



HP P4000 LeftHand SAN Solutions with VMware vSphere Best Practices

Technical whitepaper

Table of contents

Executive summary	2
New Feature Challenge	3
Initial iSCSI setup of vSphere 5	4
Networking for the iSCSI Software Adapter	4
Two network ports	5
Four network ports	5
Six network ports	6
More than six network ports	7
10GbE options	7
Enabling the iSCSI software adapter	7
HBA connectivity and networking	7
Multi-pathing iSCSI	8
Connecting and using iSCSI volumes	8
Creating the first iSCSI volume on the SAN	8
Enabling Load Balancing for performance	9
Discovery of the first iSCSI volume	9
Discovering additional volumes	9
Disconnecting iSCSI volumes from ESXi hosts	9
Troubleshooting volume connectivity	9
Creating a new datastore on the iSCSI volume	10
Expanding an iSCSI LUN in P4000	10
Expanding a volume and extending the datastore	10
Expanding an iSCSI LUN in CMC	10
Expanding an iSCSI LUN in vSphere	10
Snapshots, Remote Copy, and HP SmartClone volumes	13
Resignaturing Snapshots	13
HP P4000 Application Managed Snapshots in ESX	14
HP P4000 snapshots of raw devices	15
HP P4000 snapshots VMFS datastores	15
HP P4000 Remote Copy volumes and SRM	15
HP P4000 Multi-Site SANs	16
HP P4000 SmartClone volumes	16
vMotion, Clustering, HA, DRS, and FT	16
Choosing datastores and volumes for virtual machines	17
VMware's vStorage API for Array Integration	17
vSphere 5 Storage Enhancements for P4000	19
VMFS-5	20
Storage Distributed Resource Scheduler (Storage DRS)	20
vSphere Storage APIs for Storage Awareness (VASA)	22
HP Insight Control Storage Module for vCenter	23
Storage I/O Control (Storage I/O Control)	24
Best practices	25
Frequently asked questions	26
For more information	28

Executive summary

This white paper provides detailed information on how to integrate VMware vSphere® 5.0 with HP P4000 LeftHand SAN Solutions. VMware vSphere is an industry leading virtualization platform and software cloud infrastructure enabling critical business applications to run with assurance and agility by virtualizing server resources. Complementing this technology, the HP P4000 LeftHand SANs address the storage demands and cost pressures associated with server virtualization, data growth, and business continuity. P4000 SANs scale capacity and performance linearly without incurring downtime, enabling it to meet the requirements of small pay-as-you-grow customers to the mission-critical applications of an enterprise.

This document presents configuration guidelines, best practices and answers to frequently asked questions that will help you accelerate a successful deployment of VMware vSphere 5.0 on HP P4000 SAN Solutions.

Target audience: VMware and SAN Administrators who are implementing VMware vSphere 5.0 on an HP P4000 LeftHand SAN Solution.

New Feature Challenges

With vSphere 5.0, VMware continues to increase both its infrastructure and application services.

- VMware ESXi™ converges on a thinner hypervisor architecture supporting a hardened 144MB disk footprint, streamlined patching and configuration model.
- New Virtual Machine (VM) virtual hardware version format v8 evolves additional graphic and USB 3.0 support. VMs can now have 32 virtual CPUs and 1TB of RAM.
- Storage Distributed Resource Scheduler (SDRS) aggregates storage into pools, simplifying scale management, ensuring best VM load balancing avoiding storage resource bottlenecks.
- Profile-Driven storage matches appropriate storage to given VMs ensuring correct placement.
- VMFS 5 enhances scalability and increased performance supporting larger 64TB datastores, a unified 1MB block size, and small file space utilization with 8KB sub-block sizes.
- Storage APIs – enhanced VAAI and new VASA storage awareness and discovery enhances performance and brings array information into DRS and Profile-Driven features.
- High Availability (HA) architecture updated to simplify deployment and increase scalability.

The HP P4000 SAN Solution increases and builds upon features found in SAN/iQ 9.0 with 9.5.

- VAAI feature extension support found in vSphere 4.1 improves storage efficiency and performance. Note that newer VAAI features extensions introduced in vSphere 5.0 are not supported in SAN/iQ 9.5.
- SAN/iQ 9.5 expands application integration of VMware integrated snapshots supporting application consistent, point in time copies for speedy and reliable VM recovery.
- SAN/iQ 9.5 features “zero to VSA” automated installation for easy deployment.
- SAN/iQ 9.5 supports new Storage Replication Adapter (SRA) for VMware vCenter™ Site Recovery Manager 5 (SRM).
- Network RAID 5 and 6 dynamic stripe size improves efficiency and is applied to the entire volume stack. Scheduled snapshots for space efficiency are no longer required.

Successfully addressing and deploying these new features across the solution is imperative if you wish to maximize the return on investment (ROI) for your HP P4000 SAN and VMware vSphere infrastructure while continuing to meet the dynamic needs of the business. To help you successfully deploy these new features and maximize the performance and scalability of your vSphere environment, this paper presents best practices for configuring, tuning, and deploying your environment.

Initial iSCSI setup of vSphere 5

Networking for the iSCSI Software Adapter

Before SAN connectivity via the vSphere 5 iSCSI software adapter is enabled, specific network configuration is recommended. To correctly enable connectivity, a VMkernel network is created with sufficient network access to the iSCSI SAN. As a best practice, use at least two 1GbE or 10GbE network adapters teamed together for performance and failover (Figure 1). Network teams might want to be created by assigning multiple network adapters to a standard switch under Configuration – Hardware Networking (Figure 2). For the iSCSI Software Adapter, Round Robin path selection on active-active I/O paths is optimal (Figure 3).

Figure 1 – Standard Virtual Switch – Two Physical Adapters

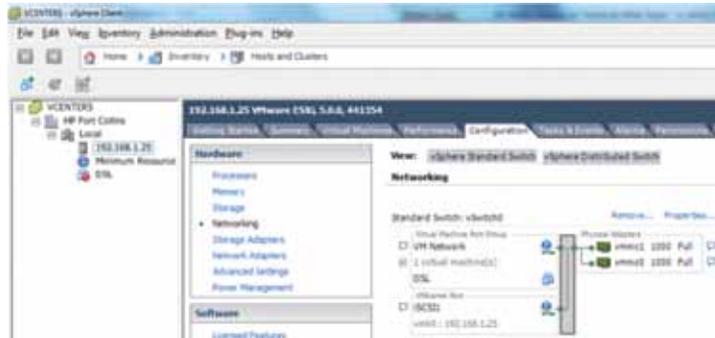
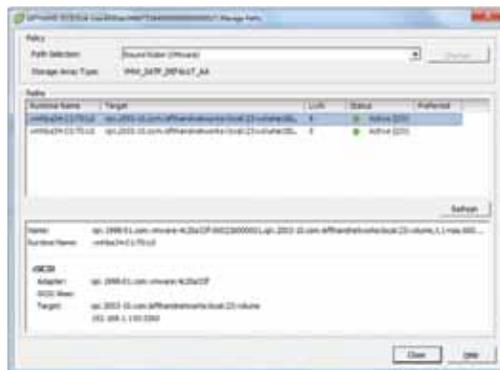


Figure 2 – iSCSI VMkernel Properties – NIC Teaming



Figure 3 – Manage Paths – Active-Active I/O



As a best performance practice, the VMkernel network for iSCSI might be separate from the host management and virtual networks used by virtual machines, however, high availability is usually preferred with minimal network hosts. If enough networks are available to provide network redundancy, it is preferred to separate vSphere vMotion and Fault Tolerance (FT) networks on unique virtual switches. These networks are placed on vSphere separate standard switches within each host or distributed switches, which handle networking traffic for all associated hosts on the datacenter. Port groups, defined with unique network labels, also limit network broadcast domains across multiple hosts.

In a vSphere distributed switch, with Network I/O Control enabled, distributed switch traffic is divided across predefined network resource pools: VMware Fault Tolerance traffic, iSCSI traffic, vMotion traffic, management traffic, NFS traffic, and virtual machine traffic. Note the similarities with standard switches and manual segregation, however, Network I/O Control enforces these rules across a datacenter, thereby easing policy enforcement and management.

The resources that you can deploy vary based upon the type and quantity of resources available. Larger performance scalable P4000 SANs will have different requirements than smaller P4000 clusters. HP BladeSystem c7000 solutions with local P4800 SANs have unique benefits with internal no latency iSCSI characteristics versus this same SAN connecting to external devices. VM network traffic requirements differ in every application. These basic guidelines help you follow best practice for connecting P4000 SANs, however understanding the approach to bandwidth tuning involves understanding the complete solution objectives and needs. HP BladeSystems with Flex-10 present tunable network interfaces as physical adapters within vSphere thereby offering competing bandwidth allocation and enforcement opportunities than with vSphere. Management and deployment options will vary. In larger multiple HP BladeSystem environments, HP CloudSystem Matrix offers additional orchestration opportunities to leverage network bandwidth automation. It is also noted that 10GbE interfaces provide increased bandwidth options over 1GbE.

Two network ports

VMware vSphere 5 hosts with only two Gigabit (1GbE) network ports are not ideal for iSCSI SANs connected by the iSCSI Software Adapter because they compete with other network traffic to ensure good iSCSI SAN performance. As shown in Figure 1, vSphere 5 hosts with only two Gigabit network ports are configured with a single standard virtual switch that comprises both Gigabit ports (teamed together) and contains:

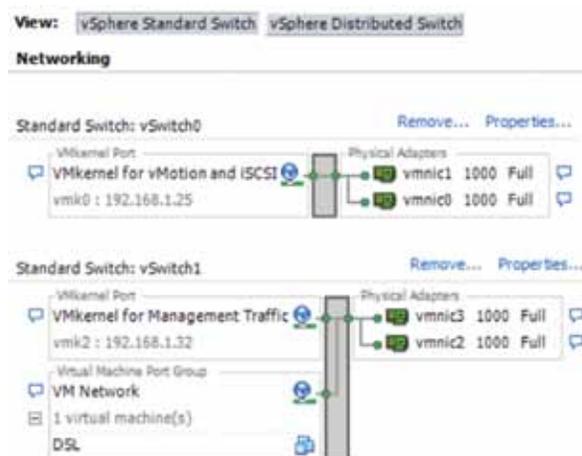
- A Virtual Machine Port Group network for virtual machines.
- A VMkernel Port network with vMotion, Management Traffic, and iSCSI Port Binding enabled.

Four network ports

VMware vSphere 5 hosts with four 1GbE network ports perform better if you separate virtual machine traffic from VMkernel vMotion, Management, and iSCSI traffic. As illustrated in Figure 4, vSphere 5 hosts with four Gigabit network ports are configured with two virtual switches, each comprising two Gigabit ports teamed together. If hosts contain multiple dual-port network adapters, use one port from each of the two dual-port adapters balancing high availability across both adapters. For example, if using two onboard Gigabit adapters and a dual-port Ethernet card, team together port 0 from the onboard adapter and port 0 from the Ethernet card, and then team together port 1 from the onboard adapter and port 1 from the Ethernet card. This configuration provides protection from some bus or network interface card failures.

- The first standard virtual switch should have:
 - A VMkernel Port network with iSCSI Port Binding, vMotion and Management Traffic enabled.
- The second standard virtual switch should have:
 - A Virtual Machine Port Group network for virtual machines.

Figure 4 – Two Standard Switches – Four Network Ports



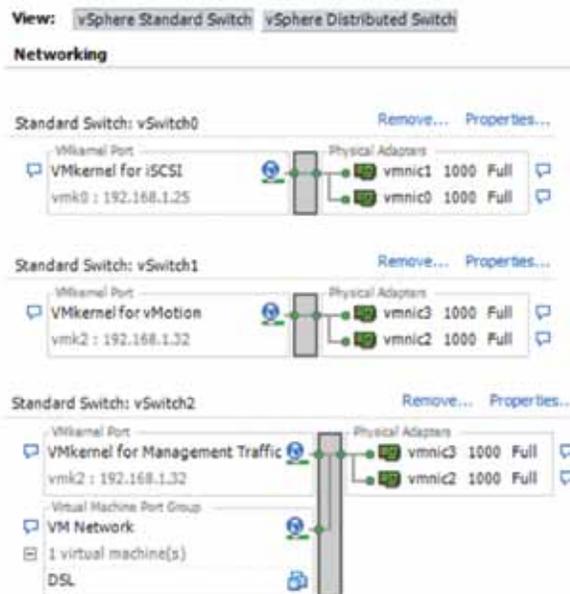
Six network ports

VMware vSphere 5 hosts with six Gigabit network ports are ideal for delivering performance with the iSCSI Software Adapter. The improvement over four ports is achieved by separating vMotion and Fault Tolerance traffic from iSCSI traffic so that they do not have to share bandwidth. Each traffic type performs better in this environment.

To configure vSphere 5 hosts with six Gigabit network ports, use three standard virtual switches, each comprising two Gigabit ports teamed together, as shown in Figure 5. If possible, one port from each of the separate Gigabit adapters are used in each team to prevent some bus or network interface card failures from affecting an entire standard virtual switch.

- The first standard virtual switch should have:
 - A VMkernel Port network with iSCSI Port Binding enabled.
- The second standard virtual switch should have:
 - A VMkernel Port network with vMotion and Management Traffic enabled.
- The third standard virtual switch should have:
 - A Virtual Machine Port Group network for virtual machines.

Figure 5 – Three Standard Switches – Six Network Ports



More than six network ports

If more than six network ports are available, add more ports to the iSCSI virtual switch to increase available bandwidth, or add more ports for any other network services desired.

10GbE options

Both vSphere 5 hosts and HP P4000 SANs support 10GbE interfaces and benefit from increased bandwidth. 10GbE network ports offer additional bandwidth over 1GbE interfaces. With P4000 SANs supporting 10GbE interface upgrades, throughput has been seen to increase 2–3x depending upon access patterns, SAN cluster size, data path and bandwidth availability to the hosts. Note that in 10GbE networks, full mesh backplanes and latency options vary and impact performance. Flexible network infrastructure tailors variable bandwidth options to need. Generally, a SAN's performance is limited by the number of disk spindles dictating IOP characteristics capabilities of the physical storage. In scalable storage architectures, this is distributed across the P4000 SAN cluster. Note that in all previous examples, 1GbE could as easily been replaced by faster network bandwidth options and will impact best practice choices. It is best practice to balance networking technologies end to end for overall solution performance characteristics. For example, HP P4000 10GbE solutions are limited with 1GbE switching infrastructure and therefore a bottleneck.

Enabling the iSCSI software adapter

If a hardware iSCSI host adapter is not used, enable the vSphere 5 iSCSI Software Adapter on each ESXi host. The iSCSI Software Adapter is managed from each host's Storage Adapters list. Here are some guidelines:

- Enable the iSCSI Software Adapter on each ESXi host.
- Copy or write down the iSCSI qualified name (IQN) that identifies each vSphere 5 host; it will be needed for authentication later in defining volume access configuring the P4000 SAN.

HBA connectivity and networking

An iSCSI host bus adapter uses dedicated hardware to connect to a 1GbE or 10GbE network mitigating the overhead of iSCSI and TCP processing of Ethernet interrupts thereby improving the performance of servers. iSCSI HBAs can include PCI option ROM to allow booting from an iSCSI target. SAN connectivity via iSCSI HBAs therefore enables both offloading of the iSCSI processing from the vSphere 5 server and another method of booting ESXi from the iSCSI SAN. HBAs do not require licensing or special networking within vSphere 5 hosts as they provide a dedicated network connection for iSCSI traffic only and are seen within the host similarly to FC host adapters. The physical network for HBAs may be a 1GbE or 10GbE network dedicated to the SAN, just as it is for iSCSI Software Adapters. As best practice, use two HBA initiators (a dual port or two single ports), each configured with a path to all iSCSI targets for failover. Configuring multiple HBA initiators to connect to the same target requires authentication for each initiator's IQN to be configured on the SAN. This is configured in the P4000 Centralized Management Console (CMC) software as two servers (one for each HBA initiator), each with authentication permissions to the same volumes on the SAN. Note that a group of hosts in the same cluster will also require authentication for each volume. Please go to HP's Single Point of Connectivity Knowledge base for tested and supported boot from SAN iSCSI HBAs with P4000 SANs. (<http://www.hp.com/storage/spock>)

Note that most 1GbE enterprise network adapters support CRC checksum offload efficiencies and make server performance improvements negligible. However, 10GbE has shown the importance of TCP processing offloads in adapters to lessen server impact. Some of the 10GbE adapters have also introduced protocol convergence of 10GbE Ethernet, offloading TCP and iSCSI processing, BIOS boot support, and iSCSI protocol support (some even with multiple protocols) providing new compelling reasons in leveraging these advantages.

Multi-pathing iSCSI

Native Round Robin iSCSI multi-pathing in vSphere 5 provides superior performance by leveraging multiple paths to the SAN. Each path to the P4000 SAN will establish a unique iSCSI session and may be distributed over different storage nodes in the P4000 cluster. Since P4000 SANs scale leveraging multiple storage nodes' compute, disk and network resources, iSCSI multi-pathing's advantage is seen in multi-threading data approach, up to the maximum throughput sizing capability of a P4000 SAN. Configuring iSCSI multi-pathing requires at least two network ports on the virtual switch. The following steps are performed on each ESXi host individually.

- From the host's Configuration – Networking, select Properties of the standard virtual switch. Select the Network Adapters tab and add the appropriate number of adapters.
- Select the Ports tab and Edit or Add the VMkernel. On the Properties, NIC Teaming tab select Load Balancing and other appropriate policy exceptions.
- From the host's Configuration – Storage, highlight the iSCSI datastore. Right click and select Properties and then select the Manage Paths option. Policy Path Selection is Round Robin (VMware). Select Change option. Note the valid multiple paths listed and the Status of Active (I/O) for each path
- From the host's Configuration – Storage, highlight the iSCSI datastore. Note the Datastore Details. Total Paths should accurately reflect the number of valid paths to the storage volume.

Once configured correctly, perform a Rescan All, scanning for new storage devices and scan for new VMFS volumes. An iSCSI session is connected for each path bound to the iSCSI Software Adapter. The example referred to in Figure 3 gives each iSCSI LUN/datastore two iSCSI paths using two separate physical network adapters.

To achieve load balancing across the multiple paths, datastores are configured with a path selection policy of round robin. This can be done manually for each datastore in the vSphere client. A reboot of the host is required for storage devices currently managed by the host. However, new storage devices do not require a host reboot.

The round robin policy ensures that the host uses an automatic path selection algorithm rotating through all available paths. This implements load balancing across all of the available physical paths. Load balancing spreads host I/O requests across all available host paths with the goal of optimizing performance storage throughput.

It is important to note that native vSphere 5 multi-pathing should not be used with HP P4000 Multi-Site SAN configurations that utilize more than one subnet and iSCSI Virtual IP.

Connecting and using iSCSI volumes

Creating the first iSCSI volume on the SAN

Before a vSphere 5 host can access a new iSCSI datastore, that volume must be created on the P4000 SAN and must have authentication configured to allow the host to access the volume. All hosts within a datacenter cluster must also have access to that same volume in order to support vMotion, High Availability and Fault Tolerance. Use the P4000 Centralized Management Console (CMC), as shown in Figure 6,

or use the command-line interface (CLI) to create a new volume. Then, using the CMC, create a new server to represent the ESXi host's iSCSI Software Adapter, using the IQN(s) copied from the vSphere host's initiator(s), or use the Challenge Handshake Authentication Protocol (CHAP) as required. SAN volumes are created and then assigned to vSphere 5 hosts within the CMC or CLI.



Figure 6 – New Volume

Enabling Load Balancing for performance

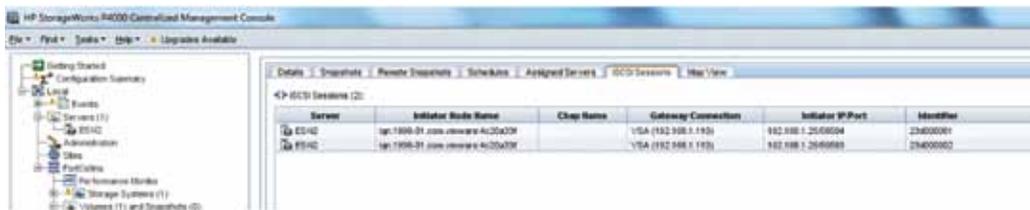
Virtual IP load balancing (VIPLB), or load balancing, is a setting on each defined server in the CMC that allows for the distribution of iSCSI sessions throughout the P4000 cluster of storage nodes in order to maximize performance and bandwidth utilization. Most iSCSI initiators, including the vSphere 5 Storage iSCSI Software Adapter and HBAs, support this feature. Virtual IP load balancing is enabled by default in the CMC. All vSphere 5 iSCSI initiators connected to the P4000 SAN should have this feature enabled.

Discovery of the first iSCSI volume

To discover volumes, the iSCSI Software Adapter of each ESXi host must have added the Virtual IP (VIP) address of the P4000 cluster which contains the volume in its Dynamic Discovery Send Targets list. New targets can then be discovered by simply choosing the Rescan All LUNS option from the vSphere Client. Alternatively, Static Discovery may also be manually entered for each target, however this is generally not considered best management practice.

The iSCSI session status of an ESXi host can be viewed in the CMC by selecting volume and selecting the iSCSI Sessions tab. Volumes that are connected show an IP address in the Gateway Connection field on the server's Volumes & Snapshots tab (Figure 7). If an IP address is not listed, look for an incorrect configuration on the ESXi host(s) or on the server in the CMC.

Figure 7 – HP P4000 Centralized Management Console – iSCSI Sessions



Server	Initiator Node Name	Chap Name	Gateway Connection	Initiator IP/Port	Identifier
ESXi01	ip:1928-01.com:esware:AC2002F		VSA (192.168.1.192)	192.168.1.254/5004	23400001
ESXi02	ip:1928-01.com:esware:AC2002F		VSA (192.168.1.192)	192.168.1.254/5004	23400002

Discovering additional volumes

Additional volumes may be dynamically discovered, if so configured, by simply choosing the Rescan All LUNS option from the vSphere Client. Static Discovery may also manually add each target for the iSCSI Software Adapter or HBAs.

Disconnecting iSCSI volumes from ESXi hosts

The vSphere 5 iSCSI Software Adapter does not have a session disconnect feature. Before an iSCSI volume can be disconnected, it must first be unassigned from the ESXi host within the CMC otherwise the next Rescan All LUNS will automatically connect to it again if Dynamic Discovery is configured. Unassigned volumes are not forcibly disconnected by the CMC. Instead the host simply will not be allowed to log in again. Rebooting an ESXi host will clear all unassigned volumes iSCSI sessions. Individual iSCSI sessions can be reset forcefully from the P4000 CLI with the reset session command. Before an iSCSI session is forcibly reset from the CLI, all VMs accessing the volume through that session should be powered off or, using vMotion, reassigned to another host or moved to a different datastore that will continue to be accessible.

Troubleshooting volume connectivity

If new volumes are not showing up as expected, try the following troubleshooting steps:

- Ping the virtual IP address from the iSCSI initiator to ensure basic network connectivity.
 - HBAs typically have their own ping utilities inside the BIOS of the HBA, so use the appropriate utility for your HBA.
 - ESXi has a network troubleshooting utility that can be used from the ESXi Management Console (or remotely through SSH) to attempt a ping to the SAN.
- Double-check all IQNs or Challenge Authentication Protocol (CHAP) entries. For iSCSI authentication to

work correctly, these must be exact. Simplifying the IQN to something shorter than the default can help with troubleshooting.

- Make sure that all hosts on the SAN have load balancing enabled in the CMC. If there is a mix—if, for example, some have it enabled and some don't—then hosts that don't have load balancing enabled may not connect to their volumes.

Creating a new datastore on the iSCSI volume

After the vSphere 5 host has an iSCSI P4000 SAN volume connected, it can be formatted as a new VMware vStorage Virtual Machine File System (VMFS) datastore or mounted as a raw device mapping (RDM) directly to virtual machines. New datastores are formatted from within the VMware vSphere client in the new VMFS-5 format. An example of the process for creating a new datastore is shown in Figure 8.

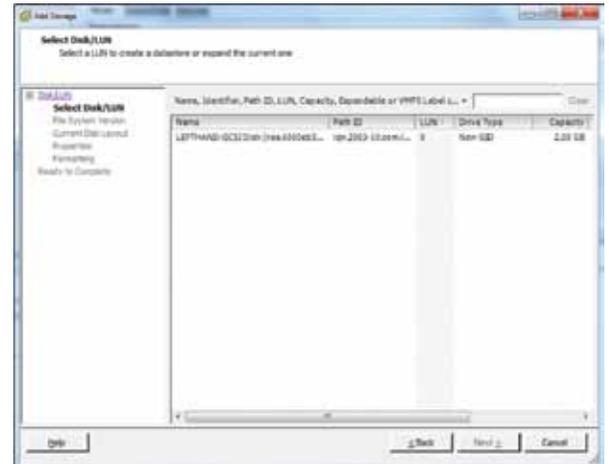


Figure 8 – Creating a new datastore – Select the iSCSI Disk. Note the HP P4000 SAN volume will be identified as a LeftHand iSCSI Disk.

Expanding an iSCSI LUN in P4000

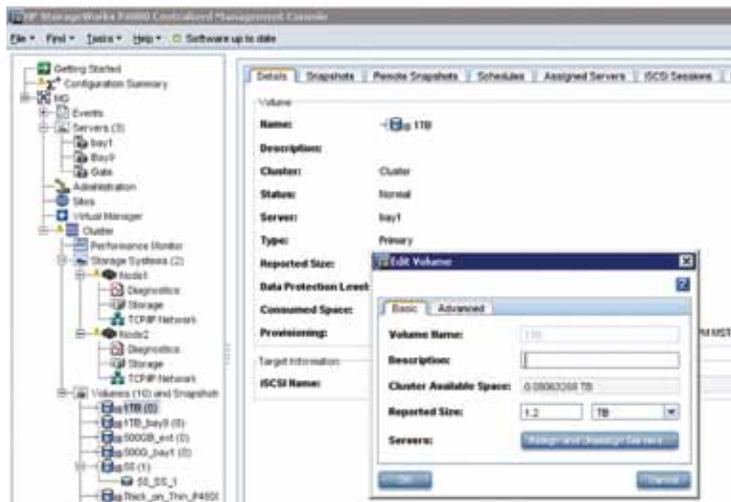
Expanding a volume and extending the datastore

Both P4000 SAN volumes and vSphere 5 datastores can be expanded dynamically. If space is running low on a datastore, you must first expand the volume on the P4000 SAN. To increase the size of the P4000 SAN volume, simply edit the volume in the CMC or CLI and increase its size. The P4000 SAN immediately changes the LUN size and lays out data across the cluster accordingly, without affecting the volumes availability. Once the volume has been expanded, you can increase the size of the VMFS datastore from within the vSphere client. Increasing the size of the datastore is preferred over adding extents.

Expanding an iSCSI LUN in the CMC

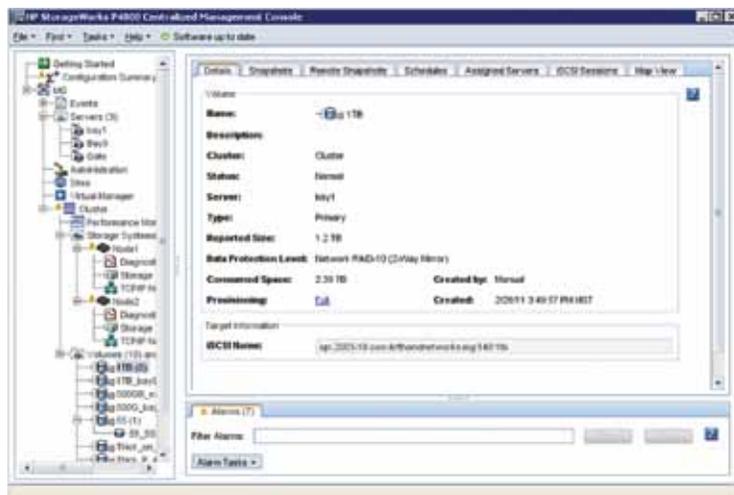
1. Go to the CMC.
2. Highlight the volume to be expanded.
3. Right click on the volume and click on “Edit volume”
4. A new window pops up showing the LUN size as 1TB.
5. Change the reported size to the desired size and click on ok (Figure 9). In this example, it is changed to 1.2 TB.

Figure 9 – Edit the Volume Size in the P4000 CMC



6. The reported size shows the expansion size immediately (Figure 10).

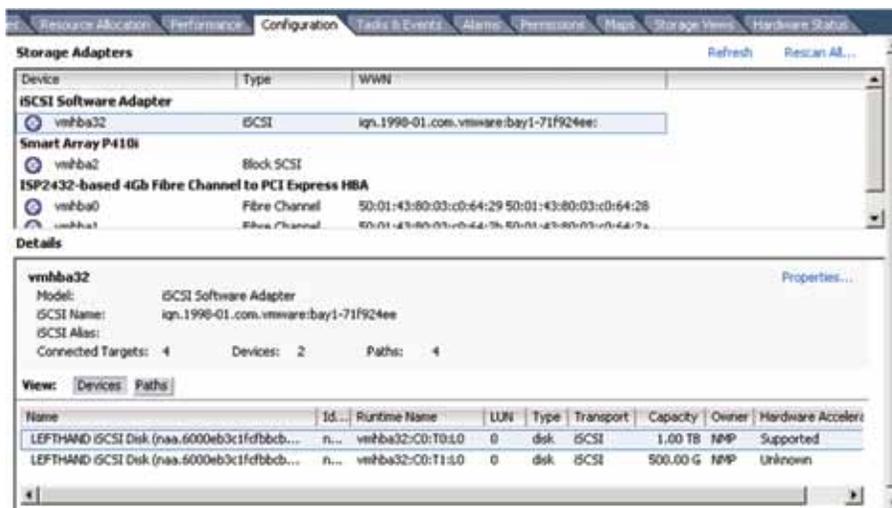
Figure 10 – The Reported Size has now changed and is immediately available



Expanding an iSCSI LUN in vSphere

1. Highlight the ESXi host where the iSCSI LUN is presented.
2. Click on the Configuration tab.
3. Highlight the iSCSI adaptor. The LUN is still showing 1TB (Figure 11).

Figure 11 – vSphere Client shows iSCSI volume still as 1TB prior to Rescan All



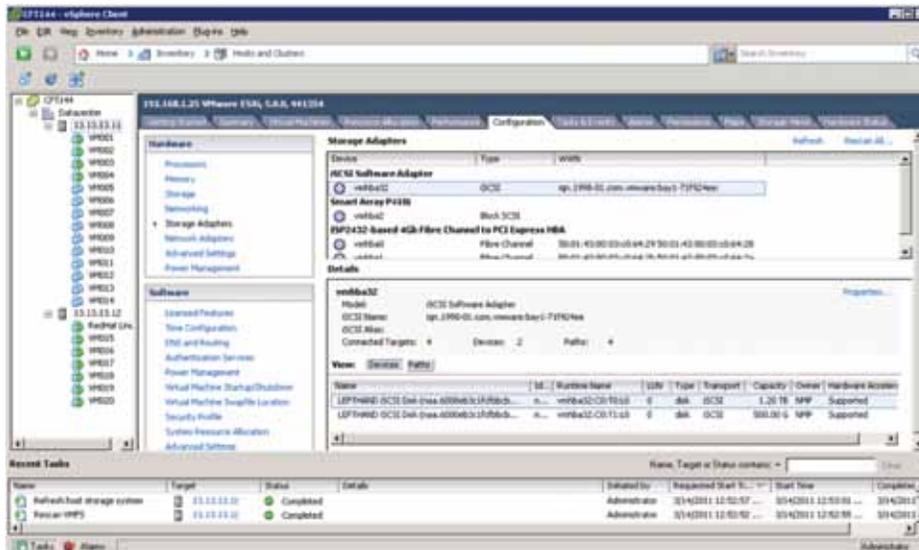
- Click on “Rescan All” and pick both “Scan for Storage devices” and Scan for New Volumes (Figure 12).

Figure 12 – Rescan for the change in the iSCSI LUN



- The new size is reported immediately in the vSphere Client (Figure 13).

Figure 13 – vSphere Client now reflects the change made to the iSCSI volume



Snapshots, Remote Copy, and HP SmartClone volumes

Resignaturing Snapshots

In order for vSphere 5 hosts to utilize SAN-based snapshots or clones, a snapshot or a clone is taken of a volume and presented to the vSphere hosts. The process of presenting the snapshot or cloned volume is similar to adding a new iSCSI volume datastore. vSphere 5 supports 3 options to mount a new SAN based snapshot or cloned volume that has VMFS data: Keep the existing signature, assign a new signature or format the disk (Figure 14).

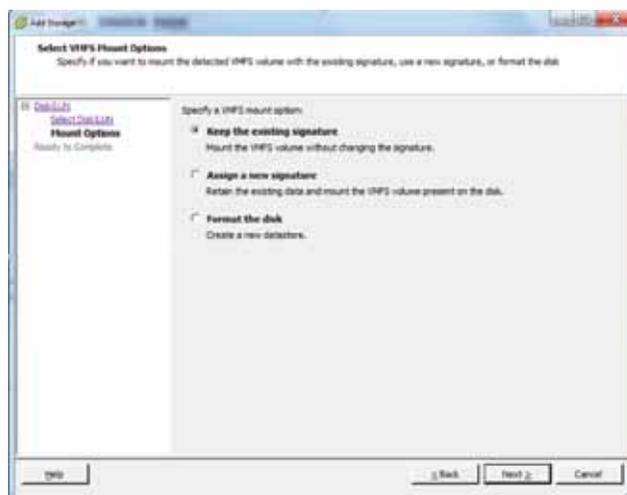
- Keep the existing signature allows you to mount the VMFS volume without changing the signature.
- Assign a new signature allows you to retain the existing data and mount the VMFS volume present on the disk. The signature will change and uniquely identify the volume without impacting data within the VMFS file system.
- Format this disk will destroy all data on the volume.

In keeping the existing signature during a mount may generate an error, HostStorageSystem.ResolveMultipleUnresolvedVmfsVolumes. In this example, a VMFS datastore can be mounted only if it does not conflict with an already mounted VMFS datastore that has the same UUID (VMFS signature). If the original LUN in a vSphere datacenter that contains the original VMFS datastore is currently mounted at the same time with its snapshot LUN, force mounting the VMFS datastore from the snapshot LUN on this same ESX host is not allowed. In this case, the only way to mount that VMFS datastore from the snapshot LUN is to assign a new signature and choose a new label for that datastore.

In assigning a new signature, references to the existing signature from virtual machine configuration files will need to be updated. Note in the Ready to Complete summary, the Original UUID signature is provided along with Assign a new UUID. The event task is resignature unresolved VMFS volume.

With a successful mount, the datastore mounts with a generate Identification "snap-XXXXXXXX-<Original label>". On volume properties, select Rename to provide a more suitable volume Identification name.

Figure 14 – Add Storage when mounting a volume clone or snapshot



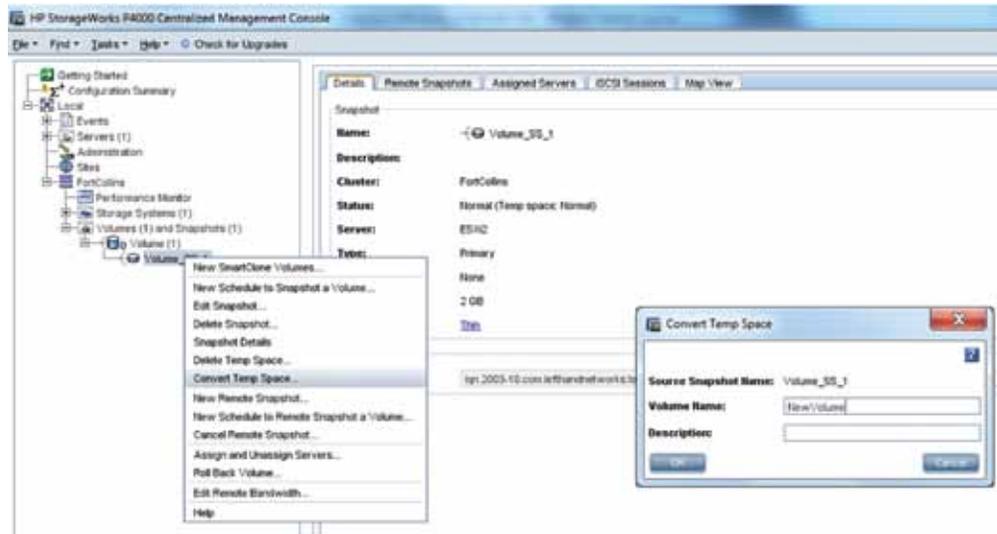
Note that P4000 snapshots are read only persistent points in time representing data on the source volume at the time it was taken. As an extension to this feature, the P4000 SAN allows read-only snapshots to be mounted and written to in a temporary area. This data area is called a temporary space and contains delta writes from the original snapshot data. By assigning a new signature in vSphere 5, we are writing data to this temporary space. From the P4000 SAN's perspective, temporary space will be lost if the SAN is rebooted and retains the same high availability characteristics of the source volume. If data written to the temporary space along with the snapshot needs to be preserved in a new volume, the CMC may be used to create a new volume copy with this temporary delta data merged with the source snapshot. In order to preserve data on the volume prior to conversion, VMs or virtual disks depending on this volume must be stopped, suspended or disconnected.

Alternatively, selecting the Delete Temp Space option will forcibly remove this data from the P4000 SAN. Once temp space is deleted from the SAN, it may not be recovered. It is a best practice to unmount the snapshot prior to deleting temp space. If you wish to remount the snapshot again, you can issue the Rescan All command in the vSphere Client prior to mounting the snapshot again.

The CMC supports the Convert Temp Space option to add this data to a new volume on the P4000 SAN. Once the new volume is created with a new name, it may be presented as a new datastore with the same Resignaturing Snapshot rules above.

Snapshots offer excellent ways at promoting data recovery.

Figure 15 – P4000 Centralized Management Console – Convert Temp Space



HP P4000 Application Managed Snapshots in ESX

New to HP P4000 SANs in SAN/iQ 9.5 expands Application Integration functionality to include VMware integrated snapshots. This new Application Aware Snapshot Manager enables VMware volume SAN based snapshots, thereby creating application consistent, point in time copies for reliable recovery of VM states. Without this integration option, snapshots previously were as on volume from the P4000 SAN state. In flight cached data may not have been fully quiesced creating an application recovery point. Now, full data recovery does not depend upon a paused or stopped VM state and flushed data to the SAN.

The configuration of the Application Aware Snapshot Manager is performed within the HP P4000 CMC. Users configure the IP address of the Virtual Center server. The Centralized Management Console communicates with the vSphere Virtual Center Server(s) during the snapshot process. Virtual Center quiesces VMs and takes a VMware snapshot. VMware tools within each guest VM quiesces applications. Once quiesced, the HP P4000 SAN performs a hardware based snapshot which contains vSphere VM snapshots when examined thru snapshot manager and represented to vSphere hosts.

HP P4000 snapshots of Raw Devices

P4000 SAN snapshots of vSphere 5 RDMs are supported in exactly the same way as for physical servers, either booting from the SAN or accessing LUNs on the SAN. Detailed information about P4000 snapshots and how they work for specific applications is available on the P4000 website: <http://www.hp.com/go/p4000>.

In vSphere 5, two compatibility modes for RDMs are available. Virtual compatibility mode allows an RDM to act exactly like a virtual disk file, including the use of VMware based snapshots. Physical compatibility mode allows direct access of the LUN device for applications that need lower level direct raw device control.

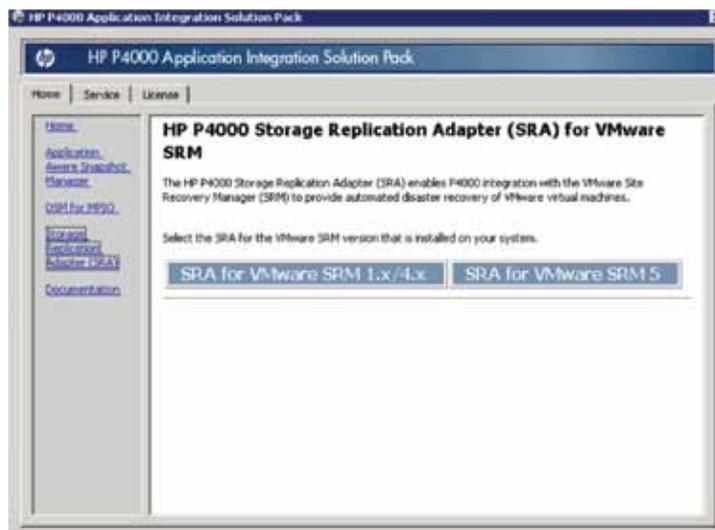
RDMs offer some benefits including:

- Dynamic Name Resolution – the identification associates each RDM with the device regardless of physical server changes, adapter hardware, path changes or device relocation.
- Distributed File Locking – RDMs are available as shared raw LUNs without losing data across multiple vSphere hosts accessing the same LUN.
- Snapshots – VM snapshots are possible on RDM LUNs in virtual compatibility mode. Combined with SAN/iQ 9.5 Application Managed snapshots, raw devices may be quiesced in context VMs. Note that this is NOT supported with physical compatibility mode RDMs.

HP P4000 snapshots on VMFS datastores

P4000 SAN snapshots are very useful in a vSphere 5 environment. All virtual machines stored on a single volume can be snapshot and rolled back to at any time. Moreover, P4000 snapshots can be mounted to any vSphere 5 host without interrupting access to their source volume. This may be used for data mining or testing real world data without affecting live corporate data.

Figure 16 – HP P4000 Application Integration Solution Pack – SAN/iQ 9.5 required for vSphere 5



HP P4000 Remote Copy volumes and SRM

Remote Copy replicates P4000 snapshots over WAN links to remote sites for disaster recovery or backup. vSphere 5 environments can be protected by Remote Copy volumes on a scheduled basis and automated by VMware Site Recovery Manager for a simple and complete disaster recovery solution. HP provides a Storage Replication Adapter (SRA) for VMware Site Recovery Manager (SRM) to integrate Remote Copy volumes seamlessly with a vSphere 5 environment. For more information on Remote Copy volumes, review the Remote Copy User Manual installed with the P4000 CMC.

vSphere 5 and VMware Site Recovery Manager 5 requires SAN/iQ 9.5's Application Integration Solution Pack. An option for SRA for VMware SRM 5 must be selected for vSphere 5 and Site Recovery Manager 5 support. Previous version supported SRM 4.1 in a single installer option. SAN/iQ 9.5 supports SRM 1.x/4.x or SRM 5 (Figure 16). A best practice is to ensure that SAN/iQ hardware and software are always up to date. Note that automated failback is now supported as part of the SRM 5 features.

HP P4000 Multi-Site SANs

HP P4000 Multi-Site SANs enable vSphere 5 clusters to be stretched across locations to provide multi-site VM vMotion, VM High Availability (HA), Distributed Resource Scheduler (DRS) including new extensions to storage, and Fault Tolerance (FT). Multi-Site SAN configurations use synchronous replication in the underlying SAN to create a single SAN that spans both locations. Ensure that Multi-Site SAN has adequate bandwidth and latency for storage data and the FT required logging as demonstrated in the P4000 Multi-Site User Guide. This allows vSphere 5 clusters to act exactly the same way they do when physically located in a single location. HP P4000 Multi-Site SANs require virtual IP addresses (VIPs) in unique subnets. When connecting ESXi hosts to a Multi-Site SAN each of the virtual IP addresses (VIPs) of the SAN from each site should be listed in the Dynamic Discovery Send Targets list of the iSCSI Software Adapter. Path selection policy for Multi-Site SAN volumes should be set to fixed (default).

HP P4000 SmartClone volumes

P4000 SmartClone volumes may also be used in a vSphere 5 environment (Figure 17). Using HP P4000 SmartClone Technology, all the virtual machines stored on a single volume can be cloned instantly without replicating data. SmartClone volumes consume space for data changes only from the time the SmartClone volume was created. SmartClone volumes are the best way to deploy small quantities of cloned golden image VMs or virtual desktops. Manageability and deployment of VDI environments are best to leverage VMware View implementations with Linked Clones. Please refer to the HP VDI Reference Architectures.

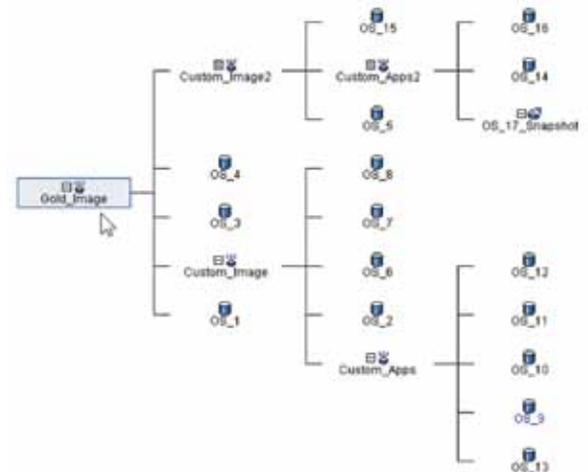


Figure 17 – Map View of HP P4000 SmartClone volumes in the CMC

SmartClone volumes can be used seamlessly with any other P4000 software feature, including snapshots or Remote Copy. SmartClone volumes are also very useful for performing tests on virtual machines by quickly reproducing them, without actually using space on the SAN to copy them. Unlike snapshots, delta data is persisted between each source snapshot. Please see the section on Resignaturing snapshots. Note that every write (including a delete) is a delta block persisted in the SmartClone delta space. If long term space efficiency is required with SmartClone volumes, minimizing writes to the SmartClone datastores include avoiding defragmentation within the guest VM. Successful approaches have also included separation of User and Application data including file redirection. Space reclamation can only be performed with a SmartClone volume by creating a new one retaining the original small source reference with minimal delta data. By removing OS dependence in separation from the User and Applications, periodic SmartClone deletion and re-creation ensures that data delta growth is minimized. Without these approaches, SmartClone volumes may eventually occupy an entire SmartClone's volume space as delta change data. SmartClone's initial value is in immediate cloning of golden image volumes. Efficient space utilization objectives need to understand use and mitigation approaches to maximize success.

vMotion, Clustering, HA, DRS, and FT

The advanced vSphere 5 features of vMotion, clustering, HA, DRS, and FT, all require multiple vSphere 5 hosts to have simultaneous access to volumes. To enable this on a P4000 volume, assign multiple servers (one for each vSphere 5 server iSCSI Software Adapter) to the same volume. To use multiple servers, simply create one server for each vSphere 5 host that connects to the SAN, using the steps outlined earlier

in the section Creating the first iSCSI volume on the SAN. When multiple servers are assigned to the same volume in the CMC, a warning indicates that this configuration is intended for clustered servers or clustered file systems such as ESXi (Figure 18).

P4000 SAN server iSCSI IQN authentication maps one vSphere 5 iSCSI Software Adapter to one or more volumes. P4000 SAN server CHAP authentication maps one or more vSphere 5 iSCSI Software Adapters to one or more volumes. CHAP authentication can also be used for increased security, either 1-way or 2-way. Each vSphere 5 host must be configured to use the correct CHAP credentials. New volumes on each vSphere 5 host must be discovered Dynamic or Static Discovery as described in the section Discovery of the first iSCSI volume.

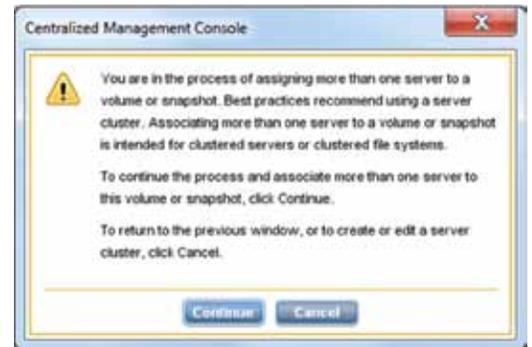


Figure 18 – Assigning multiple servers to a single volume

Choosing datastores and volumes for virtual machines

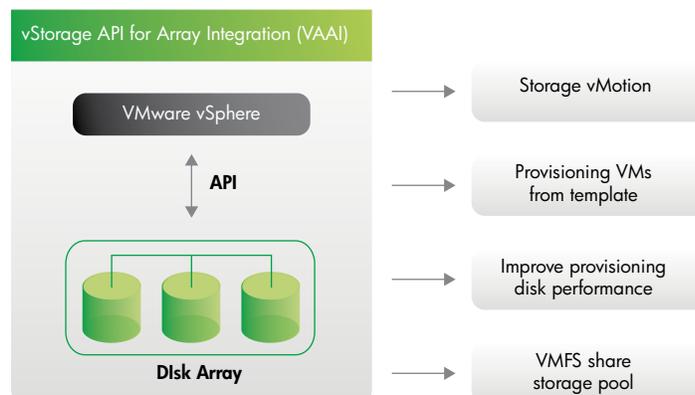
More than one virtual machine can—and generally should—be stored on each datastore or volume. The choice of where to put virtual machines should be driven by space considerations, current datastore load, the relationships of the VMs to each other, and performance. In general, virtual machines that do not have a relationship to one another should not be mixed on the same P4000 volume; however, there are always tradeoffs. Consolidating VMs becomes mandatory as their numbers increase and larger configurations are demanded. P4000 features such as Snapshots, Remote Copy, and SmartClone volumes are very useful with virtual machines, but will always affect all virtual machines on the same volume simultaneously. If virtual machines that have no relationship are mixed on a single volume, those virtual machines will have to be snapshot, rolled back, remotely copied, and cloned together.

Performance of virtual machines could also be affected if too many virtual machines are located on a single volume. The more virtual machines on a volume, the more I/O and contention there is for that volume, particularly when multiple ESXi hosts run VMs sourced on the same datastore. Refer to the section on Storage Distributed Resource Scheduler (DRS) for the new vSphere 5 feature.

VMware’s vSphere Storage APIs for Array Integration

The VMware vSphere Storage APIs for Array Integration (VAAI) enables the P4000 SAN, running SAN/iQ 9.5, to provide offloading from the vSphere 5 hosts for specific storage operations. These functions increase the performance and efficiencies for each vSphere host and offloads vSphere host storage processing to the P4000 SAN, where it is most efficient. Originally introduced in vSphere 4.1, vSphere 5.0 further extends these features. P4000 SANs only support the original functional primitives at this time in SAN/iQ 9.0 and 9.5.

Figure 19 – vStorage API for Array Integration (VAAI) Framework



Currently supported VAAI offloading capabilities:

- Full copy – Enables the P4000 SAN to make full copies of data within the array. Previously, data movement would traverse from the SAN, occupying bandwidth to the vSphere host, utilize CPU resources processing the data movement, and again traversing back to the SAN.
- Block zeroing – Enables the P4000 SAN to zero out large numbers of blocks with a single command. Previously repetitive zero's would traverse from the vSphere host to the SAN, occupying bandwidth and utilizing CPU resources processing the data movement. This feature only speeds up the process of provisioning of VMs created with the eager-zeroed thick disks option which is not the default.
- Hardware-assisted locking – Provides a new efficient mechanism to protect metadata for VMFS cluster file systems, improves scalability of large vSphere host farms sharing a datastore, and is more efficient than the previous SCSI reserve and release mechanism. iSCSI volumes are shared across multiple vSphere hosts in a cluster. Improved efficiency in sharing volumes across the cluster directly impacts the number of VMs a volume can hold, increases the size of the volume to support the increased VM count, and increases the number of hosts within a cluster sharing access to the same LUN.

Figure 20 – VAAI performance improvements with HP P4000 SAN

- Full copy clone Microsoft® Windows® 2008 R2 40GB VM
– **7x faster and 94% less load**
- Block zero a 256GB vmdk
– **21x faster and 92% less load**
- Hardware Assisted Locking
– **6x more VMs per volume/LUN**

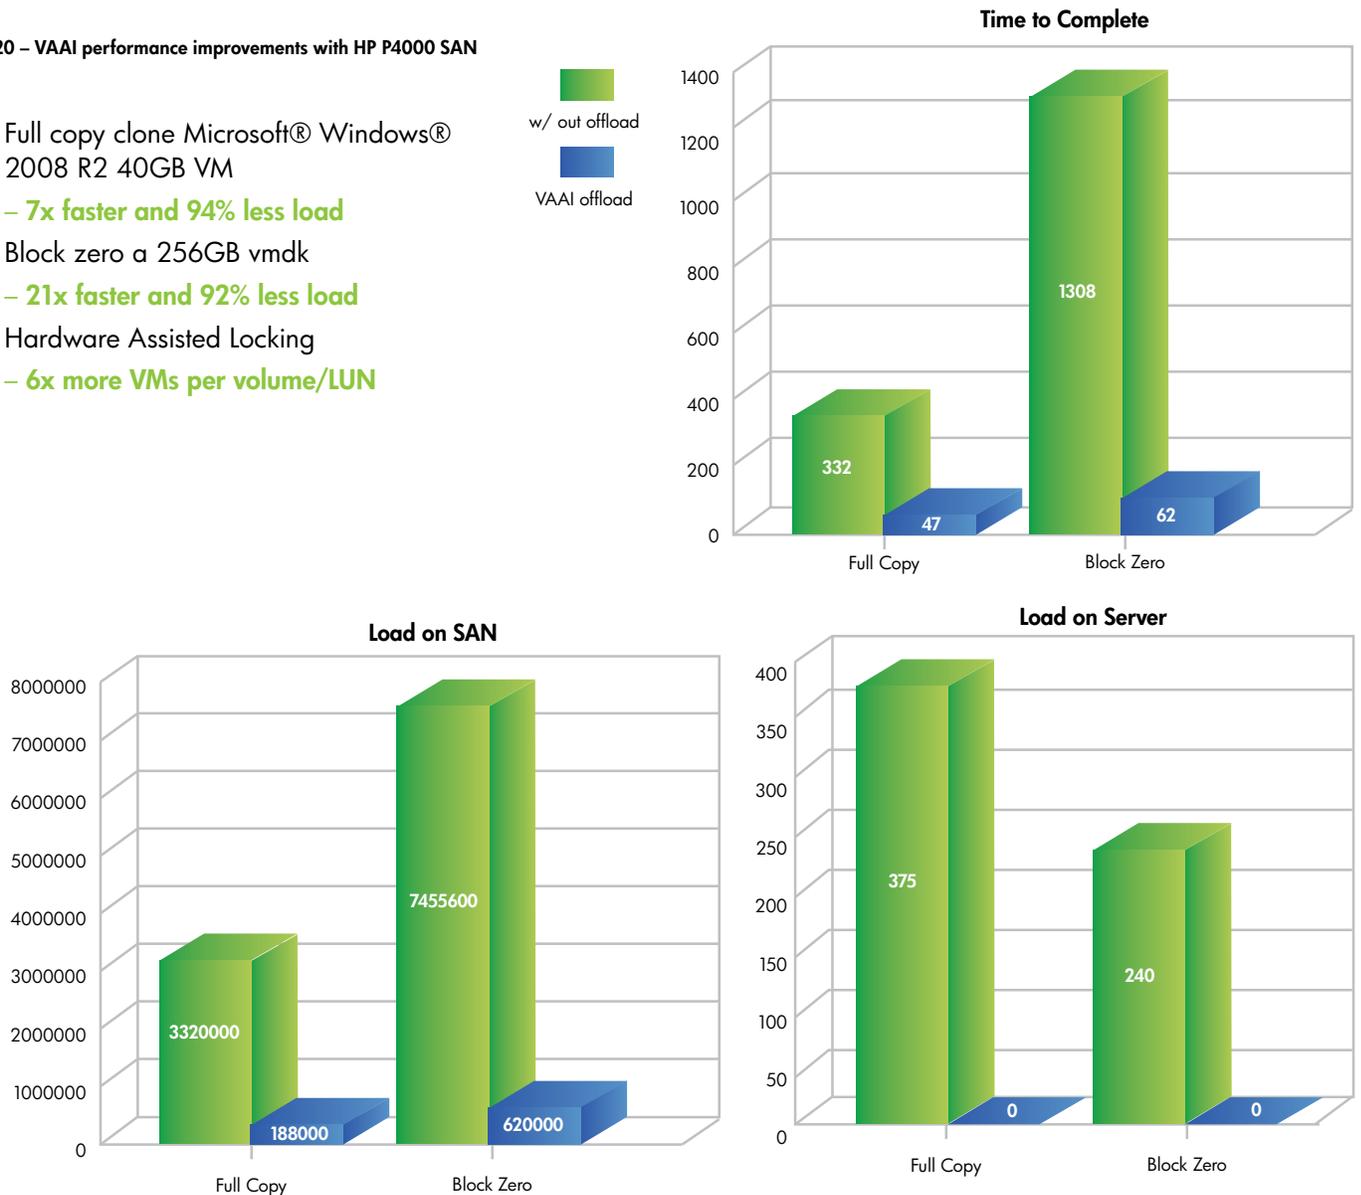


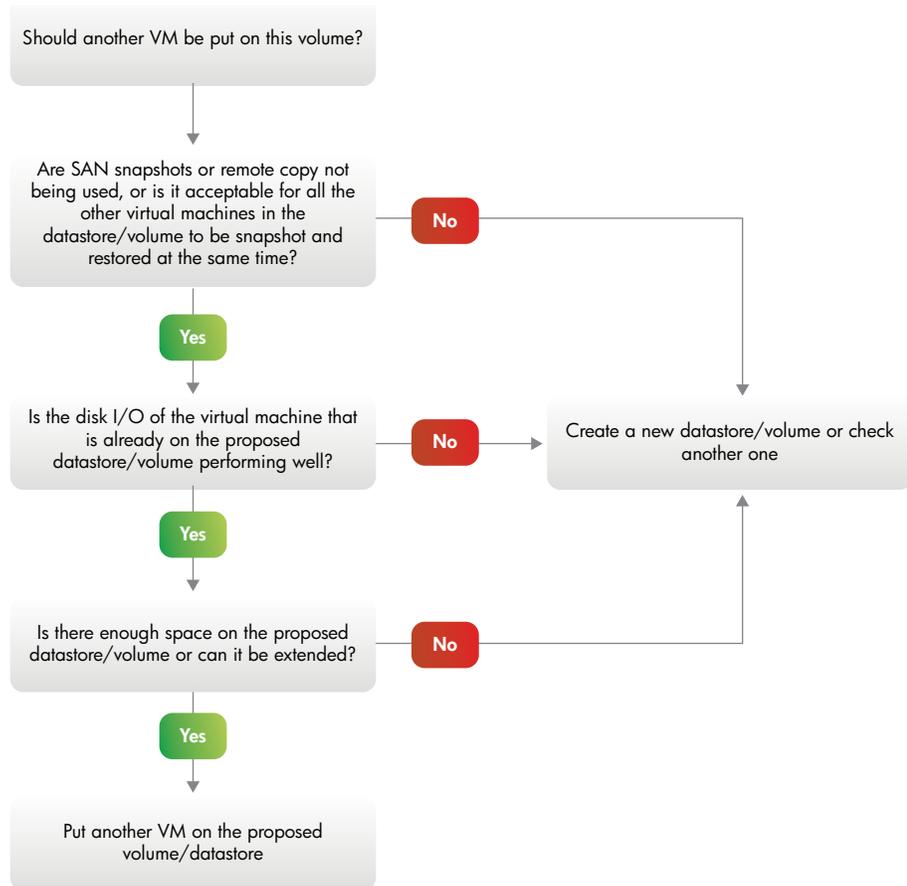
Figure 20 shows the level of improvements attributed to VMware's vSphere Storage APIs for Array Integration (VAAI) support with the P4000 SAN. Best practices previously suggested that up to sixteen (16) virtual machines will function on a single volume—but they might experience degraded performance, depending on the hardware configuration, and if they are all booted at the same time. Volumes with four to eight virtual machines are less likely to affect performance.

With VMware's vSphere Storage APIs for Array Integration (VAAI) locking Atomic Test and Set (ATS) new best practice recommendations suggest that up to ninety (96) virtual machines will function on a single volume. Volumes with twenty-four (24) to forty-eight (48) virtual machines are less likely to impact VM performance. These VMs are highly dependent upon I/O workload patterns as well as other variables

Datastore details will show if hardware acceleration is enabled.

A decision tree for VM placement is shown in Figure 21. With or without VAAI, virtual machine placement will impact storage I/O.

Figure 21 – Manual decision tree for determination of VM placement on LUN



vSphere 5 Storage Enhancements for P4000

VMware vSphere 5 offers many new capabilities improving and extending storage specific benefits with HP P4000 SAN integration over vSphere 4.1. These new features and enhancements provide increased performance optimization, easier provisioning, monitoring and troubleshooting thereby impacting utilization and operational efficiency. The major vSphere 5 storage enhancements impacting the P4000 SAN with SAN/iQ 9.5, will be presented.

VMFS-5

vSphere VMFS-5 introduces core architectural changes enabling greater scalability and performance. Major features introduced include 64TB device support, unified block size and improved subblocking. The unified block size is now 1MB. Note that upgrading VMFS-3 to VMFS-5 is non-disruptive and immediate, however, volumes will still retain their original block size. Therefore, modifying block size to fully recognize new VMFS-5 benefits will require a reformat of the volume. This is particularly problematic if the data needs to be retained or VMs depending upon the volume need to stay online. It is therefore recommended to create new volumes from the P4000 SAN as a purposed upgrade replacement. Storage vMotion would then be leveraged for moving current VMs from VMFS-3 to VMFS-5 volumes. Additional increased efficiency is achieved with subblocking, reducing the storage overhead associated with smaller files. Small files need not occupy an entire 1MB unified block as subblocks of 8KB and 1KB are now possible. Pass-through RDMs larger than 2TB are now also supported.

Note that VMFS-5 has finer grained SCSI locking controls supporting VAAI – atomic test and set (ATS). In order to take full advantage of this accelerated P4000 feature found in SAN/iQ 9.0 and above, ensure that the VMFS volumes are using the VMFS 5.0 format. VAAI ATS improves scalability and performance. The lock manager for VMFS-5 scales linearly with the number of vSphere hosts in the cluster resource pool which is particularly important with the P4000. This impacts scalability of storage resources in the same manner fully realizing benefits of scalable, clustered storage architecture core to P4000 SAN. Please refer to vSphere Storage APIs for Array Integration.

Improvements to the lock manager have also shown up to a 40% improvement during concurrent updates of shared resources in a cluster. This impacts the VMFS-5 clustered file system not only sharing volumes, but also when performing metadata state updates to the volume. This efficiency is seen during volume creation, vMotion, VM creation and deletion, VM Snapshot creation and deletion, and VM power on and off transition states.

VMFS-5 New versus Upgraded Datastore Similarities

- The maximum size of datastore device for either VMFS-5 or physical compatibility mode RDMs is 64TB.
- The maximum size of virtual compatibility mode RDMs on VMFS-5 is 2TB–512 bytes.
- The maximum size of a VMDK virtual disk on VMFS-5 is 2TB-512 bytes.
- Conversion to VMFS-5 can not be undone.

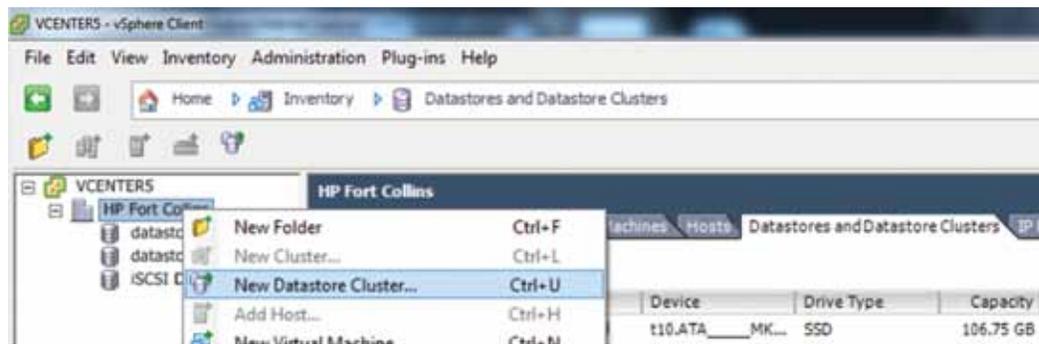
VMFS-5 New versus Upgraded Datastore Differences

- Upgraded volumes will continue to use the previous file block size which may be larger than the unified block size of 1MB. Data movement copy operations between datastores with different block sizes do not leverage VAAI storage offload acceleration.
- Upgraded volumes will continue to use 64KB instead of 8KB subblock sizes.
- Upgraded volumes will continue to have a 30,720 file limit.
- Upgraded to VMFS-5 volumes will continue to use Master Boot Record (MBR) partition types until the volume is grown above 2TB–512 bytes when it will automatically convert to GUID partition types (GPT) without impact to online VMs.
- Upgraded volumes will continue to have a partition starting on sector 128 versus newly created volumes starting on sector 2,048. This may impact best practice disk alignment benefiting P4000 SAN performance.

Storage Distributed Resource Scheduler (Storage DRS)

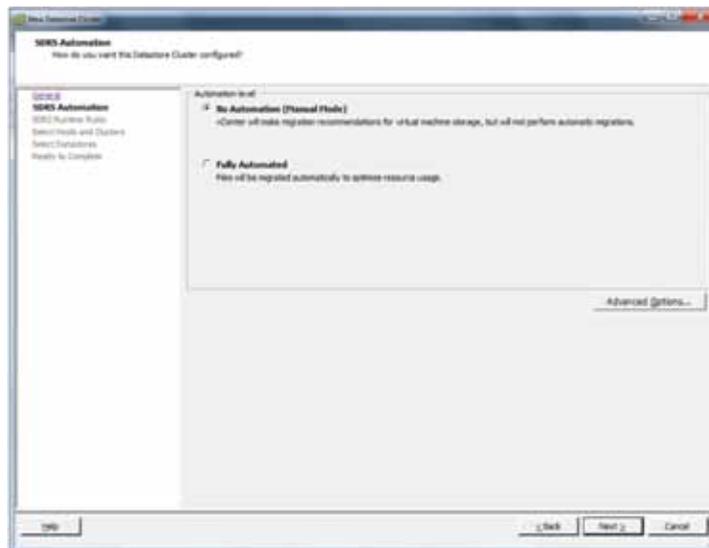
Storage DRS is a new feature providing smart VM placement across storage by making load balancing decisions which are based upon current I/O performance and latency characteristics and space capacity assessments. Essentially, day to day operational effort is decreased either by automatically or recommending best practice implementation to configurable tolerance levels. A new datastore cluster object (Figure 22) aggregates a collection of datastores into a single object.

Figure 22 – New Datastore Cluster object



Storage DRS is managed in the same manner as a compute resource cluster. This enables consistently the smartest and best placement for new VMs and virtual disks, as well as ongoing evaluation, thereby enforcing or suggesting best practice load balancing based upon existing workloads. Note that the configuration may be automatic or manual (Figure 23).

Figure 23 – SDRS Automation Level



Grouping of datastores allows for more consistent higher level abstraction and less intimate dependency for knowledge of the HP P4000 SAN datastore volumes. Storage DRS thresholds govern and trigger rules for enabling expert action (Figure 18). Storage tiering is the process of moving data to different types of storage based upon the value of the data. This tiering can be based upon performance, capacity, high availability or data segregation requirements. By leveraging multiple HP P4000 SAN clusters (SATA, MDL SAS or SAS based) and exposing different LUNs with different characteristics, such as volumes with different network RAID levels, datastore clusters can simplify management for new VM placement decisions on new or current VMs including during HA movements.

Figure 24 – SDRS Runtime rules for Datastore Clusters



Ongoing balancing recommendations are made when a datastore in a Datastore Cluster exceeds user-configurable space utilization or I/O latency thresholds. These thresholds are defined during the Datastore Cluster configuration (Figure 24). I/O load is evaluated every 8 hours by default which is also modifiable. When configured maximum space utilization or I/O latency exceeds thresholds, Storage DRS calculates all possible moves to balance VMs accordingly with a cost/benefit analysis. If the Datastore Cluster is configured for full automation, VMs will move and rebalance accordingly, if possible.

Storage DRS affinity rules enable control over which virtual disks should or should not be placed on the same datastore in a datastore cluster:

- **VMDK Anti-affinity** – virtual disks of a VM with multiple virtual disks are placed on different datastores. This ensures maximizing multiple datastore throughput over unique virtual disks. This will generally complicate P4000 SAN based snapshot and restoration.
- **VMDK Affinity** – virtual disks are kept together on the same datastore. This is usually best practice and ensures that datastore snapshots fully contain all VM data.
- **VM Anti-Affinity** – two specified VMs, including associated disks are placed on different datastores. This is recommended when known high storage resource VMs should be separated to ensure best VM SLAs.

vSphere Storage APIs for Storage Awareness (VASA)

vSphere Storage APIs for Storage Awareness (VASA) detects the capabilities of the storage array volumes and exposes them to the vCenter environment providing a foundation for quality of service based profile management. VASA support is enabled for the P4000 SAN through the HP Insight Control for VMware vCenter Server version 6.3 software. Key capabilities integrated into the vCenter console with HP Insight Control integration (with VASA):

- **Combined physical and virtual view.** From the vCenter console, monitor status and performance of VMs and the P4000 SAN.
- **Integrated troubleshooting.** Receive pre-failure and failure alerts from the P4000 SAN.
- **Provision P4000 SAN Storage.** Add new datastores, deleting or expanding an existing datastore, creating new VMs, or cloning existing VMs.

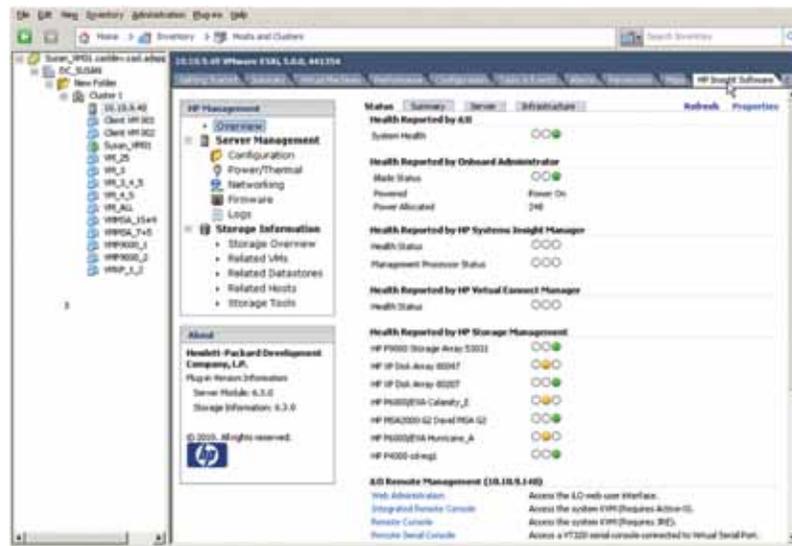
Storage capability, RAID level, thin or thick LUN provisioning, and replication states are now all visible within vCenter. This enables profile driven storage allowing rapid and intelligent placement of VMs based upon a service level agreement (SLA), availability, or performance. With profile-driven storage, characteristics defined as a tier can be requested in a defined VM storage profile. These profiles are then used during provisioning, cloning and vMotion to ensure that storage stays in compliance. VM compliance status enables the checking of all VMs and disks in a single view. VM Administrators that may not have access to all layers can now validate virtual machine compliance easily. Profile-Driven storage also removes the need for maintaining complex mapping and validating compliance during every migration, creation of VMs or virtual disks.

HP Insight Control Storage Module for vCenter

The HP Insight Control for VMware vCenter server 6.3 Storage Module is a plug-in module to VMware's vCenter Server which enables VMware administrators to monitor and manage their HP servers, networks and storage. Administrators may clearly view the relationships between virtual machines, datastores and the storage arrays, and manage these arrays directly from the vCenter console.

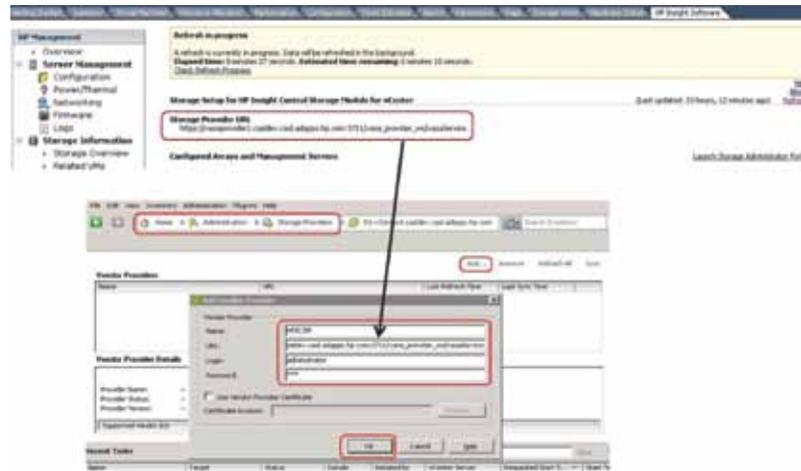
Instead of provisioning storage through the CMC, Insight Control supports creation of datastores and VMs from within the vCenter Server Console. Also note, with version 6.3 now supporting VASA, the P4000 VASA provider HPICSM is installed through the 6.3 installation process in registering the VASA Provider (Figure 26). HP Insight Control registers with vSphere and has the core framework and functionality to integrate the server, network and storage environment. The HPICSM VASA provider allows users to map the VMware objects (VMs, Datastores, RDMs) to the virtual disks on the P4000 SAN. Additionally, storage provisioning operations – create datastore, expand datastore, create VM from template and clone VM are also supported.

Figure 25 – HP Insight Control for VMware vCenter



The VASA interface provides information about the SDRS management capability for LUNs or filesystems. A VASA Provider can indicate that the storage device supports SDRS migration for a particular LUN or filesystem. A Provider can also indicate whether VMDK file migration between LUNs is recommended.

Figure 26 – Registering the HP P4000 SAN VASA Provider



For additional information, best practices and software, please go to the HP Insight Control site <http://h18000.www1.hp.com/products/servers/management/integration.html>.

Storage I/O Control (Storage I/O Control)

Storage I/O Control was introduced in vSphere 4.1 to provide I/O prioritization of VMs. Dynamic allocation of I/O queue slots across vSphere hosts. Storage I/O control automatically throttles a VM that is consuming disparate amount of I/O bandwidth when configured latency thresholds are reached. Storage DRS and Storage I/O Control both prevent deprecation of service level agreements (SLAs) while providing fair storage I/O distribution across available resources.

Best practices

- Use at least two Gigabit network adapters teamed together for better performance and for failover of the iSCSI connection. 10GbE network adapters will also increase performance.
- Teaming network adapters provides redundancy for networking components such as adapters, cables, and switches. An added benefit to teaming is an increase in available I/O bandwidth. Network teams on HP storage nodes are easily configured in the CMC by selecting two active links and enabling a bond. Adaptive load balancing (ALB) is the most common teaming method used on P4000 storage nodes. From the storage node, ALB bonding supports receiving data into the storage from one adapter, while sending out through both. The best bonding option is to leverage Link Aggregation Control Protocol (LACP) 802.3AD if the switching infrastructure supports it. From the storage node, LACP bonding supports both sending and receiving data from both adapters. Network teams on vSphere 5 servers are configured at the virtual switch level. As an added improvement over virtual switch bonding enable native MPIO round robin for iSCSI. Network teams should show Active I/O across both interfaces. Properties on the vSwitch will show both adapters as active. Properties on the datastore, manage paths will also verify Active I/O on all paths.
- The VMkernel network for iSCSI should be separate from the management and virtual networks used by virtual machines. If enough networks are available, vMotion and FT should also use a separate network. Separating networks by functionality, iSCSI, vMotion, FT and virtual machines provides higher reliability and performance of those functions.
- For improved performance, enable VIP load-balancing for iSCSI in the CMC.
- Hardware based flow control is recommended for P4000 SANs on all network interfaces and switch interfaces.
- The load-balancing feature of the P4000 software allows iSCSI connections during volume login to be redirected to the cluster's least busy storage node. This keeps the load on storage nodes throughout the cluster as even as possible and improves overall performance of the SAN. This setting is enabled by default in P4000 SANs.
- Virtual machines that can be backed up and restored together should share the same datastore volume.
- HP recommends using HP network cards that are capable of offloading TCP/IP in iSCSI network. Network cards capable of offloading processes from the CPU to the network card ASIC enhances the performance of ESX servers and balances the processing in the solution. For a list of vSphere 5 supported network interfaces, refer to the VMware Compatibility Guide for network I/O devices <http://www.vmware.com/resources/compatibility>.
- Because P4000 snapshots, Remote Copy volumes, and HP SmartClone volumes work on a per-volume basis, it is best to group virtual machines on volumes based on their backup and restore relationships. For example, a test environment made up of a domain controller and a few application servers would be a good candidate for being put on the same volume. These elements could be snapshot, cloned, and restored as one unit.
- HP recommends synchronizing P4000 nodes and vSphere hosts with the same time server. Time synchronizing your environment enhances your ability to troubleshoot problems.

Frequently asked questions

I added the P4000 cluster virtual IP address to vSphere 5 host's Dynamic Discovery Send Targets list; why don't I see a new target?

Most likely, either you need to perform a Rescan All on the storage adapters, or else you have not configured authentication on the SAN correctly. Also, you should confirm all network configurations. Verify on the iSCSI SAN that IP addresses can be pinged. Please refer to Discovery of the first iSCSI volume.

Is virtual IP load balancing supported in vSphere 5?

The vSphere 5 Software Adapter and hardware adapters are supported by P4000 Virtual IP Load Balancing. As a best practice, load balancing should be enabled in the P4000 CMC on the servers defined for assigning authentication of all vSphere 5 hosts.

What initiator should I use for a virtual machine boot volume/LUN?

There are many ways to connect and present an iSCSI volume to a virtual machine on a vSphere 5 host. These include using the vSphere 5 Software Adapter, a hardware adapter (HBA), or the software iSCSI initiator provided by some guest operating systems.

For VMFS datastores containing a virtual machine's definitions and virtual disk files, you need to use the vSphere host's hardware or software adapter. You may use either adapter. Both give you full vSphere 5 functionality, and both are supported equally. The HBA adapter has the advantages of supporting boot from SAN (prior to ESXi) and offloading iSCSI processing from the vSphere 5 host. Note that although ESXi supports iBFT, it is not currently supported or listed in HP's Single Point of Connectivity Knowledge base or the P4000 Compatibility Matrix. VMware's Auto Deploy feature may also be another stateless booting mechanism. If boot from SAN is not necessary, then the iSCSI Software Adapter is a good choice. The impact of iSCSI processing on modern processors is minimal, so with either adapter, performance will be related primarily to the quality of the physical network, disk quantity, and disk rotation speed of the attached SAN. In addition, most enterprise network interfaces offer offloading of CRC check sums further improving the comparison between network interfaces and dedicated iSCSI HBAs. Most vSphere deployments on P4000 SANs are done with the iSCSI Software Adapter included with vSphere.

What initiator should I use for additional raw device mappings or virtual disks?

Aside from the boot LUN/volume, additional volumes should be used for storing application data. Specifically, as a best practice, database and logs volumes for many applications require separate volumes. These should be presented as either raw devices (RDMs) through your chosen vSphere 5 host or connected as iSCSI disks directly through the software initiator of the virtual machine's guest operating system. Using RDMs or direct iSCSI allows these application volumes to be transported seamlessly between physical and virtual servers, as they are formatted in the native file system (NTFS, EXT3, etc.) of the guest VM OS. To use the guest operating system initiator successfully, follow these guidelines:

- Make sure the guest operating system initiator is supported by P4000 SAN technology. Please consult the HP Single Point of Connectivity Knowledge base (SPOCK) <http://www.hp.com/storage/spock>.
- For optimal performance and failover, the guest network used by the initiator should have at least dual Gigabit links and be separate from other virtual networks, including VMkernel, vMotion, FT, host management, and virtual machine public networks.
- Make sure that the guest operating system is using the fastest paravirtualized network driver vmxnet3 NIC driver from VMware tools.
- Make sure that the virtual machine will not be used in conjunction with VMware Site Recovery Manager (SRM). SRM does not work with volumes connected by guest initiators.

How many virtual machines should be on a volume or datastore?

Refer to Choosing datastores and volumes for virtual machines.

Are jumbo frames supported?

Jumbo frames (IP packets configured larger than the typical 1,500 bytes—i.e., up to 9,000 bytes) are supported by all HP P4000 SANs. In order for jumbo frames to be effective, they must be enabled end to end, including in network adapters and switches. For most Gigabit iSCSI implementations jumbo frames are not required for good performance. 10 Gigabit iSCSI implementations do see significant benefits from jumbo frames.

What is the maximum size iSCSI volume supported in vSphere 5 and what size vmdk?

The maximum volume size supported by vSphere 5 is 64TB. Note that VMFS-5 is required. If upgrading, ensure that datastores are running the newer VMFS-5 format. The maximum VMDK size has not changes in vSphere 5 and is still 2TB. Note that VSAs based upon SAN/iQ 9.5 will still support a maximum of 10TB with 5 virtual disks of 2TB each.

What version of P4000 Software supports VMware vSphere 5?

All HP P4000 releases, HP SAN/iQ Software versions 9 .0 and above, support vSphere 5, but does not fully support all of the new VAAI primitives. For SRM 5, SAN/iQ 9.5 is required with the new SRA 5.0 plugin.

For more information

For more information on HP P4000 SANs, visit: www.hp.com/go/p4000

For more information on the HP P4000 Virtual SAN Appliance, visit: www.hp.com/go/vsa

HP single point of connectivity, www.hp.com/storage/spock

HP Virtualization with VMware, www.hp.com/go/vmware

HP Insight Control management software, www.hp.com/go/insightcontrol

HP BladeSystem, www.hp.com/go/bladesystem

HP BladeSystem interconnects, www.hp.com/go/bladesystem/interconnects

HP Systems Insight Manager (HP SIM), www.hp.com/go/hpsim

HP Virtual Connect Technology, www.hp.com/go/virtualconnect

VMware fault tolerance demo using P4000 SAN, http://h30431.www3.hp.com/?fr_story=217cb39a5a037d7a3050746dc4838b02d5700d31&rf=bn

SAN/iQ 9 P4000 VMware VAAI Whiteboard, <http://www.youtube.com/watch?v=pkHJRwW-jXc&playnext=1&list=PL4B2F1EE84DAFE90B>

To help us improve our documents, please provide feedback at http://h20219.www2.hp.com/ActiveAnswers/us/en/solutions/technical_tools_feedback.html.

vmware®

Share with colleagues



Become a fan on  >

Follow on  >



Get connected

www.hp.com/go/getconnected

Current HP driver, support, and security alerts delivered directly to your desktop

© Copyright 2011 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

4AA3-6918ENW, Created September 2011

