Technical white paper

HP LeftHand Storage with VMware vSphere: Design considerations and best practices

Table of contents

Executive summary

This white paper provides detailed information on how to integrate VMware vSphere 5.0 with HP LeftHand Storage. VMware vSphere is an industry-leading virtualization platform and software cloud infrastructure that enables critical business applications to run with assurance and agility by virtualizing server resources. Complementing this technology, HP LeftHand Storage (formerly known as HP P4000 LeftHand SAN solutions) addresses the storage demands and cost pressures associated with server virtualization, data growth, and business continuity. HP LeftHand Storage scales capacity and performance linearly without incurring downtime, allowing 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 vSphere 5.0 on HP LeftHand Storage.

Target audience:
VMware and Storage administrators who are implementing vSphere 5.0 on HP LeftHand Storage.

New Storage feature

With vSphere 5.0, VMware continues to increase both its infrastructure and application services. This also influences storage deployments with VMware vSphere:

- VMware ESXi 5.0 converges on a thinner hypervisor architecture supporting a hardened 144 MB disk footprint, streamlined patching, and configuration model.
- Storage Distributed Resource Scheduler (Storage DRS) aggregates storage into pools, simplifies scale management, ensures best VM load balancing, and helps avoid storage resource bottlenecks.
- Profile-driven storage matches appropriate storage to given VMs, thus ensuring correct placement.
- Virtual Machine File System version 5 (VMFS-5) enhances scalability and increased performance supporting larger 64 TB datastores, a unified 1 MB block size, and small file space utilization with 8 KB sub-block sizes.
- Improved vSphere Storage APIs for Array Integration (VAAI) and new vSphere Storage APIs for Storage Awareness (VASA) enhance performance and bring array information into DRS and profile-driven features.

HP LeftHand Storage increases and builds upon features found in HP LeftHand operating system (HP LeftHand OS; formerly known as SAN/iQ) version 9.0 with version 9.5:

- VAAI introduced in vSphere 4.1 improves storage efficiency and performance.
  Note that newer VAAI feature extensions introduced in vSphere 5.0 are not supported in HP LeftHand OS 9.5.
- HP LeftHand OS 9.5 expands application integration with Application-managed Snapshots for VMware. Application-consistent point-in-time copies allow speedy and reliable Virtual Machine (VM) recovery.
- HP LeftHand OS 9.5 features "Zero to VSA" assisted and automated installation for easy deployment of HP LeftHand Storage Virtual SAN Appliance (VSA).
- HP LeftHand OS 9.5 supports new Storage Replication Adapter (SRA) for VMware vCenter Site Recovery Manager 5.0 Site Recovery Manager 5 (SRM) offering integrated storage failover and failback.
- 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 HP LeftHand Storage and VMware vSphere 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 and tuning.

vSphere 5 storage enhancements for HP LeftHand Storage

vSphere 5 offers many new capabilities over vSphere 4.1—improving and extending storage-specific benefits with the integration of HP LeftHand Storage. These new features and enhancements provide increased performance optimization, easier provisioning, monitoring, and troubleshooting thereby impacting utilization and operational efficiency.

VMFS-5

VMFS-5 introduces core architectural changes enabling greater scalability and performance. Major features introduced include 64 TB device support, unified block size and improved sub-blocking. The unified block size is now 1 MB. 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 on HP LeftHand Storage as a proposed upgrade replacement. Storage vMotion would then be leveraged for moving current VMs from VMFS-3 to VMFS-5 datastores. Additional increased efficiency is achieved with sub-block addressing, reducing the storage overhead associated with smaller files. Small files need not occupy an entire 1 MB unified block as sub-blocks of 8 KB and 1 KB are now possible. Raw Device Mappings (RDMs) larger than 2 TB are now also supported.

Note that VMFS-5 has finer grained SCSI locking controls supporting hardware-assisted locking in VAAI. In order to take full advantage of this feature found in HP LeftHand OS version 9.0 and above, ensure that the VMFS datastores are using VMFS-5. Hardware-assisted locking 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 useful in conjunction with the scalability characteristics of HP LeftHand Storage. This impacts the number of virtual machines per VMFS datastore, the overall storage scalability and helps to fully realize in the same manner, fully realizing the benefits of scalable, clustered storage architecture core offered by HP LeftHand Storage.

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

Storage DRS

Storage DRS is a new feature that provides 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 (figure 1) aggregates a collection of datastores into a single entity.

Figure 1. Defining a new Datastore Cluster in VMware vCenter
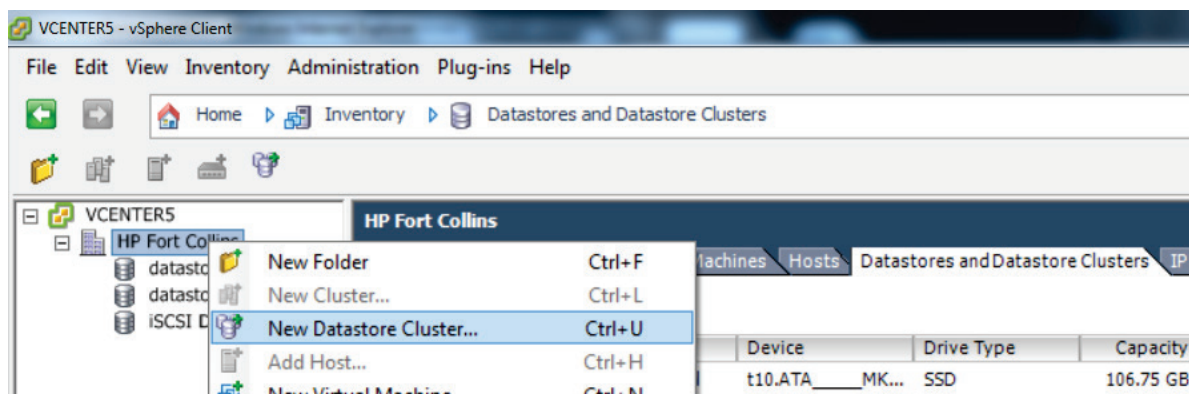
Storage DRS is managed in the same manner as DRS in compute resource cluster. This consistently enables the smartest and best placement for new VMs and virtual disks, as well as ongoing evaluation, thereby enforcing or suggesting best practice load balancing based upon existing workloads. Note that the configuration may be automatic or manual (figure 2).

Figure 2. Choosing the level of Storage DRS automation
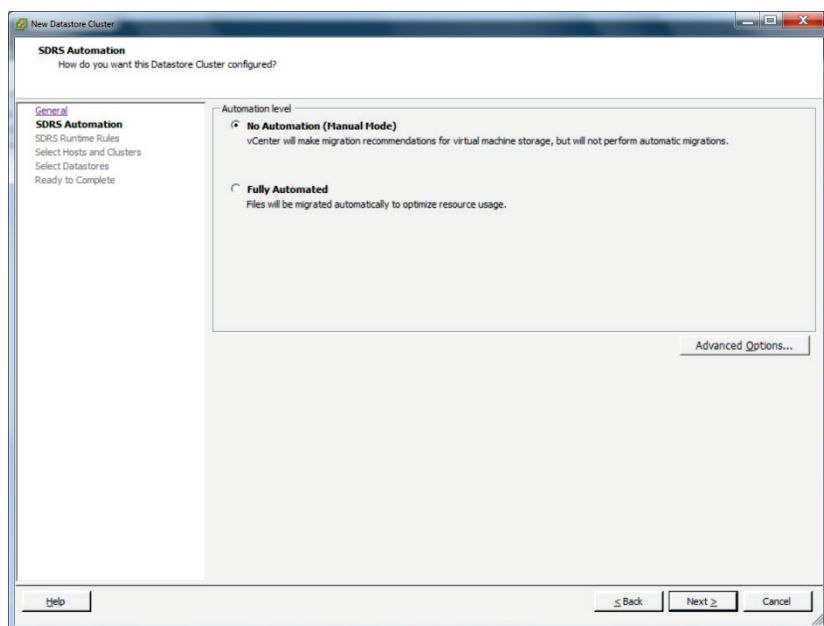
Clustering of datastores allows for more consistent higher level abstraction and less intimate dependency for knowledge of the HP LeftHand volumes. Storage DRS thresholds govern and trigger rules for enabling expert action (figure 3). Storage tiering is the process of moving data to different types of 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 clusters of HP LeftHand Storage based on different drive technology (MDL SAS, SAS, or SSD-based) and exposing volumes with different performance and capacity characteristics, such as volumes from different HP LeftHand clusters 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 3. Storage DRS settings for thresholds and timers
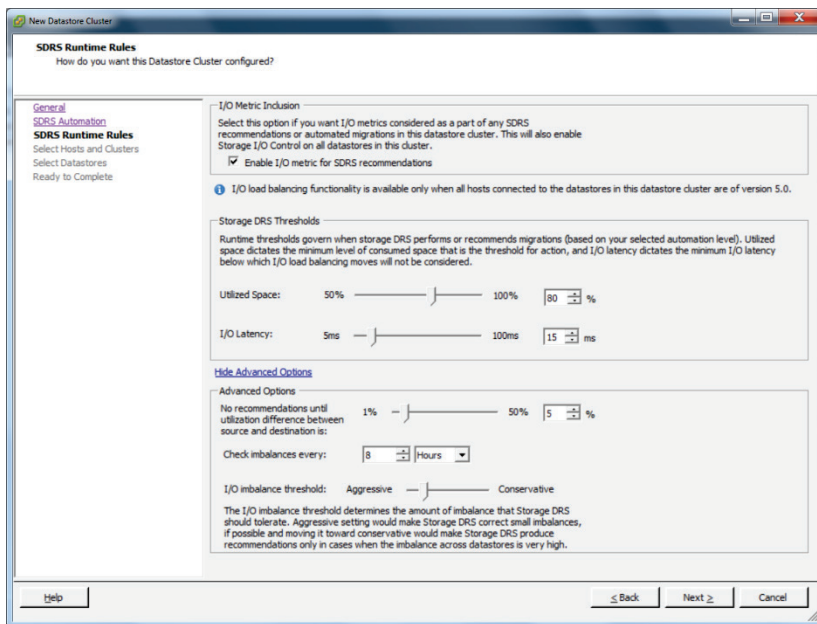
Ongoing balancing recommendations are made when a datastore in a Datastore Cluster exceeds user-configurable space utilization or I/O latency thresholds. These thresholds are defined during the datastore cluster configuration (figure 3). I/O load is evaluated every eight hours by default which is also configurable. 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 helps maximize multiple datastore throughput over unique virtual disks. This will generally complicate HP LeftHand Snapshot and rollback.

• VMDK Affinity
   Virtual disks are kept together on the same datastore. This is usually a best practice and enables 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 promote best VM SLAs.

HP Insight Control Storage Module for vCenter and vSphere Storage APIs for Storage Awareness (VASA)

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

Instead of provisioning storage through the HP LeftHand Centralized Management Console, the Storage Module of Insight Control for VMware vCenter Server supports creation of VMFS datastores and VMs from within vCenter. With version 6.3 and above, the HP Insight Control Storage Module (HPICSM) also supports VASA, the HP LeftHand VASA provider is part of HPICSM and is installed through the 6.3 installation process in registering the VASA provider (figure 5). The HPICSM VASA provider allows users to map the VMware objects (VMs, datastores, RDMs) to the volumes on HP LeftHand Storage.

Figure 4. HP Insight Control for VMware vCenter Server with monitoring for servers, storage and Virtual Connect networking
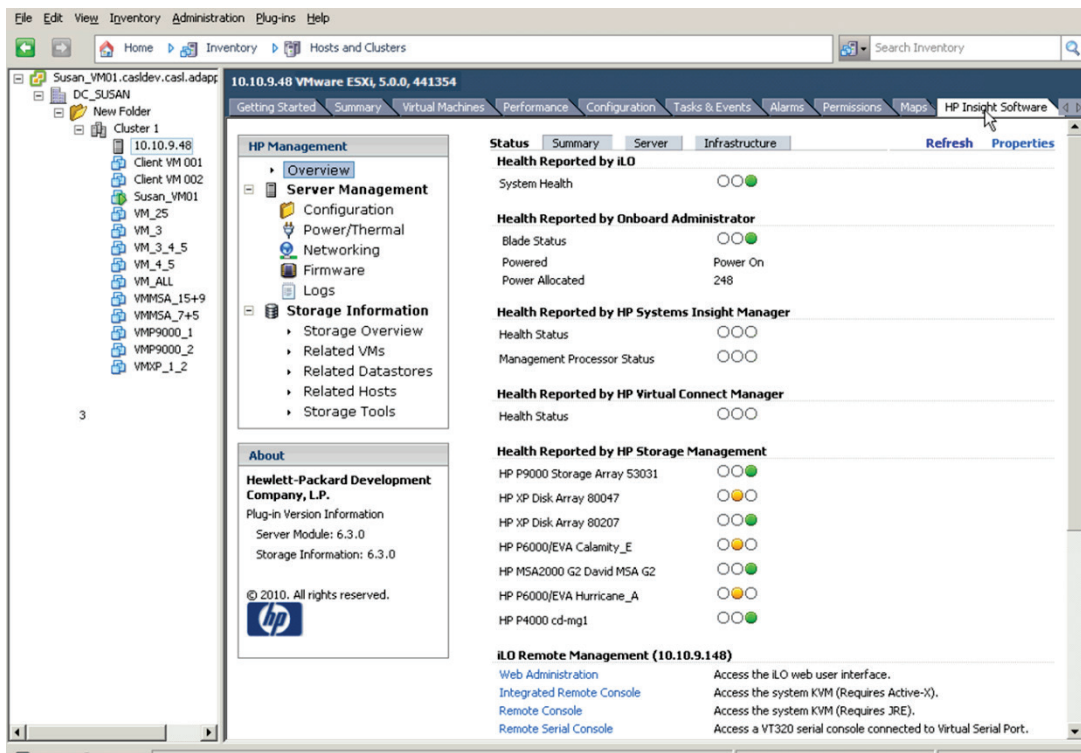
The VASA interface provides information about the Storage DRS management capability for VMFS datastores and volumes. A VASA provider can indicate that the storage device supports Storage DRS migration and whether migration between VMFS datastores is recommended.

**Figure 5.** Registering the VASA provider for HP LeftHand Storage
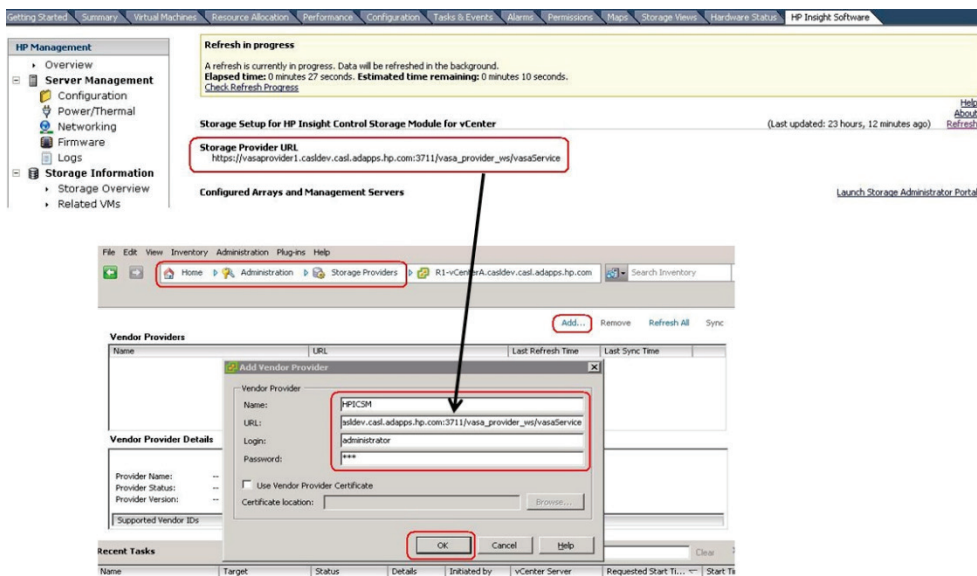
For additional information, best practices and software, please go to the HP Insight Control site
http://h18000.www1.hp.com/products/servers/management/integration.html.

VMware High-Availability and Distributed Resource Scheduler

The advanced vSphere 5 features, such as vMotion, Clusters, High-Availability (HA), Fault Tolerance (FT) and Distributed Resource Scheduler (DRS), require multiple vSphere 5 hosts to have simultaneous access to the same set of volumes on shared storage. To enable this on a HP LeftHand 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 connection for each vSphere 5 host that connects to HP LeftHand Storage, and create a server cluster using the steps outlined in the section Creating the first iSCSI volume. When multiple server connections 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 vSphere 5 (figure 6). When assigning a server cluster rather than individual server connections, this warning will not appear.

A server connection with iSCSI qualified name (IQN) authentication maps one vSphere 5 iSCSI software adapter to one or more volumes. In addition, Challenge Authentication Protocol (CHAP) authentication can also be used for increased security, either one-way (target secret) or two-way (target and initiator secret). 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 6. Warning when assigning a volume to multiple server connections
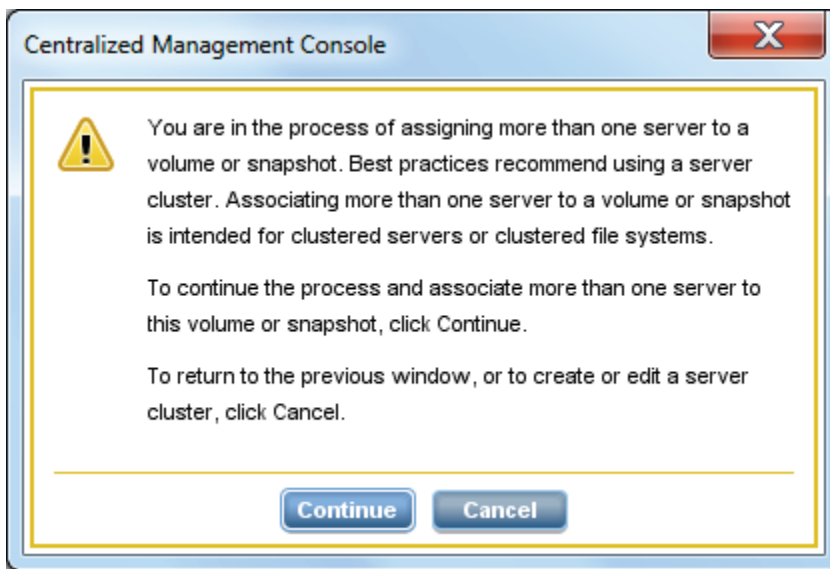
# Initial iSCSI setup of vSphere 5

**Networking for the iSCSI software adapter**

Before storage connectivity via the iSCSI software adapter is enabled in vSphere 5, a specific network configuration is recommended. For best connectivity, use at least two 1GbE or 10GbE network adapters for performance and failover (figure 7) on the iSCSI network which connects hosts running vSphere 5.0 to HP LeftHand Storage.

Figure 7. vSphere standard switch with two physical adapters and two VMkernel ports for iSCSI
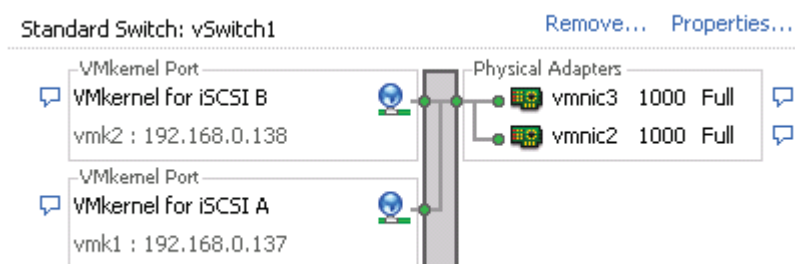
Built-in storage multi-pathing capability uses multiple VMkernel ports for iSCSI mapped each to one network adapter which are assigned to one or multiple standard vSwitches under Configuration – Networking (figure 8).

Figure 8. Mapping of VMkernel ports to network adapters
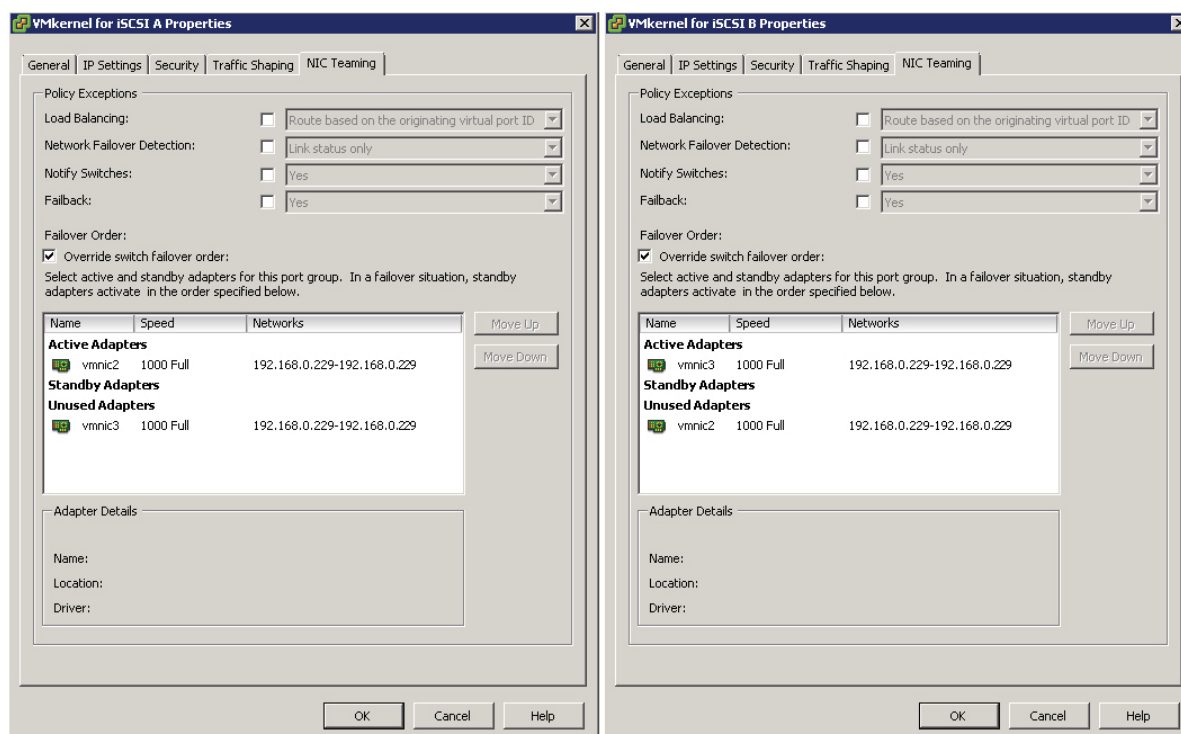
As a best performance practice, the VMkernel ports for iSCSI should be separate from the host management and virtual networks used by virtual machines. However, high availability is usually preferred over splitting networks on hosts with a small number of network adapters. If enough networks are available to provide network redundancy, it is ideal to separate vSphere vMotion and Fault Tolerance (FT) networks on unique virtual switches. These networks are placed on separate vSphere 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 help to organize and separate traffic. In a vSphere distributed switch, with network I/O control enabled, distributed switch traffic is divided across predefined network resource pools: VMware FT 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 HP LeftHand Storage clusters can accommodate different requirements than smaller HP LeftHand clusters. HP BladeSystem c7000 solutions with HP LeftHand P4800 G2 in the same enclosure have unique benefits with internal extremely low latency network characteristics for iSCSI through HP Virtual Connect technology vs. this same cluster connecting to external devices. A VM's network traffic requirements differ in every application.

These basic guidelines help you follow best practice for connecting HP LeftHand Storage. However, understanding the approach to bandwidth tuning involves understanding the complete solution objectives and needs. HP BladeSystem with HP Virtual Connect Flex-10 or HP Virtual Connect FlexFabric 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

vSphere 5 hosts with only two Gigabit (1GbE) network ports are not ideal for iSCSI storage connected by the iSCSI software adapter because they compete with other network traffic and do not offer good iSCSI performance. As shown in figure 9, vSphere 5 hosts with only two Gigabit network ports are configured with a single standard virtual switch that comprises both Gigabit ports and contains:

- A virtual machine port group network for virtual machines, with the two network adapters for load balancing (default),
- A VMkernel port network with vMotion and management traffic, with the two network adapters for load balancing (default), and
- Two VMkernel ports with iSCSI port binding enabled (on a different VLAN), each exclusively mapped to one network adapter as an active adapter (figure 8).

Figure 9. Two network adapters
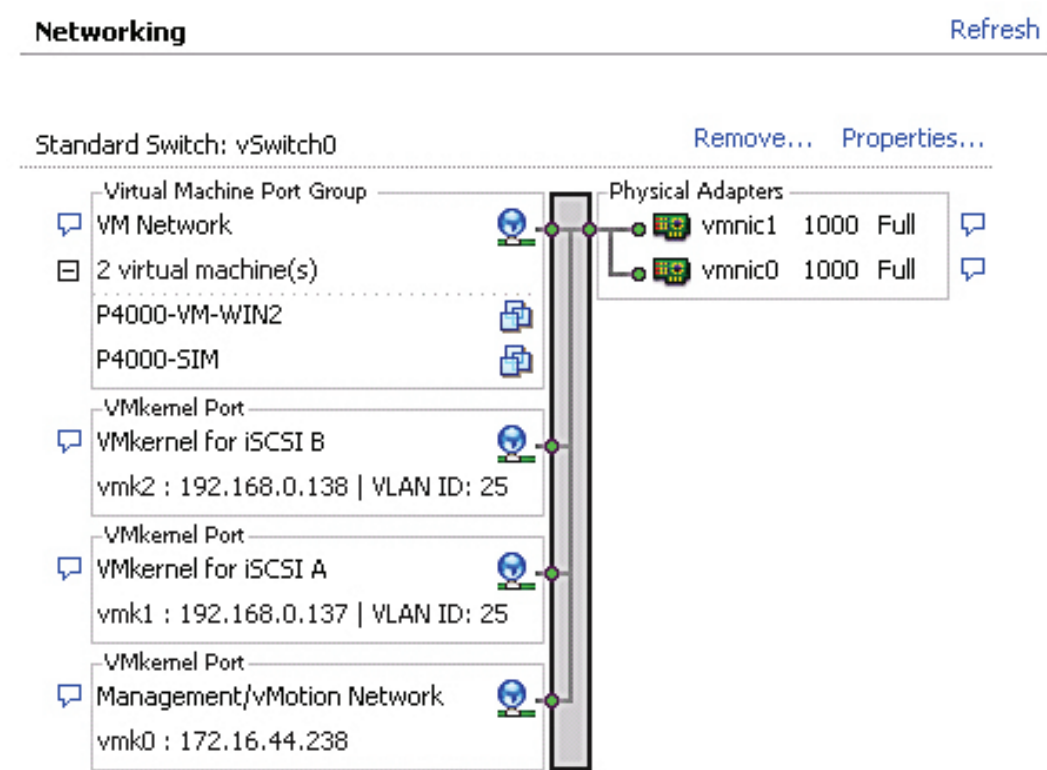
Four network ports

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 10, vSphere 5 hosts with four Gigabit network ports are configured with two virtual switches, each comprising two Gigabit ports for load balancing on virtual machine networks, management, and vMotion.

If hosts contain two dual-port network adapters, use one port from each of the dual-port adapters balancing high availability across both adapters. For example, if using two on-board Gigabit adapters and a dual-port Ethernet card, assign port 0 from the on-board adapter and port 0 from the Ethernet card to one vSphere switch, and configure port 1 from the on-board adapter and port 1 from the Ethernet card on a second vSphere switch. This configuration provides protection from some bus or network interface card failures.

- The first standard virtual switch should have:
  - A VMkernel port network with vMotion and management traffic enabled, with the two network adapters for load balancing (default), and
  - a virtual machine port group network for virtual machines, with the two network adapters for load balancing (default).

The second standard virtual switch should have:

  - Two VMkernel ports with iSCSI port binding; each exclusively mapped to one network adapter as active adapter (refer to figure 8).

Figure 10. Four network adapters (two dedicated to iSCSI)
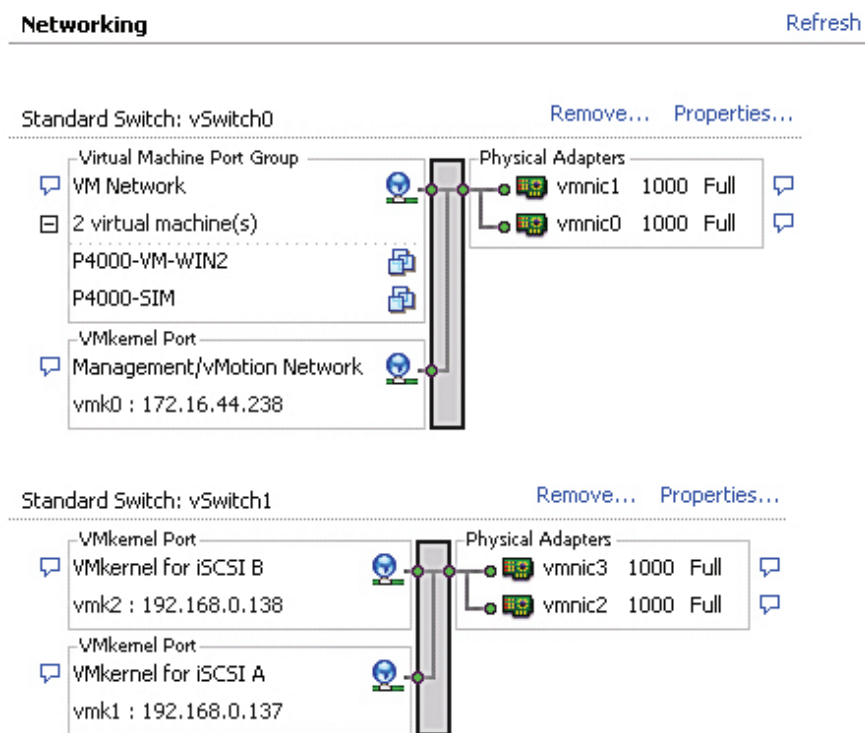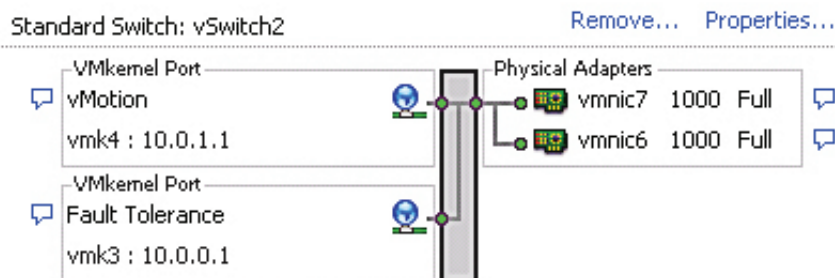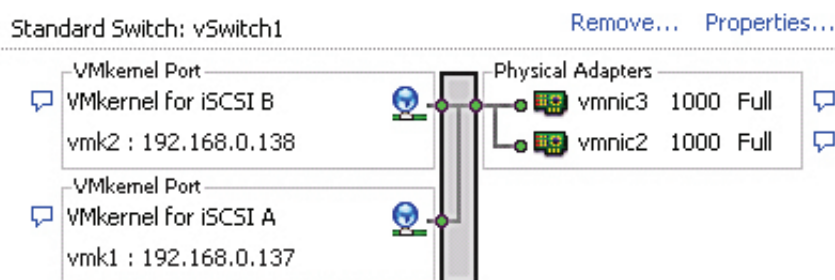

Six network ports

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 FT traffic from management 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, as shown in figure 11. 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 management traffic enabled, with the two network adapters for load balancing (default), and
  - a virtual machine port group network for virtual machines, with the two network adapters for load balancing (default).
- The second standard virtual switch should have:
  - Two VMkernel ports network with iSCSI Port binding enabled, each exclusively mapped to one network adapter as an active adapter.
- The third standard virtual switch should have:
  - A virtual machine port group network for vMotion, with the two network adapters for load balancing (default), and
  - A virtual machine port group network for Fault Tolerance, with the two network adapters for load balancing (default)

Figure 11. Six adapters
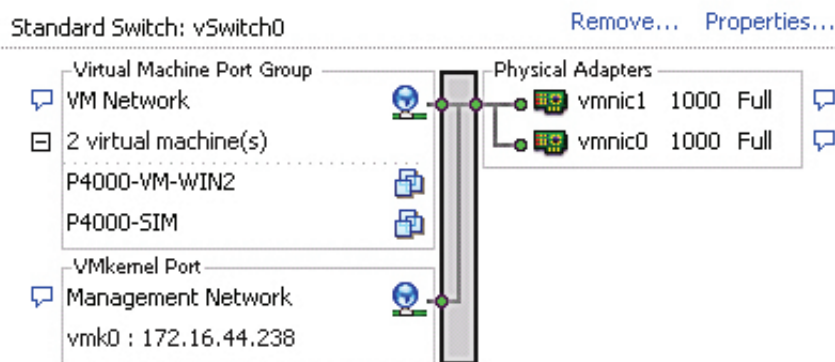
More than six network ports

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, such as virtual machine networking, as required.

When adding more ports for iSCSI, be sure to create additional VMkernel ports for iSCSI, map them to one physical adapter on the vSwitch and add them to the iSCSI software initiator.

10GbE options

Both vSphere 5 hosts and HP LeftHand Storage support 10GbE interfaces and benefit from increased bandwidth. 10GbE network ports offer additional bandwidth over 1GbE interfaces. With HP LeftHand supporting 10GbE interface upgrades, throughput has been seen to increase 2–3x depending upon workload patterns, HP LeftHand cluster size, and networking infrastructure 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, storage performance for small block, random workloads is limited by the number of disk spindles dictating input/output operations per second (IOPS) characteristics capabilities of the physical storage. In scalable storage architectures, like HP LeftHand Storage, workload is distributed across the storage cluster.

Note that in all previous examples, 1GbE could have easily been replaced by faster network bandwidth options—doing so will impact best practice choices. It is best practice to balance networking technologies end to end for overall solution performance characteristics. For example, HP LeftHand Storage with 10GbE connectivity can be limited by using a 1GbE switching infrastructure.

Enabling the iSCSI software adapter

The vSphere 5 iSCSI software adapter is the most common way to connect a vSphere 5 host to HP LeftHand Storage. It has to be enabled on each vSphere 5 host. To enable the iSCSI software adapter go to Configuration – Storage Adapters to add it to the host's list of storage adapters. Configuration of the iSCSI software initiator is done through each host's storage adapters list. Here are some guidelines when using the software iSCSI software adapter:

• Enable the iSCSI software adapter on each vSphere 5 host
• Make sure all VMkernel instances which are connected to the iSCSI network are assigned to the iSCSI software adapter (figure 12)
• Copy or write down the IQN that identifies each vSphere 5 host; it will be needed for the server connection on the HP LeftHand management group later to define volume access on the storage cluster.

Figure 12. Two VMkernel ports bound to the iSCSI software initiator
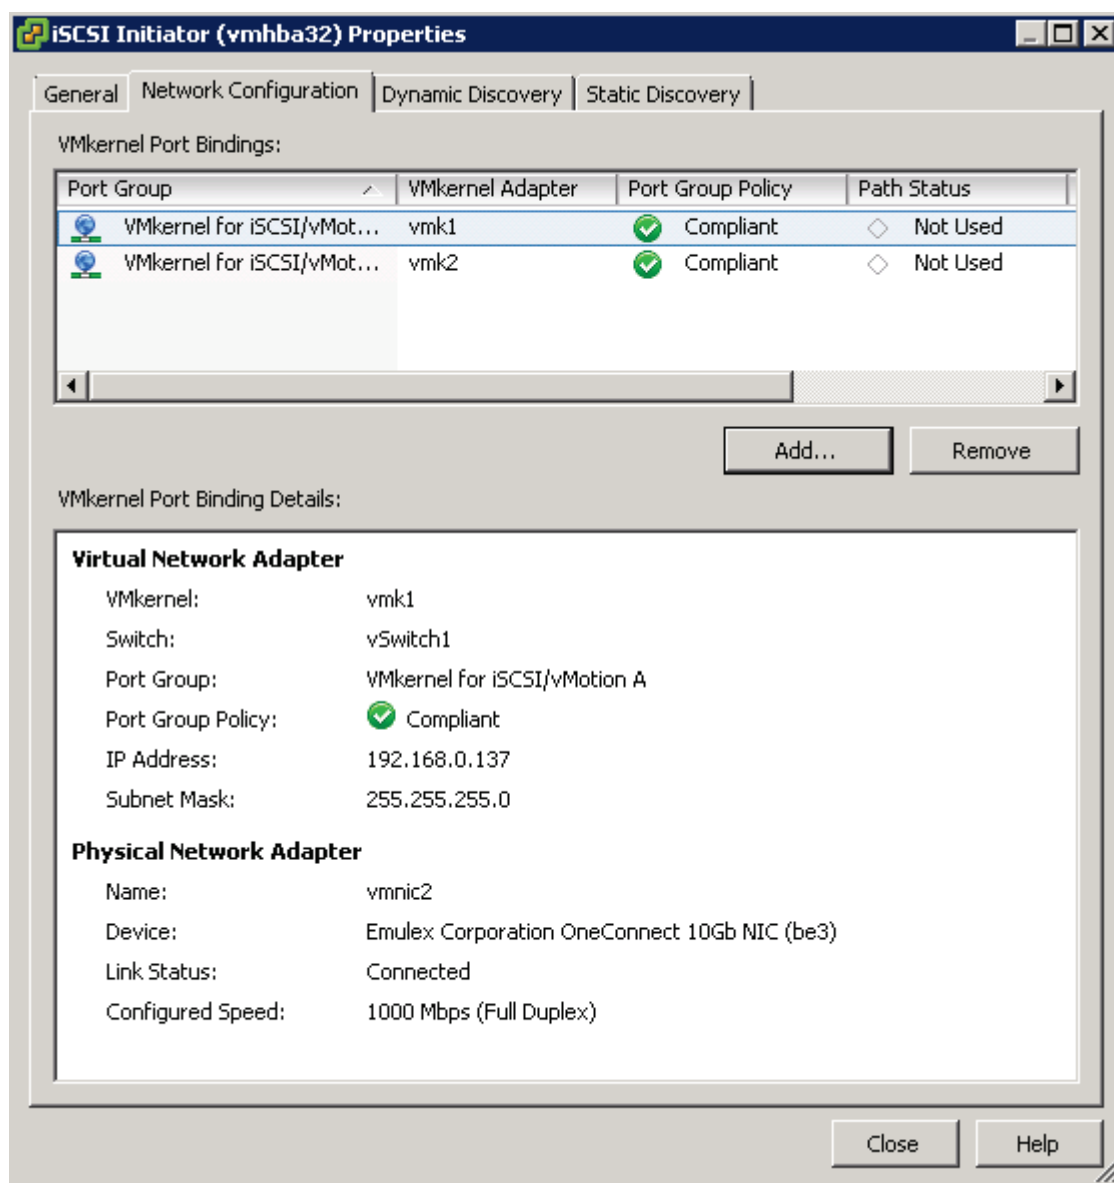
iSCSI HBA connectivity

An iSCSI host bus adapter (iSCSI HBA) uses dedicated hardware to connect to a 1GbE or 10GbE network while mitigating the overhead of iSCSI and TCP processing of Ethernet interrupts and thereby improves the performance of servers. iSCSI HBAs can include support for booting from an iSCSI target. Storage connectivity via iSCSI HBAs therefore enables both offloading of the iSCSI processing from the server and booting from HP LeftHand Storage.

iSCSI HBAs do not require licensing or special networking within vSphere 5 hosts as they provide a dedicated storage connection for iSCSI targets only and are managed within vSphere 5 similarly to Fibre Channel HBAs. As best practice, use two iSCSI HBA initiators (a dual port adapter or two single port adapters), each configured with a path to all iSCSI targets for failover. Configuring multiple iSCSI HBA initiators to connect to the same target requires each iSCSI HBA initiator's IQN to be configured as server connection on the HP LeftHand Management Group. This is configured in the HP LeftHand Centralized Management Console as two server connections (one for each HBA initiator), each with permissions to the same volumes. Note that a group of hosts in the same vSphere cluster will also require authentication for each volume. For easier handling a server cluster on the HP LeftHand Management Group can be created in the CMC to assign volumes to all hosts and all their HBAs at once.

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, BIOS boot support, and iSCSI protocol support (some even with multiple protocols) providing new compelling reasons in leveraging these advantages. Using offloading, iSCSI acceleration or boot from iSCSI targets requires a tested and supported network adapter.

Please go to the HP Single Point of Connectivity Knowledge for tested and supported iSCSI HBAs with HP LeftHand Storage. (http://www.hp.com/storage/spock)

Multi-pathing for iSCSI

For the iSCSI software adapter which uses multiple VMkernel ports, round robin path selection for active I/O on all paths is optimal. Native round robin iSCSI multi-pathing in vSphere 5 provides superior performance and resiliency by leveraging multiple paths to HP LeftHand Storage. Each path to an HP LeftHand volume will establish a unique iSCSI session. Since HP LeftHand Storage scales by 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 the HP LeftHand cluster.

Configuring iSCSI multi-pathing requires at least two network adapters on the iSCSI network. The following steps are performed on each vSphere 5 host individually:

• From the host's Configuration – Networking, select Properties of the standard virtual switch. Select the Network Adapters tab and add the appropriate adapters on the iSCSI network.

• Select the Ports tab and Edit or Add VMkernel Ports. Configure as many VMkernel ports as network adapters used for iSCSI. On the VMkernel Properties, NIC Teaming tab select one physical adapter (vmnic) to be listed as Active Adapter; all other physical adapters should be listed as Unused Adapter (not Standby Adapter). Repeat for remaining VMkernel ports, and select another physical adapter to be the active adapter, such that each VMkernel port has only one exclusive physical adapter assigned to it. See the following example for two NICs and two VMkernel ports (figure 8).

| Port group (VMkernel port name)	Active adapter	Unused adapter
iSCSI A	vmnic2	vmnic3
iSCSI B	vmnic3	vmnic2

• From the host's Configuration – Storage Adapter, highlight the iSCSI Software Initiator. In the Network Configuration tab, add all VMkernel ports which are connected to the iSCSI network.
Note that all added VMkernel ports should be listed with Port Group Policy Compliant. If not, check your network configuration, VMkernel ports for iSCSI and their NIC teaming settings.

• In the iSCSI Initiator Properties, add the HP LeftHand cluster's virtual IP address to the Dynamic Discovery (figure 13) tab and leave the port of 3260 unchanged. Close the dialog and Rescan All storage adapter to look for new volumes and scan for new VMFS datastores. All HP LeftHand volumes should appear in the iSCSI software adapter's devices section.

Figure 13. Dynamic Discovery with a single Virtual IP address of a HP LeftHand Cluster
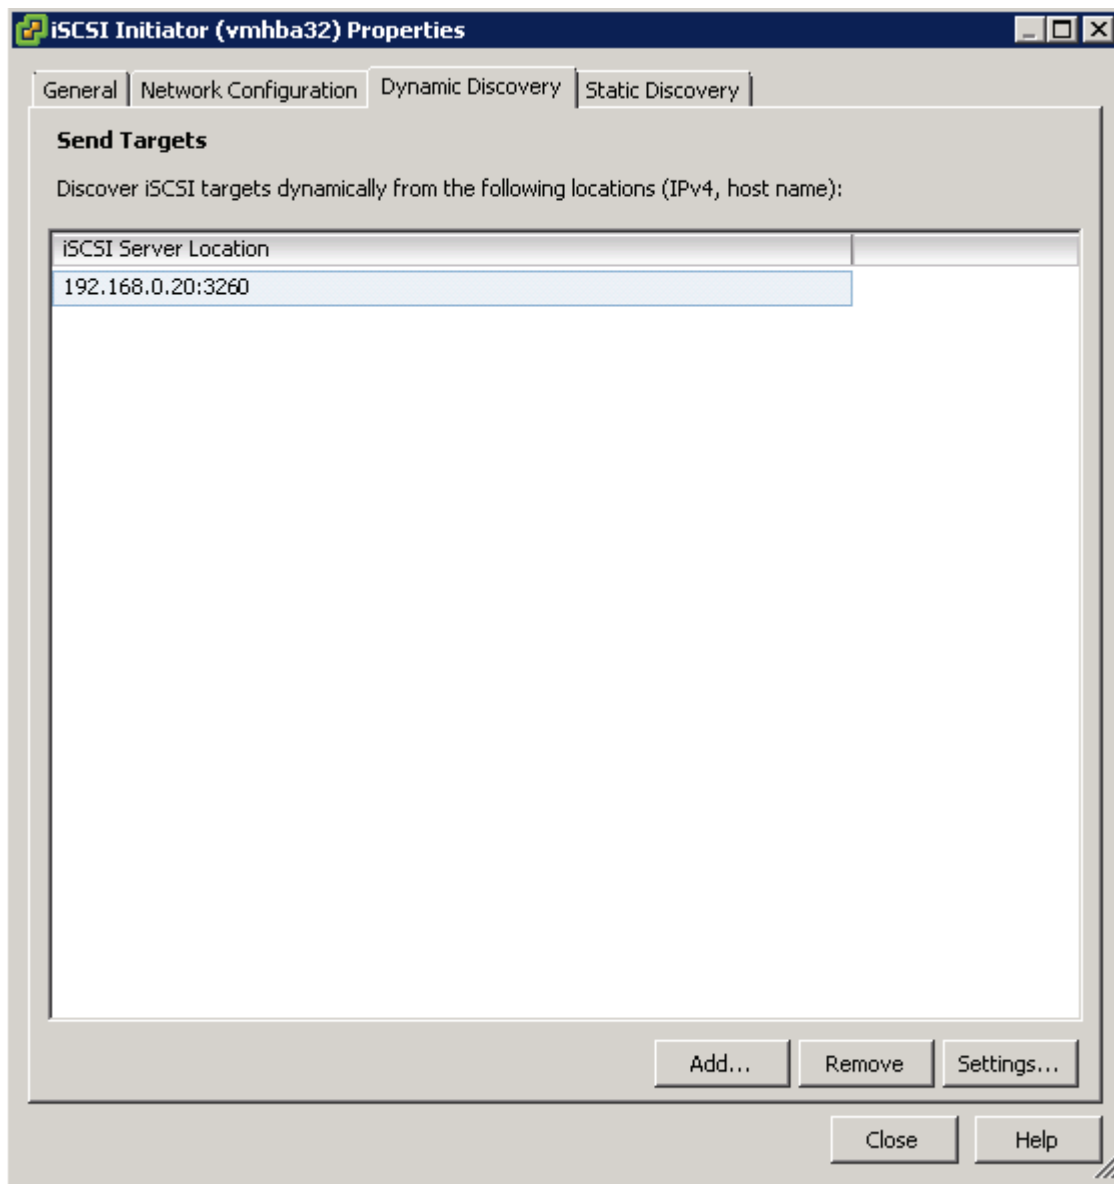
- From the host's Configuration – Storage Adapters, highlight the iSCSI volumes on HP LeftHand Storage from the iSCSI software initiator's list of devices. Right click and select the Manage Paths option. Select the Policy Path Selection to be Round Robin (VMware). Select Change to apply the change in the configuration (figure 14). Note that all paths list the status of Active (I/O).

Figure 14. Manage paths and select round robin policy
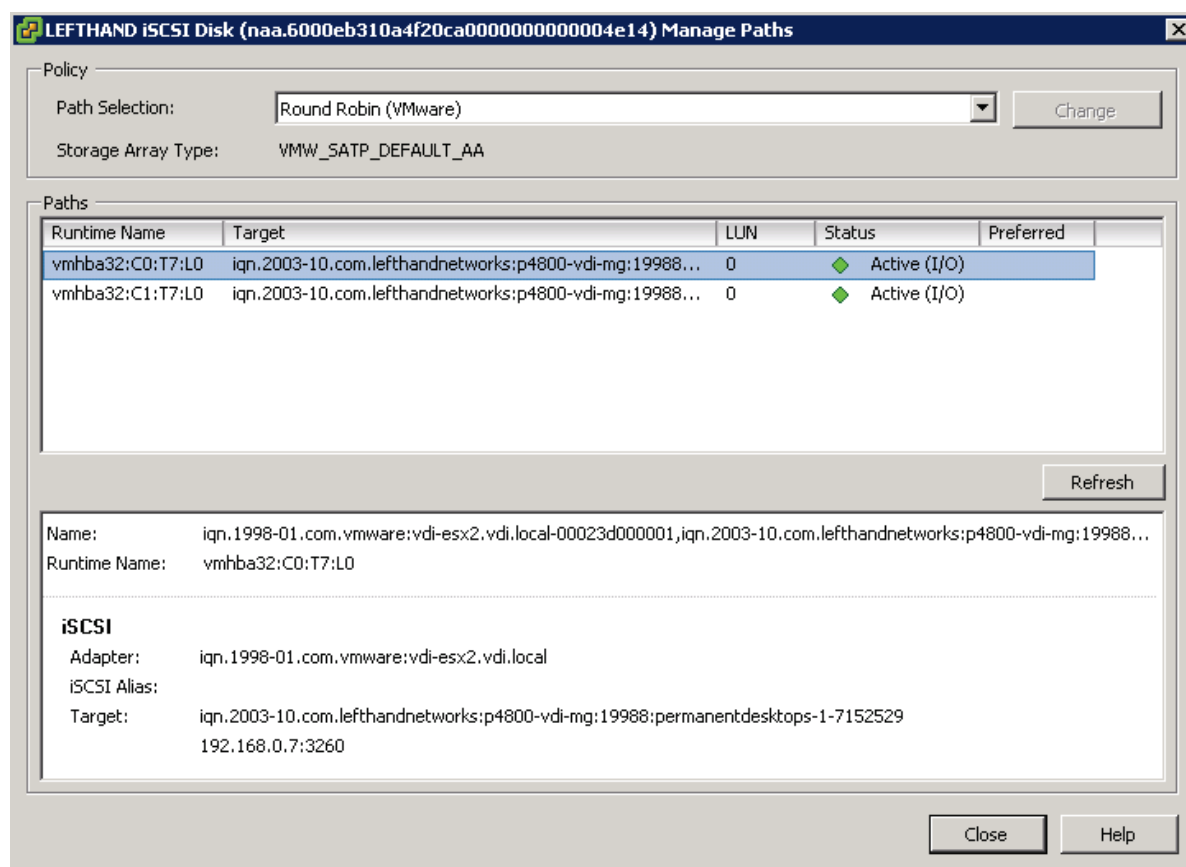
Alternative: In the host's Configuration – Storage, highlight the VMFS datastore for which you want to manage the paths, right-click, select Properties and select the Manage Paths option.

- From the host's Configuration – Storage, highlight the VMFS datastore. Note in Datastore Details, the number of Total Paths should accurately reflect the number of valid paths to the storage volume.

For each network adapter used by the iSCSI software adapter an individual iSCSI session is established to HP LeftHand Storage. The example referred to in figure 14 gives each HP LeftHand volume two iSCSI paths using two separate physical network adapters.

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.
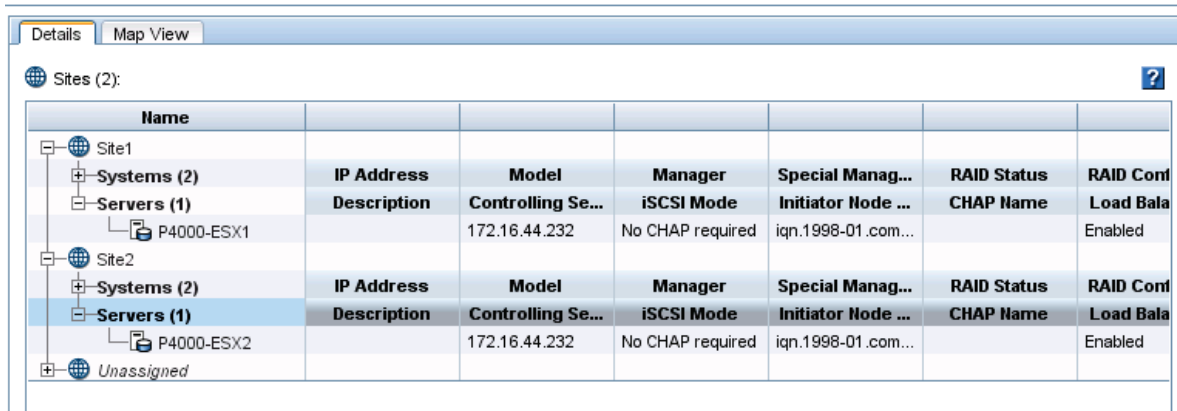
Connecting and using iSCSI volumes

Enabling load balancing

Virtual IP load balancing (VIP-LB) is a setting on each defined server connection (or server cluster) in the CMC that allows the distribution of iSCSI sessions over the storage nodes in the HP LeftHand cluster in order to maximize performance and bandwidth utilization.

Most iSCSI initiators, including the iSCSI software adapter in vSphere 5 and most iSCSI HBAs, support this feature. Virtual IP load balancing is enabled by default in the CMC and all vSphere 5 server connections on HP LeftHand Storage should have this feature enabled.

HP LeftHand multi-site deployments should use site-preferencing to allow for best effort data locality. This is achieved by assigning server connections to sites (figure 15).

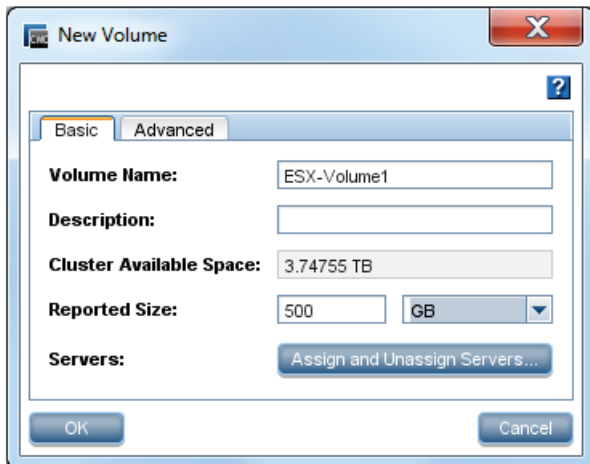**Figure 15.** Assigning servers to Sites in a HP LeftHand Management Group

Creating the first iSCSI volume

Before a vSphere 5 host can create a new VMFS datastore on an HP LeftHand volume, that volume must be created on the HP LeftHand cluster and must have been configured to allow the host to access the volume. All hosts within a vSphere 5 cluster must also have access to that same volume in order to support vMotion, High Availability, and Fault Tolerance.

Use the CMC, as shown in figure 16, or use the command-line interface (CLI) to create a new volume. Then, using the CMC, create a new server connection with the vSphere 5 host's IQN, as configured for the iSCSI software adapter; optionally use CHAP secrets as required. Please refer to the HP LeftHand User Guide for more information.

Figure 16. Creating a new volume on HP LeftHand Storage

Discovery of the iSCSI volumes

To discover volumes, the iSCSI software adapter of each vSphere 5 host must have added the Virtual IP (VIP) address of the HP LeftHand cluster which contains the volume in its Dynamic Discovery Send Targets list (see figure 13). New targets can then be discovered by simply choosing the Rescan All option from the vSphere 5 host's Configuration – Storage Adapters. 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 a vSphere 5 host can be viewed in the CMC by selecting volume and selecting the iSCSI Sessions tab (figure 17). All hosts and all their network adapters which are used for iSCSI should be listed with their individual iSCSI session to the selected volume.

Figure 17. iSCSI sessions of vSphere cluster

| Details | Snapshots | Remote Snapshots | Schedules | Assigned Servers | iSCSI Sessions | Map View |

iSCSI Sessions (4):

Server ⌃	Initiator Node Name	Chap Name	Gateway Connection	Initiator IP/Port	Identifier
vdi-esx1	iqn.1998-01.com.vmwa…		P4000-EMEA4 (192.168…	192.168.0.223/60220	23d000001
vdi-esx1	iqn.1998-01.com.vmwa…		P4000-EMEA4 (192.168…	192.168.0.123/55389	23d000002
vdi-esx2	iqn.1998-01.com.vmwa…		P4000-EMEA4 (192.168…	192.168.0.222/58834	23d000001
vdi-esx2	iqn.1998-01.com.vmwa…		P4000-EMEA4 (192.168…	192.168.0.122/52210	23d000002

Disconnecting iSCSI volumes from vSphere 5 hosts

iSCSI software adapter in vSphere 5 offers options to remove iSCSI volumes from the vSphere 5 hosts. It is considered best practice to gracefully

1. Unmount the VMFS datastore (option is available on Configuration – Storage),
2. Detach the volume (in the volume list from Configuration – Storage Adapters) on all vSphere 5 hosts, and
3. Unassign the volume for all vSphere 5 hosts via the CMC.

If all steps above are completed, the volume will no longer be displayed when rescanning the iSCSI software adapter for volumes.

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.
  - iSCSI HBAs typically have their own ping utilities inside the BIOS of the HBA, so use the appropriate utility for your HBA.
  - vSphere 5 has a network troubleshooting utility that can be used from the Management Console (or remotely through SSH) to attempt a ping to HP LeftHand Storage.
- Ping the VMkernel port from one of the HP LeftHand storage nodes by going into the network configuration on the CMC and selecting Ping from the Network Tasks.
- Double-check all IQNs or CHAP entries. For iSCSI authentication to work correctly, these must be exact.
- Make sure that all server connections (and server clusters) in the CMC have load balancing enabled. If there is a mix—if, for example, some have it enabled and some do not—then hosts that do not have load balancing enabled may not correctly connect to their volumes.

If Pings do not work, check the switch and VLAN configuration if the VMkernel ports with iSCSI binding and HP LeftHand Storage are on the same network segment. If in an HP BladeSystem environment, check HP Virtual Connect for proper mapping/tunneling of networks (i.e. VLANs) to uplink sets and server bays. Please refer to HP Virtual Connect documentation for more information.

Creating a new datastore on a volume

After the vSphere 5 host has connected to a HP LeftHand volume, the volume can be formatted as VMFS datastore or mounted as a raw device mapping (RDM) directly to virtual machines. New VMFS 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 18.

Figure 18. Creating a new VMFS datastore using wizards in vCenter
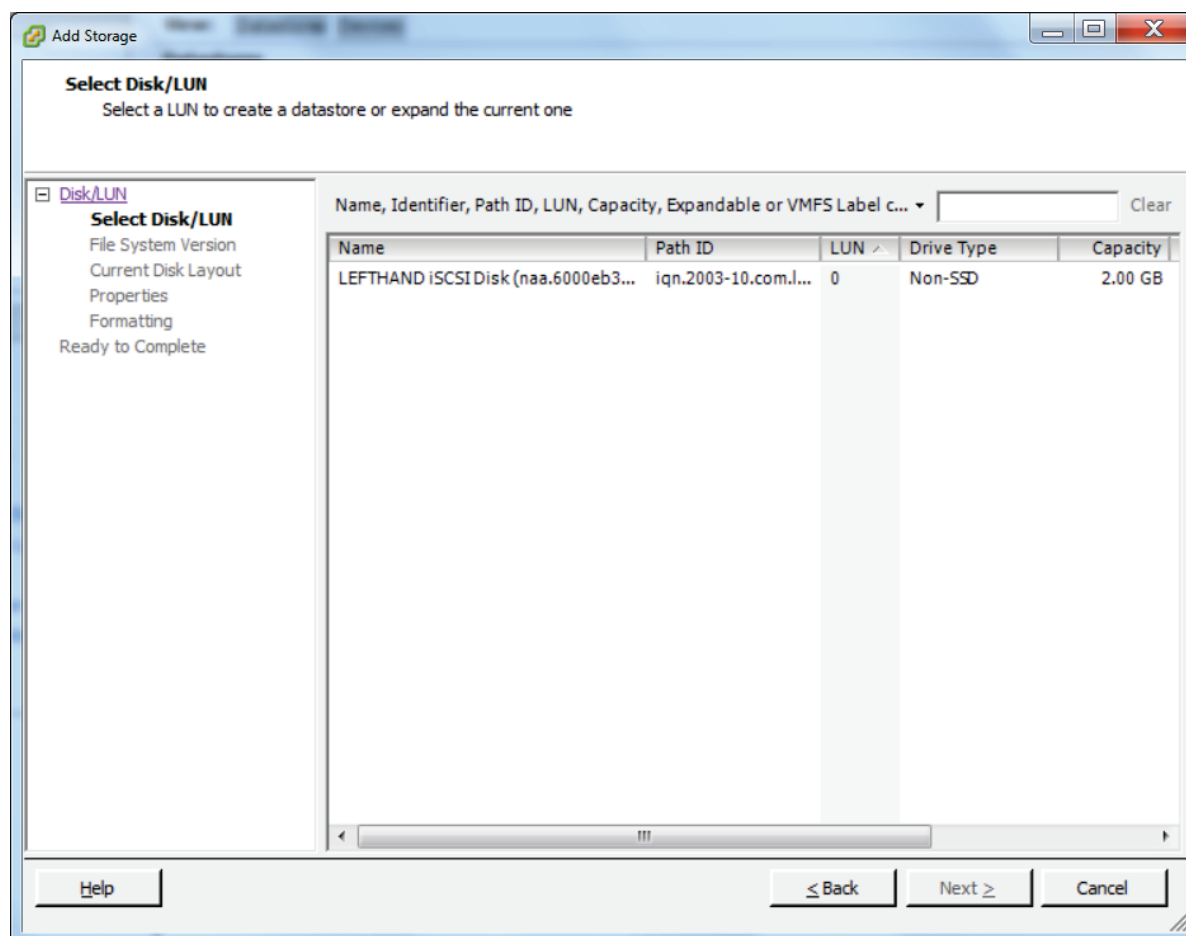
Note:
The HP LeftHand volume will be identified as an HP LeftHand iSCSI Disk.

Both HP LeftHand volumes and VMFS datastores can be expanded dynamically. If space is running low on a VMFS datastore, you must first expand the volume on the HP LeftHand cluster. To increase the size of the HP LeftHand volume, simply edit the volume in the CMC or CLI and increase its size. HP LeftHand Storage 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 on one vSphere 5 host via the vSphere client. Increasing the size of the VMFS datastore is preferred over adding extents on other HP LeftHand volumes to it.

Expanding an iSCSI LUN in the CMC

1. Go to the CMC, log in to the HP LeftHand management group, and locate the cluster's volumes.
2. Highlight the volume to be expanded.
3. Right-click on the volume and click on Edit Volume.
4. The Edit Volume window pops up showing the LUN size as 1 TB.

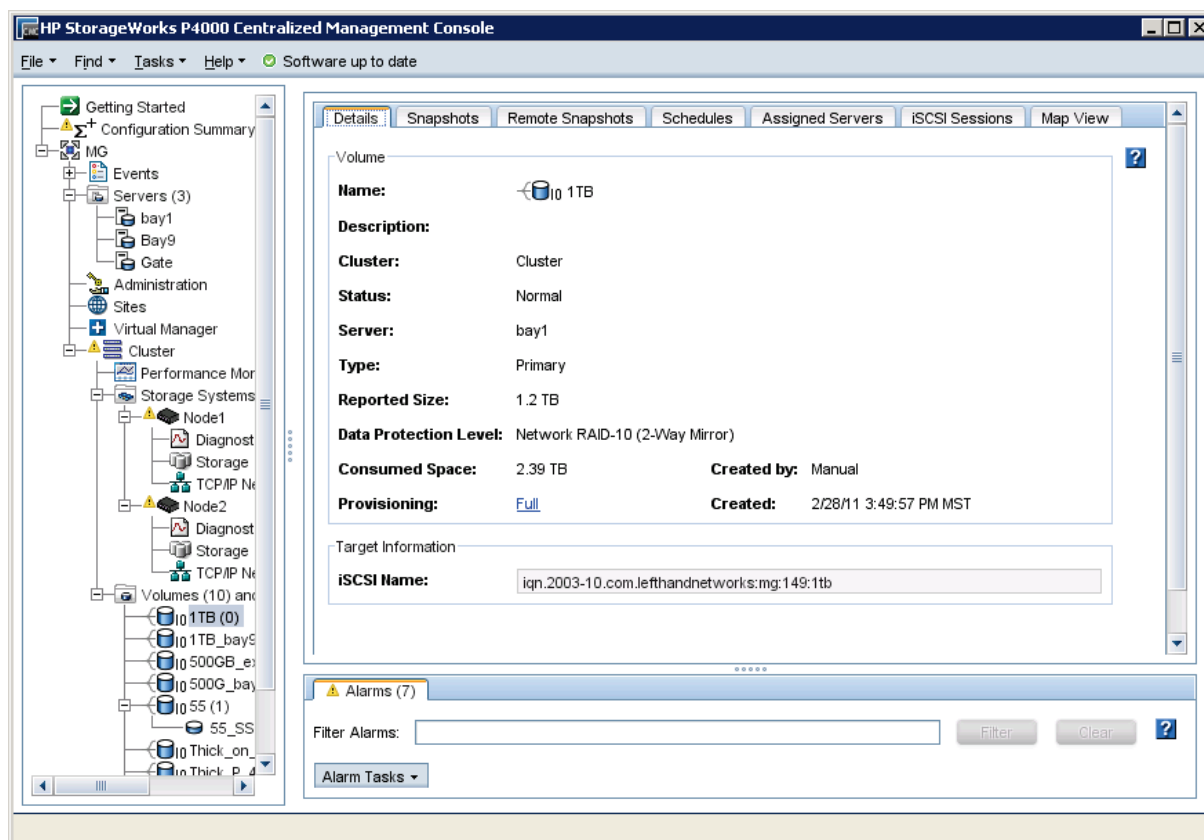5. Change the reported size to the desired size and click on OK (figure 19).

In this example, it is changed to 1.2 TB.

Figure 19. Changing the volume size in the CMC
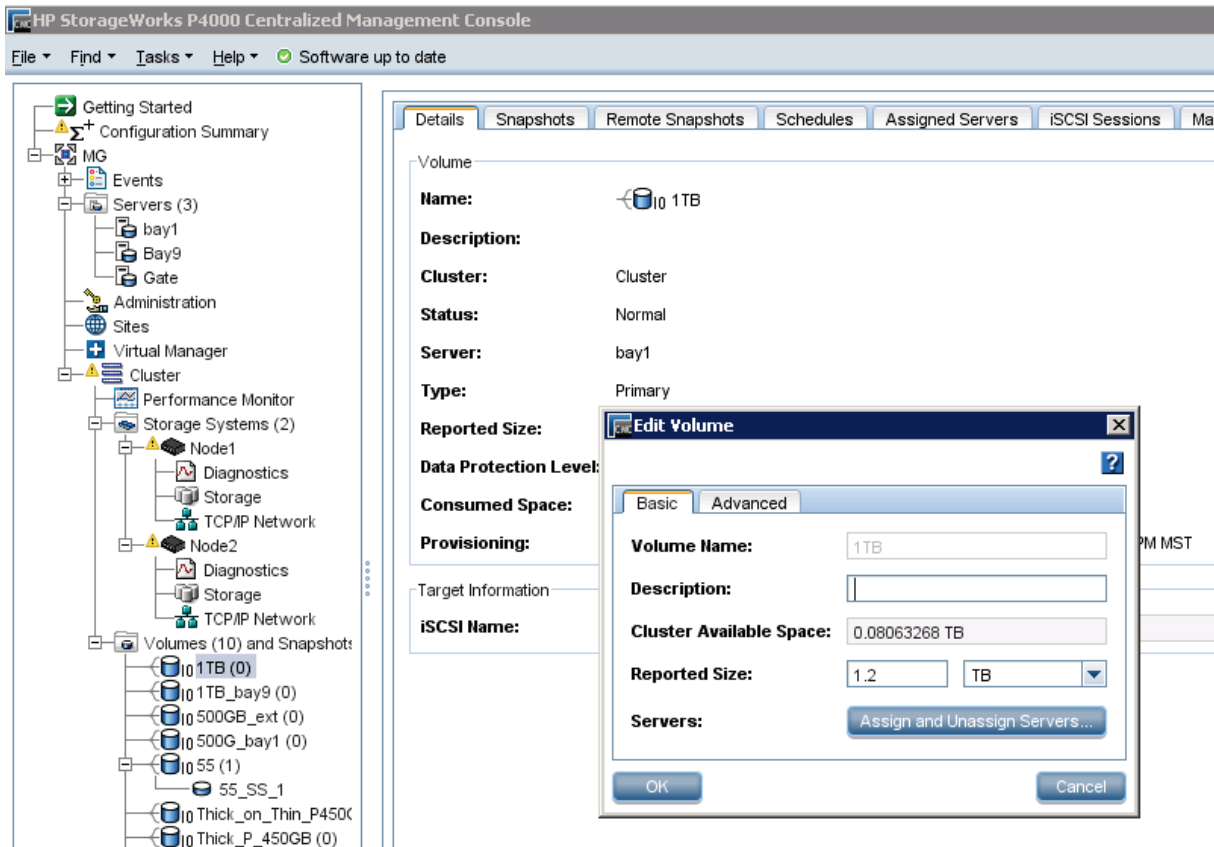
6. The reported size shows the expansion size immediately (figure 20).

Figure 20. The reported size has now changed and is immediately available


Expanding a VFMS datastore

1. Connect to vCenter and select one vSphere 5 host to which the iSCSI LUN is presented.
2. Click on the Configuration tab and select Storage Adapters.
3. Highlight the iSCSI software adapter and locate the volume in the Details section. The LUN is still showing 1 TB (figure 21).

Figure 21. vSphere Client shows iSCSI volume still as 1 TB prior to Rescan All
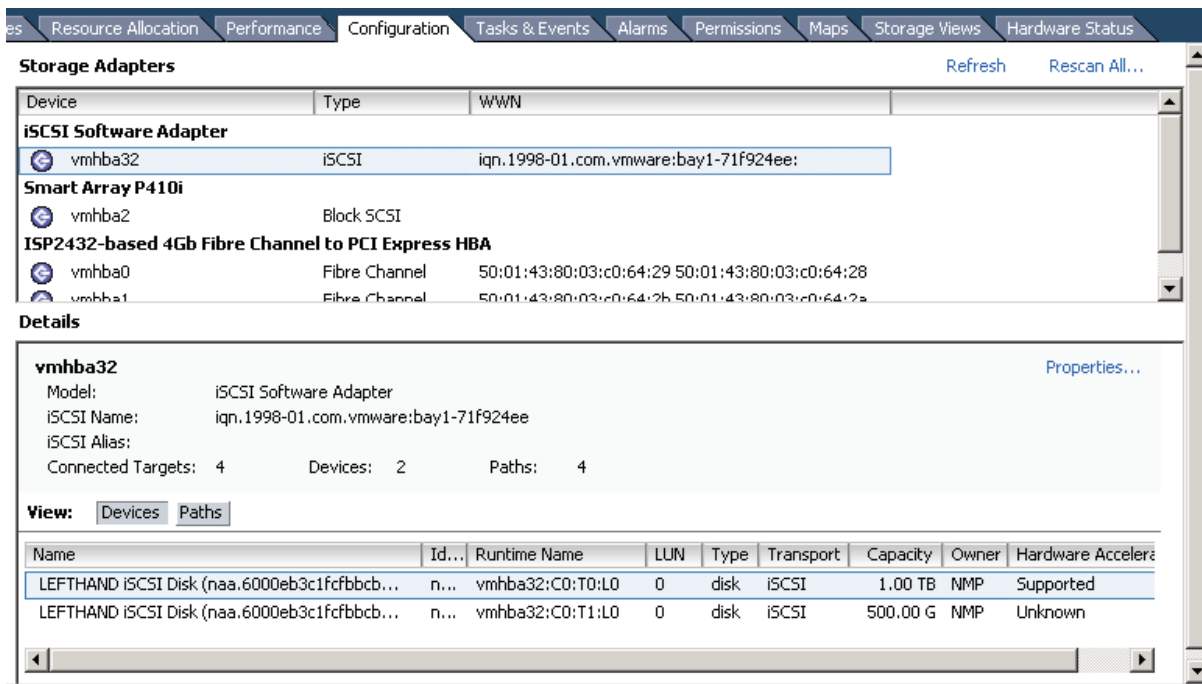
4. Right-click on the iSCSI software adapter and select Rescan.
5. The new size is reported immediately (figure 22).
6. Go to Storage and select the VMFS datastore which should be expanded.
7. Click on Properties and choose increase.
8. Select the same HP LeftHand Storage volume which reports the changed size.
9. The VMFS datastore will be increase by using the free space on the volume.

Figure 22. vSphere Client now reflects the change made to the iSCSI volume
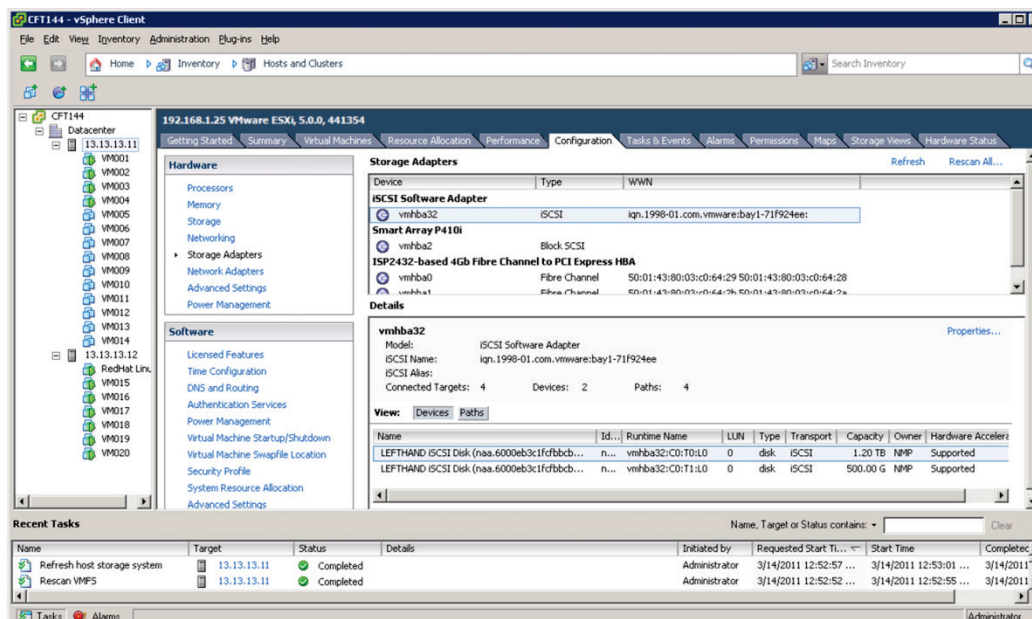
Snapshots, Remote Copy, and HP SmartClone volumes

HP LeftHand Application-managed Snapshots in vSphere 5

HP LeftHand OS version 9.5 expands application integration functionality to include volumes with VMware Virtual Machines. HP LeftHand Application-aware Snapshot Manager enables HP LeftHand snapshots for VMFS datastores, thereby creating application-consistent, point-in-time states for reliable recovery of VM states.

Without this integration option, snapshots are crash-consistent only. In flight cached data may not have been fully quiesced and may not represent an application recovery point. Now, with Application-managed Snapshots for vSphere 5 rolling back an entire snapshot or recovering single VMs from a snapshot typically restores to point when the VMs was quiesced by a VMware snapshots. The creation of this snapshot is fully integrated into the HP LeftHand Application-aware Snapshot Manager.

To use Application-managed Snapshots, the HP LeftHand Application-aware Snapshot Manager is installed on the vCenter Server and configured with the HP LeftHand Management Group's credentials to communicate with HP LeftHand Storage. Additional configuration is performed within the HP LeftHand Centralized Management Console: The server connection (or all server connections in a server cluster) is configured with the IP address of the vCenter Server as Controlling Server (figure 23).

Figure 23. Controlling Server IP Address setting in server connection

When creating the Application-managed Snapshot of a volume that contains a VMFS datastore, the CMC communicates with vCenter Server during the snapshot process. vCenter Server quiesces VMs and takes VMware Snapshots. VMware Tools

within each guest VM may quiesce applications inside the VMs during this process. Once all VMware Snapshots are taken, HP LeftHand Storage performs a snapshot of the volume which now contains VMware Snapshots on the HP LeftHand storage cluster. This can be seen when mounting the snapshot to a vSphere 5 host. After the snapshot of the VMFS datastore volume has been created on HP LeftHand Storage, the vSphere VM snapshots are removed.

Resignaturing snapshots

In order for vSphere 5 hosts to utilize SAN-based snapshots or SmartClones, a snapshot or a SmartClone of a volume is taken and presented to the vSphere hosts. The process of presenting a snapshot or SmartClone volume is the same as assigning volumes to a server connection in the CMC. To connect to the newly assigned snapshots or SmartClones, rescan the iSCSI software adapters on the vSphere hosts.

vSphere 5 supports three options when mounting a HP LeftHand snapshot or SmartClone volume that has VMFS data: Keep the existing signature, assign a new signature, or format the disk (figure 24).

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

Keeping the existing signature during a mount operation may generate an error, when there is conflict with an already mounted VMFS datastore on the host or cluster that has the same UUID (VMFS signature). If the original volume that contains the original VMFS datastore and its snapshot is mounted at the same time in a vSphere datacenter, force mounting the VMFS datastore from the snapshot LUN on this same vSphere 5 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 VMFS datastore.

By 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 during the VMFS datastore mount process, the Original UUID signature is provided along with Assign a new UUID. The resulting event task is resignature unresolved VMFS datastore.

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

**Figure 24.** Add Storage wizard when mounting a volume clone or snapshot
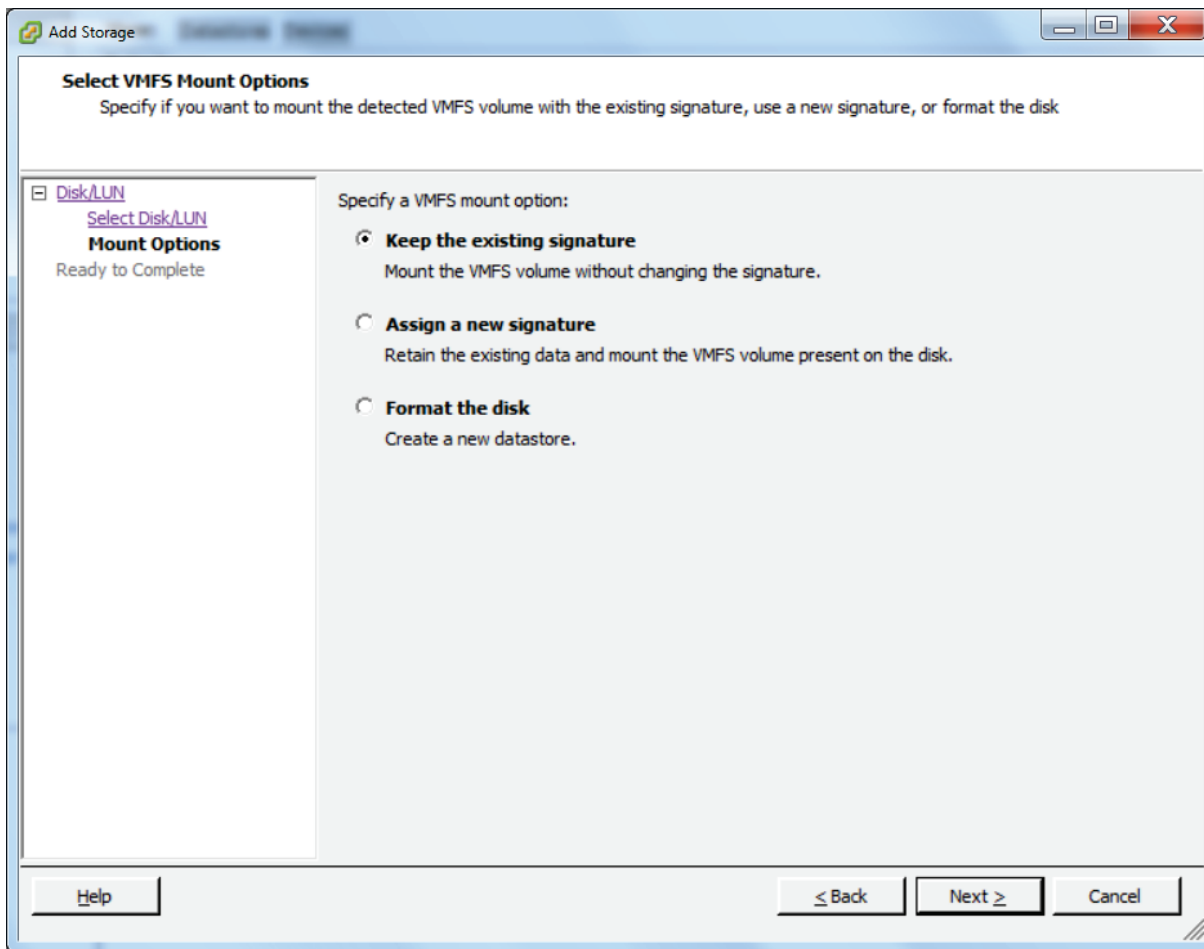
Note that HP LeftHand snapshots are persistent points in time representing data on the source volume at the time it was taken. As an extension to this feature, HP LeftHand allows snapshots to be mounted and any changes are written to in a temporary area. This data area is called a temporary space and 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 HP LeftHand storage cluster perspective, temporary space will be lost if HP LeftHand Storage is rebooted or manually removed. It 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 space-efficient new volume under a new name and new access permissions (figure 25). The resulting SmartClone volume contains the changes from the temporary space and point to the source snapshot. In order to preserve data on the volume which was written prior to conversion a SmartClone, the volume needs to be disconnected from the vSphere 5 host(s).

Alternatively, selecting the Delete Temp Space option from a context menu of the snapshot in the CMC will forcibly remove this data from HP LeftHand Storage. Once temporary space is deleted, it may not be recovered. It is a best practice to unmount the snapshot from any vSphere 5 hosts prior to deleting temporary space. If you wish to remount the snapshot, you can issue the Rescan All command in the vSphere Client prior to mounting the snapshot again.

Snapshots offer excellent ways for promoting data recovery.

Figure 25. Converting Temp Space into a SmartClone volume using the CMC
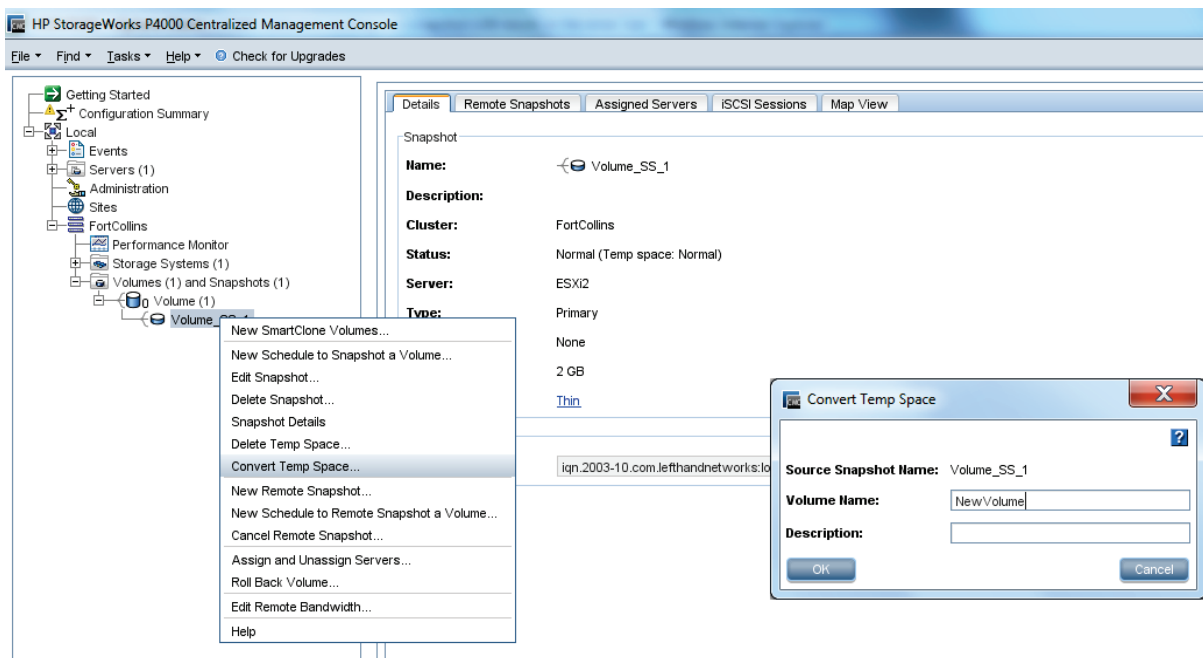
HP LeftHand Snapshots of Raw Devices

HP LeftHand Snapshots of vSphere 5 RDMs are supported in exactly the same way as for physical servers which either boot from SAN or access volumes on the HP LeftHand cluster.

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 to the volume 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 volumes without losing data across multiple vSphere hosts accessing the same volume.
- Snapshots
  VM snapshots are possible on RDM volumes in virtual compatibility mode. Combined with HP LeftHand Application-managed Snapshots, raw devices may be quiesced in the context of VMs using the RDM. Note that this is not supported with physical compatibility mode RDMs.

HP LeftHand Snapshots on VMFS datastores

HP LeftHand Snapshots are very useful in a vSphere 5 environment. All virtual machines stored on a single VMFS datastore (i.e., on the same volume) can be snapshot and rolled back to at any time. Moreover, HP LeftHand 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 26. HP LeftHand Application Integration Solution Pack–HP LeftHand OS version 9.5 required for vSphere 5
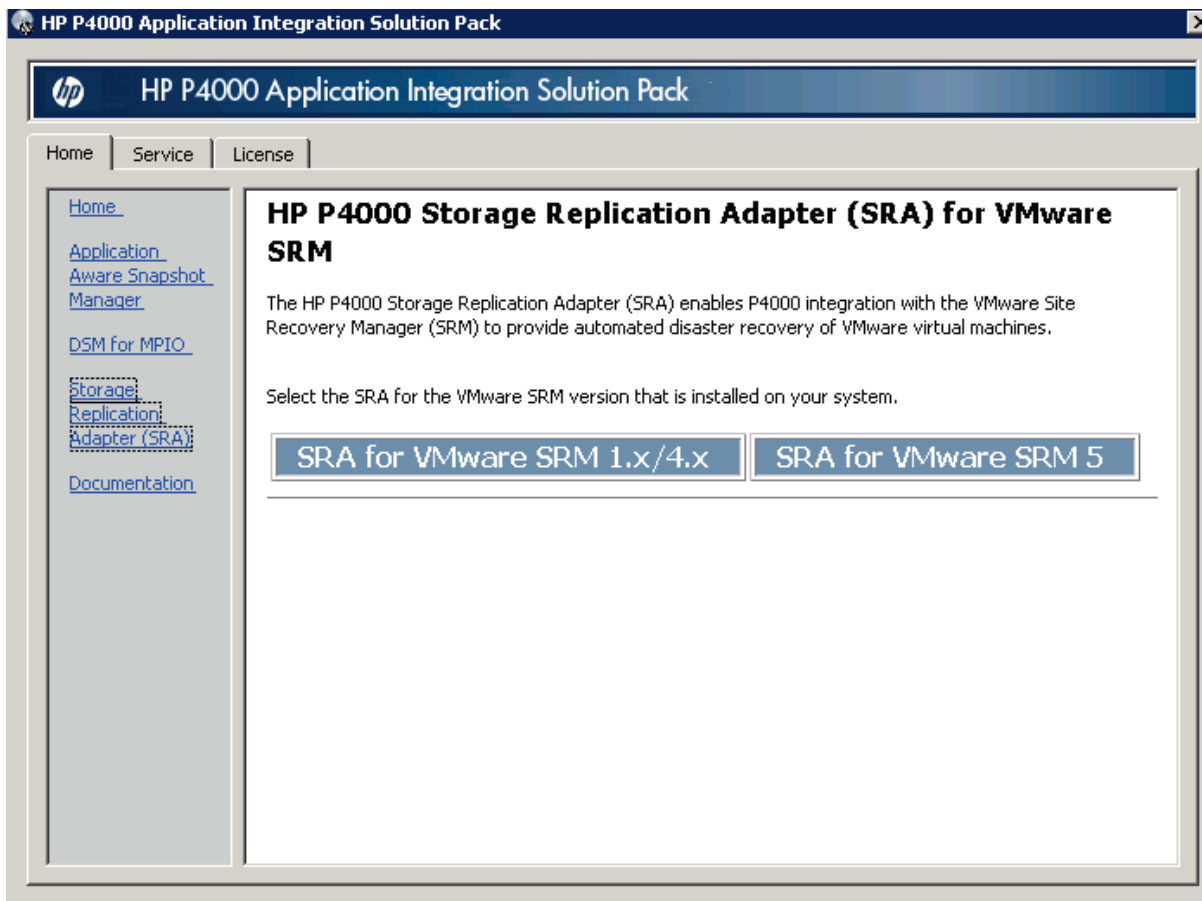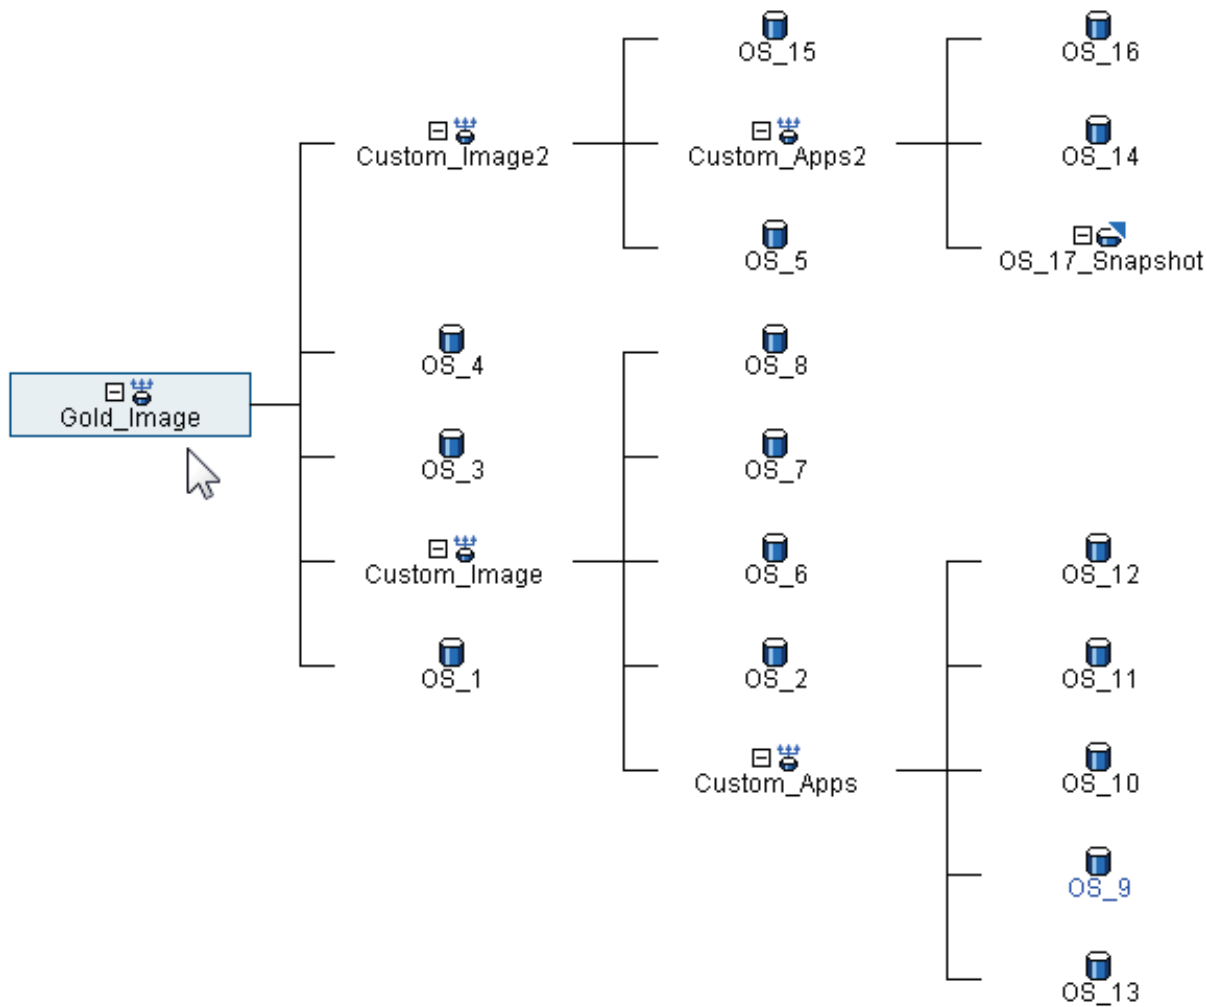

HP LeftHand Remote Copy volumes and SRM

Remote Copy replicates HP LeftHand 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 an SRA for VMware 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 HP LeftHand Centralized Management Console.

vSphere 5 and VMware SRM 5 requires HP LeftHand Application Integration Solution Pack. An option for SRA for VMware SRM 5 must be selected for vSphere 5 and SRM 5 support. Previous version supported SRM 4.1 in a single installer option. HP LeftHand OS 9.5 supports SRM 1.x/4.x or SRM 5 (figure 26). A best practice is to ensure that HP LeftHand software components are always up to date. Note that automated failback is now supported as part of the SRM 5 features.

HP LeftHand SmartClones

HP LeftHand SmartClone volumes may also be used in a vSphere 5 environment. Using HP LeftHand SmartClone Technology, all the virtual machines stored on a single volume can be cloned instantly without replicating data. Instead of a full copy of data for each SmartClone volume, each SmartClone volume is merely a set of pointers back to the original volume, thus consuming only a few MB of space, rather than the entire size of a volume. SmartClone volumes consume space for data that is written to the SmartClone volume 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 27. Map view of HP LeftHand SmartClone volumes in the CMC

SmartClone volumes can be used seamlessly with any other HP LeftHand 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 HP LeftHand cluster to copy them. Unlike temporarily mounted snapshots, SmartClones are meant to stay on the HP LeftHand cluster for a longer time and any delta data is persisted in the SmartClone volume. 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 stick to the same best practices as for thinly provisioned volumes, such as minimizing writes to VMFS datastores on a SmartClone volume, including 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 helps minimize data delta growth. 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.

Virtual machine placement and VAAI

Choosing datastores for virtual machines

More than one virtual machine can—and generally should—be stored on each VMFS datastore or volume. The choice of where to put virtual machines should be driven by space considerations, current VMFS 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 HP LeftHand volume; however, there are always tradeoffs. Consolidating VMs becomes mandatory as their numbers increase and larger configurations are demanded.

HP LeftHand 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 VMFS datastore. The more virtual machines on a VMFS datastore, the more I/O and contention there is for that HP LeftHand volume, particularly when multiple vSphere 5 hosts run VMs sourced on the same VMFS datastore. Refer to the section on Storage Distributed Resource Scheduler (Storage DRS) for the new vSphere 5 feature.

vSphere Storage APIs for Array Integration (VAAI)

VMware's VAAI enables HP LeftHand Storage, running HP LeftHand OS version 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 HP LeftHand Storage, where it is most efficient. Originally introduced in vSphere 4.1, vSphere 5.0 further extends these features. HP LeftHand Storage supports the original functional primitives at this time in HP LeftHand OS version 9.0 and 9.5.

**Figure 28.** VAAI framework
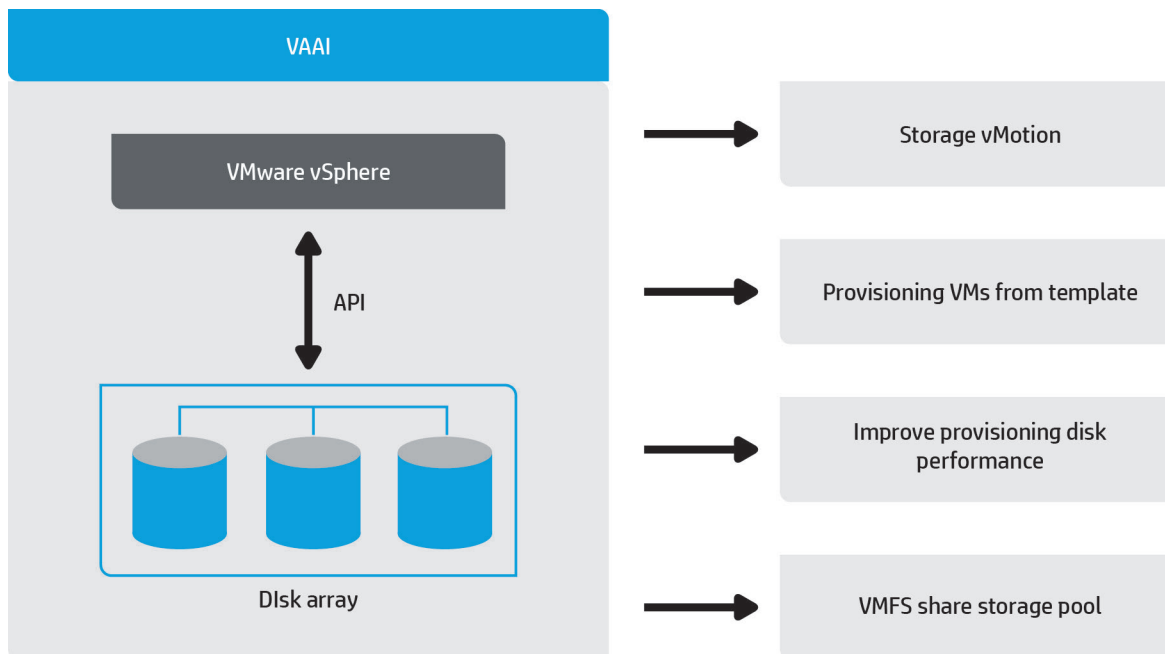
Currently supported VAAI offloading capabilities:

- Full copy
  Enables HP LeftHand Storage to make full copies of data within the HP LeftHand storage clusters. Previously, data movement would traverse from the storage arrays, occupying bandwidth to the vSphere host, utilize CPU resources processing the data movement, and again traverse back to the storage array.

- Block zeroing
  Enables HP LeftHand Storage to zero out large numbers of blocks with a single command. Previously repetitive zeros would traverse from the vSphere host to the storage arrays, 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 but required for Fault Tolerance.
- Hardware assisted locking (Atomic Test and Set)
  Provides a new efficient mechanism to protect metadata for VMFS in cluster deployments, improves scalability of large vSphere host farms sharing a VMFS datastore, and is more efficient than the previous SCSI reserve and release mechanism. HP LeftHand volumes are shared across multiple clustered vSphere hosts. 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 29. VAAI performance improvements with HP LeftHand Storage

Time to complete

Load on SAN

Load on Server
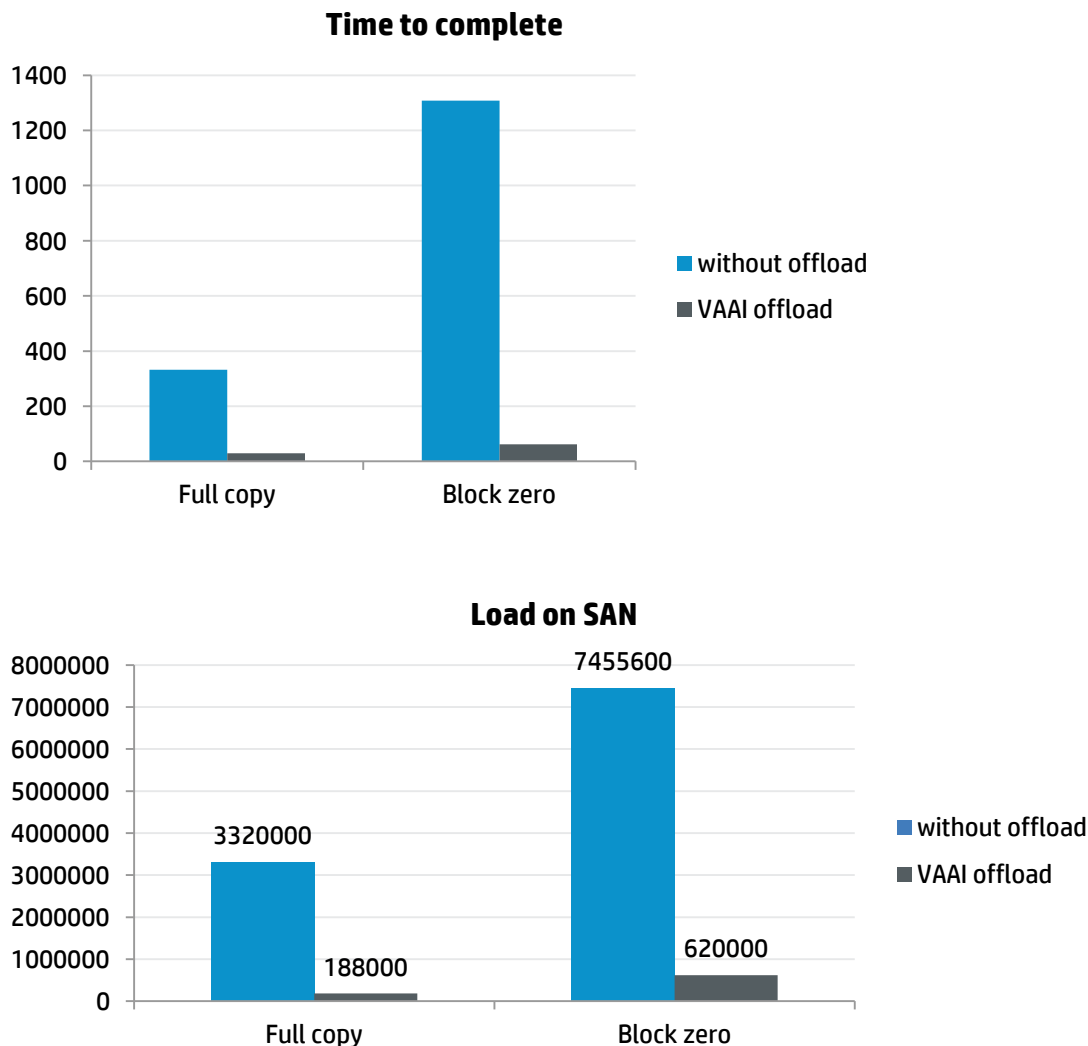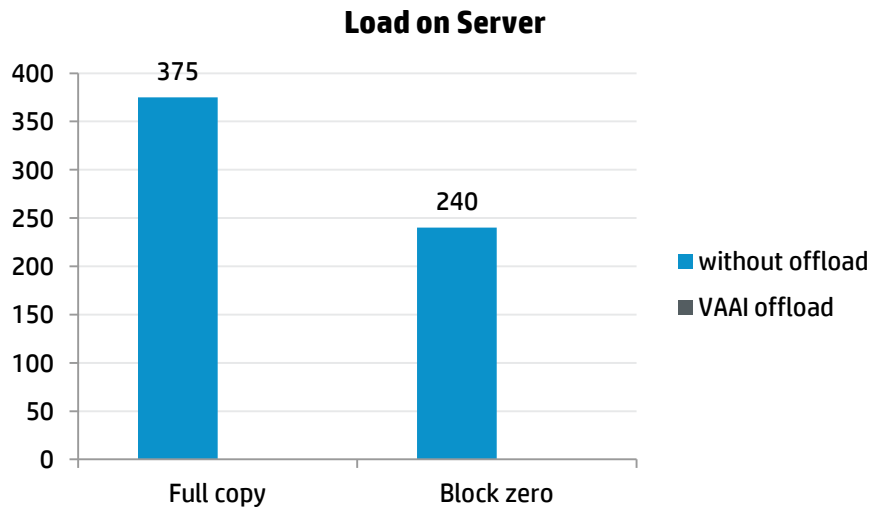
- Full copy clone Microsoft® Windows® 2008 R2 40 GB VM
  - 7x faster and 94 percent less load
- Block zero a 256 GB Virtual Machine Disk
  - 21x faster and 92 percent less load
- Hardware-assisted locking
  - 6x more VMs per volume/LUN

Figure 29 shows the level of improvements attributed to VAAI support with HP LeftHand Storage. Best practices previously suggested that up to 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 VAAI's hardware-assisted locking new best practice recommendations suggest that up to 48 virtual machines will function on a single volume. These figures are highly dependent upon I/O workload patterns as well as other variables. No separate software component or setting is needed to enable VAAI, VMFS datastore details will show if hardware acceleration is enabled.

Figure 30. Manual decision tree for determination of VM placement on HP LeftHand volumes
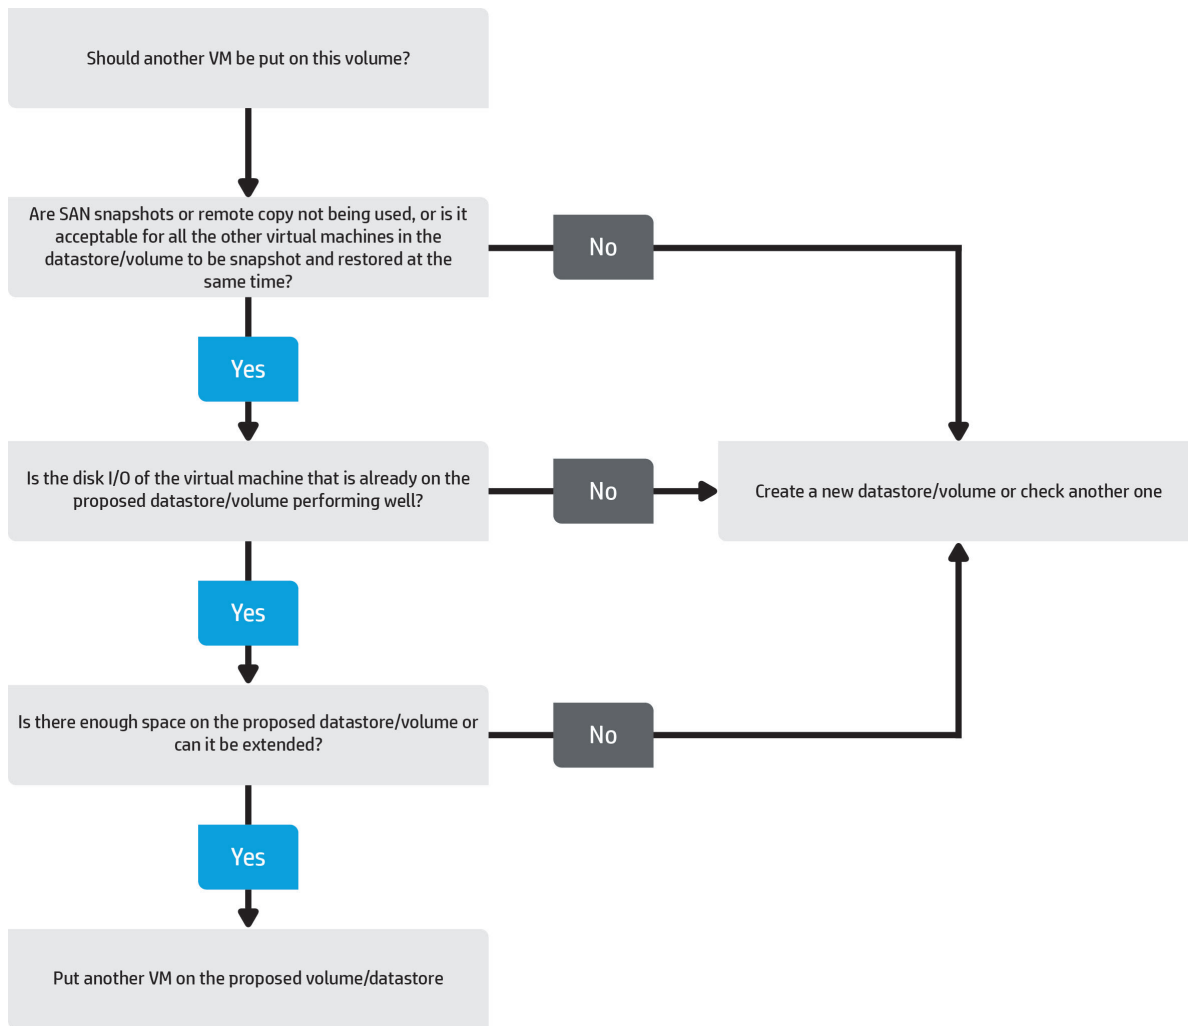
A decision tree for VM placement is shown in figure 30. With or without VAAI, virtual machine placement will impact storage I/O.

Best practices

• Use at least two Gigabit network adapters on vSphere 5 hosts for improved performance and for failover of the iSCSI connections. 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 bond on HP storage nodes is easily configured in the CMC by selecting two active links and enabling a bond: Adaptive load balancing (ALB) is the most common and easiest method used on HP LeftHand storage nodes. From the storage node, ALB bonding supports sending and receiving (for more than one iSCSI Initiator accessing the HP LeftHand storage node) through both interfaces. 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 sending and receiving data from both adapters.

On vSphere 5 hosts, native storage multi-pathing (MPIO) is using multiple VMkernel ports (each mapped to a single active adapter) to achieve increased bandwidth and resiliency. The preferred multi-path policy is round robin. All paths to volumes on HP LeftHand Storage should show up as active (I/O).

• VMkernel ports with iSCSI bindings should be separate from the management and virtual networks used by virtual machines. If enough network adapters 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.

• Hardware-based flow control is recommended for HP LeftHand Storage on all network interfaces and switches.

• The VIP load-balancing feature of HP LeftHand Storage allows iSCSI sessions during volume login to be redirected to the storage 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 HP LeftHand Storage. This setting is enabled by default; leaving it unchanged is recommended.

• 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 enhance the performance of vSphere 5 servers and balance 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 and HP SPOCK.

• Because HP LeftHand 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. Virtual machines that can be backed up and restored together should share the same VFMS datastore.

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 HP LeftHand 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 HP LeftHand 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 you have not correctly configured server connections and volume access in HP LeftHand Storage. Also, you should verify all network configurations. Verify on the iSCSI network that IP addresses can be pinged. Please refer to Discovery of the first iSCSI volume.

Is virtual IP load balancing supported in vSphere 5?

In vSphere 5, iSCSI software adapter and iSCSI HBAs are supported by HP LeftHand Virtual IP Load Balancing. As a best practice, load balancing should be enabled in the CMC on the server connections (or server clusters) of all vSphere 5 hosts. This is enabled by default.

What initiator should I use for a virtual machine's boot volume?

There are many ways to connect and present an iSCSI volume to a virtual machine on a vSphere 5 host. These include using the iSCSI software adapter, iSCSI HBAs, or the iSCSI initiator provided by guest operating systems. Please check the compatibility listing for HP LeftHand Storage on HP SPOCK for information on iSCSI initiator support.

For VMFS datastores containing a virtual machine's definitions and virtual disk files, you need to use the iSCSI software adapter or an iSCSI HBA on the vSphere 5 host. You may use either adapter. Both give you full vSphere 5 functionality, and both are supported equally. Using iSCSI HBAs has the advantage of offloading iSCSI processing from the vSphere 5 host.

Note that although vSphere 5 supports iSCSI Boot Firmware Table (iBFT), it is currently not supported or listed in HP SPOCK or the HP LeftHand Compatibility Matrix.

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

Additional volumes are used for storing application data; typically for separation from a virtual machine's boot volume. Specifically, as a best practice, database and log volumes for many applications require separate volumes. These should be presented as either RDMs through your vSphere 5 hosts or connected as iSCSI disks directly through the iSCSI initiator software of the virtual machine's guest operating system (direct iSCSI).

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 operating system. To use the guest operating system initiator successfully, follow these guidelines:

- Make sure the guest operating system initiator is supported by HP LeftHand Storage. Please check the compatibility listing for HP LeftHand Storage on HP SPOCK.
- For optimal performance and failover, the virtual machine network used for iSCSI by the guest's iSCSI initiator should have at least dual Gigabit links and be separate from other virtual networks, including VMkernel instances for vMotion, Fault Tolerance logging, host management, and virtual machine public networks.
- Make sure that the guest operating system is using the para-virtualized VMXNET 3 network adapter driver which is typically installed through VMware Tools on supported guests.
- When using direct iSCSI, make sure that the virtual machine will not be used in combination with VMware SRM. SRM does not work with volumes connected by guest iSCSI Initiators.

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

Refer to the section Choosing datastores and volumes for virtual machines in this document.

Are jumbo frames supported?

Jumbo frames (Ethernet frames configured larger than the typical 1,500 bytes—i.e., up to 9,000 bytes) are supported by HP LeftHand Storage. In order for jumbo frames to be effective, they must be enabled and configured end to end with the same frame size, including all network adapters and switches.

For most Gigabit iSCSI implementations jumbo frames are not required and do not increase performance, whereas 10 Gigabit iSCSI implementations may see benefits from jumbo frames for sequential workload.

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

The maximum volume size supported by vSphere 5 is 64 TB. Note that VMFS-5 is required. If upgrading, ensure that VMFS datastores are running the newer VMFS-5 format. The maximum VMDK size has not changed in vSphere 5 and is still 2 TB.

What version of HP LeftHand OS supports vSphere 5?

All HP LeftHand Storage platforms (running HP LeftHand OS version 9 .0 and above) support vSphere 5.

For SRM 5, HP LeftHand OS version 9.5 is required with the new SRA 5.0 plugin.

For more information

HP LeftHand VSA, hp.com/go/vsa

HP Single Point of Connectivity Knowledge, hp.com/storage/spock

HP LeftHand Compatibility Matrix, hp.com/go/lefthandcompatibility

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

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

HP BladeSystem, hp.com/go/bladesystem

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

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

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

VMware fault tolerance demo using HP LeftHand Storage,
http://h30431.www3.hp.com/?fr_story=217cb39a5a037d7a3050746dc4838b02d5700d31&rf=bm

P4000 VMware VAAI Whiteboard, http://www.youtube.com/watch?v=pkhJRwW-jXc&

To learn more about HP LeftHand Storage, visit hp.com/go/lefthand

Get connected