3PAR Utility Storage with VMware vSphere

Optimize, simplify, and save in VMware vSphere environments with HP 3PAR Storage Systems

Technical white paper

Table of contents

Executive summary

This paper discusses the benefits of deploying VMware vSphere with HP 3PAR Utility Storage including: greater virtual machine (VM) density, simplifying administration, and realizing significant cost savings in virtualized server environments. The paper also includes best practices for an integrated VMware vSphere and HP 3PAR Utility Storage solution.

Target audience: The information is intended to assist solution architects and system integrators designing or implementing a VMware virtualized environment utilizing HP 3PAR Storage Systems.

Introduction

Server virtualization from VMware provides a powerful environment for consolidating a large number of servers, thereby delivering cost savings and enhanced flexibility to the data center. HP 3PAR Utility Storage provides an ideal complement to VMware vSphere deployments by providing a highly virtualized storage platform designed to meet the specific needs of virtual server environments.

This paper presents the benefits of using HP 3PAR Storage Systems for server virtualization with VMware vSphere and documents best practices for this type of deployment.

Overview

When supported with the correct underlying storage platform, server virtualization delivers greater consolidation, administrative efficiency, and cost savings. As a result, server virtualization is not only transforming the data center, but also the businesses that those data centers fuel. However, these transformative results depend on enterprise storage to deliver the performance, availability, and flexibility to keep up with the dynamic and consolidated nature of virtualized server environments.

HP 3PAR Utility Storage was built from the ground up to exceed the economic and operational requirements of even the most demanding and dynamic IT environments, and to support a converged infrastructure by providing the SAN performance, scalability, and availability that clients need to transform the data center. The next generation of federated Tier 1 storage, HP 3PAR Utility Storage delivers 100% of the agility and efficiency demanded by virtual data centers and cloud computing environments as part of an HP Converged Infrastructure. It does this through an innovative system architecture that offers storage federation, secure multi-tenancy, built-in thin processing capabilities, and autonomic management and storage tiering features that are unique in the industry.

When deployed together, VMware vSphere and HP 3PAR Utility Storage deliver a compelling virtual data center solution that increases overall resource utilization, provisioning agility, administrative efficiency, and both capital and operating costs.

HP 3PAR Utility Storage uniquely enhances the flexibility and return on investment (ROI) of VMware vSphere deployments in the following ways:

Greater virtual machine (VM) density: The unique HP 3PAR Architecture stripes volumes widely across all drives to deliver maximum I/O throughput and minimum latencies, which mitigates server memory bottlenecks and traditional storage constraints. Increased array performance can not only boost VM-based application performance, but when paired with the superior reliability of the HP 3PAR Storage System and advanced support of VMware's vSphere Storage APIs – Array Integration (VAAI) capabilities, can result in higher VM density. This benefit enables organizations to double virtual machine density on physical servers, placing twice as many VMs on physical servers as with traditional storage platforms.

Simplified storage administration: HP 3PAR Utility Storage reduces storage administration time by up to 90%. Rapid provisioning, autonomic load balancing, and software management products such as HP 3PAR Autonomic Groups automate repetitive storage administration tasks. Autonomic storage tiering tools like HP 3PAR Dynamic Optimization Software allow administrators to tune storage allocations without interruption to servers. Storage federation capability enables data and workloads to be simply and fluidly shifted between storage arrays without disruption.

Integration with VMware vCenter Server: The HP 3PAR Management Software Plug-In for VMware vCenter allows administrators to monitor HP 3PAR storage volumes from within the vSphere console. HP 3PAR Recovery Manager Software for VMware vSphere uses HP 3PAR Virtual Copy Software to create point-in-time, VM- and application-aware, disk-based snapshots giving VMware administrators a simple process for recovering whole Virtual Machine Disks (VMDKs), individual VMs, or even individual files.

Simplified disaster recovery: HP 3PAR Replication Adapter Software for VMware vCenter SRM provides simple and reliable replication and disaster recovery for VMware vSphere environments based on HP 3PAR Remote Copy Software.

Broader VMware vSphere deployment through cost-effective HP 3PAR Thin Provisioning Software: HP 3PAR Utility Storage enhances flexibility and reduces physical storage capacity needs in VMware vSphere environments by leveraging the cost-effective usable capacity delivered by Fast RAID 5 and RAID 6 (also called RAID Multi-Parity or RAID MP) and HP 3PAR Thin Provisioning Software. Thin Provisioning allows clients to safely over-allocate capacity upfront and then purchase actual physical capacity incrementally and only for written data. The result is a savings of 50% or more on purchased capacity as compared to traditional storage platform – guaranteed[1].

HP 3PAR support for VMware end-user computing initiatives: HP 3PAR Utility Storage provides rapid provisioning of both storage and virtual desktops for VMware View and VMware View Composer deployments. Together, these solutions provide storage efficiency for boot images as well as user data and deliver the performance, scalability, and high availability demanded by enterprise end-user computing deployments.

Greater virtual machine density

Consolidating workloads in virtual server environments is key to driving increased savings. The number of virtual machines per physical server – known as VM density – is a leading indicator of consolidation success and is a function of high performance and reliability. HP 3PAR Utility Storage, wide striping, mixed workload support, and integration with the VMware Adaptive Queue Depth Throttling algorithm and Storage APIs provide the performance necessary to increase VM density per server by at least twofold when compared to VMware deployments using traditional storage arrays.

Since a VMware ESX/ESXi server can host many different VMs, each with its own I/O patterns, I/O is generally random in nature. Random I/O means the server derives less benefit from storage caching than a server with a more predictable access pattern, forcing more requests to access disk. Thus, performance is determined by the number of drives (IOPS density) that make up the LUN. Wide striping spreads even modestly sized volumes across all drives in the HP 3PAR Storage System, resulting in unmatched I/O performance, even when caching is minimally effective.

In addition to using storage for files and similar data, VMware ESX implements aggressive memory management techniques, including higher rates of paging and swapping to/from disk than are typical of contemporary, non-virtualized servers. VMware therefore recommends placing swap files on high-performance storage. In ESX deployments, the amount of memory that can be installed in the server often limits the number of virtual machines that can be placed on a single physical server. HP 3PAR clients have found that the improved swapfile performance afforded by using HP 3PAR Storage

[1] Get Thin Guarantee: http://www.hp.com/storage/getthin

Systems allows a twofold to fourfold increase in the number of VMs per physical server, which is otherwise constrained by server memory. For even faster performance, SSD devices can be utilized in conjunction with the 3PAR Storage Systems and HP 3PAR Adaptive Optimization Software for autonomic tiering that optimizes both cost and performance.

The increased reliability of HP 3PAR arrays is a function of architectural features such as redundant hardware, fast RAID rebuild capability, and RAID isolation as well as HP 3PAR InForm Operating System Software features such as HP 3PAR Persistent Cache. These resiliency features collectively minimize service outages, allowing greater consolidation in virtualized server environments without undue risk.

The following sections provide an overview of HP 3PAR Utility Storage features that contribute to higher performance and reliability in a VMware environment, and provide a number of best practices for maximizing performance with ESX and HP 3PAR Storage Systems.

Performance and reliability features

As mentioned previously, the HP 3PAR Architecture includes a number of features that drive higher levels of performance and reliability across a VMware vSphere deployment:

Wide striping distributes each virtual volume across all drives and other resources of the array. This applies to both VMware Virtual Machine File System (VMFS) and Raw Device Mapping (RDM) volumes, including those which use HP 3PAR Thin Provisioning Software. Wide striping allows even small volumes to gain the performance benefits of many drives and provides for consistent levels of performance, even at increased levels of capacity utilization.

Mesh-Active clustering of controllers ensures that all array resources are available to all LUNs at all times – providing the host with symmetric, active-active access to all LUNs. Unlike traditional arrays, where controllers have preferred drives and access to "foreign" drives within the array is expensive, with HP 3PAR arrays, the high performance, Mesh-Active backplane enables each LUN to have uniformly high performance access to all resources within the array – regardless of which controller nodes are used to connect to the array.

Tech Tip
ESX/ESXi servers gain the full performance benefits of the Mesh-Active design by changing from the default "Fixed" to "Round Robin" multipathing.

Adaptive cache provides dynamic and autonomic balancing of read and write cache levels to increase performance by adjusting caching without the need for manually setting or modifying policies or schedules.

Mixed workload support enables different types of applications (involving both transaction-based and throughput-intensive workloads) to run without contention on a single HP 3PAR array. Mixed workload support is especially important in consolidated and virtualized environments, where the same array must support a wide mix of application types across the cluster.

Persistent Cache provides resiliency in the event of a node failure by using the high-performance, Mesh-Active backplane to quickly re-mirror write-back cache to other nodes in the system. In the event of a node failure, Persistent Cache enables the HP 3PAR array to maintain the substantial performance benefits of write-back caching without risk of data loss. This feature allows "always on" application and virtual server environments to gracefully handle an unplanned controller failure without the substantial performance penalties associated with traditional arrays and "write-through" mode.

RAID isolation preserves access to data even if an entire drive chassis (up to 40 drives) is lost. Traditional RAID merely guarantees continued data access in the event of the failure of a single drive, or two drives with RAID 6.

Queue depth throttling

Each port on the HP 3PAR Storage System has a finite queue depth that depends on the host bus adapter (HBA) model; each server attached to a port shares that port's queue. If a host sends an I/O request to a port with a full queue, the host receives a "queue full" SCSI response from the HP 3PAR array. I/O commands sent to a port in an HP 3PAR array that has reached its maximum queue depth are not processed beyond the "queue full" SCSI response.

Historically, an ESX host's default reaction to this response would be to recognize it as a valid command and to continue sending requests to that port. Lack of I/O responses can result in VMs becoming unresponsive and can lead to a crash of the ESX server. ESX 3.5 Update 4 and later include an adaptive queue depth throttling algorithm which adjusts the LUN queue depth in the VMkernel I/O stack. This algorithm is activated when the storage array indicates I/O congestion by returning a "queue full" SCSI status. When congestion is detected, the VMkernel throttles the LUN queue depth and attempts to gradually restore the queue depth when congestion conditions subside.

Without adaptive queue depth throttling, administrators are forced to limit the number of VMs per physical server so as to reduce the risk associated with any particular VM overrunning I/O queues. Administrators are also forced to manually tune the number of VMs when they detect congestion –a reactive, slow, and error-prone process. By automating congestion control, administrators can confidently create a higher number of VMs per physical server without the need for manual congestion control.

For more information, refer to the VMware Knowledge Base entry on how to enable the VMware adaptive queue depth throttling algorithm: http://kb.vmware.com/kb/1008113. If it is not possible to enable VMware's adaptive queue depth throttling algorithm, then the *HP 3PAR VMware ESX Implementation Guide* contains details to tune/throttle the queue depth on the ESX host. To access this document, go to http://www.hp.com/go/3PAR, navigate to the "HP 3PAR InForm Operating System Software" product page under "3PAR Software", click "Support for your product", and then click "Manuals".

Tech Tip

With either the VMware adaptive queue depth algorithm or target throttling, no more than 16 ESX hosts should be attached to any one 2 Gb port, 32 hosts for 4 or 8 Gb ports. Even with the adaptive algorithm, it is still important to pay attention to the number of VMs that are hosted on each ESX server for performance.

You can monitor the "Qlen" values on the system (using System Reporter or the command `statvlun -ni -rw -host <ESX host>`) to make sure you are not exceeding these values.

Metadata locking with SCSI reservations

Metadata updates by ESX must ensure exclusive access when multiple ESX hosts are accessing shared storage. Locking is necessary to prevent multiple hosts from concurrently writing to the metadata, which would lead to data corruption. Examples of operations that require metadata updates include: powering on/off a VM, adding/removing a VMFS datastore, creating a new VM, using VMware vMotion, and growing a VMFS datastore.

Prior to vSphere 4.1, SCSI reservations were used to perform this locking. A SCSI reservation (in essence, is a lock on the LUN) was made before the metadata was updated, and released after the metadata update operation had completed, thus allowing other operations to continue. Under this system, too many reservations being made at once could lead to I/O failures if a host was unable to make a reservation because another host already had a lock on the LUN. When a host is unable to make a reservation due to a conflict with another host, it retries at random intervals until it is successful. If too many unsuccessful retries are made, the operation fails.

With SCSI reservations, it is best to limit the number of operations that can cause reservation conflicts and stagger them so that relatively few reservations are attempted at the same time. For example, it is not advisable to power on/off multiple VMs at the same time or to run multiple instances of VMware vMotion in parallel.

With vSphere 4.1 and later, the vSphere Storage APIs – Array Integration (VAAI) provide a means to perform this locking on the hardware at a much lower performance cost.

vSphere Storage APIs – Array Integration (VAAI)

The vSphere Storage APIs are a set of technologies and interfaces that enable vSphere to leverage storage resources to deliver the efficiency, control, and ease of customization that clients demand of their IT environment. The vSphere Storage APIs – Array Integration (VAAI) is one of these technologies. Under the VAAI initiative, APIs have been introduced to improve performance and scalability by leveraging efficient array-based operations.

HP has developed the HP 3PAR Management Software Plug-In for VMware VAAI to deliver enhanced performance, agility, and scalability using vSphere commands introduced in vSphere 4.1. Initial support of the SCSI (T10) standard was introduced in vSphere 4.1 through the use of block level commands. These standard commands were enabled by a standard VMware plug-in which enabled Hardware Assisted Locking, Fast Copy, and Block Zeroing.

vSphere 5.0 provides enhanced support for the T10 standards without the need to install a plug-in, enabling vSphere to directly utilize more advanced features of the storage array. HP 3PAR InForm OS version 2.3.1 does not offer native T10 support and will still require the installation of the 3PAR VAAI plug-in. However, InForm OS version 3.1.1 and later include native T10 support and no longer require the plug-in to take advantage of the enhanced VAAI functionality within vSphere 5.0.

Some of the important hardware commands that VAAI enables are documented below:

Hardware Assisted Locking eliminates SCSI reservation contention by providing a fast, fine-grained locking mechanism. The ATS ("Atomic Test and Set") command verifies that a block of metadata is what is expected (test) and then replaces it with an updated block (set) in a single, atomic operation. Using this command, the ESX server can lock a portion of a LUN related to a single VM instead of locking the whole LUN as described in the previous section, thereby allowing other VMs on the same LUN to continue operating normally. The implementation of ATS on HP 3PAR Utility Storage arrays uses the HP 3PAR ASIC to further improve performance. The combination of ATS and the HP 3PAR ASIC allows an increase in VM density per LUN and greater scalability for vSphere deployments.

Fast Copy uses the XCOPY command to improve the performance of common operations like VM cloning and Storage vMotion by performing large data movement operations directly within the storage array. By not requiring each block to make a round-trip to the host, the time required for these operations is significantly reduced and storage network traffic minimized. When combined with HP 3PAR Thin Persistence Software, drive I/O and storage capacity can also be reduced since blocks of zeros are not written due to the array's zero-detection capability, which is integrated into the HP 3PAR ASIC.

Block Zeroing uses the standard SCSI command WRITE_SAME to offload large, block-level write operations of zeros from the host to the storage array. Block zeroing improves host performance and efficiency when allocating or extending Eager Zeroed Thick (EZT) virtual disks, or on initial access to a block on a non-EZT virtual disk. When combined with built-in zero-detection and EZT virtual disks, storage array bandwidth, disk I/O bandwidth, and disk consumption is minimized. Initialization of EZT virtual disks in seconds rather than minutes eliminates the tradeoff between fast VM creation and fast run-time performance.

Space Reclamation is new in vSphere 5.0 and uses the standard T10 command UNMAP for space reclamation with thin provisioned volumes. Using the UNMAP command, vSphere can automatically issue a command to the storage array to unallocate blocks of storage freed by vSphere, for example from a deletion of a VM, and return it to the storage array's resource pool without the involvement of a storage administrator. Prior to vSphere 5.0 and UNMAP, space would remain assigned to the thin provisioned LUN until a storage administrator manually reclaimed the storage associated with it.

This enhancement improves the effectiveness of HP 3PAR thin provisioned volumes by minimizing the physical storage required. Minimizing the physical storage required increases the ROI through more efficient use of installed storage capacity and reduction of associated the power and cooling costs.

Tech Tip
If you have upgraded a legacy solution to ESXi 5.0/VMFS5, you can manually initiate a reclaim of space from a thin provisioned volume using `vmkfstools -y 60% <datastore>`.

Storage I/O Controls

Storage I/O Controls (SIOC), added in vSphere 4.1, is a feature which protects the storage performance of high-priority virtual machines. It does this by detecting high storage I/O latency (in excess of 30 ms response time by default). Once excess latency is detected, SIOC throttles the host-side queue depths of lower-priority VMs (based on administrator-defined shares of storage I/O bandwidth) to prevent low-priority VMs from gaining more than their share of I/O bandwidth. Integration of HP 3PAR Utility Storage with SIOC minimizes the impact of I/O congestion. This allows clients to increase the number of VMs per physical server and add more higher-performing applications to physical servers when attached to HP 3PAR arrays.

vSphere Storage APIs – Storage Awareness (VASA)

The vSphere Storage APIs – Storage Awareness (VASA) is a new set of APIs introduced with vSphere 5.0 that enables VMware vCenter Server to detect the capabilities of the storage array LUNs and their datastores. This visibility into the array's configuration of its datastores and their capabilities, simplifies vSphere administration with HP 3PAR Utility Storage. Capabilities such as RAID level, thin or thick provisioned, device type (SSD, Fibre Channel, or Nearline) and replication state can now be made visible within the standard vCenter disk management interface. This allows vSphere administrators to select the appropriate disk for virtual machine placement based on its needs. VASA eliminates the need for maintaining complex spreadsheets detailing the storage capabilities of each LUN previously required to guarantee the correct SLA.

The concept of a storage profile, introduced in vSphere 5.0, extends the base VASA functionality. These profiles are used in conjunction with the capabilities of the LUN to determine which LUNs meet the needs of a VM. vSphere 5.0 can use this information to migrate between LUNs for load balancing while maintaining the needs (RAID level, etc) of the virtual machine. These profiles also allow vSphere to make placement decisions automatically based on the needs of the VM and the available datastores, further reducing the administration impact.

VMFS versus RDM

VMware Virtual Machine File System (VMFS) is a clustered, hierarchical file system that is designed and optimized by VMware specifically for virtualized server environments. It is the container for the virtual disks used by VMs. VMFS increases resource utilization by allowing multiple ESX servers to access shared storage concurrently. It provides distributed locking for VMs so all VMs can operate safely in a Storage Area Network (SAN). In addition, VMFS uses an optimized architecture to enable users to achieve higher VM-to-physical server density. A VM is represented as a Virtual Machine Disk (VMDK) in a VMFS datastore.

Raw Device Mapping (RDM) is an alternative to VMFS. RDMs provide VMs with direct I/O access to raw storage volumes. RDMs can be used for providing shared storage between multiple VMs, for example a data or quorum disk in a cluster.

Cost savings

Today's IT administrator faces potentially spiraling storage costs as a result of the explosive growth of data and the new storage performance challenges created by technologies such as server virtualization. The massively parallel HP 3PAR Architecture uniquely addresses these performance challenges while HP 3PAR Thin Technologies – including HP 3PAR Thin Provisioning Software, HP 3PAR Thin Conversion Software, and HP 3PAR Thin Persistence Software – enable organizations to meet storage needs with 50% less capacity, energy, and floor space.

Virtual volumes: thin or thick?

In general, a "large" LUN should be configured as a VMFS datastore for use by multiple VMs. A large LUN gives VMware vSphere users the most flexibility by not requiring storage administrators to provision new storage every time a new VM is created. A large LUN also minimizes rescans for new LUNs on the ESX server, which can impact VMFS I/O. LUNs presented to ESX servers should be striped across many physical drives.

With its massively parallel architecture, the HP 3PAR Storage System uses all available drives of the same type behind its virtual volumes, which are presented as LUNs to the ESX servers. Workloads are automatically spread across all internal resources (nodes, drives, ports, etc.), delivering high application service levels and predictable levels of performance (high IOPS and low latencies). With wide and autonomic distribution of workloads, storage administrators need not spend time figuring out what drives are available in an array or how best to layout a LUN. Instead, creating a virtual volume on an HP 3PAR Storage System takes less than a minute and is a matter of a few clicks with the HP 3PAR InForm Management Console or a simple command from the command line interface.

With traditional arrays, storage administrators generally create LUNs in the 500 GB to 2 TB range for VMFS volumes, with a 1 TB LUN being a very common size. This seems to be the acceptable tradeoff between creating a large enough LUN to satisfy the performance requirements of VMware vSphere without wasting too much unused space in the LUN. However, HP 3PAR Thin Provisioning Software eliminates this tradeoff altogether.

With HP 3PAR Thin Provisioning Software, one can allocate as much logical capacity to a VMFS volume as is needed over the lifetime of that volume without actually dedicating any physical capacity up front. Physical capacity is allocated seamlessly on an as needed basis and is consumed only when vSphere writes to the thin provisioned virtual volumes. There is no wasted capacity in the LUN that is presented to the ESX host. Figure 1 and Figure 2 illustrate that the same 2 TB of physical storage can be used to host better than 2x the servers because of thin provisioning.

Figure 1. Traditional Provisioning

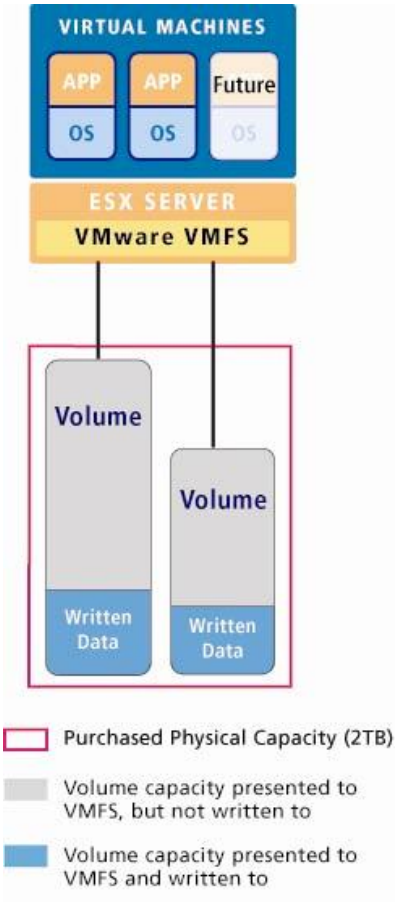

Figure 2. Thin Provisioning

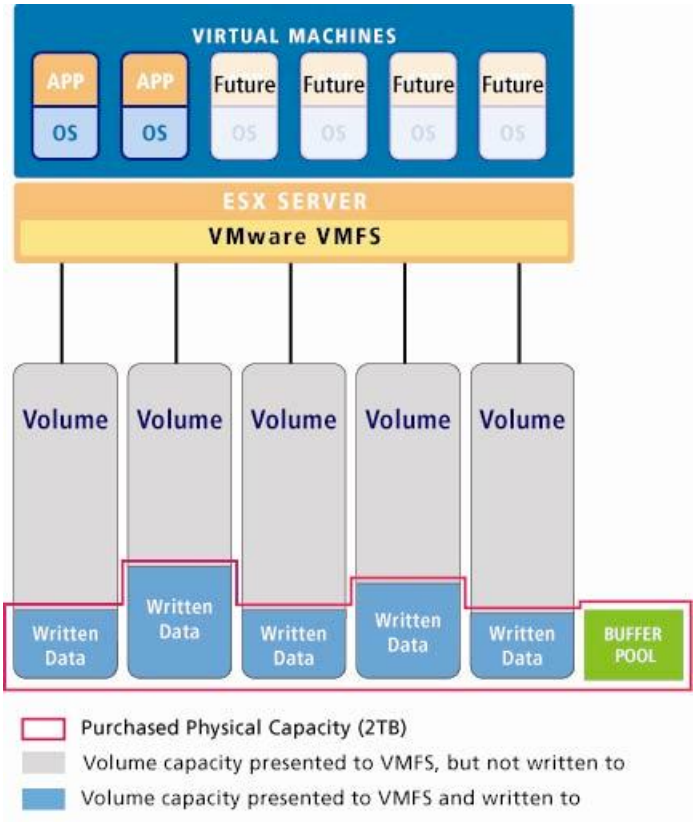

For this reason, compared to traditional arrays, storage utilization is very high when using HP 3PAR Thin Provisioning Software. Aside from the storage utilization differences between a thin provisioned LUN and a traditional "thick" provisioned LUN, the advantage provided by the HP 3PAR Architecture in this situation is that both LUN types enjoy the same wide striping across all the drives in the system. Unlike many traditional approaches to thin provisioning, HP 3PAR Thin Provisioning Software does not require the creation or use of separate pools of drives within the system.

Virtual disk format options

When creating VMs, there are a number of options that are available for the VMDK files, not all of which are "thin-friendly" (compatible with HP 3PAR Thin Provisioning Software). By default, VMware vSphere creates VMs using the *"Lazy Zeroed Thick"* option, which is thin-friendly. With this option, when a new VM is created, the full size of the VMDK is not immediately zeroed. It returns zeros when asked to read from unwritten areas but does not actually write the zeros to disk.

For performance-intensive environments, VMware recommends using "Eager Zeroed Thick" (EZT) virtual disks. EZT disks are also the only disk format supported with VMware Fault Tolerance (FT). EZT disks have the smallest overhead[2] but require zeros to be written across all of the capacity of the VMDK at the time of creation. Thus, this VMDK format negates all the benefits of traditional thinly provisioned LUNs since all of the physical storage is allocated when the volume is zeroed-out during

[2] VMworld 2009 session TA2942: vSphere 4.0 Performance Best Practices

creation. However, as described in the next section, using HP 3PAR Thin Persistence Software allows clients to retain thin provisioning benefits when using Eager Zeroed Thick VMDKs without sacrificing any of the performance benefits offered by this VMDK option.

HP 3PAR Thin Conversion and Thin Persistence Software

In a perfect world, all storage would start thin by using HP 3PAR Thin Provisioning Software. However, it is never too late to become more efficient, which is why the HP 3PAR storage platform also gives clients the ability to "thin" existing volumes using HP 3PAR Thin Conversion Software, and to keep them thin over time using HP 3PAR Thin Persistence Software. HP's unique 3PAR Thin Built In technology is a hardware capability built into the HP 3PAR ASIC to provide an efficient, silicon-based zero-detection mechanism for converting traditional "thick" volumes to more efficient "thin" volumes on the HP 3PAR array without performance impact.

This hardware capability leverages HP 3PAR's own "Thin Engine" – a software-based virtualization mapping engine for space reclamation – which gives HP 3PAR arrays the power to remove allocated but unused space in existing volumes.

Figure 3. Using HP 3PAR Thin Persistence Software to "thin" a volume

Together with VMware Storage vMotion or any standard data migration utility, HP 3PAR Thin Conversion Software uses the zero-detection capabilities of the HP 3PAR ASIC to convert "thick" volumes to "thin" volumes by identifying unused space (that is, zeros) in the volume. The conversion relies on the virtualization mapping capabilities of the HP 3PAR Thin Engine to eliminate the need to store any of the volume's unused space. Through this process, users can migrate from "thick" volumes on any other storage array to "thin" volumes on HP 3PAR Utility Storage to gain the benefits of HP 3PAR Thin Provisioning Software.

For volumes that start thin (or "get thin" with HP 3PAR Thin Conversion Software), the next challenge is to stay thin over time. This is where HP 3PAR Thin Persistence Software comes in. Like Thin Conversion, Thin Persistence uses Thin Built In and the HP 3PAR Thin Engine to reclaim unused space associated with deleted data within HP 3PAR storage volumes without disruption (Figure 4).

Figure 4. Using 3PAR Thin Persistence to stay thin over time

The process starts by activating HP 3PAR Thin Persistence Software (if necessary), then using standard file system tools (*sdelete* in Microsoft® Windows® or *dd* in Linux) to write zeros across deleted space in a VM's file system. The zeros are autonomically detected by the HP 3PAR ASIC and the disk space they were consuming is freed up and returned to the thin provisioned volume. Thus, as files are deleted within VMs, Thin Persistence ensures that the underlying volume stays thin over time.

Another benefit of Thin Persistence is when deleting a VM. The HP 3PAR Management Software Plug-In for VMware VAAI or T10 support enables the use of WRITE_SAME or UNMAP commands (depending on the version of vSphere/VMFS in use). These commands can enable the release of the allocated storage back to the array rather than keeping it assigned to the LUN. With older VMware solutions, a similar benefit can be obtained by writing zeros to the VMDK and allowing the HP 3PAR zero detection to release the storage. However, with vSphere 5.0 and either the HP 3PAR Management Software Plug-In for VMware VAAI or HP 3PAR InForm OS version 3.1.1 or higher, this space reclamation takes place autonomically.

Thin Persistence also ensures that when a new, fully formatted volume is created, the entire volume is not allocated from physical storage since only zeros have been written. This situation is encountered with the Eager Zeroed Thick VMDK format since this VMDK format first zeroes out the full VMDK file and thus consumes the full size of the VMDK file on the array even before the VM has been used.

With HP 3PAR Thin Persistence Software and the built-in zero-detection capability of the HP 3PAR ASIC, as the ESX host writes zeros to the VMDK file, the zeros are detected in-line by the ASIC, and no space is allocated for the VMDK in the thin provisioned volume. Contrast this with arrays that do not offer this capability: on those arrays, an Eager Zeroed Thick VMDK negates any thin provisioning benefits because it fills up the thin provisioned volume as it is being initialized.

Increased administrative efficiency

Virtual server environments enable increased administrative efficiency for servers, but these savings are often diluted by archaic and inflexible storage. HP 3PAR Utility Storage includes a number of products and features to assist the VMware administrator. HP 3PAR Autonomic Groups is a feature of the HP 3PAR InForm Operating System Software that enables rapid provisioning to reduce storage administration time in virtualized environments. The HP 3PAR Management Software Plug-In for VMware vCenter enhances visibility into storage resources for the VMware administrator by integrating HP 3PAR-specific information into vCenter.

HP 3PAR Management Software Plug-In for VMware vCenter

The HP 3PAR Management Software Plug-In for VMware vCenter is a vSphere management console plug-in that allows easy identification of HP 3PAR virtual volumes used by VMs and datastores. It provides an integrated view of the VMs and associated storage resources. Properties such as volume type (Thick or Thin Provisioned Virtual Volume (TPVV)), device type (FC disk, NL disk, or SSD), RAID level, etc. are displayed via the 3PAR tab in the vSphere management console.

The HP 3PAR Management Software Plug-In for VMware vCenter provides the VMware administrator a view into the system via a single pane (Figure 5). There is no need to login to the system to identify space consumption or determine how a volume maps to a datastore. This information is easily visible via the plug-in, as is capacity usage and other data.

Figure 5. Viewing virtual volume mapping information via the HP 3PAR plug-in

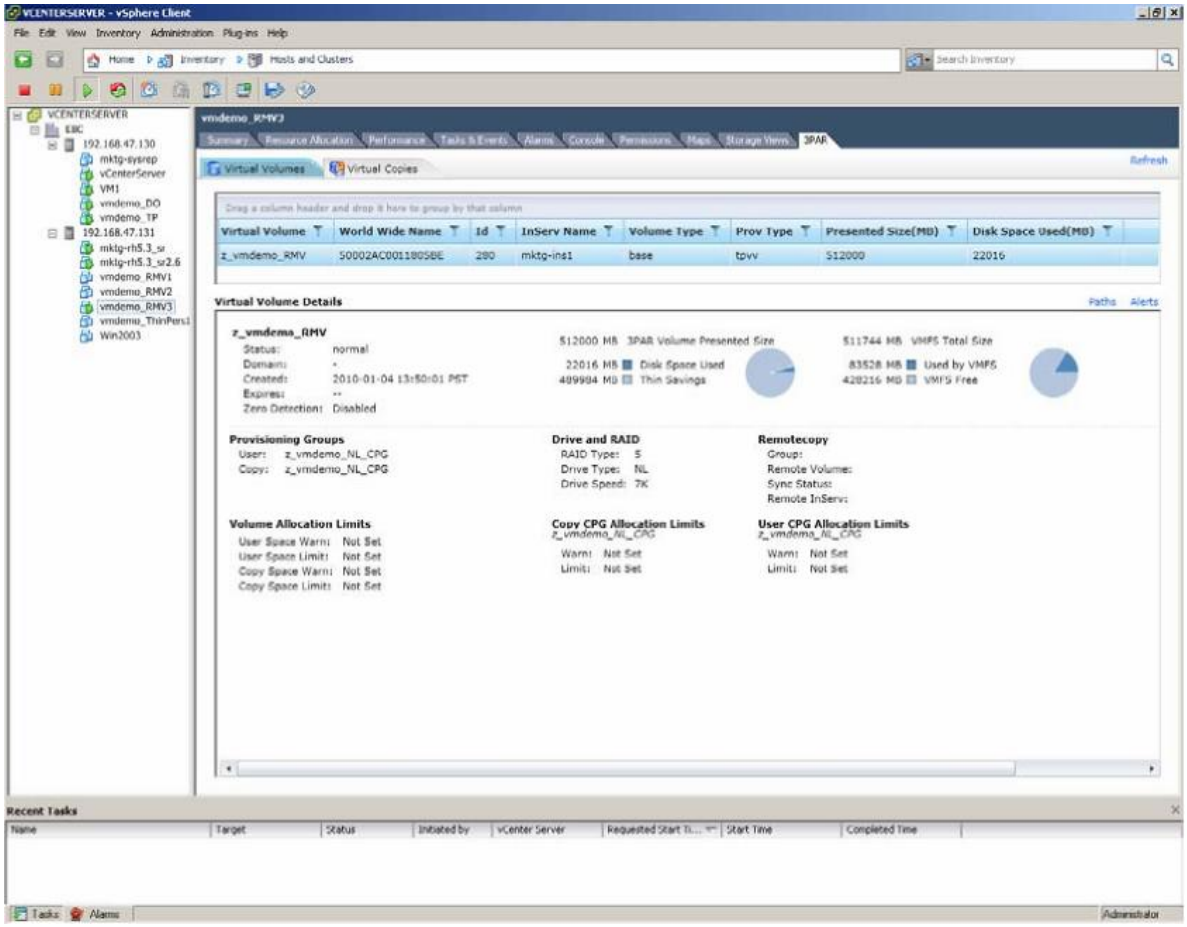

Beginning with vSphere 5.0, the HP 3PAR Management Software Plug-In for VMware vCenter also includes a VASA plug-in which allows vSphere to display detailed (device type, RAID level, etc.) information on the 3PAR LUNs directly in the vSphere disk management interface. This information can then be used to define storage profiles ensuring the storage meets the needs of the virtual machine.

HP 3PAR Recovery Manager Software for VMware vSphere

HP 3PAR Recovery Manager Software for VMware vSphere enables the protection and rapid recovery of VMs and datastores (Figure 6). It provides virtual copy management and allows the administrator to take LUN-level snapshots of VMs and datastores via the vSphere management console. HP 3PAR Recovery Manager Software provides array-based snapshots that are fast, space-efficient, and VM-aware. This plug-in solves the issues associated with traditional, agent-based backup schemes that are typically slow, complex, and fail to offer the flexibility and granularity that array-based snapshots can provide.

HP 3PAR Recovery Manager Software makes it possible to create hundreds of virtual copies. The number of virtual copies to retain and the retention period for each virtual copy can easily be specified. Once a virtual copy has been created, this plug-in allows the flexibility of granular restores at the VMFS layer, the VM layer, or the individual file level.

For more detailed information on the HP 3PAR Management Software Plug-In for VMware vCenter or HP 3PAR Recovery Manager Software for VMware vSphere, see http://www.hp.com/go/3PAR.

Figure 6. Virtual copy of a VM as displayed from inside the vSphere Management console

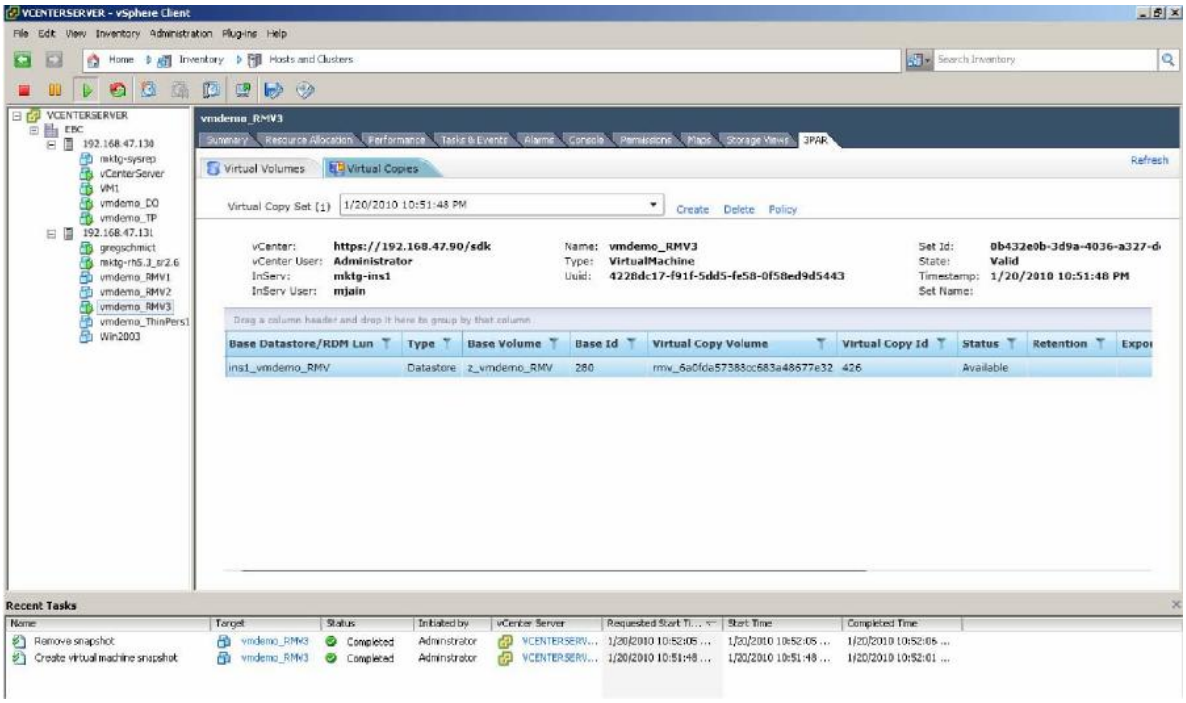

HP 3PAR Replication Adapter Software for VMware vCenter SRM

VMware vCenter Site Recovery Manager (SRM) provides end-to-end management of array-based replication, virtual machine failover, and automated disaster recovery management for environments that use VMware vCenter Server. HP 3PAR Replication Adapter Software for VMware vCenter SRM was developed to provide integration between VMware vCenter Site Recovery Manager (SRM) and HP 3PAR Remote Copy Software.

HP 3PAR Dynamic Optimization Software and VMware vSphere

HP 3PAR Dynamic Optimization Software is a software product that offers an online and non-disruptive way to make changes to volumes on the HP 3PAR Storage System. Storage administrators can move volumes between different drive types (Fibre Channel, Nearline, SSD), convert between RAID levels (RAID 1, RAID 5, or RAID 6/RAID MP), and/or rebalance volumes as new drives are added, all without impacting any hosts that the system is busy serving.

In VMware vSphere environments, HP 3PAR Dynamic Optimization Software can be used to move running VMs between different drive tiers without impacting what the VMs are doing. Similarly, as new drives are added to the array, the LUN that ESX is using can be striped across the new drives on the fly without taking an outage at the ESX server level.

VMware vMotion offers somewhat similar functionality, but at the hosting layer. HP 3PAR Dynamic Optimization Software works at the storage layer, bypassing any application-level overheads. Therefore, it can be used to optimize data service levels while VMware vMotion can be used to optimize CPU utilization across multiple hosts.

Boot from SAN environments

In a boot from SAN environment (Figure 7), each ESX server's operating system is installed on a volume on the array, not on the server's internal disk. In this situation, administrators should create a separate virtual volume for each ESX server to be used for the boot image. The "boot from SAN" approach brings a number of benefits to a VMware vSphere environment. For example, HP 3PAR Virtual Copy Software can be used to take snapshots of the boot image and provide rollback capabilities when upgrading ESX. Users can toggle between different ESX versions by simply pointing the ESX server's HBA to the different LUNs containing the respective ESX versions. In addition, replacing an ESX server simply means pointing the new server to the boot LUN.

Note: For detailed information on configuring your server for boot from SAN see http://h18000.www1.hp.com/storage/networking/bootsan.html.

Figure 7. A "Boot from SAN" configuration

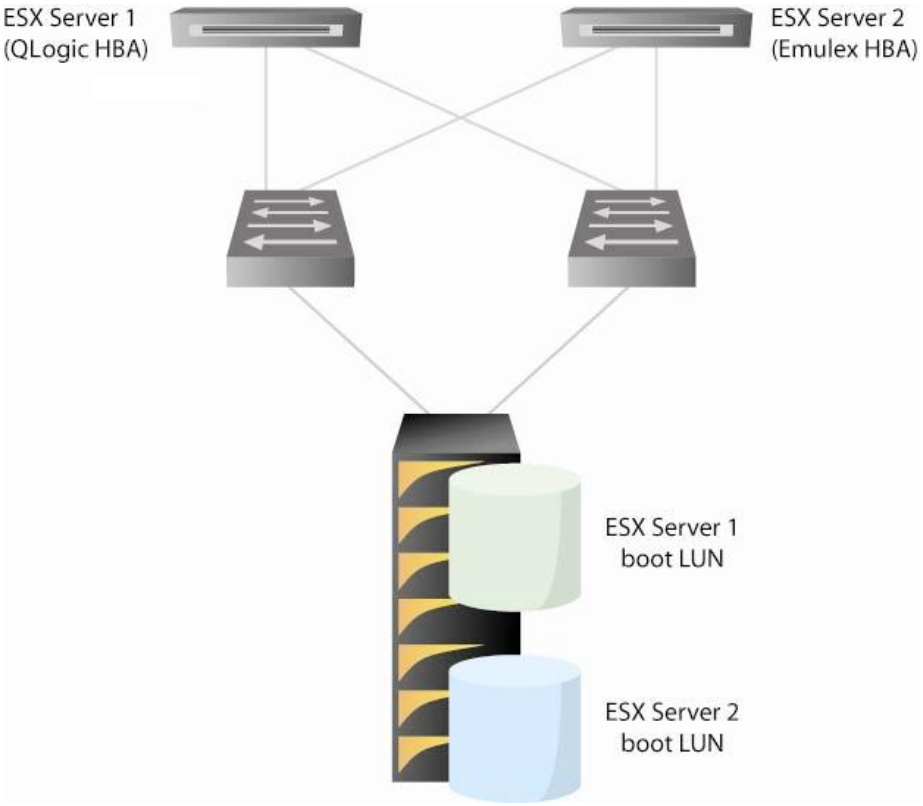

Conclusions

Infrastructure convergence holds the keys to enabling organizations to overcome the inflexibility and high costs created by IT sprawl in order to have the freedom to shift resources away from operations in favor of fostering innovation and driving strategic initiatives that will grow the business. A fundamental element of this strategy is the deployment of a storage infrastructure that addresses the specific needs of virtual and cloud data centers with the fundamental flexibility to handle not only today's demands, but to serve as the foundation for a data center transformation with the power to poise that data center for what comes next. While there is indeed no way to "future-proof" the data center, there are certainly actions that can be taken today to maximize infrastructure efficiency and to build in the agility necessary to meet even rapidly changing business demands.

Deploying HP 3PAR Utility Storage in VMware vSphere environments helps remove the management, provisioning, and monitoring headaches associated with traditional storage platforms. Traditional SAN storage struggles to meet the performance and agility needs of a virtualized server environment, whereas HP 3PAR Utility Storage meets or even exceeds the heavy demands that server virtualization places on storage by using massive parallelization for exceptional performance and high availability features for superior resilience, thus enabling clients to consolidate with confidence.

Support for the latest version of VMware vSphere and integration with vSphere Storage APIs deliver enhanced performance, agility, and scalability in vSphere environments while HP 3PAR Thin Provisioning Software allows physical storage to be consumed only when required for actual written data, rather than when allocated. HP 3PAR Dynamic Optimization Software can be used to tailor storage performance without disruption to VMware vSphere and contribute new autonomic space reclamation functionality. These are among the unique advantages that make HP 3PAR Utility Storage the ideal foundation for building or expanding a virtualized server environment with VMware vSphere as part of a converged infrastructure to meet the needs of the Instant-on Enterprise.

Appendix: frequently asked questions

The following questions commonly arise when integrating HP 3PAR Storage Systems with VMware vSphere:

Q: Should a Thin Provisioned Virtual Volume (TPVV) be used for ESX datastores?

A: Thin provisioning makes sense as long as you are not going to fill up the volume right away. If the volume is going to be full in a short time, there is no benefit to be gained from thin provisioning. If VMware snapshots are being created and are not being cleaned up, the benefits of thin provisioning will be negated. To the ESX host, it makes no difference if the host sees a thin provisioned LUN or a traditional "thick" provisioned LUN.

Q: Is there any overhead to using Thin Provisioned Virtual Volume (TPVV)?

A: The additional overhead of TPVVs as compared to traditional volumes is negligible.

Q: If using a 2 TB Thin Provisioned Virtual Volume (TPVV), what will vSphere see?

A: vSphere will see a 2 TB LUN available for its use. Without the use of VAAI and/or T10, vSphere cannot determine if a volume is a thin provisioned virtual volume and a traditional virtual volume. With vSphere 4 or vSphere 5.0 on InForm OS 2.3.1 we recommend installing the VAAI plug-in. For vSphere 5.0 and InForm OS 3.1.1 or higher, no plug-in is required.

Q: What size Virtual Volume (VV) should be created?

A: The volume size is not as important to HP 3PAR Storage Systems, as VVs are widely striped across as many drives as possible within the array. If using HP 3PAR Thin Provisioning Software, actual storage capacity is only consumed upon write. However, a 2 TB VV will be able to accommodate many more VMs than a 500 GB VV. ESX 4.1 or later has improved the way it performs metadata locking, meaning that you can now create a VV that is as large as you are comfortable with. For ESX 4.0 and prior, in order to minimize the impact of SCSI reservations and keep the environment well balanced, it is best to create 500 GB volumes (thin provisioned or "thick" provisioned).

Q: How many VMs can be put on a single Virtual Volume (VV)?

A: It depends. There is no one answer that will work for every situation. A number of factors such as server hardware, number of CPUs, amount of memory, type of VMs, applications running in the VMs, etc. will determine how many VMs can be comfortably hosted on a LUN.

Q: What type of path policy (Fixed, MRU, Round Robin) should be used with HP 3PAR Storage Systems?

A: With ESX 3.5, use the default policy (Fixed). With ESX 4.0 and later, change the default policy to the Round Robin path policy. To do this, log in to the service console for each ESX 4.0 host and run the following command line:

```
esxcli nmp satp setdefaultpsp --satp VMW_SATP_DEFAULT_AA --psp VMW_PSP_RR
```

For vSphere 5 you can change the path selection algorithm using the Manage Paths dialog box either from the Datastores or Devices view or from the command line on each host:

```
esxcli storage nmp satp set --satp=VMW_SATP_DEFAULT_AA --default-psp=VMW_PSP_RR
```

Q: VMware supports thin VMs. Which thin provisioning should be used: VMware, HP 3PAR, or both?

A: VMware thin provisioning only applies to VMs at the VMFS level. It allows one to over-allocate VMs to maximize VMFS usage. If the goal is to reduce storage costs and maximize storage utilization, then use HP 3PAR Thin Provisioning Software to provision large VMFS volumes with minimal upfront storage costs. There are no additional storage savings to be realized by using VMware thin provisioning. VMware thin provisioning does consume some CPU cycles on the ESX host as it is performed at the software layer (as compared to HP 3PAR Thin Provisioning Software, which is performed on the array). It is perfectly fine to place VMware thin VMs on HP 3PAR Thin Provisioning Software volumes so long as you are prepared to manage thin provisioning at both the VMware level and the array level.

Q: Is the UNMAP primitive of VAAI supported across the board?

A: No, UNMAP is not supported prior to vSphere 5.0 with InForm OS 3.1.1.

Q: When do I need to install the VAAI plug-in?

A:

- VAAI is not supported on HP 3PAR InForm OS 2.3.1 MU1 or earlier.
- On vSphere 4.1, the VAAI 1.1 plug-in is required.
- When running vSphere 5.0 with HP 3PAR InForm OS 2.31, the 3PAR VAAI 2.2 plug-in is required.
- For vSphere 5.0 with HP 3PAR InForm OS 3.1.1 or higher, no plug-in is needed as all of the VAAI primitives are supported natively.

For more information

Visit www.hp.com/go/VMware and www.hp.com/go/3PAR.

To help us improve our documents, please provide feedback at http://h71019.www7.hp.com/ActiveAnswers/us/en/solutions/technical_tools_feedback.html.

4AA3-4023ENW, Created May 2011; Updated August 2011, Rev. 1