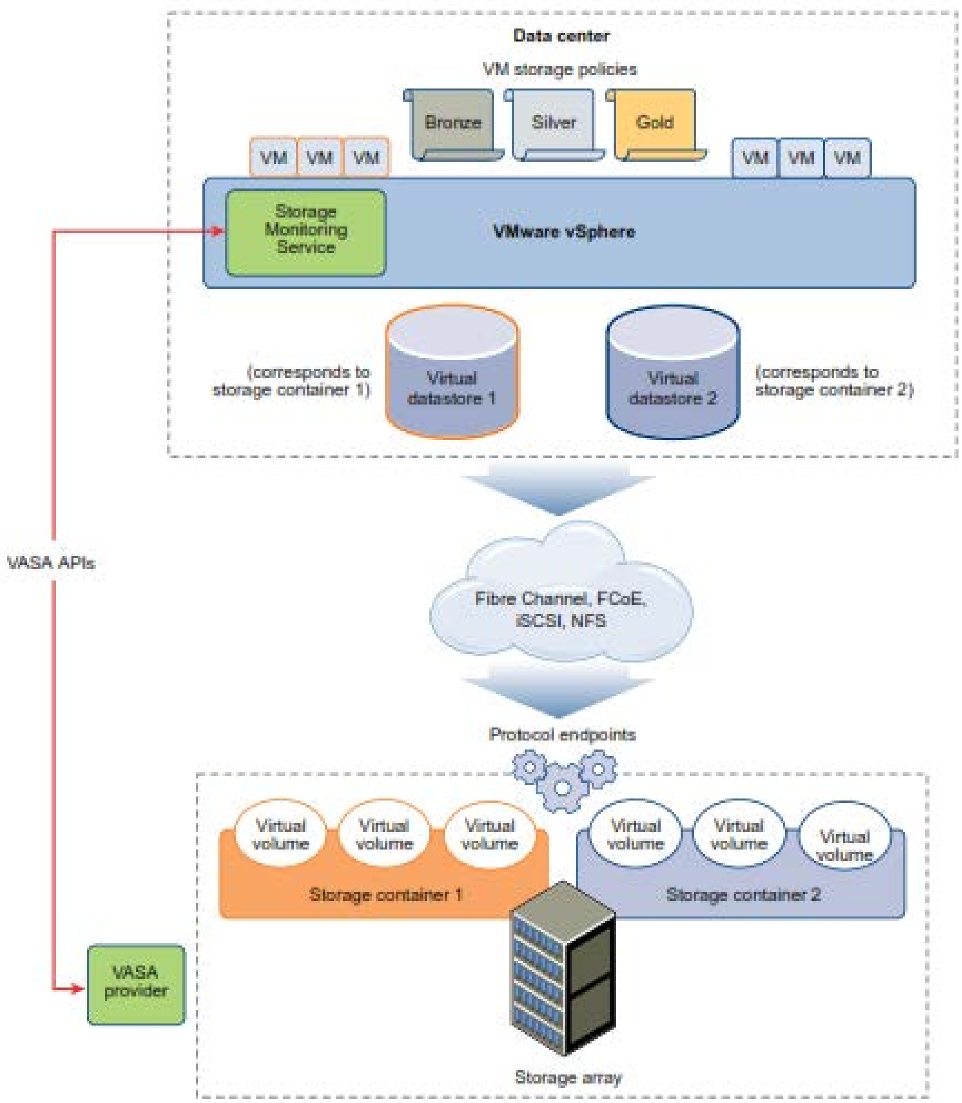

Figure 1

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com

Protocol Endpoint

A PE allows the data to flow from the hypervisor to the storage array, and acts as a proxy that allows the data to pass through to its Virtual Volume. It does not store any data and is simply a demultiplexer that passes data to a Virtual Volume. On a SAN-based storage system, the PE appears as a small LUN, and on a NAS-based storage system it appears as a mount point. Each storage container (SC) must have one or more PEs associated with it. A PE can only be associated with a single storage SC. However, multiple ESXi hosts can be attached to the same PE. Unified arrays (arrays that support both SAN and NAS) can have multiple PEs that point to the same storage container.

We have not seen any differentiators between LUN and Volume based PEs. There may be a performance impact based on queue depth as the number of PEs per SC increases, but we have not seen definitive proof of this.

Storage Container

The storage administrator carves out a pool of raw storage from their arrays into SCs. Each SC has capabilities associated with it such as capacity, performance, replication, encryption and protection. SCs can be used to partition an array's capacity for multi-tenancy. For vendors without complex rule-sets or sophisticated QoS features, SCs can be used to create tiers of performance or capabilities.

SCs are mapped to vSphere via a datastore. The storage container shown in Figure 2 is a datastore with a type of "VVOL."

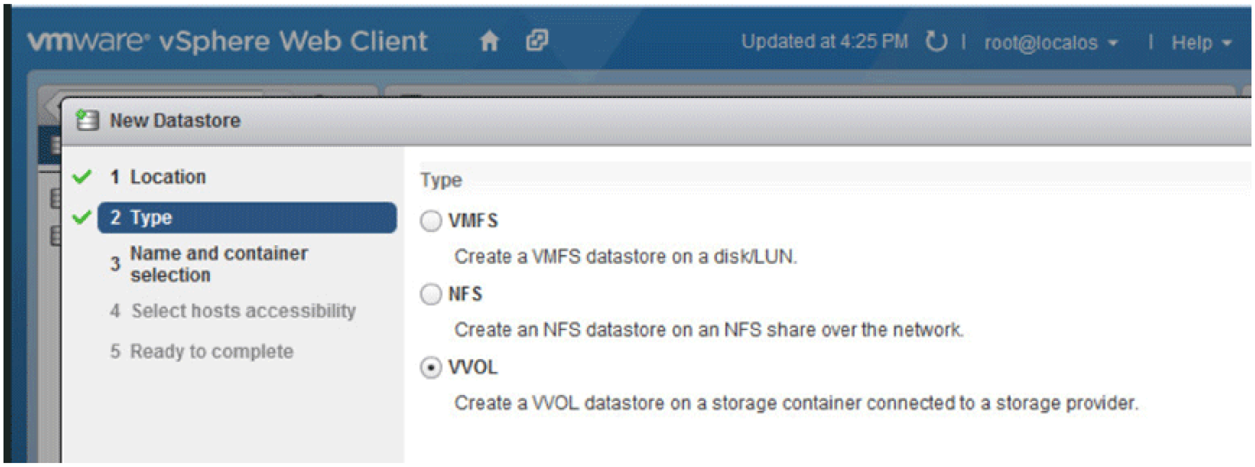

Figure 2

VASA Provider

The VASA Provider (VP) provides a bidirectional communication conduit between an array and vCenter. The VP is deployed either as a VM or as a feature on the storage controller. Both deployment models have their advantages and disadvantages. For example, a VP deployed on a storage controller may require a firmware upgrade but will not require installation. A VP deployed as a VM may be able to be more easily updated, but will be dependent on the hypervisor in order to insure uptime.

Capabilities and Rule-Sets

An array's capabilities are exposed to vSphere via the vendor's VASA provider in the form of "rule-sets," which are used to create storage policies. Figure 3 shows a policy being built up from a rule-set. Each policy can have one or more rule-sets included in it. When a VM is created, a storage policy is selected and all of the VVol datastores that can meet the requirements of the policy are displayed,

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

one of which can be selected for use. Throughout the life of a VM, the storage requirements of a VM will often change; in order to accommodate these changes, the VM can be reassigned a different storage policy.

We found that most vendors are able to surface many of their array's underlying capabilities in the form of rule-sets. We also found that some vendors are able to create pre-defined policies that group the capabilities into predefined sets and then expose these up as a single rule-set. Predefined sets of rules take advantage of the tribal knowledge of a particular vendor and make it extremely easy to categorize and assign policies to VMs.

Figure 3 shows how a storage vendor would surface up their array's capabilities as rule-sets that can be used to build a policy.

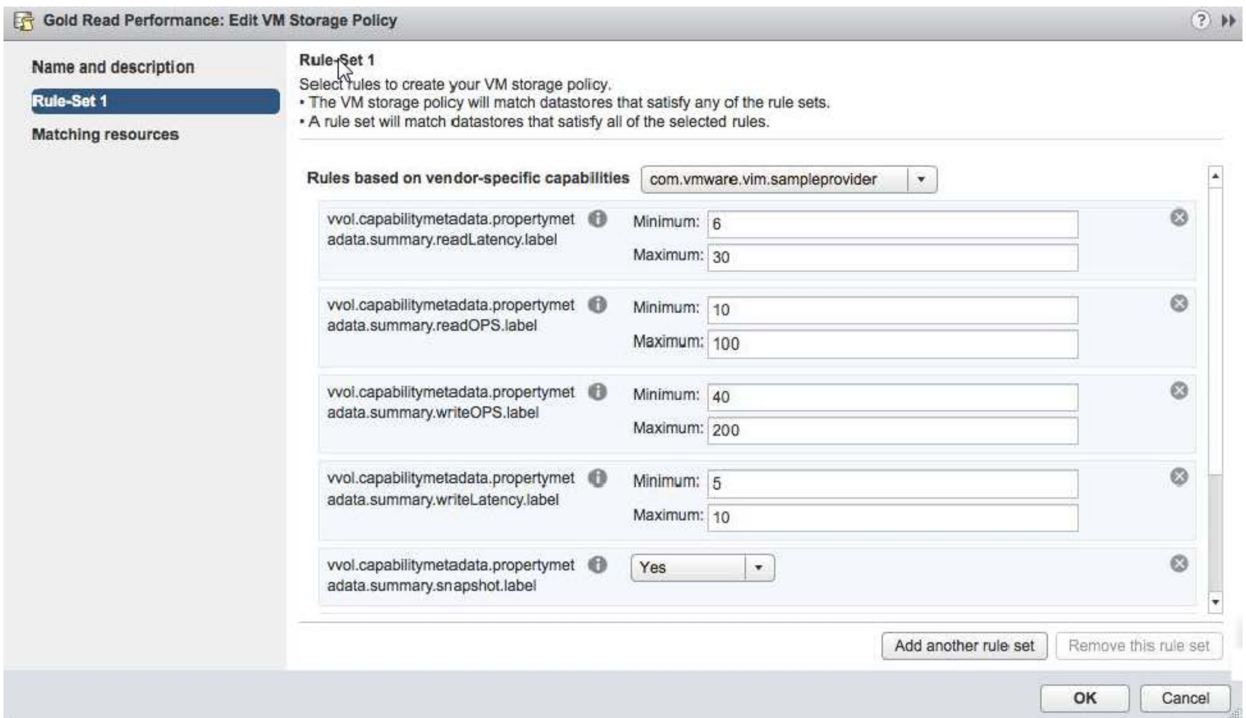

Figure 3

Virtual Volumes

Virtual volumes are not only the name of the product, but also a construct in a storage container. A virtual volume is the object that actually holds the data and metadata about a VM. A minimum of three virtual volumes (data, config and swap) are needed for each VM, but other constructs such as snapshots will also require a virtual volume. Conceptually, think of the VVOL as the container data object for a VMDK file. The number of VVols an array supports is directly related to the number of VMs and snapshots it supports.

Snapshots

Snapshots have always been problematic in vSphere. VMware KB article 1025279 provides best-practices information about using snapshots with vSphere. These best practices include many caveats such as the advice to use 2-3 snapshots per VMDK and to not use a snapshot for more than

72 hours. The use of snapshots can clearly cause performance issues. VVol has mitigated many of these issues as snapshots are offloaded to the array. With VVol, there are two types of snapshots: "Managed Snapshots," which are managed by vSphere with a maximum of 32 snapshots per VMDK; and "Unmanaged Snapshots," which are managed by the storage array with the maximum number of snapshots dictated by the storage array. Managed snapshots are initiated, managed and deleted using vSphere while unmanaged snapshots are managed solely by the array.

As VVol snapshots are done on the array, the limitations noted above are no longer applicable. However, some array vendors may have their own limitations for snapshots.. Some storage vendors include snapshot policies with their rule-sets and will take unmanaged snapshots periodically. Some vendors state that their snapshots (both managed and unmanaged) can be kept indefinitely.

Replication

VVol is NOT interoperable with SRM or "array-based replication" (VMware KB article 2112039). VMs stored on VVols can use "vSphere replication," albeit with limitations. While many storage vendors support VM replication, their replication cannot be integrated with vSphere and products like SRM cannot use VMs stored using VVols. All the vendors that we talked to about this have many caveats about replicating VMs stored using VVol.

After much thought, we decided that vendor replication with VVol objects is too immature to be included in this paper. We do know that IT users have a keen interest in replication with VVol objects and we are planning a companion piece to this paper to look more deeply into it.

What we asked the Vendors

We asked the vendors 32 questions about their VVol implementation. Below is a sampling of the types of questions that we asked:

- Are the PEs NAS or SAN-based, how many PEs can be associated with a storage container, and how many PEs can each SC support?
- Are a SC's capabilities surfaced individually or collectively? What capabilities are surfaced up to vSphere?
- How many SCs can a single array support? Is the SC static in size or can it be expanded and contracted in size?
- How many virtual volumes can an individual VM as well as the array support? Is that limit hard or soft-enforced?
- How many snapshots and clones can an individual VM have and how many total for the array?
- What Quality of Service (QoS) features are available for your VVol integration?
- Provide any additional comments about your implementation of VVol, including the features and capabilities that set them apart from other VVol offerings.

THREE TYPES OF VVOL INTEGRATION

VVol is just a framework. VMware does not specify the capabilities that can be surfaced up from the underlying array; VVols simply supply the mechanism from which the capabilities can be surfaced up. To help categorize arrays, we have organized them into three types using VM-level storage capabilities as a way to distinguish between them.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com

Type 1

This category is the most basic VVol integration on an array. A Type 1 array has one or more SCs to store VVols, one or more PEs per SC to communicate with vSphere, and surfaces up its capabilities as rule-sets. A Type 1 array does array-side snapshots and clones and supports thin provisioning of VVols. Type 1 arrays enable storage to be automatically provisioned at a VM level of granularity.

Type 1 arrays do not have the ability to set performance metrics for volumes and workload prioritization. Vendors with Type 1 integration may set up separate SCs based on the performance characteristics of the underlying storage (SSD, 15K HDD, 10K HDD, etc.) in order to offer different levels of performance, but they do not offer any QoS features (Min/Max performance targets) for the volumes stored on the storage. Array that have multi-tiering capabilities but do not surface up QoS capabilities are still classified as Type 1 VVol arrays.

Figure 4 shows a Type 1 array that has two storage containers; one uses SSD for its underlying storage and the other uses hybrid storage. Both SCs allow unmanaged snapshots, replication, deduplication and encryption to be surfaced up as rule-sets and selected on a per VM basis. The VM on the SSD based storage container has policies for replication and encryption while a VM stored the Hybrid SC has policies for snapshots, replication, deduplication and encryption. With Type 1 integration, the performance capabilities are static to the storage container and the VI administrator cannot choose the performance targets for a VM that is stored on the SC. With a type 1 integration if a VM needs to use higher performance media, it must be moved using vMotion to a different SC.

Implementing VVol on an array is not a trivial task. Some vendors have found that they are able to implement VVol rather quickly by using Type 1 integration to give their users basic VVol benefits. Vendors that offer arrays with a single performance level (e.g. all flash) or automated tiering capabilities may also decide to do Type 1 integration, since they do not necessarily need to surface up QoS features to meet their customers' requirements.

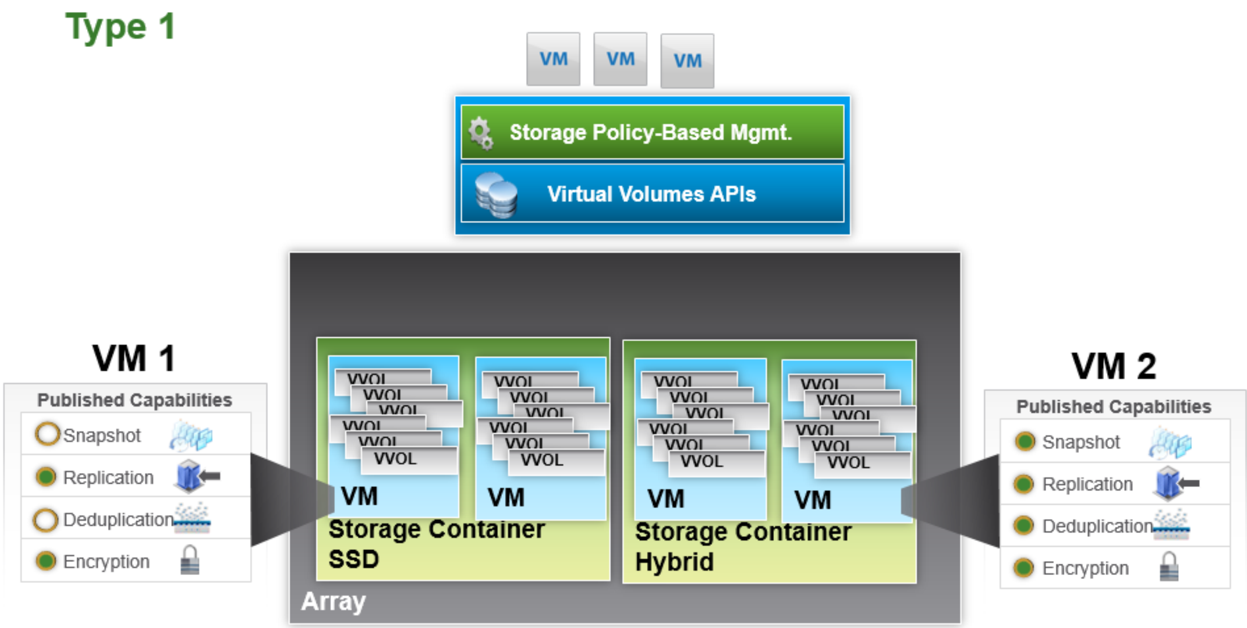

Figure 4

Type 2

Type 2 integration differs from Type 1 in being able to surface up QoS features in addition to Type 1 capabilities. By allowing QoS, a VM's performance can be exactly matched to the performance needed to meet SLAs. The QoS capability may be tied to the amount of IOPS or bandwidth or to the latency of a VVol. If the performance characteristic of a VM needs to be modified, the policy of the VM can be changed and as long as the underlying SC can accommodate the policy, the VM will not need to vMotion'ed to another SC. Type 2 integration allows for the consolidation of SCs as a single SC can accommodate different performance characteristics.

Figure 5 below shows a Type 2 array that has a single storage container. This single container can surface up any capability that the array is capable of, including QoS capabilities. For multi-tenancy reasons an enterprise may wish to create multiple SCs.

VMware Storage IO Control

vSphere 4.0 introduced a new feature: Storage IO Control (SIOC). SIOC allowed for storage I/O prioritization by using shares and limits. SIOC was only available with the Enterprise Plus edition of vSphere and arrays with automated storage tiering capabilities needed to be on a special VMware compatibility list.

VMware does not seem to be pushing SIOC and we have not seen the community adopt vSphere-based SIOC in great numbers. SIOC is a valuable feature and with the advent of VVols, we hope that more arrays will begin to integrate I/O prioritization into their storage and allow more end users to embrace this technology. The net-net is that the Taneja Group feels that storage prioritization, like many other storage features, is best done at the array level rather than the vSphere level.

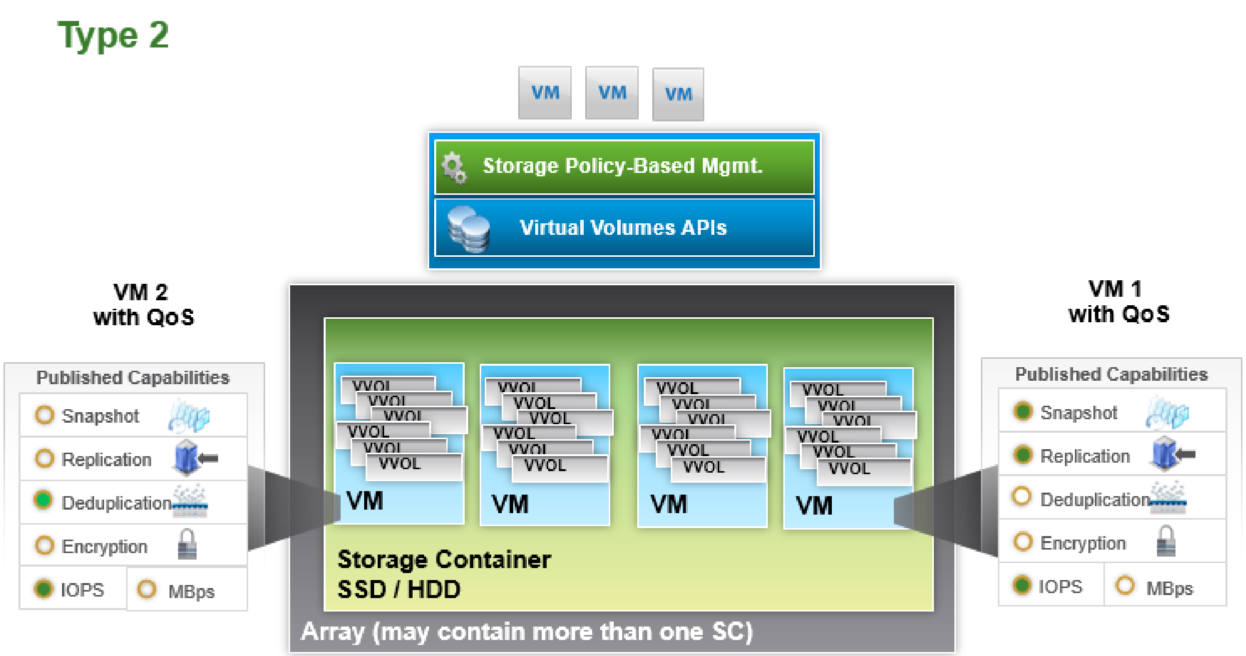

Figure 5

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557
www.tanejagroup.com

Type 3

This category has all the features of a Type 2 integration, along with the addition of workload prioritization. QoS with performance targets works well until the storage system comes under contention, at which point the array needs to prioritize the workload. Workload prioritization allows the VI administrator to allocate storage resources based on the relative value of a VM. vSphere for years has allowed VI administrators to do workload prioritization with compute (CPU) and memory using Shares, Limits and Reservations, and enterprises have found these features invaluable. VVols will allow customers to extend these same advantages to storage. Workload prioritization may be implemented as a discrete rule or bundled up as with other QoS features as a rule.

We included “Pre-Defined Performance Policies” as a separate “nice to have” capability for Type 3 categorization, since the addition of such policies will enhance the effectiveness of an array and simplify the deployment of VMs on it.

Figure 6 below shows a Type 3 array that has a single SC with two VMs, both with the same QoS policy but one (VM 1) having a higher workload priority. When storage resources come under contention, the VM with the higher priority will have more resources allocated to it.

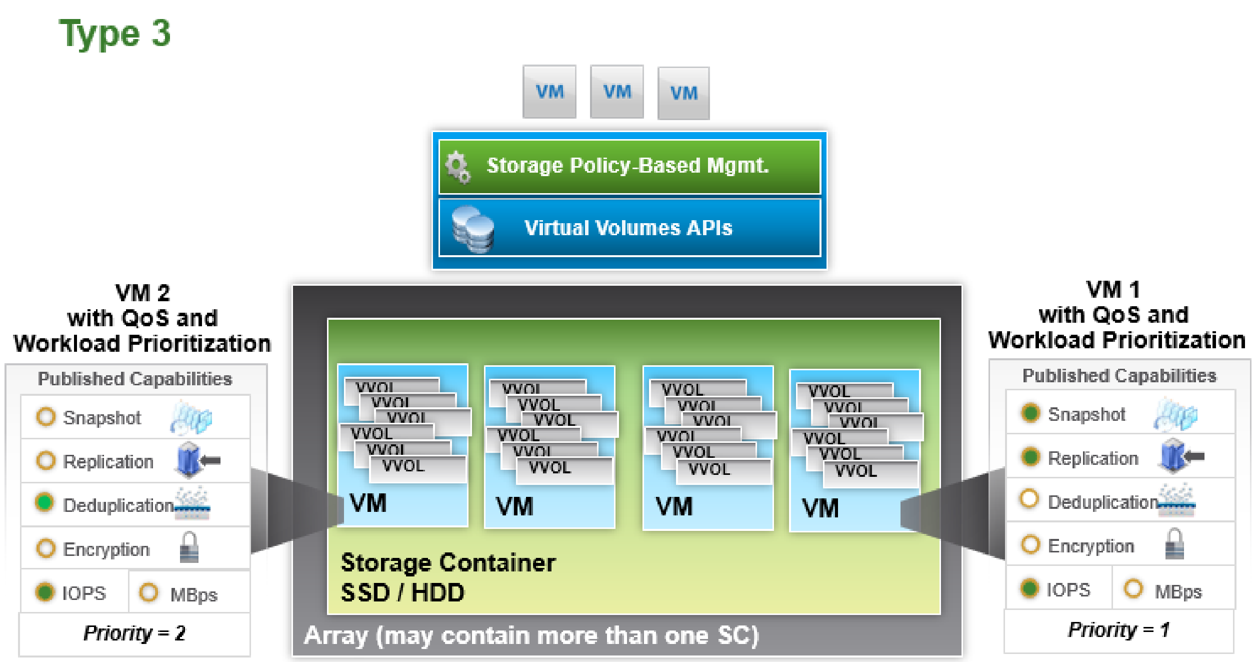

Figure 6

Summary of Classification

	Capabilities	**Benefits**
Type 1	Protocol Endpoints	Automated Per VM Storage Provisioning VM-level Storage Mgmt Granularity
	Storage Containers	
	Policies/Rule Sets	
Type 2	Protocol Endpoints	Automated Per VM Storage Provisioning VM-level Storage Mgmt Granularity
	Storage Containers	
	Policies/Rule Sets	
	Performance Metric Min/Max Targets	**Change Performance, no Storage vMotion; Consolidate Storage Containers**
Type 3	Protocol Endpoints	Automated Per VM Storage Provisioning VM-level Storage Mgmt Granularity
	Storage Containers	
	Policies/Rule Sets	
	Performance Metric Min/Max Targets	Change Performance, no Storage vMotion; Consolidate Storage Containers
	Workload Prioritization	**Guaranteed Performance during Contention**
	Pre-Defined Performance Policies (nice to have)	**Scalable Performance Management**

VVOL IMPLEMENTATION BY VENDOR

In April 2015, we conducted a survey of eleven companies that currently have, or will shortly have, VVol-capable arrays. We took their answers at face value but followed up with vendors when we had questions or concerns.

Below is a chart that shows how the vendors that we surveyed fit into our Type 1 – 3 categorization. More details on the VVol implementation of these vendors' products are included in the General Market Landscape Report which is available from Taneja Group. For information on this report, please send an email to info@tanejagroup.com

Disclaimer: The vendors that supplied us the information below make no representation and undertake no obligation with regard to product planning information, anticipated product characteristics, performance specifications, or anticipated release dates (collectively, "Roadmap Information"). Roadmap Information is provided by the vendors as an accommodation to the recipient solely for purposes of discussion and without intending to be bound thereby. As VVol is a new product, we suggest you follow up with a particular vendor to ensure that the information is current.

SUMMARY TABLE

	Availability	**Type**	**Deployment**	**Management**
Dell EQ	Today	Type 1	VM	vSphere Plug In
HDS	Today	Type 2	VM	Array UI
HP 3PAR	Today	Type 1	Array F/W	vSphere Plug In
IBM	XIV: Today, others: 2H 15	Type 1	VM	vSphere Plug In
Kaminario	1H 16	Type 2	Array F/W	vSphere Plug In
NetApp	Currently shipping	Type 2	VM	vSphere Plug In
Nexenta	2H 16	Unknown	Array F/W	vSphere Plug In
NexGen	Mid 15	Type 3	Array F/W	vSphere Plug In

Nimble	2H 15	Type 1	Array F/W	Array UI
Pure	2H 15	Type 1	Array F/W	vSphere Plug In
SolidFire	Unknown	Type 2	Array F/W	vSphere Plug In

SPOTLIGHT ON NEXGEN STORAGE - TYPE 3

Availability	Type	Deployment	Management
Mid 2015	Type 3	Array F/W	vSphere Plug-In

NexGen was a VMware VVol development partner. NexGen worked with VMware to deliver integration between Storage QoS and VVols.

All models of the NexGen N5 arrays will support VVols. Support is planned for Mid 2015.

VVol support will require a firmware upgrade. The upgrade will be non-disruptive and is free to customers with valid support contracts.

The VASA Provider is implemented via the storage controller on the array. All management can be done through a vSphere plug-in or a web based UI. NexGen solutions are designed to deliver predictable, consistent VM performance in a shared storage environment. Integrating Storage QoS and VVol will allow customers to change VM performance without a Storage vMotion, avoid the creation of multiple storage containers, and guarantee VM performance even when the storage system experiences contention via its workload prioritization scheme.

NexGen + VVol – The Numbers

Release Date	iOControl 3.5 (mid 2015)
Number of PEs per SC	2 Minimum
Number of SCs per array	2 Minimum; more can be created as needed
Number of VVol based VM's per array	300 – recommended, not enforced
Total number of virtual volumes per array	1000 – recommended, not enforced
VVol snapshots per VM	50 – recommended, not enforced
Total number of snapshots per array	1000 – recommended, not enforced
Clones per VM	50 – recommended, not enforced
Total number of clones per array	1000 – recommended, not enforced
Discrete QoS capabilities per VM via policies based on rule-sets	IOPs, BW, Latency controls Min, Automated Max Targets Multi-Variant Workload Prioritization Pre-Defined QoS Policies (5)
Other capabilities surfaced via rule-sets	Performance QoS, 5 Pre-Defined Policies
Data-efficiency features on array	Inline deduplication

NexGen + VVol – The Details

NexGen offers hybrid flash arrays that use PCIe flash and use Storage QoS to ensure predictable performance. They were one of VMware's original VVol development partners and have taken this knowledge to create a Type 3 VVol implementation of their arrays, which they expect to release mid-year 2015.

NexGen combines storage QoS with management of VMs at the granularity of VVols thereby providing customer value beyond standard VVol implementations. These benefits include VM level storage management granularity and automated storage provisioning.

NexGen differentiates itself from other arrays by giving customers the ability to prioritize VM workloads using three categories of service levels: Mission Critical, Business Critical, and Non Critical. These categories enforce workload priorities during storage contention states, often caused by unexpected workload spikes; and degraded performance from device rebuilds and firmware updates.

NexGen has done an excellent job of creating five pre-defined Storage QoS policies that can be applied to individual VMs. QoS functionality allows VI administrators to set performance metrics targets. Examples include minimum and maximum performance targets on a per-VM basis within a single storage container. This allows the VI administrator to change the performance of a VM without having to execute a storage vMotion, which reduces the need to implement multiple storage containers and keeps infrastructure management as simple as possible. NexGen implements the new policy and adjusts VM performance in a matter of seconds.

The NexGen VVol integration allows the administrator to set a Storage QoS policy on a per VM basis. As the mission of a VM changes over its lifecycle the VM's policy can be changed on the fly. This results in more predictable performance of the Mission Critical VMs and prevents the miss-allocation of storage resources.

Perhaps NexGen's most unique feature is workload prioritization. During periods of storage contention the ability to dictate what VMs are the most valuable is a feature that that sets a NexGen's array apart from others.

NexGen exposes five different rule-sets (Figure 7) to vSphere that map to the QoS policies on the array. These five policies cover all but the most niche uses cases. In the future NexGen plans to expose individual bandwidth and latency rules to vSphere, so administrators could create additional custom policies. Coincidentally, or not, this management paradigm aligns perfectly with VMware's Storage Policy Based Management (SPBM) approach.

Each of these policies has a different set of IOPS, bandwidth, and latency targets. VMs can be easily assigned to the appropriate policy; no need to define specific QoS targets to each VM. The five policies are:

- Mission-Critical Policy-1
- Business-Critical Policy-2
- Business-Critical Policy-3
- Non-Critical Policy-4
- Non-Critical Policy-5

Users can make informed decisions about VM workload priorities using detailed real time metrics that help them visualize the VM workload performance with NexGen's UI.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com

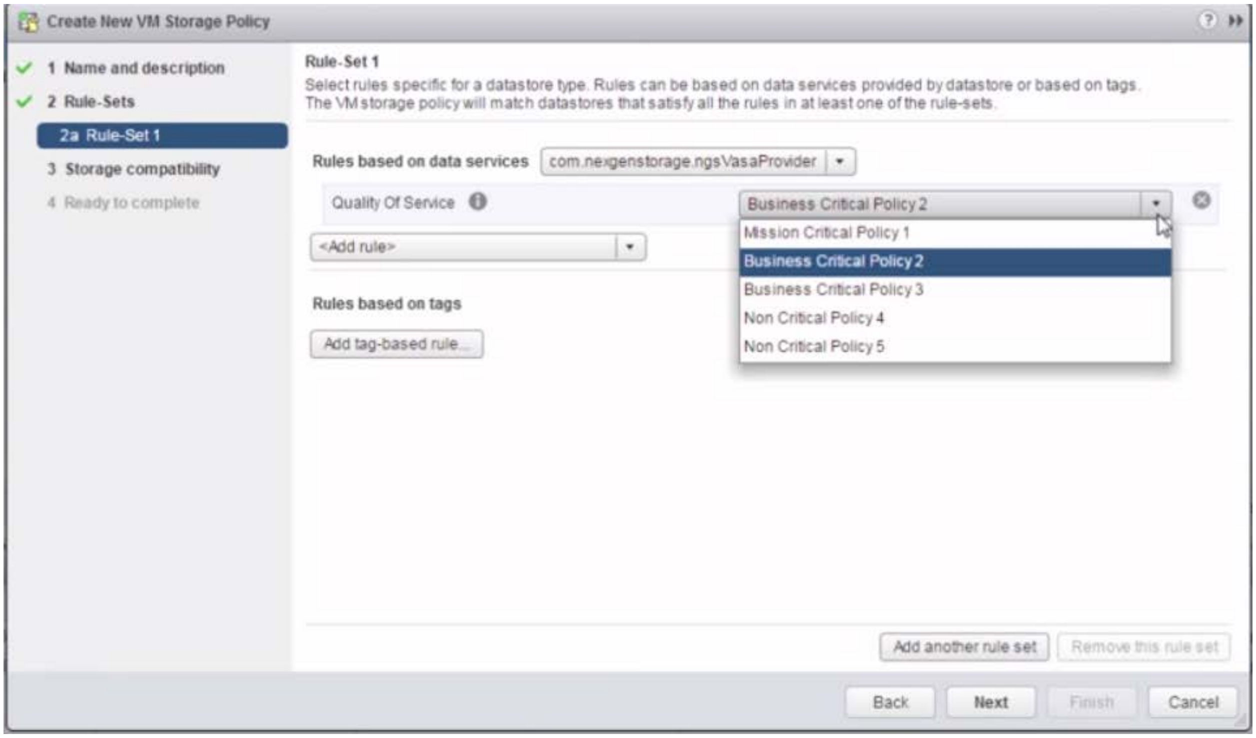

Figure 7

NexGen Arrays

NexGen currently offers four different arrays: the N5-200, N5-300, N5-500 and N5-1000. NexGen arrays are unique from all the other arrays in this report as they use a mixture of RAM, PCIe Flash, and HDD drives to maximize capacity and performance. All NexGen arrays come standard with dual active-active processors, PCIe Flash, RAID 6 disk protection, hot-swap drives, 4 x 10GbE interfaces for data and 4 x 1GbE for array management.

Depending on the model, NexGen arrays offer 2-15TB of flash, 32-256TB of raw disk capacity and deliver 150k-250K of sustained IOPS.

TANEJA GROUP’S OPINION ON NEXGEN VVOL IMPLEMENTATION

NexGen was industry’s first storage array to build in true QoS and thereby allow storage administrators to provide certain applications with a minimum set of resources while ensuring others got no more than a set amount. While it worked at a LUN level, NexGen’s QoS could deal with resource contention as well—it was able to systematically allocate resources to ensure the most important applications were not starved.

This functionality was a barn burner at that time and while some other competitors have caught on since then, few know about true QoS as NexGen does. So the fact that they are supporting QoS in the VVols implementation is no surprise. What is also not a surprise is that their QoS implementation mirrors their non-VVol implementation in how it deals with resource contention. This is what differentiates them from several other vendors who are adding QoS to their VVol-based arrays.

By providing a storage container that can host many different performance levels, NexGen has eliminated the need for storage silos. This results in a very efficient switchover when VMs need to be

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com

moved from one policy to another. Some of the other arrays that we looked at required separate storage containers to be created in order to provide different levels of performance based on the underlying storage hardware (All-Flash, 15K disk, near line disks).

This combination of functionality results in an array with Type 3 VVol integration that eliminates storage silos, provides VM-centricity across many data services and incorporates industry-leading QoS capability. No other storage array attained a Type 3 ranking in this survey. NexGen was a trend setter when it initially came to market with QoS and we consider them to be a trendsetter now in their VVol support.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com