

TECHNICAL WHITE PAPER

Top Reasons to Use Pure Storage FlashArray with vSphere Virtual Volumes

Optimize and simplify your VMware vSphere storage

Contents

- VMware Changed How Storage Is Used** 3
 - VMware: No Longer Just Virtualizing Compute 3
- Top Reasons to Use Pure Storage and Virtual Volumes** 4
 - The Best Storage Platforms for Your VMware Hybrid Cloud 4
 - Snapshot Overhead and Performance 4
 - Storage Policy-based Management 5
 - Data Mobility and Presentation 5
 - Simplified Business Continuity and Disaster Recovery Workflows 6
 - VM Aware Storage Insights and Operations 7
 - Virtualize Reluctant Tier 1 Workloads 8
 - VMware Storage as Code 8
- Summary** 8
- How to Get Started with Pure Storage and Virtual Volumes** 8



VMware Changed How Storage Is Used

With the introduction of server virtualization and the impact VMware has had on the compute and storage industry, how companies deployed and maintained their line of business applications drastically changed—and for the better. With the ability to collapse operating systems and applications into an isolated environment, virtual machines (VM)—and running multiple VMs on a hypervisor, specifically ESX—delivered far greater return on investments (ROI) for server purchases due to higher utilization and deep densities.

Compute resources wasn't the only resource VMware influenced: With their Virtual Machine File System (VMFS) VMware was able to change the storage landscape as well. VMFS is a clustered file system developed to store the needed files and images for virtual machines and additionally allowed for multiple servers and ESX hosts to read and write to the same file system simultaneously while individual virtual machine files are locked.




FIGURE 1 VMFS datastore and ESXi hosts

However, VMFS does have some drawbacks and limitations:

- It can be challenging to fully understand how storage is provisioned and what is consuming it.
- It is harder to identify and track down performance challenges. With multiple hosts and multiple VM accessing and running on the file system, it is a lot harder to isolate issues. A VM running on host A could be directly impacting the performance of a VM running on host B as they share the same storage.
- VMs don't have access to the array storage features as the underlying storage has been abstracted away. The VMs aren't aware of the storage, and the storage isn't aware of the VMs.

VMware: No Longer Just Virtualizing Compute

As environments grow and scale in size, the above challenges become even more magnified. For these reasons, along with others not mentioned, VMware developed and released vSphere Virtual Volumes (vVols) in 2015. vVols is a framework that virtualizes storage resources to make them more efficient and optimized for virtual workloads, centering on the application itself and not the infrastructure.



Top Reasons to Use Pure Storage and Virtual Volumes

With some of the challenges out of the way, how does vSphere Virtual Volume with Pure Storage® not only solve those problems, but improve on the administration and management of your storage environment? Below are eight examples of “How” vVols can make an impact, in no particular order.

The Best Storage Platforms for Your VMware Hybrid Cloud

Pure Storage has a long-standing partnership with VMware as a Technology Alliance Partner (TAP) - Advanced. Pure Storage not only supports VMware features and technologies, but is heavily involved with their inception and advancement. This allows us to build solutions and integrations that take advantage of the latest VMware technologies and make VI administrators jobs easier. That could be with the joint development of vVols outlined here, or our vSphere Client Plugin for FlashArray™ that makes administering storage from a single interface simple and streamlined. If you need a vVol replication that is easy to configure and manage, have a look at our Replication Manager feature. Using Aria Operations, you can monitor FlashArray with our Management Pack. If you are considering or using VMware Cloud on AWS (VMC on AWS) or Azure VMware Solution (AVS), Pure Cloud Block Store™ can dramatically reduce your compute and storage total cost of ownership (TCO). Whatever you are doing with VMware, you can make it better with Pure Storage.

Snapshot Overhead and Performance

Snapshots, and their use in vSphere environments, is second nature to storage and virtualization administrators. Whether that be array-based snapshots or vSphere based snapshots. But whichever you choose, there are tradeoffs, with some of them being more impactful than others. Array-based snapshots with FlashArray are lightweight with zero performance impact. However, they are not granular; it is an all-or-nothing proposition as you are snapshotting at the volume level. Their vSphere counterparts, while quick and easy to use, have a direct impact on your virtual workloads' performance. And it is significant. By VMware's own testing, a VM with a single snapshot can lose up to 70% of its IO performance when using VMFS or vSAN! That number only gets worse as you add additional vSphere snapshots to the chain




FIGURE 2 I/O loss

Virtual volumes with Pure Storage gives you the best of both worlds. You get the lightweight and zero performance impact to your workloads that array-based snapshots provide with the VM and VMDK granularity level that can be accomplished with a vSphere based snapshot without any tradeoffs. Again, there is zero performance impact to the virtual workload and zero consolidation when removed, tested and validated by [VMware](#). Add to it the automation and scheduling capabilities provided by storage policy based management (SPBM), and you can massively simplify your data protection strategy.




FIGURE 3 I/O loss with virtual volumes



Storage Policy-based Management

Administering your storage and virtual machines via policies delivers a control plane that simplifies the presentation of storage services to your virtual workloads and validates that settings are applied. Historically data services have been delivered at the VMware Datastore or volume level with the workloads being unaware, but this method has brought a monolithic approach to data services and how they were consumed. As mentioned before, it lacks granularity.

The use of SPBM in tandem with vVols gives you the power to associate the required data services at the virtual machine or VMDK level of a workload. Configure a “Gold” policy that protects your Tier 1 workloads with snapshots every hour and replicates those snapshots to a secondary location every hour. Assign the gold policy to the wrong VM, easily remove it or change it to a different policy, live, with zero impact to the virtual machine. Want to limit the IOPS or bandwidth of a test VM that is impacting production workloads, create a “Test/Dev” policy to set limits to these workloads and eliminate noisy neighbor scenarios




FIGURE 4 Policy levels

Data Mobility and Presentation

Sharing is a good thing, but when working with physical or virtual servers sharing can quickly become problematic. VMDKs and virtual machine file systems (VMFS) are designed and created for virtual workloads: physical servers can't access this data. Storage native volumes are designed for direct server access and are presented as such, whether that be bare metal hosts or ESXi servers. Virtual machines can access them via raw device mappings (RDM), but that brings additional layers of complexity and drawbacks. Add to that the use of public cloud storage, and how you move and use your data in a hybrid cloud model amplifies these data mobility challenges.

Using Pure Storage FlashArray snapshots, volume cloning, and replication functionality along with vVols remediates these challenges. No longer do you need complex scripting, clunky workarounds, or time consuming network copies to get your data to the workloads that need it. With a vVol being the single object of record for workload or a database, you can bridge the gap between the physical, virtual, and cloud worlds.

As a practical example, say you are running a production database on a virtual machine backed with vVols and you want to provide the database data to your bare metal data warehouse environment, your virtual Test/Dev environment, and a DR copy to the cloud. This is simple with FlashArray and vVols. For the data warehouse and Test/Dev machines the process is

1. Take a snapshot of the vVol holding the database files
2. Clone the snapshot to a new volume for each environment
3. Present and mount the new volume to the bare-metal server and to the virtual Test/Dev server.



That's it. For the DR copy, the scenario is even simpler. Utilizing Pure Cloud Block Store in either AWS or Azure, configure a snapshot/replication schedule with SPBM, or natively in the array, to synchronize a copy of the database data in the cloud.




FIGURE 5 Data Mobility with virtual volumes

Simplified Business Continuity and Disaster Recovery Workflows

Mapping out and preparing a Business Continuity/Disaster Recovery (BC/DR) plan is a complex and time consuming endeavor. There are lots of variables to account for: application start order, network connectivity, authentication and security, and of course getting the data from point A to point B. Automation and orchestration is paramount to the successful configuration and implementation of these plans.

Whether you are using VMware Site Recovery Manager (SRM)—which Pure Storage FlashArray has full integration and support for—or another BC/DR orchestration engine to solve for the application, networking, etc. The movement of, and scheduling for, the data with vVols and FlashArray is simple. First, the requirement for a storage array adapter (SRA) with SRM is removed so that is one less thing to install and manage. Second, with the granularity of vVols, the configuration of protection groups and the applications that are being protected can be fine tuned. You're no longer protecting, replicating, and failing over at the datastore. You have the ability to be selective—down to the individual VMDK—of what you want to protect and how you want to protect it. Paired with SPBM, you can configure and set the replication schedule along with the retention period between the paired sites in a single and simple concise interface, the vSphere client.




FIGURE 6 Policy levels for snapshots and replication



VM Aware Storage Insights and Operations

One of the hardest, if not the hardest, items to manage in a vSphere environment is storage, whether that be administration, capacity planning, operations, or troubleshooting. VMFS is and has been great, but it does bring some challenges when it comes to each of those various points that both VI administrators and Storage Administrators have had to overcome. While being able to provision 25 VMs on a single datastore/volume is easy enough, managing those decisions is a different story. vVols can and will change this for the better.

One way vVols can change things for the better is VM provisioning. With vVols you'll no longer lookin vCenter Server and have to guess which datastore you should place a new workload on. Is it the one with the most amount of free space? Or do you go with the one with the least amount of VMs on it? Making either of these decisions without really knowing what is going on with a given datastores performance, noisy neighbors, IO conflicts, etc, can lead to disastrous results. Using vVols removes this guessing game as it eliminates the datastore provisioning model. When provisioning a workload, you know to select the vVols-enabled FlashArray and that's it—the VM and associated VMDKs are provisioned.

Speaking of noisy neighbors and IO conflicts, the next advantage of vVols is how easily you can see and troubleshoot performance issues. Instead of looking at a multitude of interfaces like the vSphere client, ESXTop, FlashArray UI, or a 3rd party monitoring tool to try to identify and isolate performance bottlenecks with vVols, that information is surfaced to you directly in the vSphere client. With vVols one-to-one mapping of VMDK to volume, you can see the exact performance characteristics of that VMDK. No longer are you trying to correlate IOPs, latency, etc with the other 24 VMs residing on a datastore: The information is right there at your fingertips.




FIGURE 7 vSphere client UI

The last advantage is capacity planning and utilization. With vVols and the Pure Storage vSphere Remote Client plugin, you can see exactly how much storage a VM or VMDK is consuming as well as the data reduction value with the dedupe and compression functionality of FlashArray. Going beyond what you can observe, vVols makes it far easier to reclaim dead or unused space via UNMAP. VMFS-6 does support UNMAP operations, but it is an asynchronous process. This means that it is a background process scheduled via ESXi to reclaim the unused space, which can take over 24 hours to complete. Utilizing vVols with UNMAP gives instant results when deleting a file in the VMDK or the VMDK itself. As vVols are direct volumes, the guest file system is actually issuing UNMAP directly to the FlashArray. There's no scheduling or waiting for a background operation.



Virtualize Reluctant Tier 1 Workloads

Believe it or not, there are still some workloads that people are nervous about virtualizing. It could be an electronic medical record (EMR) system in healthcare, high-performance Microsoft SQL or Oracle databases for financial organizations, or even an in-house developed line-of-business application. Whatever the specific application is, vVols can help with the testing and prove whether an application can or can't be virtualized. As mentioned before, with vVols allowing ease of data mobility and access to both physical and virtual servers, testing the given workload(s) can be accomplished with little to no effort.

VMware Storage as Code

Whether your cloud is private, public, or hybrid, automation is a core tenant in cloud computing. The ability to rapidly provision, configure, or remove objects through code or scripting just isn't a requirement. It is a necessity. vVols with Pure Storage make this a reality when working with storage in a vSphere environment. Since virtual volumes are a first class disk (FCD) in vSphere and on FlashArray, you can create new virtual disks independently of a virtual machine. This opens the door for a wide range of capabilities that are difficult—if not impossible—when working with VMFS. With PowerCLI and the provided APIs, you can automate the cloning, copying, replication, recovery, and snapshots of these FCDs. Note that because these processes are being completed on the FlashArray, they complete nearly instantaneously.

Summary

Virtualizing compute resources has simplified the management and operations of workloads for application owners and systems administrators. Why not do the same for your virtualization and storage administrators when it comes to your storage? Migrating to or implementing vVols with Pure Storage FlashArray in your vSphere environment can decrease the time to resolution of performance issues, decrease the overhead of workload provisioning, deliver a pragmatic approach to managing data services for your workloads, increase your ease of data mobility, and simplify your BC/DR planning and workflows.

How to Get Started with Pure Storage and Virtual Volumes

- Pure Storage FlashArray//X™ [solution page](#) Pure Storage vSphere virtual volumes [YouTube series](#)
- Pure Storage vSphere virtual volumes: [Quick Reference Page](#)
- Pure Storage: [Implementing vSphere virtual volumes with FlashArray](#)
- VMware virtual volumes: [Solution Page](#), [Tech Zone](#)