

C o m m u n i t i e s

All Communities ▾ Blogs ▾

Communities > Blogs

VMware Storage Blog

February 04, 2010

Introduction to vStorage APIs for Data Protection

[vStorage APIs for Data Protection](#) were introduced in vSphere 4.0 to facilitate seamless virtual machine backup and recovery. Protecting virtual machines in the VMware environment with traditional data protection methodology had its challenges. vStorage APIs for Data Protection enables partners to integrate their backup/ recovery products; hence, giving customers more choices.

The vStorage APIs for Data Protection enable backup software to perform backups without interrupting or affecting the ESX servers with the processing of backup tasks. A backup product using these APIs can backup and restore virtual machine data from its central backup server or virtual machine without requiring that backup agents and backup processing be done inside of each virtual machine on the ESX server. This allows you to offload backup processing from ESX servers to reduce costs by allowing each ESX server to run more virtual machines.

The vStorage APIs for Data Protection leverages the snapshot capabilities of [VMware vStorage VMFS](#) to enable backup without requiring downtime for virtual machines. As a result, backups can be performed non-disruptively at any time of the day without requiring extended backup windows and the downtime to applications and users associated with backup windows. Native integration to data protection product facilitates user experience that is similar to physical backup. This also results in simpler management and restore.

vStorage APIs for Data Protection support:

- Windows and Linux guests
- All storage architectures (FC, iSCSI and NAS.)
- Faster incremental, differential and full VM image backup
- File level backup and restore

Some people think this is the new VMware Consolidated Backup, but it is not (Table 1.) Unlike VCB (VMware Consolidated Backup), VMware vStorage APIs for Data Protection (VADP) doesn't require Proxy server and is fully integrated with the Data Protection software such that backup application will leverage it seamlessly out of the box. Though, VMware continues to offer VMware Consolidated Backup to support the existing install base and on vSphere 4.0.

	VADP	VCB
Requires additional download & install	No, built into the data protection software	Yes
Full VM image backup	Yes, single step copy – source to target	Yes, with two step copy – source to VCB proxy and VCB proxy to the target
Incremental VM image backup	Yes	No
File level backup	Yes, both Windows and Linux	Yes, Windows only
Full VM image restore	Yes	Yes, by using VMware Converter
Incremental VM image restore	Yes	No
File level restore	Yes, using restore agents	Yes, using restore agents

Table 1: Differences between VADP and VCB

Although there are significant benefits for leveraging vStorage APIs for Data Protection compared to

VMware Blogs

- [The Console \(VMware Mgmt\)](#)
- [VMTN Blog \(Technical News\)](#)
- [Developer Center Blog](#)
- [ESXi Chronicles](#)
- [Team Fusion](#)
- [Uptime \(Business Continuity\)](#)
- [VI Team Blog](#)
- [Virtualization for SAP Solutions](#)
- [VIXAPI Blog](#)
- [VMware Code Central Blog](#)
- [VMware Communities Blog](#)
- [VMware Knowledge Base Blog](#)
- [VMware Knowledge Base Weekly Digest](#)
- [VMware Networking Blog](#)
- [VMware PhD \(Education\)](#)
- [VMware Security Blog](#)
- [VMware Storage Blog](#)
- [VMware ThinApp Blog](#)
- [VMware vApp Developer Blog](#)
- [VMware: Virtual Reality](#)
- [VMware View Blog](#)
- [VMware View-Point](#)
- [VMware Virtual Disk Development Kit Blog](#)
- [VMware vSphere Blog](#)
- [VMworld Team](#)
- [VROOM \(Performance\)](#)
- [vSphere PowerCLI Blog](#)
- [Workstation Zealot](#)

Recent Comments

- [Backup on Introduction to vStorage APIs for](#)
- [Giri on Introduction to vStorage APIs for](#)
- [Storage on Introduction to vStorage APIs for](#)
- [Ito 5 tape on VMFS vs. NFS for VMware Infrastructure?](#)
- [seslichat on Introduction to vStorage APIs for](#)

Categories

[Weblogs](#)

VMware Consolidated Backup, I would like to emphasize that it depends on the customers' IT environment and their business requirements. VMware gives the customer the choice to choose the right data protection solution for their environment.

vStorage APIs from Data Protection also provides the framework for [VMware Data Recovery](#) product that was introduced in vSphere 4.0. VMware Data Recovery is a disk-based backup and recovery solution that enables quick, simple and complete data protection for your virtual machines. It's fully integrated with [VMware vCenter Server](#) to enable centralized and efficient management of backup jobs and also includes data de-duplication to save on disk storage for your backups.

My next post will be on the vStorage APIs for Array Integration. In the meantime if you would like to read more on the other vStorage API feature, read the blog on [vStorage APIs for Multipathing](#)

Author: Ibrahim "Ibby" Rahmani
Sr Product Marketing Manager - Storage, VMware

Posted by Ibby Rahmani on February 4, 2010 | [Permalink](#) | [Comments \(7\)](#) | [TrackBack \(0\)](#) |  [BOOKMARK](#)   

October 19, 2009

vStorage Multi Paths Options in vSphere

Multi Path challenges

In most SAN deployments it is considered best practice to have redundant connections configured between the storage and the server. This often includes redundant host based adaptors, Fibre Channel switches, and controller ports to the storage array. This results in having four or more separate paths connecting the server to the same storage target. These multiple paths offers protection against a single point of failure and permit load balancing. In a VMware virtualization environment a single datastore can have active IO on only a single path to the underlying storage target (LUN or NFS mountpoint) at a given time. The storage stack in an ESX server can be configured to use any of the available paths but only one at a time prior to vSphere. In the event of the active path failing, the ESX server will detect the failure and failover to an alternate path.

The two challenges that VMware Native Multi Path addresses has been the 1) presentation of a single path when several are available (aggregation) and 2) handling of failover and failback when the active path is compromised (Failover). In addition the ability to alternate which path is active was supported through the introduction of round robin path selection in ESX release 3.5. Through the command line one can change the default parameters to have the path switched after a certain number of blocks or a set number of transactions going across an active path. These values are not currently available in the vCenter interface and can only be adjusted via the command line. Another issue which the ESX server storage stack has to address was how to treat a given storage device as some offered and active/active (A/A) controllers and others had active/passive controllers. A given LUN can be addressed by more than one storage controller at the same time in an active/active array. Where as the active/passive (A/P) array may access a LUN from one controller at a time and requires the transfer control to the other controller when access thru that controller is needed. To deal with this difference in array capabilities, VMware had to place logic in the code for each type of storage array supported. That logic defined what capabilities the array had with regards to active/active or active/passive as well as other attributes.

Path selection and failback options included: Fixed, Most Recently Used and Round Robin. The Fixed path policy is recommended for A/A arrays. It allows the use to configure static load balancing of paths across ESX hosts. The MRU policy is recommended for A/P arrays. It prevents the possibility of path thrashing due to a partitioned SAN environment. Round robin may be configured for either A/A or A/P arrays, but care must be taken to insure that A/P arrays are not configured to automatically switch controllers. If A/P arrays are configured to automatically switch controllers (sometimes called pseudo A/A), performance may be degraded due to path ping ponging.

vStorage API for Multipathing

Pluggable storage architecture (PSA) was introduced in vSphere 4.0 to improve code modularisation, code encapsulation, and to enable 3rd party storage vendors multipath software to be leveraged in the ESX server. It provides support for the storage devices on the existing ESX storage HCL through the use of Storage Array Type Plugins (SATP) for the new Native Multipathing Plugins (NMP) module. It also introduces an API for 3rd party storage vendors to build and certify multipath modules to be plugged into the VMware storage stack.

In addition to the VMware NMP, 3rd party storage vendors can create their own pluggable storage modules to either replace the SATP that VMware offered. Or they could write a Path Selection Module (PSP) or MultiPath Plug (MPP) in that could leverage their own storage array specific intelligence for increased performance and availability.

As with earlier versions of ESX server, the VI admin can set certain path selection preferences via the vCenter path selection preferences. The path selection options of Fixed, MRU and Round Robin still exist as options in vSphere. A default value, set by the SATP that matched the storage array could be changed even if not advised as best practice. But with the new 3rd party SATP the vendor could set what defaults were optimized for their array. Through this interface, ALUA became an option that could be supported. The VMware Native MultiPath module has sub-plugins for failover and load balancing.

- **Storage Array Type Plug-in (SATP)** to handle failover and the
- **Path Selection Plug-in (PSP)** to handle load-balancing.

- NMP "associates" a SATP with the set of paths from a given type of array.
- NMP "associates" a PSP with a logical device.
- NMP specifies a default PSP for every logical device based on the SATP associated with the physical paths for that device.
- NMP allows the default PSP for a device to be overridden.

Rules can be associated with NMP as well as 3rd party provided MPP modules. Those rules stored in the /etc/vmware/esx.conf file.

Archives

[February 2010](#)

[October 2009](#)

[September 2009](#)

[June 2009](#)

[September 2008](#)

[July 2008](#)

[Subscribe to this blog's feed](#)

Rules govern the operation of both the NMP and the MPP modules.
In ESX 4.0 rules are configurable only thru the CLI and not through vCenter.
To see the list of rules defined use the following command:

```
# esxcli corestorage claimrule list
```

To see the list of rules associated with a certain SATP use:

```
# esxcli nmp satp listrules -s <specific SATP>
```

The primary functions of an SATP are to:

- o Implements the **switching of physical paths** to the array when a path has failed.
- o Determines when a **hardware component** of a physical path **has failed**.
- o **Monitors the hardware state** of the physical paths to the storage array.

VMware provides a **default SATP** for each supported array as well as a **generic SATP** (an active/active version and an active/passive version) for non-specified storage arrays.

To see the complete list of defined **Storage Array Type Plugins (SATP)** for the **VMware Native Multipath Plugin (NMP)**, use the following commands:

```
# esxcli nmp satp list
```

Name	Default PSP	Description
VMW_SATP_ALUA_CX	VMW_PSP_FIXED	Supports EMC CX that use the ALUA protocol
VMW_SATP_SVC	VMW_PSP_FIXED	Supports IBM SVC
VMW_SATP_MSA	VMW_PSP_MRU	Supports HP MSA
VMW_SATP_EQL	VMW_PSP_FIXED	Supports EqualLogic arrays
VMW_SATP_INV	VMW_PSP_FIXED	Supports EMC Invista
VMW_SATP_SYMM	VMW_PSP_FIXED	Supports EMC Symmetrix
VMW_SATP_LSI	VMW_PSP_MRU	Supports LSI and other arrays compatible with the SIS 6.10 in non-AVT mode
VMW_SATP_EVA	VMW_PSP_FIXED	Supports HP EVA
VMW_SATP_DEFAULT_AP	VMW_PSP_MRU	Supports non-specific active/passive arrays
VMW_SATP_CX	VMW_PSP_MRU	Supports EMC CX that do not use the ALUA protocol
VMW_SATP_ALUA	VMW_PSP_MRU	Supports non-specific arrays that use the ALUA protocol
VMW_SATP_DEFAULT_AA	VMW_PSP_FIXED	Supports non-specific active/active arrays
VMW_SATP_LOCAL	VMW_PSP_FIXED	Supports direct attached devices

To take advantage of certain storage specific characteristics of an array, one can install a 3rd party SATP provided by the vendor of the storage array, or by a software company specializing in optimizing the use of your storage array.

Path Selection Plug-in (PSP).

The primary function of the PSP module is to determine which physical path is to be used for an I/O request to a storage device. The PSP is a sub-plug-in to the NMP and handles load balancing. VMware offers three PSPs:

To see a complete list of Path Selection Plugins (PSP) use the following command:

```
# esxcli nmp psp list
```

Name	Description
VMW_PSP_MRU	Most Recently Used Path Selection
VMW_PSP_RR	Round Robin Path Selection
VMW_PSP_FIXED	Fixed Path Selection

Fixed — Uses the designated **preferred path**, if it has been configured. Otherwise, it uses the first working path discovered at system boot time. If the ESX host cannot use the preferred path, it selects a random alternative available path. The ESX host automatically reverts back to the preferred path as soon as the path becomes available.

Most Recently Used (MRU) — Uses the first working path discovered at system boot time. If this path becomes unavailable, the ESX host switches to an alternative path and continues to use the new path while it is available.

Round Robin (RR) — Uses an automatic path selection rotating through all available paths and enabling load balancing across the paths. It only uses active paths and is of most use on Active/Active arrays. However in Active/Passive arrays, it will load balance between on ports to the same Storage Processor. However, Round Robin policy is not supported in MSCS environments.

3rd party vendors can provide their own PSP to take advantage of more complex I/O load balancing algorithms that their storage array might offer. However, 3rd party storage vendors are not required to create plug-ins. All the storage array specific code supported in ESX release 3.x has been ported to ESX release 4 as a SATP plug-in.

MultiPath Plugins (MPP) or Multipath Enhancement Module (MEM)

Multipath Plugin (also referred to as Multipath Enhancement Modules (MEM)) enable storage partners to leverage intelligence within their array to present one connection that is backed by an aggregation of several connections. Using the intelligence in the array with their own vendor module in the ESX server, these modules can increase the performance with load balancing and increase availability of the multiple connections. These modules offer the most resilient and highly available multipathing for VMware ESX servers by providing coordinated path management between both the ESX server and the storage.

A MPP "claims" a physical path and "manages" or "exports" a logical device. The MPP is the only code that can associate a physical path with a logical device. So that path can not be managed by both a NMP and MPP at the same time.

The PSA introduces support for two more features of the SCSI 3 (SPC-3) protocol: TPGS and ALUA.

However, although SCSI-3 is supported for some VMFS functions (LSI emulation for MSCS environments) the volume manager functions within VMFS are still SCSI-2 compliant. As such, the 2TB per LUN limit still applies.

Target Port Group Support (TPGS) - Is a mechanism that gives storage devices the ability to specify path performance, and other characteristics to a host, like an ESX server. A Target Port Group (TPG) is defined as a set of target ports that are in the same target port Asymmetric Access State (AAS) at all times with respect to a given LUN.

Hosts can use AAS of a TPG to prioritize paths and make failover and load balancing decisions.

Control and management of ALUA AASs can operate in 1 of 4 defined modes:

1. **Not Supported** (Report and Set TPGs commands invalid).
2. **Implicit** (TPG AASs are set and managed by the array only, and reported with the Report TPGs command)
3. **Explicit** (TPG AASs are set and managed by the host only with the Set TPGs command, and reported

with the Report TPGs command)

4. **Both** (TPG AASs can be set and managed by either the array or the host)

Asymmetric Logical Unit Access (ALUA) is a standard-based method for discovering and managing multiple paths to a SCSI LUN. It is described in the T10 SCSI-3 specification SPC-3, section 5.8. It provides a standard way to allow devices to report the states of their respective target ports to hosts. Hosts can prioritize paths and make failover/load balancing decisions.

Since target ports could be on different physical units, ALUA allows different levels of access for target ports to each LUN. ALUA will route I/O to a particular port to achieve best performance.

ALUA Follow-over Feature:

A Follow-over scheme implemented to minimize path thrashing in cluster environments with shared LUNs. Follow-over supported for explicit ALUA only and is not part of the ALUA standard.

When an ESX host detects a TPG AAS change that it did not cause:

- It will not try to revert this change even if it only has access to non-optimized paths
- Thus, it follows the TPG of the array

Posted by Paul Manning on October 19, 2009 | [Permalink](#) | [Comments \(0\)](#) | [TrackBack \(0\)](#) |  [BOOKMARK](#) 

September 29, 2009

VMware vStorage reduces storage challenges

VMware introduced VMware vStorage at VMworld 2008 to highlight key capabilities and interfaces that VMware has been developing to enable vSphere administrators to rapidly manage, plan and provision the storage infrastructure to meet the requirements of their virtualized applications. vStorage reduces the complexity of managing the underlying storage and enables administrators to integrate with existing storage administration tools. vStorage helps the vSphere administrator and storage administrator to work together more efficiently. vStorage is a set of features and APIs offered by VMware to reduce storage challenges to meet application requirements. vStorage includes both technologies delivered by VMware as well as technologies from storage ecosystem partners that are integrated with the VMware platform. VMware vStorage includes features like VMFS, Storage VMotion, Thin Provisioning, and the vStorage APIs (which currently include APIs for Multipathing, Data Protection and Site Recovery Manager). vStorage fulfills the need of all customers from small to large businesses. Also, VMware vStorage will continue to add new features and APIs with future vSphere releases to provide better integration and user experience in a vSphere environment.

VMware vStorage is not a product that you buy as an add-on to vSphere. vStorage is part of vSphere and its components are packaged to include functionality targeted to specific needs. More information on the various editions can be found [here](#).

vStorage technologies and interfaces help customers reduce the complexity of managing the underlying storage, meet the service levels required by their applications, and leverage their choice of storage architectures and storage management tools.

My next post will be on the vStorage APIs. In the meantime if you would like to read more on the new vStorage features in vSphere see the [What's new in vSphere - Storage paper](#).

Authors: Jon Bock, Group Product Marketing Manager - BC/ DR &

Ibrahim "Ibby" Rahmani, Sr Product Marketing Manager - Storage

Posted by Ibby Rahmani on September 29, 2009 | [Permalink](#) | [Comments \(0\)](#) | [TrackBack \(0\)](#) |  [BOOKMARK](#) 

June 16, 2009

Summary of what is new for storage in vSphere

With the launch of vSphere 4.0, VMware introduced a new version of the ESX Server and the management server, vCenter Server. In addition vSphere enables greater resource efficiency, management control, and freedom of choice for virtual datacenter storage resources and connectivity options. A brief summary of these new storage management capabilities in vSphere 4.0 are as follows:

Virtual Disk Thin Provisioning

Virtual disk thin provisioning increases the efficiency of storage utilization for virtualization environments by dynamically allocating storage and using only the amount of underlying

storage resource needed to store the contents of the virtual disk. On average customers feedback is that they expect a reduction of 50% of their storage space.

Improved iSCSI Software Initiator Efficiency

The new vSphere iSCSI SW initiator has been optimized for virtualization IO and provides significant improvements in CPU efficiency and throughput when compared to the earlier 3.5 version. VMware has released a new iSCSI software initiator that is far more efficient in its use of the ESX CPU cycles to drive storage IO. The entire iSCSI software initiator stack has been tuned to optimize cache affinity, enhance VMkernel TCP/IP stack, and to make better use of internal locks.

New vCenter Server Storage Enhancements

To help manage storage space in vSphere, each VM and ESX in the inventory now has its own storage tab that shows information about storage resources for those objects. The datastore can now be also managed as an object within vCenter in a manner that allows the vSphere administrator to view, group, and set permissions for each datastore. vSphere also provides a detailed view of all the components in the storage layout. These topology maps provide key information to administrators about which paths are available, as well as the grouping of objects sharing storage resources and new alerts and alarms are a key capability for to effectively managing datastores in which thin provisioned disks might create an over allocation of storage resources.

Enhanced Storage VMotion

Management of Storage VMotion is now fully integrated into vCenter Server. It offers full support to move a VM home from one datastore to another to facilitate storage migration. Some of the use case for the migration can be changing the storage protocol (FC, iSCSI or NFS) and/or change from thick to thin virtual disk format without impacting the VM. In vSphere, Storage VMotion is also significantly more efficient because it leverages a change block tracking method that moves the VM home over to the new datastore much faster than in snapshot method used in the ESX Server release 3.5, and it no longer requires twice the CPU and memory of a give VM.

Dynamic Expansion of VMFS Volumes

VMFS Volume Grow offers a new way to increase the size of a datastore that resides on a LUN that has been recently increased in capacity, via the vSphere Client GUI without disrupting the running virtual machine. It complements the dynamic LUN expansion capability that exists in many storage array offerings today. If a LUN is increased in size, then VMFS Volume Grow enables the VMFS volume extent to dynamically increase in size as well.

vStorage APIs for Multipathing

In vSphere, leveraging third party storage vendor multi-path software capabilities has been introduced through a modular storage architecture that allows storage partners to write a plug-in for their module to leverage. These modules can communicate with the intelligence running in the storage array to determine the best path selection, as well as leveraging parallel paths to increase performance and reliability of the IO from the ESX to the storage array. By default the native multipath driver (NMP) supplied by VMware will be used. It can be configured to support round-robin multipath as well. However, if storage vendor module is available, it can be configured to manage the connections between the ESX and the storage. There are three types of plugins, Path Selection Plug-in (PSP), Storage Array Type Plugin (SATP) and Multi-Pathing Plug-in (MPP.)

Conclusion

VMware is now providing more storage options and features that provide greater control, choice and efficiency to be leveraged in a vSphere environment. The new storage features, as well as previously existing ones (like VMFS) are grouped into an infrastructure category named vStorage. vSphere 4.0 enables the vSphere administrator to do more with less. Extending the benefits provided by virtualization reach new levels of storage efficiency, management control and additional choices in storage connectivity. To read more on the new features in vSphere see the [What's new in vSphere - Storage paper](#), for detailed information on configuring see the [Configuration Guide](#).

September 03, 2008

VMFS vs. NFS for VMware Infrastructure?

The dynamic, flexible environment that we call VMware Infrastructure requires shared, coordinated storage between ESX servers. There are two families of storage technologies that can meet this requirement today, SAN-based block storage (e.g. Fibrechannel or iSCSI) and NAS. VMware supports both forms of storage access for our customers.

VMFS is VMware's technology for managing block storage. It is a clustered file system used by multiple ESX Servers to coordinate and share access to the underlying block storage. It was explicitly designed to enhance the performance and reduce the complexity of block storage as it relates to virtual infrastructure. VMFS is responsible for two broad functions in the ESX storage subsystem.

1. **Structured data layout** – VMFS defines the structure and layout of data on the block devices. VMFS is optimized for virtual disk storage, but as is common to most file systems, VMFS brings a file and directory structure to the on-disk data. Most importantly, VMFS handles the layout of individual vmdk files (these are the virtual disks that are eventually presented to the VMs). VMFS also has some basic layout or volume management functions – presenting and managing combinations of multiple block storage devices (LUNs) as a single VMFS volume.
2. **Clustering** - The second and vitally important feature of VMFS is support for sharing or clustering. VMFS is responsible for the coordination of multiple independent ESX servers safely reading and writing to the same block device. It is VMFS which handles the basic safe transfer of ownership of data from server to server during high level VM activities such as vmotion, DRS, storage vmotion, HA, etc.

NFS is a NAS (Network Attached Storage) or file based protocol and is used to establish a client-server relationship between the ESX hosts and the NAS device. As opposed to block storage, a NAS system itself is responsible for managing the layout and structure of the files and directories on the physical storage. It also handles some of the aspects of shared access and provides primitives that can be used for coordinated access from multiple servers. ESX uses these primitives to provide similar shared access for VMware Infrastructure as is provided by VMFS in the block storage case.

So which to use? The first criteria is to continue to use the type of storage infrastructure you are familiar with. If your organization uses block based storage – use VMFS. If NAS is in use, it may make more sense to deploy VMware Infrastructure with NFS. Other aspects of storage management, such as the basic virtualization of storage on behalf of the VM or the internal structure of the virtual disk files (VMDK) are handled independently of this choice. You get the same high level VM functionality regardless.

For new deployments, there are the traditional storage tradeoffs. SAN or block based storage is a well understood infrastructure choice, has a great deal of existing physical management tools and may offer optimal performance. NAS-based storage has a reputation for less complexity and is in many ways easier to manage and setup than block based storage. And for some VM workloads, the performance of NAS is comparable to SAN. Individual NAS and SAN products may offer advanced data management features that become even more valuable with Virtual Infrastructure such as data de-duplication and thin provisioning.

As with many of VMware's virtualization features, we place a great deal of emphasis on the encapsulated VM abstraction. Consistent with this philosophy, the virtual disks that are presented to virtual machines are an idealized and highly portable view of storage which is independent of the backing storage implementation. ESX handles both storage technologies in a fashion that continues to maintain this portability of the VM. With ESX and VMware Infrastructure, one can move VMs between storage technologies quite easily and now with Storage vMotion, that portability even applies to live running workloads, offering customers the freedom to readily switch between NAS or block storage as they see fit.

July 21, 2008

VMware's "Proprietary" Clustered File System

There has been some discussion lately from some of our competitors on their storage architectures or lack thereof. Their boast is that VMware has a "proprietary clustered file system" that forces you to change your storage management practices while they preserve your physical storage management. They're right, they do completely preserve the complexity of physical storage management. What they fail to mention is that this is a problem, not a solution. Physical SAN management is overly complex and a headache for most customers. And in the virtual world, the status quo is compounded! Instead of managing dedicated LUNs for every physical server, if you buy their argument you'll be managing separate dedicated LUNs for every virtual machine!

VMware doesn't have a vested interest in preserving physical complexity. That's why VMware invented VMFS and our hypervisor-based storage stack. Virtualization is all about abstracting the physical world to ease management burdens and reduce complexity. VMFS is very special purpose clustered file system that is largely transparent to our customers. It is used by VMware's ESX Hypervisor to abstract and share LUNs used to store Virtual Machines and their Virtual Disks.

The result? Great customer benefits, such as:

1. VMware's instant one click provisioning, including storage. Quick, easy provisioning of a new VM, OS and application that does not require physical storage LUN provisioning.
2. Mobility/Portability, i.e. vMotion and storage vMotion. In a virtual world, workloads should be abstracted from, not beholden to, physical storage. Just like they should be abstracted from physical servers.

3. Encapsulation and HW Independence. VMs should be entirely encapsulated from the physical world. This simple, but critically important facet of virtualization unleashes the power of virtual infrastructure. For an example, look at VMware's new Site Recovery Manager that enables DR solutions that no longer require identical hardware (and software) configurations at each site.
4. Reduced complexity. SAN management is hard, complicated work. Why shouldn't it be simplified?

Now is VMFS good for absolutely every use case? No, sometimes it's important to preserve a physical LUN format for a given workload. That's why VMware supports RDMs (Raw Disk Maps) which allows a customer to directly map an entire LUN into the guest OS while still maintaining virtual infrastructure benefits. The trade off with RDMs is more LUNs to provision and manage, but direct access from the VM to a native physical LUN.

What our competitors fail to mention is that not only does VMware support all the use cases they claim to provide, but we also bring all the benefits of virtualization to the entire data center, storage included. And that's something worth touting!

Posted by Scott Davis on July 21, 2008 in [Weblogs](#) | [Permalink](#) | [Comments \(5\)](#) | [TrackBack \(0\)](#) |  [BOOKMARK](#) 

Welcome

Welcome to the VMware Storage Blog, a blog dedicated to all aspects of storage in the virtual environment. We get a lot of questions about Storage Architecture in VMware Virtual Infrastructure. It is a complex topic with many options and design trade-offs to consider. In this blog we will attempt to demystify the topic. Please let us know how we are doing! Thanks.

Posted by Scott Davis on July 21, 2008 in [Weblogs](#) | [Permalink](#) | [Comments \(3\)](#) | [TrackBack \(0\)](#) |  [BOOKMARK](#) 