



Peer Motion on HP LeftHand Storage

Table of contents

An Introduction to HP LeftHand Storage.....	2
HP Delivers Federated Storage Technology with Peer Motion.....	3
Peer Motion in HP LeftHand Storage	3
Peer Motion Benefits.....	3
What is Volume Migration?	4
Benefits of Volume Migration.....	4
When to Use Volume Migrations	4
How to Perform a Volume Migration	5
What is Cluster Swap?	8
Benefits of Cluster Swap	8
When to use Cluster Swap.....	8
How to Perform a Cluster Swap.....	9
Cluster Swap Tasks	9
Working Example of Cluster Swap	10
Cluster Swap Example.....	10
Peer Motion and LeftHand VSA Software.....	13
VSA annexing-use case	13
Benefits of VSA annexing.....	13
Best Practices for Peer Motion	14
Summary	15
Frequently asked questions	16

An Introduction to HP LeftHand Storage

HP LeftHand Storage is a highly scalable, highly available iSCSI storage solution. LeftHand's scale-out architecture allows customers to consolidate storage nodes into pools of shared storage (clusters). All available capacity and performance is aggregated and available to every volume in the cluster.

The LeftHand operating system (LeftHand OS) has built-in disaster recovery and high availability clustering features. HP combines the intelligence of the LeftHand OS with industry-leading ProLiant servers to deliver a full-featured all-inclusive enterprise iSCSI SAN solution.

Enterprise features of the LeftHand OS are:

- Storage Clustering
- Network RAID
- Thin Provisioning
- Snapshots and Clones
- Disaster Recovery
- Centralized Management
- Multi-Site Stretched Cluster Replication

The LeftHand OS provides all features and management services to:

- Aggregate physical storage nodes into a scalable pool of clustered resources
- Turn raw disk capacity into a fully virtualized, iSCSI SAN solution
- Provide seamless data mobility services
- Scale all available capacity and performance linearly within a cluster

LeftHand storage clustering technology provides a grow-on-demand architecture where storage systems can be added to the cluster as needed without application disruption. Clustering features within the LeftHand OS include, but are not limited to:

- Addition and removal of storage systems within the cluster
- Swapping of all storage systems within a cluster (cluster swap)
- Movement of volumes between clusters
- High availability and fault tolerant characteristics of the storage cluster
- Automatic restriping of data within storage volumes
- Automatic management of iSCSI storage sessions

An industry first feature of the LeftHand OS is its Network RAID technology. Network RAID delivers volume-level data protection, which supports high availability and fault tolerance of iSCSI storage volumes. This feature also provides built-in synchronous data protection for standard and multi-site storage clusters.

The HP LeftHand Centralized Management Console (CMC) provides SAN administrators with a "single pane of glass" dashboard to control and monitor all SAN activities including real-time execution of Peer Motion features such as volume migration and cluster swap.

HP Delivers Federated Storage Technology with Peer Motion

The concept of federated storage expands the benefits of shared, virtualized storage beyond a single disk array to span the entire datacenter and beyond. Storage federation eliminates the constraints associated with traditional architectures where data is trapped within the silo of a single, monolithic storage array. HP LeftHand federated storage allows administrators to manage workloads dynamically, providing the flexibility demanded by today's virtualized environments. Compared to traditional architectures, HP LeftHand federated storage offers the agility to handle performance bottlenecks and capacity crunches associated with the rapidly changing challenges of today's IT infrastructures. Storage federation from HP allows administrators to pool resources across the datacenter to optimize ROI from storage technology and drive maximum efficiency from storage resources.

Peer Motion in HP LeftHand Storage

Peer Motion is a data migration technology which offers LeftHand customers two simple and flexible methods of moving data volumes within their LeftHand environment. The two methods are:

- Volume migration
- Cluster swap

Volume migration and cluster swap employ peer-to-peer storage federation concepts to seamlessly move application workloads between tiers of data storage or storage clusters within a LeftHand environment. Peer Motion in HP LeftHand supports all LeftHand platforms ranging from the VSA to the midrange and high-performance storage systems.

Peer Motion enables dynamic resource leveling of LeftHand storage volumes without storage downtime or maintenance window outage. Resource leveling for LeftHand systems is quick, easy, and executed within a few simple steps thus helping to reduce the need for long-term storage planning and complicated budget forecasting.

Peer Motion Benefits

- Offers data and workload movement between storage clusters to address performance and capacity requirements
- Enables seamless LeftHand storage technology refreshes, eliminating downtime and/or service interruption during migration/upgrade activities
- Promotes federated asset management—by moving data from retiring storage arrays to new storage systems non disruptively
- Allows flexible migration of data volumes between multi-site clusters to align with the growing demand for cloud based storage
- Enables clients to transparently move application workloads between storage tiers in virtualized and cloud computing environments
- Increases return on investment by combining HP thin provisioning and volume migration technology

Peer Motion supports volume migration for “full” and “thinly” provisioned LeftHand volumes. Peer Motion preserves the provisioning method (full or thin) associated with existing volumes during data migration services and supports migration of data within:

- Volumes
- Snapshots
- SmartClone volumes

What is Volume Migration?

HP's Peer Motion volume migration service is the online movement of any LeftHand volume between different tiers of storage systems and/or clusters. Volume migration enables the movement of data volumes between LeftHand storage clusters while preserving application access to the storage. iSCSI sessions are automatically redirected to the new cluster when the migration is initiated and application access is maintained throughout the process. The data "restripe rate" (rate at which data is moved to the new cluster) can be managed by the administrator to control the resources allocated to data migration and limit any potential impact on performance. Peer Motion volume migration provides customers a scalable tiered solution to align virtualized and cloud-based compute requirements with customer's most appropriate and cost-effective storage resources.

Benefits of Volume Migration

Peer Motion volume migration allows administrators to rapidly respond to changing workload and capacity requirements by seamlessly moving an application to a tier of storage with a higher performance profile or a cluster with more available capacity. In contrast to legacy storage designs, these migrations can be performed online, without downtime or complex planning activities.

Combined with the LeftHand scale-out architecture, volume migration offers small, midsize, and enterprise customers dynamic resource leveling capabilities across LeftHand platforms. These capabilities offer users a scalable means of matching shifting application demands to the appropriate performance and capacity resources of federated LeftHand storage.

Volume Migration is an HP Peer Motion technology built in to every LeftHand storage system to deliver:

- Ease-of-use—dynamic volume tiering
- Flexibility—on-line movement of volumes between clusters
- Scalability—rebalancing of data volumes to higher performance or capacity storage tier and vice versa
- Agility—resilient storage infrastructure; as easy as choosing the site, cluster and platform to tier to
- High availability—offers continuous online volume access for both standard and multi-site LeftHand clusters

When to Use Volume Migrations

Practical use cases for resource leveling utilizing volume migration features are described and categorized below.

Capacity-use case

- Movement of volumes from an over-provisioned LeftHand storage cluster to a higher capacity LeftHand peer cluster
- Movement of data volumes to larger capacity storage systems, such as LeftHand P4800 storage; thus enabling linear scale-out of compute and storage within HP BladeSystem environment

Performance-use case

- Alignment of data volumes to accommodate the input/output (I/O) usage requirements of applications; movement of data to any LeftHand cluster capable of higher performance as needed
- Movement of data volumes from mid-line or serial attached SCSI (SAS) to solid-state disk (SSD) storage systems; yielding higher input/output per second (IOPS) and smaller data center footprint—lower power, cooling, and real estate—while delivering higher I/O storage system performance

How to Perform a Volume Migration

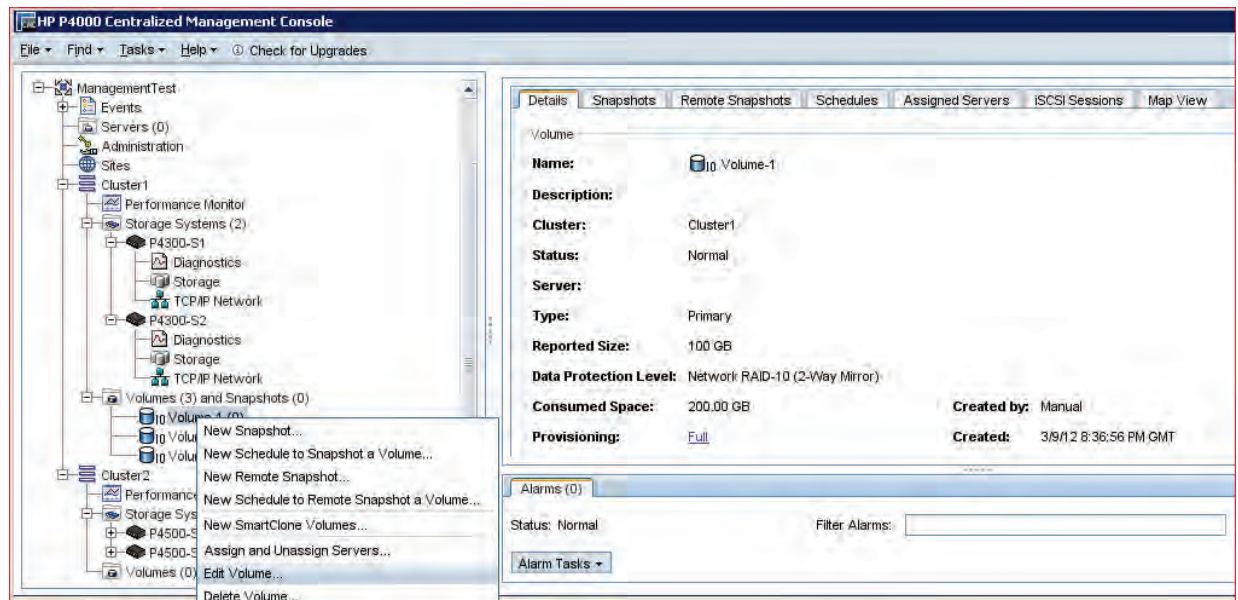
Volume migration in HP LeftHand Storage is the process of moving a volume from one LeftHand cluster to another LeftHand cluster—a Peer Motion federated concept. The movement of a volume between clusters in a LeftHand system requires adding the virtual IP address (VIP) of the target cluster to the list of target portals on each iSCSI initiator with an iSCSI connection to the volume. For Microsoft Windows, the virtual IP address should be added to the Microsoft iSCSI Software Initiator portal list prior to moving volumes between clusters to make certain volume migration services are transparent and non-disruptive to applications on to-be-migrated volumes.

The target portal list on the iSCSI initiator should include at least two VIP addresses:

- Virtual IP address of source cluster that currently owns the volume(s) (Cluster1)
- Virtual IP address of the target cluster where the volumes are migrated too (Cluster2)

LeftHand volume migration feature is initiated via the “Edit Volume” menu task in the CMC. This feature can be executed with just a few clicks after selecting the volume that will be migrated. Figure 1 illustrates the ease of using the graphical user interface of the CMC to select a volume (Volume-1) to migrate from Cluster1 to Cluster2 within the same management group.

Figure 1: Edit Volume Menu

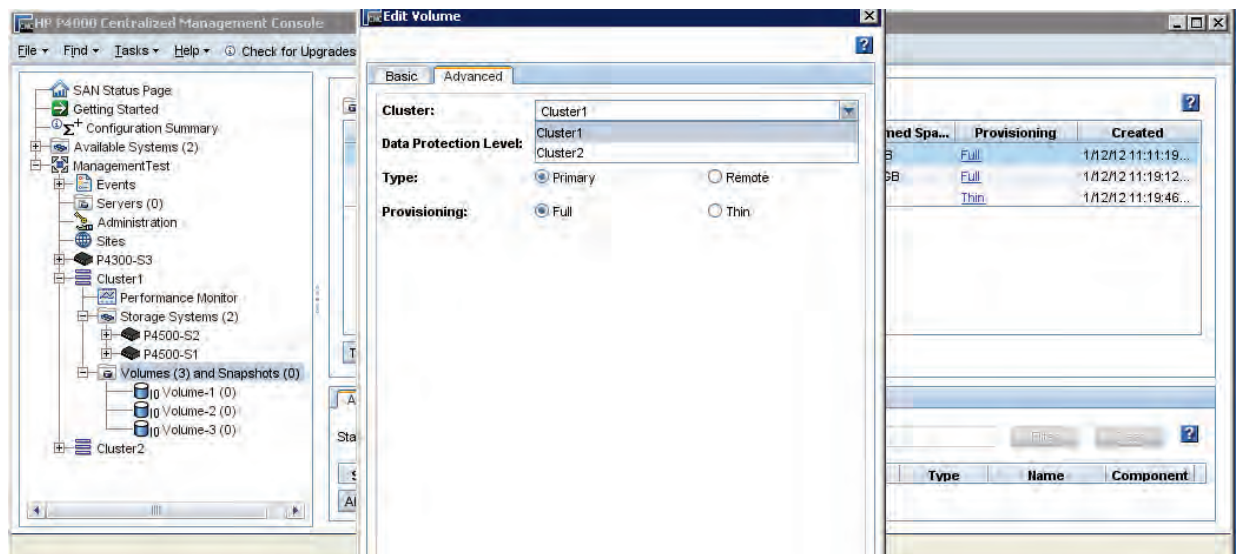


Listed below are the necessary steps to move a volume from one cluster (representing one tier of storage) to another cluster using the LeftHand CMC. Moving a volume between clusters is executed in as few as six easy steps.

Step 1: Select the volume of choice from the right navigation window as illustrated in Figure 1 above.

Step 2: Right click, and then select Edit Volume from the menu option. The Edit Volume menu popup window is displayed as shown in Figure 2 below.

Figure 2: Edit Volume Menu

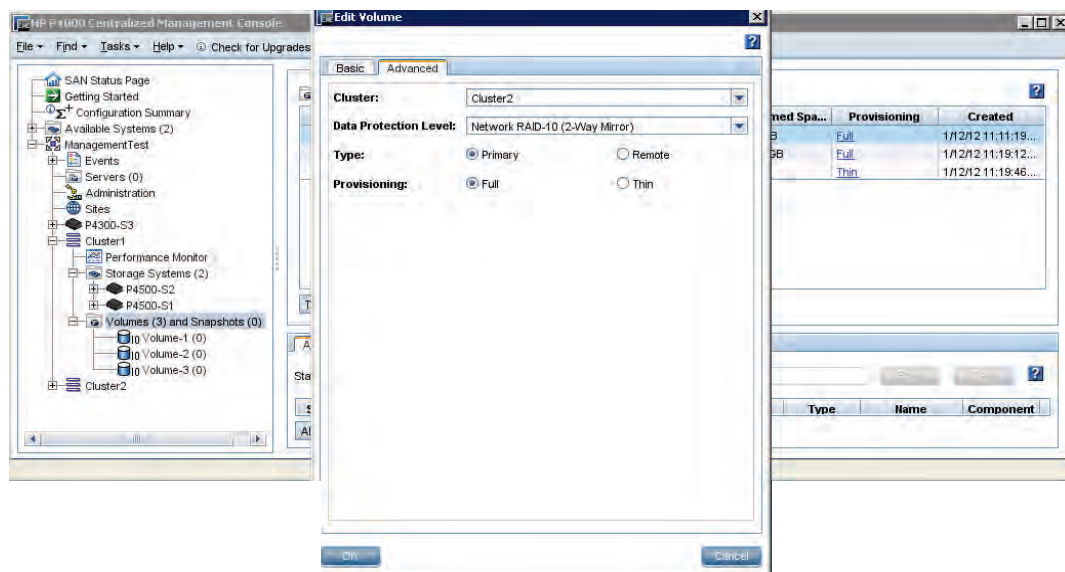


Step 3: Select the tab titled Advanced from the Edit Volume popup window as shown Figure 2.

Step 4: Next select the Cluster drop-down list. The list includes all available storage clusters within the same management group as shown in Figure 2.

Step 5: Select the name of the cluster (i.e. Cluster2) where the volume will be migrated to, then click OK as illustrated in Figure 3.

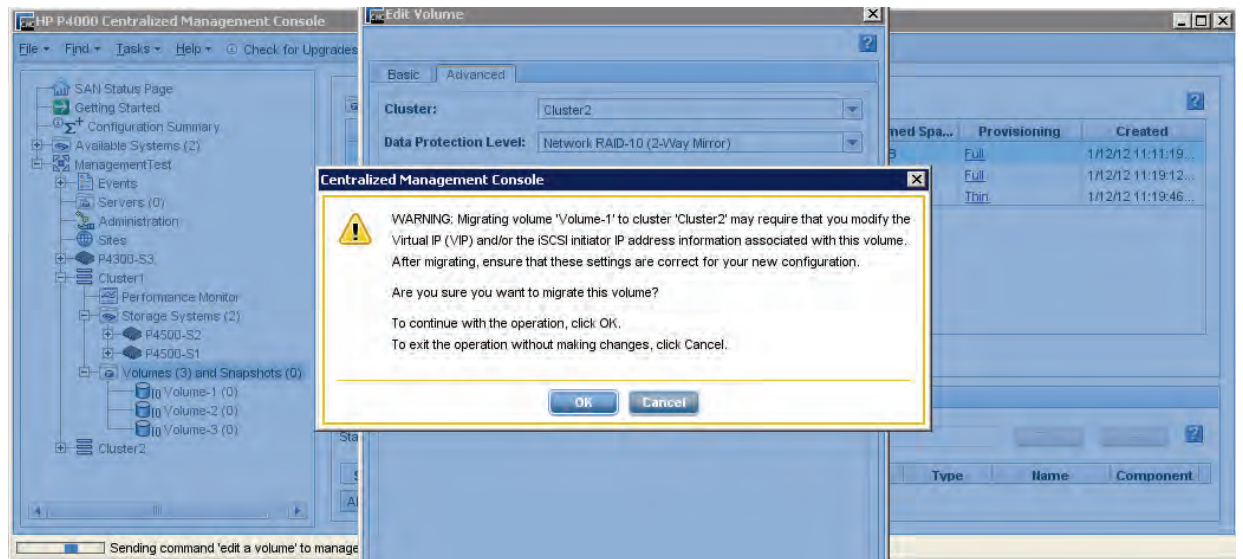
Figure 3: Edit Volume Menu Advance Tab



Step 6: A popup window with a warning message¹ appears; click OK again within this window to complete the last CMC task for this volume migration process. If necessary, repeat steps 1–6 above to migrate additional volumes to the target cluster.

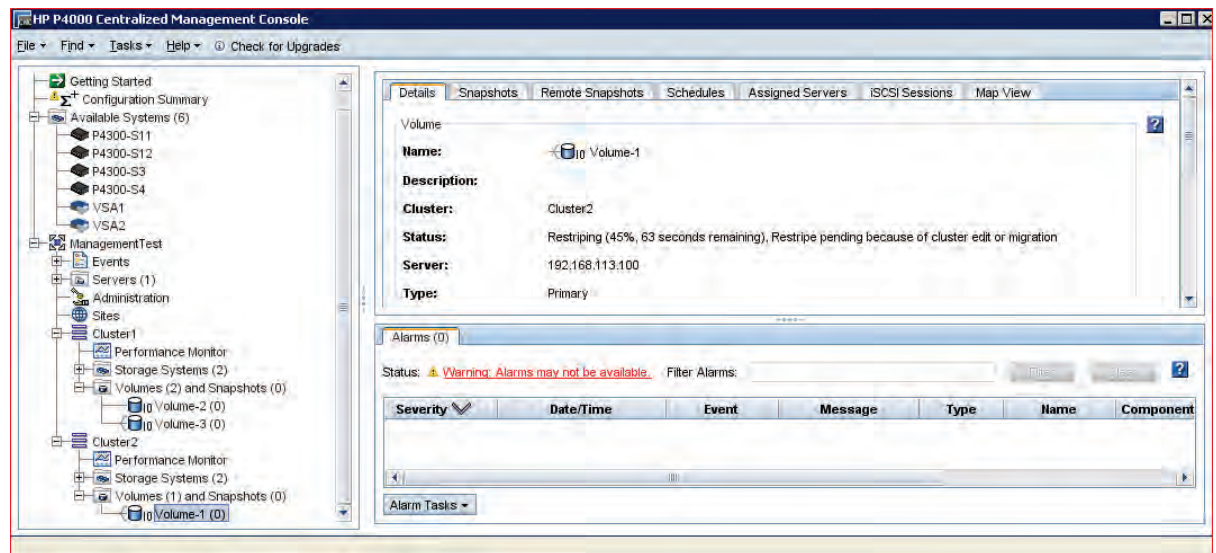
¹ **Warning Message:** After, the volume migration has completed; check client iSCSI setting to make certain the data volumes are properly mapped to appropriate cluster that owns the migrated volume.

Figure 4: Edit Volume Menu with Warning



After selecting OK, in Step 6, the program returns to the main CMC. The restriping status of the migrated volume is shown in Figure 5 below. The Details Tab of a selected volume provides status information.

Figure 5: Volume Property Menu



What is Cluster Swap?

Cluster swap enables the on-line replacement of all storage systems in a LeftHand cluster with another set of storage systems. A cluster swap replaces the physical components of a LeftHand cluster while maintaining the logical cluster and the data therein. Cluster swap offers LeftHand customers non-disruptive, built-in, and flexible data mobility services. With Cluster Swap, technology refresh of LeftHand storage can be accomplished with just a few simple steps and no downtime, maintenance windows or complex planning.

Benefits of Cluster Swap

A highlighted benefit of cluster swap is the seamless movement of data volumes without disruption to users and/or application processing. Volume migrations are performed without required changes to network mount points or changes to logical drive mappings, eliminating the need for end-user reconfiguration.

Other benefits are:

- Migration of data volumes from any LeftHand system with a few simple clicks within the CMC
- Reduce complex planning and eliminate application I/O downtime during storage system consolidations and technology refresh initiatives
- Right size LeftHand storage clusters at-will—while preserving application service levels

When to use Cluster Swap

Use the cluster swap feature to seamlessly swap out all LeftHand nodes in a cluster at once. Common use cases for cluster swap are:

- Refresh storage technology of any LeftHand platform
- Upgrade all storage systems in a cluster to a higher performing or capacity storage tier
- Downgrade all storage systems in a cluster to a less expensive storage tier

How to Perform a Cluster Swap

Cluster swap, a Peer Motion feature, can be executed in as few as six clicks starting with the Edit Cluster menu option within the CMC, after selecting the cluster of choice within the appropriate management group. The cluster swap operation should be used by customers when they want to upgrade and/or replace all storage systems within a cluster with a new set of systems. Volumes striped across all storage systems within the selected cluster are seamlessly migrated to new storage systems added during the cluster swap operation. The LeftHand OS maintains data volume consistency while restriping data blocks from current storage systems to the new.

At the time the cluster swap is initiated, all iSCSI sessions as well as the cluster Virtual IP address (VIP) are migrated to the new nodes. Client read and writes operations are handled transparently by the LeftHand OS during the swap operation. Read I/Os are processed by the old or new storage nodes depending on whether the particular data blocks have been migrated; write I/Os are processed by newly added storage nodes. The cluster swap operation automatically removes the “old” storage systems from the cluster, thus allowing these systems to be repurposed immediately after the cluster swap operation finishes.

Steps below describe CMC tasks required to execute a cluster swap.

Cluster Swap Tasks

1. Select the cluster of choice from the navigation window as shown earlier.
2. Click Cluster Tasks, next select Edit Cluster and then select Swap Storage Systems. A Swap Cluster Storage Systems popup window appears.
3. Select all the storage systems listed in the Swap Cluster Storage Systems popup window; next click the Remove Storage Systems button. The selected storage systems are removed from the list.
4. Next click Add Storage Systems. The Add Storage Systems window opens with all the available storage systems in the management group.
5. Select the storage systems to swap into the cluster and click OK.
6. Next click OK again within this window to complete the swap.

The cluster swap operation varies in time and has a dependency upon but not limited to:

- Network RAID level of each data volume migrated
- Amount of data required to be restriped
- I/O workload—number of users actively reading and writing data from volumes
- Network throughput—available bandwidth for iSCSI traffic within the local storage network as well as interconnects for multi-site cluster swaps
- Bandwidth priority value set within the management group. This setting is used to determine the maximum bandwidth allocated to the cluster swap versus application processes; this may be adjusted at any time to speed up or slow down the cluster swap operation.

Adding storage systems using cluster swap feature is not the same as adding or removing new storage systems to expand or reduce the storage capacity of a cluster. Cluster swap is used as a means to replace all LeftHand storage systems residing in the cluster at one time using the remove and add functions of the swap storage systems menu option.

NOTE: Adding new storage systems to a cluster swap causes a restripe of all data volumes in that cluster. The restripe process is transparent to the end user and may take several hours or days depending on the amount of data required to restripe. Swapped storage systems cannot be removed from the management group until the restripe process is complete

Working Example of Cluster Swap

To illustrate how Cluster Swap works we will walk through an example of moving all data volumes from a LeftHand P4300 storage cluster to a higher performance cluster. All storage nodes in this example are running the latest version of LeftHand software and reside in the same management group. The name of the management group is "ManagementTest".

In this example, the two P4300 storage systems are members of a standard cluster configuration. The cluster's name is "Cluster1". A failover manager is used in this 2-node cluster and has been added to the "ManagementTest" group to serve the role of quorum tie-breaking manager.

The following figures illustrate the replacement of two P4300 storage systems in "Cluster1" with two P4500 systems added to "Cluster1" (Swapped). This example uses two nodes; however, up to 16 nodes can be swapped within a single cluster. For information about creating clusters and volumes, see the HP LeftHand Storage Solution User Manual.

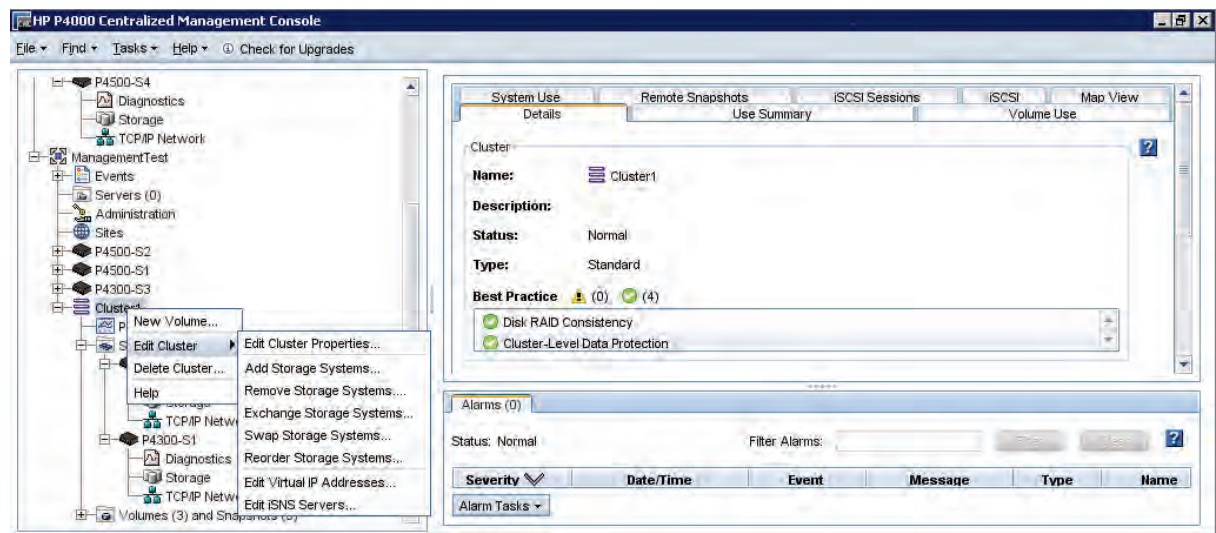
The example below assumes the HP LeftHand SAN administrator is knowledgeable of LeftHand software and design standards and has experience using the CMC.

First step is to login to the management group and make certain all storage nodes used during the cluster swap reside in the "ManagementTest" group as depicted in the following use case.

Cluster Swap Example

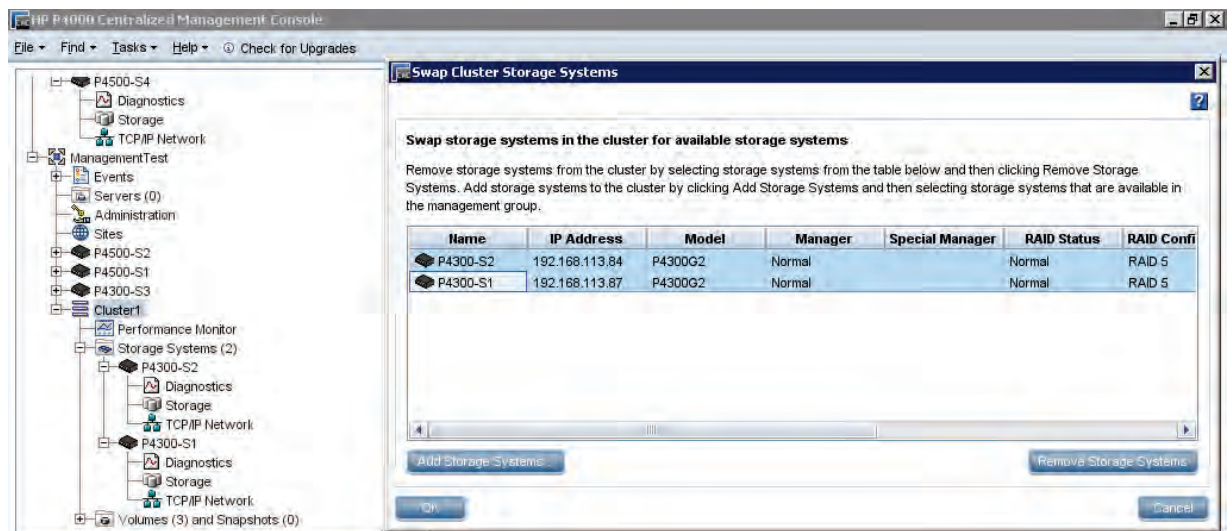
Step 1: Using the CMC, expand the tree and browse down the left-hand side of the CMC to the cluster named "Cluster1". Select Cluster1, Right click on Cluster1, Select "Edit Cluster" from the pop-up window. After selecting Edit Cluster, select "Swap Storage Systems" from the popup menu as shown in Figure 6 below.

Figure 6: Edit Cluster Task Window



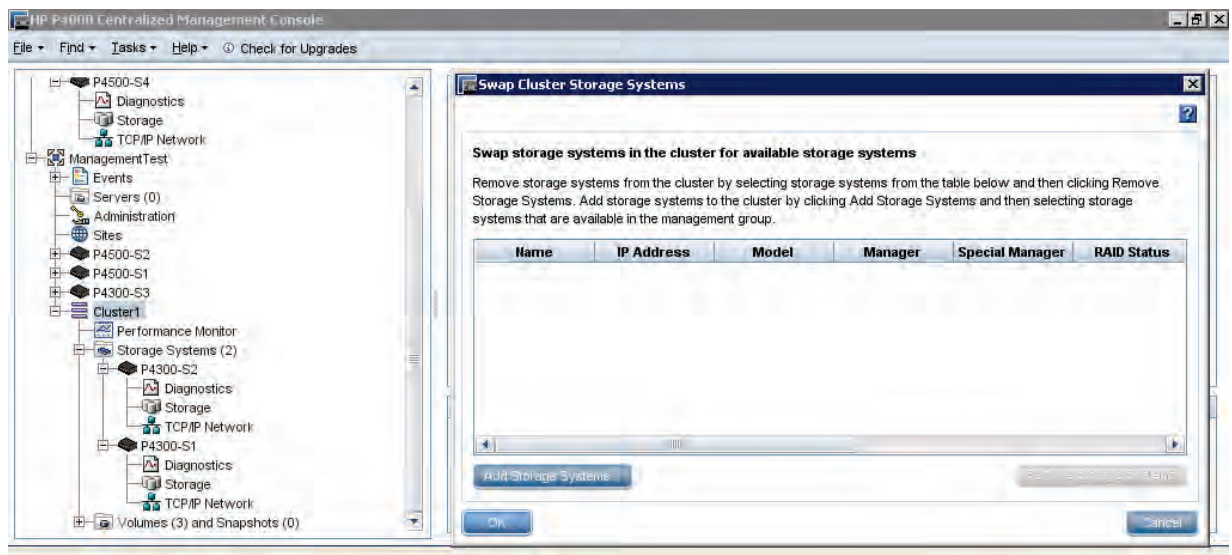
Step 2: Using the "Swap Cluster Storage Systems" pop up window, select all HP LeftHand P4300 storage systems in the list. After selecting all nodes in the list, click the "Remove Storage Systems" button.

Figure 7: Swap Cluster Storage Systems Window



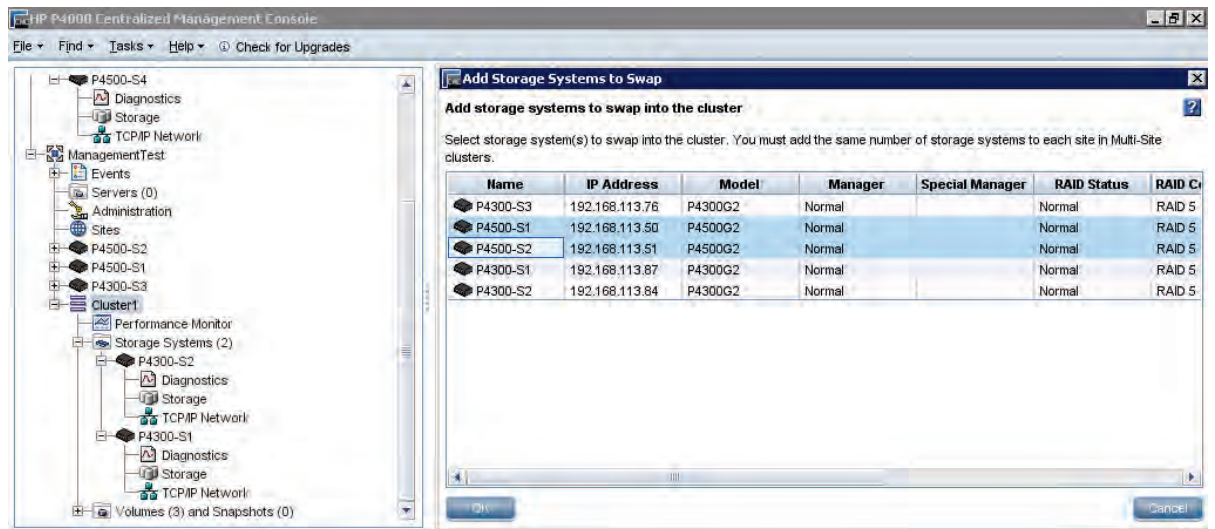
Step 3: After clicking the "Remove Storage Systems" button, a pop-up window displays with an empty list. Next select the "Add Storage Systems" button.

Figure 8: Swap Cluster Storage Systems Window



Step 4: The "Add Storage Systems to Swap", pop-up window appears with a list of all storage systems available to be swapped into "Cluster1". In this example, we selected two nodes (P4500-S1 and P4500-S2) and then click the OK button.

Figure 9: Add Storage Systems to Swap Window

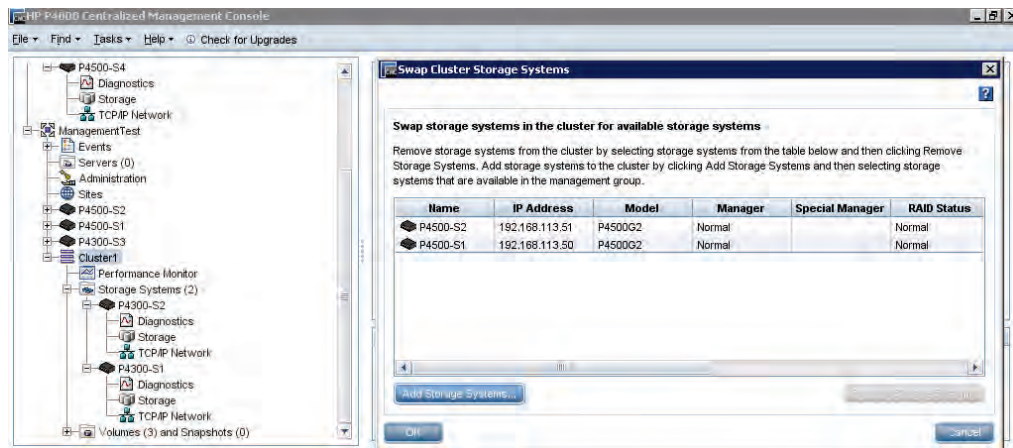


Note: The list above includes P4300 and P4500 systems. A HP LeftHand best practice is to group systems within the same P4000 platform (same capacity and disk technology) when adding systems to a cluster. In our example we selected two P4500 systems.

The P4300 nodes removed from "Cluster1" in Step 2 are available systems in the "Add Storage Systems to Swap" list. The P4300 nodes removed during Step 2 should not be selected during the Cluster Swap operation.

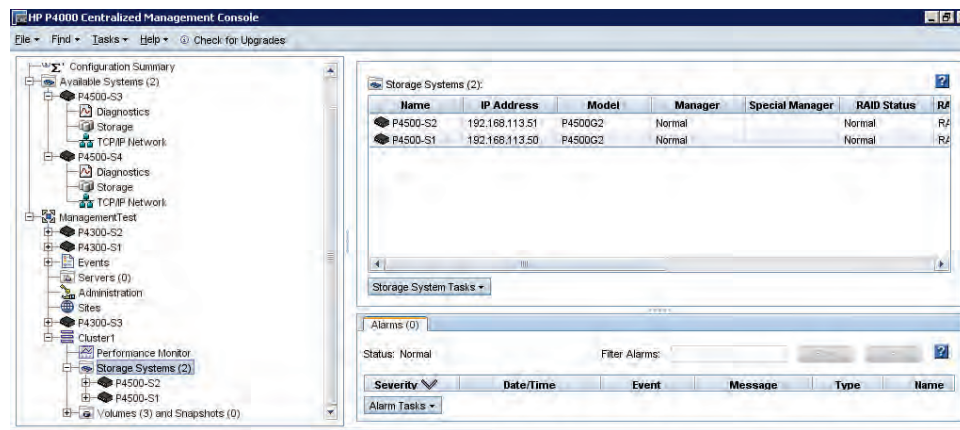
Step 5: The pop-up window below is asking for confirmation of the nodes selected in the window below. Select the "OK" button to complete the Cluster Swap activity or select the "Cancel" button to cancel the Cluster Swap activity and return to the last known state of Cluster1.

Figure 10: Swap Cluster Storage Systems Window



After selecting the "OK" button in Step 5, the program returns to the main CMC. The left side of the CMC navigation window now shows "Cluster1" with two P4500 nodes. The P4300 nodes (P4300-S1 and P4300-S2) still reside in the "ManagementTest" group but are not associated with "Cluster1" as shown in Figure 11 below.

Figure 11: Storage Systems Window



Peer Motion and LeftHand VSA Software

In federated storage environments, infrastructures can be comprised of multiple pools of storage resources and administrators have the agility to move data and workloads across these pools seamlessly. With HP Peer Motion, LeftHand users can create federated storage environments containing a variety of LeftHand platforms each offering different levels of performance and capacity and have the flexibility to migrate data between all of these resources. In many cases, this capability can be taken a step further to include other storage platforms by incorporating the HP LeftHand VSA software. The VSA software is the LeftHand Virtual SAN Appliance—a fully featured version of the LeftHand OS storage software that runs in a VMware vSphere or Microsoft Hyper-V virtual machine. With the LeftHand VSA, administrators can “annex” non-LeftHand storage arrays (any block storage device supported by the hypervisor) and add these resources to their LeftHand storage federation. Most often this is done to allow administrators to repurpose older storage capacity for use as lower tier storage within a LeftHand federation. Once a storage resource has been “annexed” Peer Motion can be used to move data to and from other resources within the environment. Similarly, cluster swap can be used if the administrator chooses to move data off a cluster of VSAs onto physical LeftHand devices at a later time. Keep in mind that data must be deleted or migrated off of the external storage before the VSA can make use of the capacity. Also, performance of a disk system or array annexed by VSA may not be equivalent to the native performance of that system or array.

VSA annexing-use case

- Repurpose legacy and/or disparate storage arrays to HP LeftHand Storage
- Consolidate disparate storage arrays under a single management console, LeftHand CMC
- Create a DR site by installing VSA on a legacy array at a DR location and replicating to it from primary LeftHand array using LeftHand Remote Copy.
- Manage and unify different storage transports (like iSCSI, FC, SAS) from legacy and third-party storage systems to the iSCSI technology of HP LeftHand Storage

Benefits of VSA annexing

An industry-leading benefit of VSA Annexing is the means to repurpose, aggregate and virtualize storage capacity of heterogeneous storage systems (arrays) together into shared pools of LeftHand storage.

Other benefits are:

- Repurpose retired, aging storage systems into new storage tier, reducing CAPEX cost
- Convert retired storage systems into a virtualized storage tier for test and development environments
- Eliminate volume migration boundaries between LeftHand and other third party storage systems
- Create federated storage, using Peer Motion on LeftHand to seamlessly move VSA connected volumes to any physical LeftHand platform, thus enabling a resilient, flexible, and agile virtualized SAN infrastructure

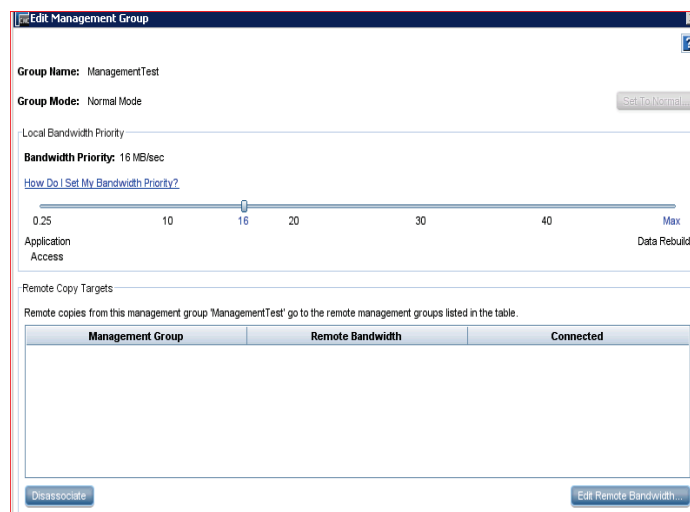
Best Practices for Peer Motion

The section to follow provides best practice considerations for planning, design and implementing a LeftHand standard or multi-site cluster swap. Considerations are:

Make certain storage nodes added to the cluster during a cluster swap adhere to HP LeftHand software and hardware prerequisites. Cluster Swap prerequisites are the same LeftHand prerequisites that one would consider when adding, removing or exchanging storage nodes within an existing standard or multisite cluster. Reference the [P4000 SAN Solution User Manual](#) to guide standard and multi-site cluster configuration prerequisites.

- Storage nodes added to a cluster during a cluster swap must be running the same version of the LeftHand OS as the storage nodes currently residing in the cluster
- Check that the storage systems (clusters) are in a good state, with no critical events.
- Run hardware diagnostics in the LeftHand CMC for all nodes that will be participating in the cluster swap.
- All storage nodes participating in a cluster swap need to be assigned to the same management group and appropriate site.
- Make certain your cluster swap adheres to the recommended number and type of running configuration managers. The end state of a cluster swap requires a minimum of two storage systems and three running configuration managers. For storage nodes less than 3, use a failover manager. The failover manager serves as a quorum tie-breaker.
- Systems within a cluster should be from the same LeftHand platform as well as populated with a homogenous group of disk drives. The storage capacity of a cluster is derived using the drive capacity of the smallest disk within the storage systems (cluster).
- Storage systems added to a management group inherit all characteristics and attributes of that management group including Domain Name Services (DNS), Simple Mail Transport Protocol (SMTP), Network Time Protocol (NTP), and more. Inheriting these attributes promotes a non-disruptive and seamless migration of data volumes without loss of user and/or application connectivity
- Make certain storage systems swapped during the cluster swap process are sized properly to accommodate the capacity and Network RAID levels of all data volumes to be migrated
- Check that all storage systems participating in the cluster swap are reachable via the IP network.
- Ensure that the hardware/disk RAID settings, network properties (frame, size and flow control) and network bonding settings are properly configured on all nodes before initiating a cluster swap.
- Keep in mind volume migration during a cluster swap requires restriping of all block data onto newly added storage systems. Consider using the local Management Group bandwidth setting within CMC to control bandwidth during data migration activities. This setting is used to determine the maximum rate per second allocated to non-application processing tasks such as data restripe. The bandwidth priority setting (See figure 12 below) controls the duration and performance impact of the restripe associated with the cluster swap. The restripe process may take several hours or days depending on the amount of data required to restripe and the amount of bandwidth allocated to restripe activity. Setting the bandwidth priority control setting greater than 40 MB (megabytes) per second may impact volume availability.

Figure 12: Bandwidth Priority Setting Control Window



- If using the HP LeftHand DSM for Windows MPIO, there needs to be at least one storage system in the management group that is not assigned to the cluster. If there are no other nodes in the management group, users may add an HP LeftHand VSA to serve this purpose. The VSA can be downloaded free of charge at www.hp.com/go/tryvsa. VSAs used for this purpose should be removed from the management group when the cluster swap is complete.

Summary

Peer Motion is a built-in data mobility feature provided at no extra cost within the all-inclusive LeftHand software suite. Key features of Peer Motion are volume migration and cluster swap. These features have been enhanced throughout HP LeftHand's product lifecycle offering customers the simplicity and agility needed to architect and design a flexible, federated SAN solution. Through features such as Peer Motion, customers have better storage resource alignment with their business and application needs while enabling data mobility to support convergence, federation, and cloud-based storage solutions and creating the building blocks of the next-generation IT-as-a-Service model. Peer Motion software delivers non-stop availability for small, mid-range and enterprise customers desiring a scalable SAN architecture that is flexible, resilient, and easy to manage.

Frequently asked questions

Can one use the cluster swap task to exchange storage nodes within a cluster?

Yes. The option can be used to add or remove individual storage systems within a cluster, however, a "Cluster Swap" implies the replacement of all nodes within the cluster. A best practice recommendation is to use the appropriate option; which is the Exchange Storage System task to add or remove (exchange) individual storage systems within a cluster.

How much time does it take for HP LeftHand Storage to complete a cluster swap operation?

The cluster swap operation can be completed in a few minutes. However, the time required to migrate all volumes from one storage system to another varies. The variance in time depends on the size and capacity of data volumes, the amount of application workload shared in concert with volume migration activities and network bandwidth. Refer, HP technical white paper titled the [Architectural Advantages of HP P4000 SANs](#) for guidance on LeftHand network best practices.

What is the minimum number of nodes supported for a cluster swap?

A two-node cluster is the minimum redundant cluster swap configuration plus a failover manager running on a separate server.

Is cluster swap supported on VSA software appliances?

Yes. It is support across all HP LeftHand Storage platforms.

Is cluster swap supported on mixed VSAs and physical LeftHand storage system clusters?

Mixing of VSAs and physical platforms is not a best practice or supported LeftHand design. A cluster swap of heterogeneous LeftHand nodes yields unpredictable performance results; in a mixed cluster scenario, performance is an average of all storage systems within the cluster.

What happens to the storage systems that were removed from the cluster during the cluster swap operation?

The prior storage systems are removed from the cluster; however, remain in the same management group. These nodes can be reused after all data volumes have been migrated to the new cluster, completing the cluster swap.

Is the HP Device Specific Module for Multipath I/O (MPIO) supported with Peer Motion?

Yes, the DSM is supported with Peer Motion. As noted in the Best Practices section above, when performing a cluster swap with DSM it is required to have at least one additional node in the management group that is not participating in the swap.

Get connected

hp.com/go/getconnected

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

© Copyright 2012 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.

4AA4-0304ENW, Created April 2012

