



# Self-provisioning storage across all of your devices with VMware Virtual Volumes and DataCore Software

By Dan Kusnetzky, Distinguished Analyst

## PROVISING STORAGE FOR VIRTUAL ENVIRONMENTS

If we visit enterprise data centers, it would quickly become apparent that most of them are storage museums. Systems and storage are assigned to support individual applications, but these assets were acquired over a long period of time – forcing IT to manage 3 to 5 year old storage. Then to overcome their shortcomings the enterprise ends up with different types of storage devices.

## CHALLENGES OF STORAGE SILOS, COST, AND STORAGE OPTIMIZATION

Enterprises are now facing pressure to optimize the use of their IT resources and to reduce their overall IT spending. This has driven the use of virtualization technology to consolidate workloads onto a smaller number of systems and to find ways to use all of their different types of storage as a unified pool. These pressures have also driven enterprises to consider how their business units can provision and manage these resources on their own.

Enterprises face a number of challenges including the need to get a handle on the storage silos they created over time. Part of the challenge is finding underutilized storage resources and putting them to work rather than just purchasing more storage when the needs of a given workload have expanded. They have learned that a given workload's storage must be able to dynamically expand and contract as needed.

Enterprises that have adopted VMware-based virtual environments learn that some of their problems have been magnified while others have been reduced. Administrators for vSphere and storage find themselves in a time-consuming, complex dance each time a change is needed. For example, each time a vSphere administrator creates a virtual machine (VM), they must ask a storage administrator for disk capacity – which is not a straightforward task.

## THE COMPLEX DANCE OF STORAGE ALLOCATION AND PROVISIONING

The storage administrator for each storage array must manually translate requests coming from the vSphere administrator into a form that their storage management software understands. This includes assigning some number of logical unit numbers (LUNs) to a class of service so that the VMware ESXi servers can access this storage over their iSCSI or Fibre Channel network. Once the storage administrator has completed this task, the vSphere administrator can get to work.



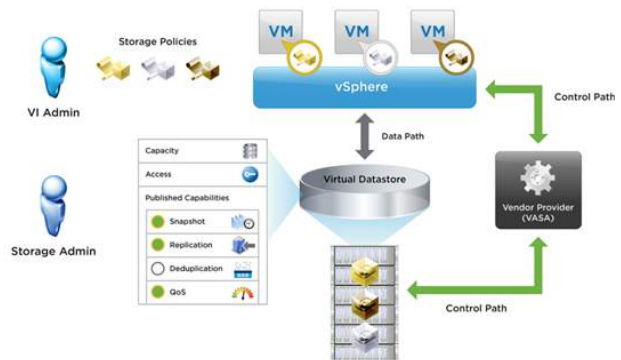
The vSphere administrator must then use vSphere administrative software to scan for the newly attached devices and create data stores on them. When something changes and the storage environment must be adjusted to address that change, and the dance begins again.

To avoid this time-consuming and labor-intensive dance, some enterprises pre-allocate storage, that is carve up storage into larger LUNs. Although this potentially wastes storage space, this approach allows the vSphere administrators to build larger data stores and quickly assign that space to multiple VMs. While addressing the problem of assigning storage to VMs, this approach creates another problem. When a vSphere administrator takes a snapshot of one of those VMs for backup, archival storage or even development purposes, the storage array must create a snapshot of the *entire LUN* and all of the VMs in it.

### ENTER VMWARE VIRTUAL VOLUMES

VMware is aware of this challenge and has introduced technology designed to simplify the environment. VMware calls this technology VMware Virtual Volumes (VVols).

VVols is a VMware integration and management framework that is designed to virtualize SAN (storage area network) and NAS (network attached storage). VMware's goal is enabling a better, more efficient virtual storage environment that will complement their virtual processing environments.



VVols focuses on the needs of the application rather than the underlying infrastructure. This change of viewpoint results in making it possible for storage to be policy driven and more agile. VM storage consumption can be dynamically adjusted in real time allowing VM storage use to grow or shrink, as needed.

VMware's VVols can be managed through a set of application programming interfaces (APIs). One of those APIs is vSphere Storage APIs for Storage Awareness (VASA).

### THE PROBLEM: MANAGING ACROSS STORAGE VENDORS

Storage providers (including array manufacturers) have to create datastores offering different classes of service to work within VMware's VASA architecture. This means developing software that makes it possible for information about their arrays and management of those arrays to be published through a VASA provider.

When the vSphere administrator creates a VM, they can simply select the datastore with the right characteristics without contacting anyone. vSphere automatically contacts the storage device requesting allocation of the desired capacity. Storage administrators can track which VM has been assigned to what available space.

The benefit of this approach is that anytime a vSphere administrator takes a snapshot, that request is automatically sent to the storage device, which sets up a separate allocation explicitly for the snapshot. This means that storage space isn't wasted.

This approach provides VM-centric and VM-granular storage allocation. But more importantly, it enables vSphere admins to self-provision capacity anytime they need it without having to involve a storage administrator.

However, IT is faced with a dilemma. While VVols makes it possible for vSphere administrators to see the available resources, assign them and manage them using VMware's administrative tools, this capability isn't available on the equipment already in place. They have to upgrade their storage to the new devices that support VVols.

At this time, only a few storage manufacturers offer VVols support on current products. They plan, instead, to support future versions. Furthermore, they don't plan to retrofit the arrays in customer sites with VVols support.

Another problem can be seen in the case of an enterprise having arrays from different manufacturers serving to vSphere. Each device could only advertise what it had control over. So each device appears as a separate datastore. This creates silos of storage and the likelihood for wasted space and higher overall costs.

#### A SOFTWARE-DEFINED STORAGE MODEL

What does this have to do with a software-defined environment? Software-defined is a catch phrase used in the IT industry to describe an environment in which basic functions, such as processing, storage, networks and the like have been implemented in a hardware-independent fashion. This typically means that the function has been placed in a virtual environment, and that its control and monitoring functions have been made available using a standard interface. DataCore has added the concept of "software-defined" to VMware's vSphere computing environment.

Enterprises are seeking a software-defined storage model that can be built upon the strong foundation of VMware's vSphere and its related virtualization technology and administrative tools.

They expect that moving to this approach would overcome some of the limitations presented by today's storage products and how enterprises have used them to serve their applications. The use of this software-defined storage environment puts administrators in control and makes it possible to extend VVols across the pools of storage created by DataCore's technology.

#### CHALLENGES PRESENTED BY TODAY'S MULTI-VENDOR STORAGE IN A VIRTUAL WORLD

Over the years, storage manufacturers have developed technology designed to allow storage access and sharing storage across many applications. Each developed a unique approach that addressed some aspects of a multi-vendor, multi-application world but created problems in other areas. Each supplier had its own approach to monitoring, management and provisioning of storage. Typically, however, their approach only worked with their equipment.

As enterprises adopted storage virtualization technology, they discovered that this "each vendor for itself" environment was getting in the way. They understood that VMware, by introducing VVols, was trying to help by simplifying the provisioning and management of storage. They wanted to make it easily possible for vSphere administrators to monitor, provision and manage the storage needed to support virtual applications. Unfortunately, storage

manufacturers have only provided VVols support for specific storage devices today while promising broader support in the future.

Enterprises find themselves in the middle. They want what VVols offers, but find that only a small percentage of their currently installed NAS and SAN storage systems work in that environment.

#### WHAT WOULD AN IDEAL SOLUTION PROVIDE ?

Enterprises have been asking the industry to provide a broad or, perhaps, an ideal solution. This solution would support storage of all types including direct attach, NAS, SAN and even Cloud storage and make it all appear to be a pool of resources. They want tools that make it possible for systems to access any and all storage as if it were a unified pool. Furthermore, these tools must make it possible for administrators to monitor, manage and provision storage in a unified way.

This ideal solution also must work directly with VMware's VVols and allow VMware's administrative tools to monitor, manage and provision any and all of the storage controlled by that tool.

#### IS THERE A TECHNOLOGICAL SOLUTION ?

Is there is a technological solution available that can address these issues today? If it existed, it would have to demonstrate the following capabilities:

- It must run on standard x86 servers, providing one set of common storage services across all storage devices.
- The combined storage capacity of all of the different devices being managed by this technology should be seen as a shared pool, regardless of the vendor, the type of device, whether it is directly attached to a system, attached to a storage server or part of a cloud storage service.
- The tool must make it possible for storage to be accessed and allocated in a way that eliminates wasted storage capacity.
- The tools must make it possible for data to be replicated between nodes eliminating a single point of failure.
- It must be possible to seamlessly scale up and scale out and allow the enterprise to select the most appropriate storage products from the most appropriate hardware manufacturer for each class of storage and servers.
- The tool must intelligently cache storage to provide the fastest access and lowest levels of storage latency.
- The tool must intelligently place data on the appropriate storage tier and storage device based upon policies and data usage.
- This capability must fit seamlessly into and support VMware's computing environment.

#### DATA CORE SOFTWARE - DEFINED STORAGE

This technology exists and is part of DataCore's SANsymphony™ and DataCore™ Hyper-converged Virtual SAN. These products provide a common set of storage services across all storage devices to make the best use of directly attached, SAN, NAS and Cloud storage services regardless of model or brand of the storage devices. It can be managed as if it were a single infrastructure-wide storage resource with no wasted space.

These unified storage resources can easily be slotted into the VMware VVols environment without facing protocol, management or communication problems.

Storage devices can be selected based upon application requirements regardless of storage vendor or storage technology.

By using this technology in a VMware VVols environment, vSphere administrators are in charge and have access to the entire global shared pool of storage. This reduces complexity and administrative costs.

This also eliminates single points of failure while also improving performance.

Enterprises may scale up via larger servers and storage devices or scale out to use more servers and storage devices depending upon the performance and capacity requirements. No longer are organizations forced to select storage from a single supplier.

DataCore makes it possible to self-provision VVols across its diverse storage pools managed, even when the underlying storage does not support VVols.

#### S U M M A R Y

Enterprises are now facing pressure to optimize the use of their IT resources and to reduce their overall IT investment. They are addressing this requirement by deploying virtualization technology to consolidate workloads onto a smaller number of systems. They are also using storage virtualization technology that allows them use all of their different types of storage as a unified pool.

They expect this technology to overcome the storage silos they created at an earlier time and allow them to find and use under-utilized storage resources rather than having to get out their checkbook and purchase more storage every time an application needs more resources.

Enterprises that have adopted VMware-based virtual environments want to get the benefits VMware's VVols offers. They also want to make it possible for vSphere administrators to be in charge of their own storage resources, thus eliminating the administrator dance.

Technology from DataCore can address these business requirements today.

Kusnetzky Group recommends visiting DataCore's website (<http://www.datacore.com>) to learn more about their unique approach of putting VVols to work for the enterprise. We recommend asking DataCore for a demonstration of their technology. It is likely to awaken many new possibilities.