



**Hewlett Packard
Enterprise**

HPE 3PAR Storage Federation

Data mobility between HPE 3PAR StoreServ Storage and
from non-HPE storage

Contents

Management summary.....	3
Intended audience	3
Terminology, abbreviations, and typography.....	4
Concepts.....	5
Zoning.....	7
Use cases.....	8
Simplicity of operation.....	8
Manage multiple systems easily.....	8
Load balancing workloads.....	9
HPE 3PAR asset management made easy.....	9
Non-HPE 3PAR asset refresh made easy.....	9
Zero footprint data mobility.....	9
Optimize system utilization.....	9
Accommodate unplanned workloads.....	10
Peer Motion throughput throttling	10
Enhanced scalability	10
Federated thin provisioning.....	11
Ultimate elasticity for cloud service providers.....	11
Empowering the storage and the server administrator	11
Save disk space	11
Save administration time.....	11
Save money	12

Technical prerequisites.....	12
Supported federation configurations.....	13
Orchestrating data mobility.....	14
Features.....	15
Federation auto-detection.....	15
Improved orchestration.....	15
Filtering federations.....	15
Unlimited data mobility.....	15
Migration resilience.....	15
Concurrent Peer Motion tasks.....	15
Import task priority.....	16
Peer links over 16 Gbps FC.....	16
Automatic host rescan.....	16
Simplified zoning.....	16
Local reads during migration.....	16
LUN ID conflict resolution.....	17
Support for migration of a subset of VVset.....	17
HPE 3PAR Storage federation advantages.....	17
Comparing HPE 3PAR Storage Federation with appliance-based storage migration tools.....	17
Ease of installation.....	17
Ease of management.....	18
Performance.....	18
Retain data services.....	18
Reduction of cost.....	18
Best practices.....	18
Logging.....	21
Licensing.....	21
Delivery model.....	21
Troubleshooting HPE 3PAR Storage Federation.....	21
Information collection when contacting support.....	22

Management summary

HPE 3PAR Storage Federation is Hewlett Packard Enterprise's scale-out strategy to turn a collection of HPE 3PAR StoreServ systems into a single entity to maximize resource utilization, balance workloads across arrays, and ease management of multiple systems. Storage Federation expands the autonomic features of Dynamic Optimization for volume level tiering and Adaptive Optimization for block level tiering into system level tiering. Within a federation, single system boundaries are removed enabling multi-array bidirectional data mobility for enterprise block storage between federation members. Concurrent migrations of volumes between any federation members in any direction offer ultimate flexibility to store application data on the system with the physical disk and volume provisioning type that suits the current situation without business disruption or complex planning. A storage federation can be extended by HPE and non-HPE storage systems, called migration sources. Volumes on one or more migration sources are imported into a federation member unidirectionally in a concurrent fashion. The federation framework is integrated in the HPE 3PAR StoreServ operating system: no external appliance in the data path is required and no additional overhead is introduced on the host resources.

Federation was designed with the customer in mind: it is easy to implement, manage, and modify using the familiar HPE 3PAR StoreServ Management Console (SSMC), eliminating the need for expensive professional services at each federation layout change. Predefined workflows present in the SSMC empower experienced and non-experienced administrators to move one, some, or all volumes associated with a host from one HPE 3PAR StoreServ Storage system to another with a single mouse click. Simplified SAN zoning and the one-time, upfront zoning setup between federated systems and their hosts allow migrations between any federation members without SAN changes.

The HPE 3PAR Storage Federation software leverages HPE 3PAR Thin Built-In technology to power the simple and rapid inline conversion of inefficient, fully provisioned volumes on source arrays to more efficient, higher-utilization thin or thin deduped volumes on the destination HPE 3PAR StoreServ system.

It expands the autonomic features of Dynamic Optimization for volume level tiering and Adaptive Optimization for block level tiering into system level tiering.

This white paper concentrates on the concepts, features, and use cases of storage federation. Text, screenshots, and tables in this paper are compatible with SSMC 3.0 and HPE 3PAR OS 3.2.2 MU4 unless mentioned otherwise. New features in 3.2.2 are discussed only when they relate to storage federation.

Intended audience

The intended audience for this white paper is storage administrators who manage HPE 3PAR StoreServ systems. Readers are assumed to have a good knowledge of HPE 3PAR StoreServ systems and the StoreServ Management Console (SSMC). Expertise on legacy HPE 3PAR StoreServ and supported non-HPE systems is required when importing data from these systems into federation members.

Terminology, abbreviations, and typography

Throughout this paper, we use the following terminology and abbreviations:

CLI	Command Line Interface
CPG	Common Provisioning Group, a template to create HPE 3PAR StoreServ volumes
Destination system	The storage system that receives the data that are migrated
FC	Fibre Channel, a physical interconnect between a server and a storage system
FIFO	First In First Out, a concept handling entries on a software queue
HBA	Host Bus Adaptor, the expansion card in a server that implements connectivity to the outside world
HDS	Hitachi Data Systems, a disk array vendor
Host	The server whose volume(s) are under migration from a source to a destination system
INQ	Acronym for Inquire, a SCSI command used to retrieve metadata of a target LUN
IOPS	I/O per second, the number of I/O operations executed in one second of time
LDEV	Logical Device, a volume on an HDS storage system
Link	The logical interconnection between the source and the destination system under a Peer Motion or Online Import migration
LIP	Loop Initialization Protocol, a method to discover devices and their characteristics on the SCSI bus
LVM	Logical Volume Manager, a method for managing space on mass-storage devices that offers striping and RAID protection levels across multiple disk drives
MAINT_IN	Acronym for MAINTENANCE_IN, a SCSI command for reporting information about a device
MC	Management Console, a graphical user interface for managing an HPE 3PAR StoreServ system
MDM	Minimally Disruptive Migration, one of the three migration types in HPE 3PAR Peer Motion and Online Import
NTP	Network Time Protocol, a framework for clock synchronization between computer systems
OIU	Online Import Utility, the console-based tool for migrating non-HPE storage systems to HPE 3PAR StoreServ ones
PMU	Peer Motion Utility, a command line oriented tool for migrating block-level storage between HPE 3PAR StoreServ systems
PROUT	Persistent Group Reservation Out, a SCSI command that sends information from a target disk back to the Initiator host
RC	Remote Copy, the replication solution for HPE 3PAR StoreServ systems
rpm	Rotations per minute, the revolution speed of the spinning platters of a disk drive
SIST	Single-initiator-single-target, a SAN zoning concept where the zone connects a single FC port on a host with a single FC port on a storage system
Source system	The storage system that contains the data to be migrated
SSMC	HPE 3PAR StoreServ Management Console, a graphical user interface for managing the HPE 3PAR StoreServ
Storage Area Network, SAN	A high-speed special-purpose network that interconnects different kinds of storage devices with servers executing applications
TUR	Test Unit Ready, a SCSI command used to determine if a device is ready for read/write operations
VLUN	A network or SAN path exporting an HPE 3PAR volume to a host
WWN	World Wide Name, a 64-bit quantity the uniquely identifies an FC Host Bus Adapter interface, an FC port on an HPE 3PAR StoreServ system or a VLUN
Zoning operation	Creates, modifies, or deletes a logical connection in the SAN between a Fibre Channel Host Bus Adapter on the server and one on the storage system

Text in the HPE Simple font is typed by a storage administrator in a computer terminal session or the result of such an activity.

Concepts

An HPE 3PAR Storage Federation is a logical construct that groups multiple HPE 3PAR StoreServ systems for concurrent, non-disruptive, bidirectional data mobility purposes. A federation can be extended by attaching a legacy HPE 3PAR StoreServ or selected HPE and non-HPE storage systems to one or more federation members. These attached systems are called “migration sources.” The grouping exists in the context of the StoreServ Management Console (SSMC), the simple, concise, and easy-to-use graphical console that manages HPE 3PAR StoreServ systems. The SSMC conducts the orchestration of an HPE 3PAR Storage Federation. It includes a graphical map of the federations it contains and their components. Figure 1 from SSMC 3.0 reveals the high-level topology of a storage federation named Brussels with three members each running HPE 3PAR OS 3.2.2. Attached to two of the federation members is a migration source running HPE 3PAR OS 3.2.1. The layout of a federation and any attached migration sources is apparent in seconds from this view.

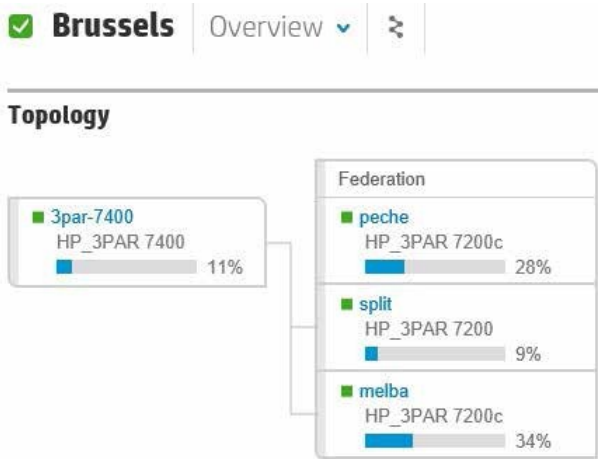


Figure 1. Topology of a three-member storage federation named Brussels with a migration source attached to two of the federation members

Figure 2 maps out the components in the “Brussels” Federation in the SSMC. We see that the hierarchy of volumes, hosts, host sets and systems is extended upwards by federations. In this composition, the three federation members and the migration source reside on the same level. With a few mouse clicks, one can explore the federation tree to the atomic level of hosts and volumes and sets of them and the relationship(s) between any two elements in the tree.

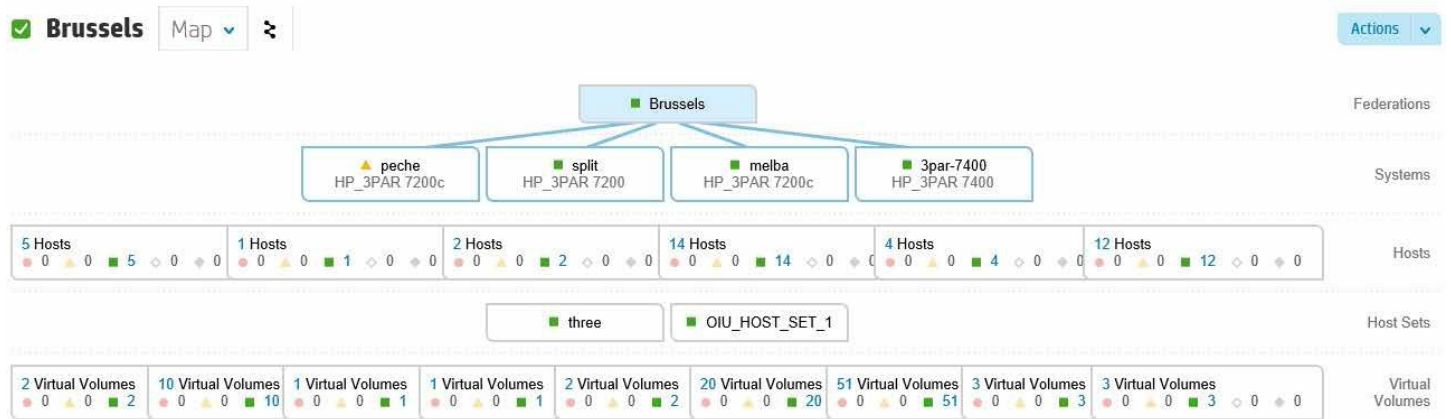


Figure 2. Map of the components of the storage federation named Brussels

The arrangement of members in a federation and migration sources attached to them is informal: adding a system to or removing it from a storage federation is live and non-disruptive for the system touched and to those already in the federation or attached to it. HPE 3PAR OS 3.2.2 is the minimum version for a system to enter a federation and execute bidirectional, multi-array migrations. SSMC 2.2 or later is needed to create and manage HPE 3PAR Storage Federations and attach migration sources. Storage federation supports mobility from and to any type of disk drive or provisioning including deduplication. The SSMC can create and manage multiple federations. However, a system can only belong to a single federation and migrations only take place between members of the same federation.

The underlying technology for data mobility between HPE 3PAR Storage Federation members in the SSMC is an enhanced version of the familiar Peer Motion software.¹ In this version, one pair of Peer links operates in each direction between any two HPE 3PAR StoreServ systems that are federation members. This way, a federation member can act simultaneously as a source and a destination system allowing concurrent, bidirectional, any-to-any member volume migration.

Data mobility between storage federation members requires no zone changes when reversing the direction of data migration or when changing the source or destination in a next migration transaction. This is thanks to the “wire once, move anywhere” concept in which the required SAN zoning for data mobility between federation members is defined and activated upfront at the creation of the federation. This obviates the need for involvement of SAN administrators for subsequent migration operations. Application data migration gets reduced to a routine activity, executable at the day and time the storage and application administrators choose, not when the SAN administrators can make the requested zoning changes. Hosts should be attached to all federation members to enable movement of any volume to any member.

Volumes on migration sources can migrate unidirectionally to a federation member. Importing data from migration sources is orchestrated by the SSMC, Peer Motion Utility, Online Import Utility, or HPE Command View EVA depending on the vendor of the source system and the HPE 3PAR OS version installed on it. With the SAN zoning of a migration source to one or more federation members done upfront, the storage administrator can move the volumes of a migration source to a federation system by a single mouse click in the SSMC. The SAN zoning between federation members and their migration sources can stay in place indefinitely to move volumes any time in the life cycle of the systems.

HPE 3PAR Storage Federation supports online, minimally disruptive and offline data mobility between its members and attached migration sources. In the minimally disruptive mode, the host undergoes a single, short outage before the data mobility takes place impacting the business minimally. The following objects can be selected in the SSMC for migration between storage federation members:

- A single host
- A single host set
- A single, some, or all volumes presented to a host, host set, or cluster
- A single virtual volume set (VVset)
- A Remote Copy group

Selecting a host or a host set includes all volumes presented to it in the migration definition. To migrate a cluster, place the cluster member in a host set. All shared and non-shared volumes exported to the host set members are selected implicitly. The host set is recreated on the destination HPE 3PAR system. Nested host sets are not supported.

If the host operating system supports ALUA and has an ALUA capable persona on both the source and destination HPE 3PAR StoreServ, a single volume, some of them or all of them can be selected for migration. VVsets on the source array can be migrated one at a time; the VVset is recreated on the destination HPE 3PAR system. An optional consistency group expedites rollback in case the migration does not complete. Per migration, one destination can be selected. When selecting a volume in a Remote Copy group for migration, all volumes of the group are included in the data transfer.

Multiple migrations can run concurrently between federation systems in the same or opposite directions. Consult the “Supported configurations” section in this paper for details on the anatomy of supported federation layouts and their migration members. The SSMC provides the option to keep the source volume(s) on the host after a successful migration; the default is to remove it. Some restrictions to this removal apply; consult the HPE 3PAR Peer Motion and HPE 3PAR Online Import User Guide.²

The HPE 3PAR SSMC can copy and synchronize storage system configuration data like user accounts, hosts, domains, and the parameters for LDAP, SNMP, NTP, and syslog between the federation member systems. Keeping these data synchronized is vital for unrestricted data migration across the federation members. Multiple volume migrations can be initiated concurrently to a single destination system by the SSMC, Peer Motion Utility, and Online Import. The data of up to nine volumes is moved in parallel into the destination system. The import task for a volume scheduled for migration beyond these first nine enters a FIFO queue on the destination system meaning the task becomes active when a running one completes.

¹ hpe.com/h20195/v2/getpdf.aspx/4AA3-5714ENW.pdf?ver=2.0

² h20564.www2.hpe.com/hpsc/doc/public/display?docId=c04783896

With multiple, concurrent data migrations possible between members of a storage federation and their migration sources, the need for a priority scheme became apparent: some migrations are more critical than others, and should get precedence although they were pushed onto the import task queue to be serviced later in time. Peer Motion in the SSMC implements a three-level priority scheme for every migration task created. The default import task priority for a volume is “medium;” “high,” and “low” are the other options. The queue with the import tasks scheduled for migrating is continuously rearranged to move tasks with the highest priority to the front and activate them at the earliest convenience. The “Features” section in this white paper adds more details to the implementation of prioritization of Peer Motion import tasks.

Peer Motion within HPE 3PAR Storage Federation is thin-aware: only the used space in a thin-provisioned volume is migrated from the source to the destination. Blocks of zeroes in fully provisioned source volumes are transferred but will not be written on the destination disks if the destination volumes are thin provisioned. HPE 3PAR Peer Motion Efficient Import is an optimization where blocks never written to on the source volume are not read from disk and hence not transferred over the Peer links. The benefit is faster volume migration, thanks to reduced SCSI read requests. It applies to thin-provisioned source volumes that are at most 75 percent full.

The integration between Remote Copy (RC) and Peer Motion has been enhanced in HPE 3PAR OS 3.2.2 MU4 to include ALUA concepts for hosts and their pathing to volumes under migration. With this, Peer Motion can migrate volumes between two HPE 3PAR StoreServ systems while a Remote Copy relationship with a secondary system is active for the Peer Motion source or destination. Synchronous and asynchronous Remote Copy is supported. During a small window of exposure, the existing Remote Copy setup is replaced by a new one incorporating the system that was the Peer Motion destination. This cutover time is minimal when HPE 3PAR OS command line scripts containing the changes are created upfront. With the snapshots of the migrating volumes taken at the end of the volume import for primary RC system Peer Motion or at the import start in the case of secondary RC system Peer Motion, the time for resynchronization of the volumes in the new Remote Copy relationship shrinks substantially. The steps for this are extensively documented in Appendix B of the HPE 3PAR Peer Motion and HPE 3PAR Online Import User Guide.³ The coalition between Peer Motion and Remote Copy requires all three HPE 3PAR StoreServ systems involved to run HPE 3PAR OS 3.2.2 or later. One or more of these systems can be a federation member.

Zoning

HPE 3PAR Storage Federation members are interconnected over preferably dual, redundant FC SAN fabrics by the Peer links. The enhanced version of Peer Motion in the SSMC supports two pairs of Peer links operating between any two HPE 3PAR StoreServ systems in the federation. This means a federation member can act in a Peer Motion setup as a source and as a destination at the same time. When interconnecting multiple federation members, a substantial number of required SAN zones in the single-initiator-single-target (SIST) concept are needed. Federation and the SSMC endorse the presence of multiple initiators and multiple targets in a single zone. Consequently, only two zones need to be constructed interconnecting the federation members. This greatly simplifies the zoning topology when creating a new federation or when a member enters or leaves an existing federation. Figures 3 and 4 show a red and a blue zone that contain the WWNs of the Peer and Host ports that make up each Peer link for use in a federation. In the figures, Hx denotes a Host port and Px denotes a Peer port.

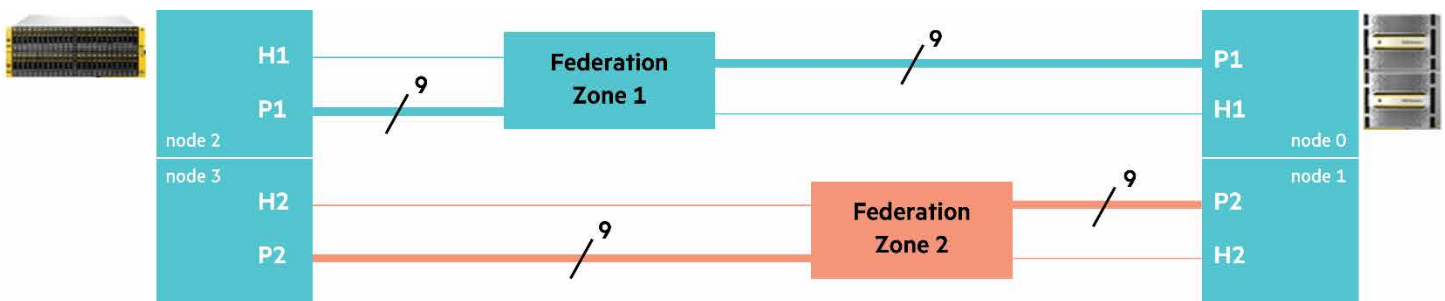


Figure 3. Peer links zoning layout for a two-member federation. Only two SAN zones are created for bidirectional data mobility.

³ h20564.www2.hp.com/hpsc/doc/public/display?docId=c04783896

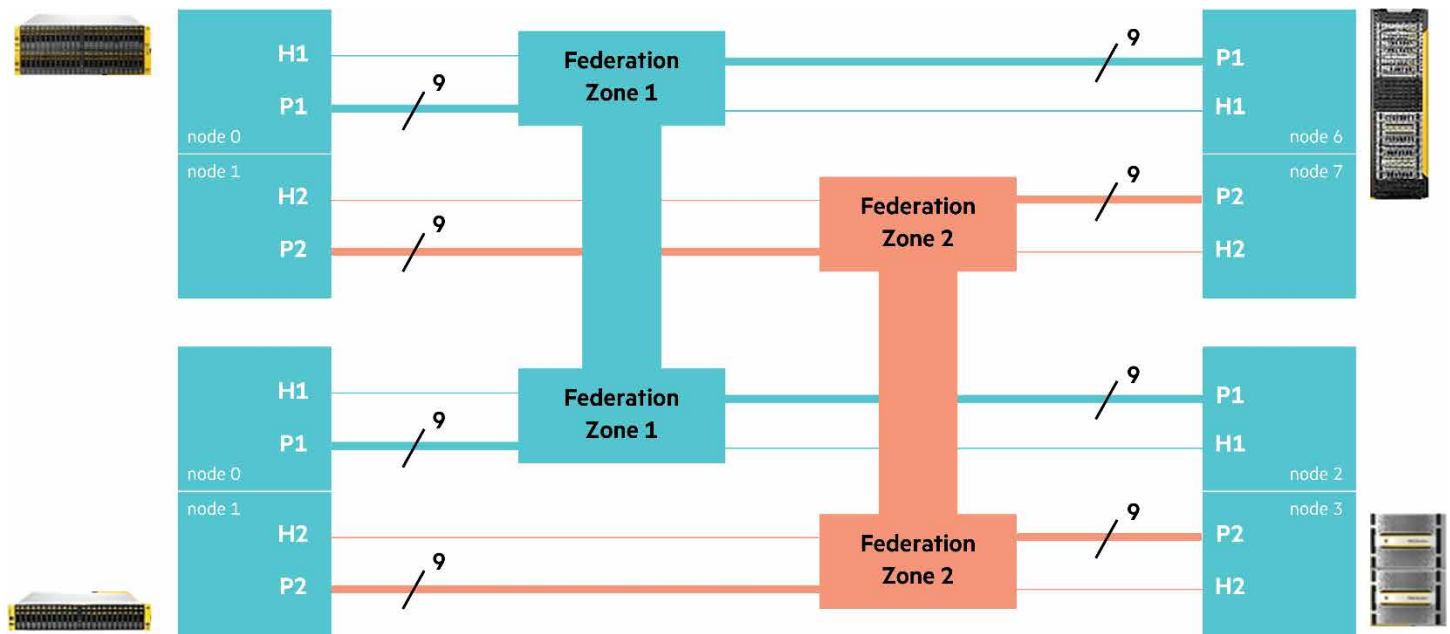


Figure 4. Peer links zoning layout for a four-member federation. Only two SAN zones are created for bidirectional data mobility.

The Peer port and the Host ports for the Peer link of the even nodes on the source and destination systems enter one zone, the other zone is composed of the Peer port and Host ports located on the odd node instances. Conveniently, the SSMC displays the list of the WWNs that should go into each of the two zones after the federation members were selected. This list includes the WWNs of the maximum count of eight virtual NPIV Peer ports to prepare for the case the volumes under migration are exported to a four-node Windows® cluster. The number of WWNs per federation zone scales with the number of federation members. In a two-member federation, every zone counts 9+1+9+1=20 WWNs as apparent in figure 3. For a three-member federation, the WWN count per zone increases to 30 and in the four-member environment shown in figure 4, we arrive at 40 WWNs per zone.

Migration sources requires their private SAN zones between the source and the federated destination array: you cannot add the Host ports of the migration source to the federation zones. The Peer ports in use for federation on the destination system are reused to build the Peer links to the migration source. No extra pair of Peer ports should be created on the federation member for the Peer links to a migration source. For migration sources, HPE prefers 1:1 single-initiator-single-target zoning for improved fault isolation and reduced bandwidth interference. At minimum, the SAN interconnection for a migration source requires two zones. The actual count depends on the number of federated systems to which the migration source will be connected. After selecting the migration source and the destination federation member, the SSMC shows the WWNs that must be included in each Peer SAN zone.

The Health section on the Federations page in the SSMC indicates the zoning status for the federation members. If the zoning was not correct and got amended, the Health state of the federation will turn to 'Normal' in seconds. You do not need re-initiate the zoning discovery in any way.

Use cases

Simplicity of operation

An HPE 3PAR Storage Federation is managed from the SSMC, HPE's advanced graphical interface for HPE 3PAR StoreServ systems. This is where simplicity of operation meets clear design: one-click data mobility permits the storage administrator to move volumes across the federation members. The SSMC conveniently lists the WWNs of all Peer and Host ports that need to enter the federation zones. The SAN zoning layout is set up once at the creation of the federation. No zoning changes are required for N:M bidirectional data mobility between any federation members.

Manage multiple systems easily

For security and fault isolation reasons, some data centers install multiple storage systems for the same application like Microsoft® SharePoint, a Web server or a database. Keeping the configuration settings of multiple storage systems serving the same application identical is a manual job that requires strict bookkeeping. In an environment like storage federation where systems are linked for data mobility, having identical configuration data across all federation members is vital for unrestricted volume movement. The HPE 3PAR SSMC copies user accounts

(username and password), hosts and sets of them, domain names with their member users and member hosts, domain sets and the parameters for LDAP, SNMP, NTP, and syslog across the federation member systems. New configuration data like a host created on one federation member are spread over the federation members by the SSMC by a single mouse click. With the SSMC as the central management console, keeping configuration data identical across multiple systems for federated data mobility has never been easier.

Load balancing workloads

Commercial workloads may consume a varying amount of resources of a storage system depending on the time of the day (day/night), the time of the year (end-of-the-quarter activities, holiday season sales) or over the lifetime of the application (natural growth of the user count or the size of a database). Because migrating a workload to a new platform is generally considered a complex task, it is common practice to size a storage system for an application's peak resource consumption possibly exhibited only once a year. This is inefficient use of array resources and the money that bought them.

With multi-array bidirectional data mobility made possible by HPE 3PAR Storage Federation without SAN zoning changes and without any downtime, moving an application to a more powerful platform once or a few times a year becomes feasible without extensive forward planning. In an alternative approach, applications on a local system can be moved over to another system temporarily to free up resources for an application that just became busier. Load balancing applications across the federation members is executed efficiently and by a single mouse click in the SSMC, resulting in reduced capital layout and improved system utilization.

HPE 3PAR asset management made easy

Customers typically write off storage equipment on 3- or 4-year cycles, so each year the data on one-quarter to one-third of the storage devices is moved out due to system retirement. At large enterprises, several projects for asset migration run in parallel every day of the year. Even when staying with the same vendor, moving data to a new system can be a challenge and require involvement of professional services because the new unit can be radically different from the decommissioned one. HPE 3PAR StoreServ systems share the same architecture, concepts, and user interface from the entry-level over the midrange to the high-end systems. Complementing this with the concept of storage federation enables the seamless replacement of retiring HPE 3PAR StoreServ with the new unit. The new HPE 3PAR StoreServ is inserted to the storage federation live and non-disruptive to itself and the existing federation members. Importing the configuration data of the federation on hosts, domains, user accounts, LDAP, NTP, etc., to the new member in a one-click operation speeds up the asset preparation of the new federation member. The migration of the data from the retiring system to the new member is non-disruptive for the applications and does not add any overhead to the hosts. Removing the decommissioned system from the federation is non-disruptive as well. Overall, access to the data is uninterrupted during the entire operation.

Non-HPE 3PAR asset refresh made easy

Compared to other vendors, HPE 3PAR StoreServ systems have decisive advantages such as default wide striping, a single OS and unified user interface across the entire portfolio from entry-level to high-end and all-flash arrays, unique scalability, industry-leading thin provisioning with hardware-assisted space reclamation, a broad software suite and superior replication options. The HPE 3PAR Online Import series of software orchestrates moving data from HPE EVA, EMC, HDS, and IBM storage systems to an HPE 3PAR StoreServ Storage system that can be standalone or a member of a storage federation. Driven from a graphical interface or the command line depending on the platform that is retired, Online Import relocates data to any current HPE 3PAR system in an online, minimally disruptive or offline fashion. HPE Online Import introduces no external appliance to manage the migration. There is no added overhead on the hosts since the physical migration path is array-to-array, a concept that reduces complexity as well. The simple and intuitive steps for performing an Online Import migration and its one-year free license with any new system purchase contribute to the decision making process to move to HPE 3PAR StoreServ storage.

Zero footprint data mobility

HPE 3PAR StoreServ competitors offer data migration solutions based on an external appliance that can nearly fill a full-height rack of floor space, includes multiple industry-standard servers complemented by custom-built hardware devices and consumes many FC and network ports in the data center. The HPE 3PAR Storage Federation subsystem comes built-in into HPE 3PAR OS 3.2.2 or later, no post-installation is needed or external appliance is needed and no client software must be installed on the migrating hosts. This saves space, power and cooling, items of increased importance in the data center today. HPE 3PAR Storage Federation is the only migration technology that delivers true zero footprint data mobility.

Optimize system utilization

Workloads can be classified by the resources they consume on the storage system connected to: some workloads are random read-intensive like Web servers, some are sequential read-intensive like video streaming, some are mixed random read-write-intensive like OLTP operations to databases, some are sequential write-intensive like filling transaction logs, disk-to-disk backup, or video recording. Sequential workloads fit

Nearline (NL) drives, random read workloads excel on SSDs while random write workloads will perform well on FC/SAS drives and SSDs. Dynamic tiering between two or three disk types can improve performance of some workloads. Choosing the right mix of drive types for a storage system is a non-trivial task when multiple workloads of a different nature execute on the system. Moreover, while the choice of drive types and number of them may match the initial set of workloads and their intended performance, this may no longer be the case after 6 to 12 months.

Setting up a storage federation of multiple HPE 3PAR StoreServ systems sporting multiple drive types and a different mix of them across the members solves this type of issue. As an example, consider an application that was brought into service one year ago. From HPE 3PAR System Reporter Region Density reports the application was found to benefit from a two-tiered SSD+FC drive setup administered by HPE 3PAR Adaptive Optimization. The workload's current storage system, being a member of a storage federation, was not equipped with SSDs however. This shortcoming was solved by moving over the application's data in hours in a non-disruptive way to another federation member that had the two disk tiers installed.

Applications executing in a storage federation framework are placed with minimal effort where they utilize best a member's resources regarding controller nodes (node count, CPU power, cache size, ASIC generation), disk type (SSD, FC, and NL) and I/O bandwidth (number and speed of FC, iSCSI, and FCoE cards). If the resource consumption modified over time or the service level objective of the application changed, moving the application to a more appropriate federation member is a routine operation that does not require downtime or an intervention from the SAN zoning administration team. By configuring HPE 3PAR Storage Federation, the storage administrator signs up for seamless data mobility on a routine basis.

Accommodate unplanned workloads

Uncertainty is more than ever the stance a data center planner works in. Planning application data placement and sizing systems has become more an art than a calculated strategy given the ever-growing number of applications to support and their voracious appetite for disk space, IOPS, and throughput. HPE 3PAR Storage Federation relaxes this organizational headache. A federation can consist of asymmetric types of HPE 3PAR StoreServ systems: any current model can enter a federation with any other type of system. This permits the customer to start with the HPE 3PAR StoreServ model of choice and add a different model of system to the federation at a later stage. As an example, a federation can be composed of an entry-level 7200c system, a midrange 4-node 8400, an all-flash 4-node 7450 model, and the new high-end 8-node 20800. This federation will handle nearly every conceivable workload in an efficient way on one of its members. As an example, the mobility of the workload to a federation member with Nearline drives when project data go into archive mode or to SSDs when an application scales up is seamless and requires no coordination with the application owners nor with the network and SAN administrators.

Peer Motion throughput throttling

HPE recommends executing volume migrations between systems during a quiet period for I/O on the migrating volumes, on the migrating host and on the hosts sharing the Host ports of the Peer links if any. This reduces the impact of the data transfer to the involved applications. To control the impact, Peer Motion migration within storage federation has the option in SSMC 2.2 and later to specify a priority for the import task of a volume. Three levels of prioritization are available directly affecting the rate of data transfer over the Peer links. A lower priority will effectively throttle the Peer Motion migration traffic redeeming I/O bandwidth to applications sharing the Host ports. The lower priority lengthens the total transfer time but that is acceptable because the migration is online. When the highest priority is selected for the import of the volume(s), the benefits of the destination HPE 3PAR StoreServ to the application will become available sooner. The prioritization option may permit the storage administrator to schedule migrations during working hours. The priority of a Peer Motion import task is specified at the creation of the migration definition but can be changed in the SSMC for an ongoing migration for optimal throughput or change of resource consumption.

Enhanced scalability

The reliability of disk arrays has grown to very high levels thanks to fault tolerance features of RAID and power distribution. Application data protection is further enhanced by efficient and reliable methods for volume replication to one or more secondary storage systems. On these grounds, it is a correct business decision to consolidate multiple applications serving tens to hundreds of hosts on a single storage system with replication enabled. HPE 3PAR StoreServ systems exhibit the required outstanding scalability for this consolidation exercise: the number of controller nodes can be upgraded from two to four in the midrange models and from two to eight in the high-end ones. CPU power, cache amount, host I/O bandwidth, and data replication capacity scale appropriately with the number of controller nodes. HPE 3PAR Storage Federation extends this massive scalability by presenting a collection of storage arrays as a single entity, manageable by the customer through a unified interface incurring no additional complexity over managing a single array. The compute, store, and I/O resources are pooled across the storage federation members. This way a storage federation offers painless scalability to a massive, pooled capacity of 60 TB and 10 million IOPS with one-click bidirectional data mobility between its members. This level of scaling is many times higher than any competitive system.

Federated thin provisioning

Thin provisioning has been a real boost to efficiency across the storage industry driving up resource utilization and cutting capacity purchases, power and cooling costs and administration time. However, thin provisioning remains managed at the level of a single storage array. Most customers happily overprovision the installed capacity of an array, but not its maximum installable capacity. The rationale behind this is that extra capacity can be added as needed to accrue a safe buffer but not more than the array can hold. With HPE 3PAR Storage Federation, this equation changes. With it, you can now quickly and easily migrate overprovisioned volumes to another array in the federation that seats enough capacity to accommodate the current and near future growth rate of the volume. Thanks to this, the thin limit is managed on the federation level allowing provisioning a single array well beyond its maximum installable capacity. This drives up utilization to much higher rates on storage federation members while enjoying the benefits and savings associated with HPE 3PAR StoreServ Thin Provisioning.

Ultimate elasticity for cloud service providers

Cloud service providers offer a large variety of server, application, and storage plans to potential customers. They always look for innovative data center deployment models to unlock new business opportunities while reducing cost. With HPE 3PAR Storage Federation, they can grow storage seamlessly to 60 TB of pooled and tiered capacity with the elastic mobility of data to a higher- or lower-level plan type by a single mouse click.

This way the service provider can offer superior flexibility to his customer for changing plans swiftly and frequently—an advantage that may generate new business on this agile market. At the same time he will save days of work for his storage administrators on larger migrations.

Empowering the storage and the server administrator

Today's trend in data center management is the fusion of the task profile of the server, networking, application, and storage administrator. For example using the HPE Insight Control Storage Module for vCenter, VMware® administrators can create, from within their familiar vCenter interface, new volumes on an HPE 3PAR StoreServ system and new datastores, assign them to new virtual machines and bring these virtual machines into an HPE 3PAR Peer Persistence relationship. This means the VMware administrator became self-sufficient for a numbers of tasks he or she has to perform on the storage side. Obviating the need to consult and coordinate with local or remote SAN and storage administration teams simplifies and speeds up the daily routine of the VMware administrator.

HPE 3PAR Storage Federation takes this independence a step further: when reversing the direction in data mobility between two federation members or when changing the source or destination in a subsequent data migration operation, no SAN zoning changes are required precluding the need to contact the SAN administration team for zoning modifications. Server, networking, and application administrators with some HPE 3PAR StoreServ background are empowered by this technology to execute migrations at the time of their choice without involvement of a storage or SAN team members. Data mobility is now reduced to a routine activity, executable by both experienced and non-experienced HPE 3PAR StoreServ administrators. This is the basis of unprecedented agile data management.

Save disk space

The storage space sold worldwide grows by double-digit numbers every quarter. Disk arrays that had abundant growth reserves when purchased hit the ceiling for expansion before reaching their predicted lifetime. Modern disk arrays support deduplication, a technique for eliminating storing duplicate copies of repeating incoming data while ensuring that data remains correct and complete. This makes it possible to store more data in less space.

HPE 3PAR Storage Federation members support hardware-accelerated inline deduplication on SSD drives. Migrating volumes from legacy systems to HPE 3PAR StoreServ All Flash arrays and arrays with SSD drives using Peer Motion or Online Import will reduce the space needed on the destination system. This saves capital cost on disk drives driving down the price of solid-state storage to the spinning disk level while increasing performance and reducing power consumption of the disk array.

Save administration time

With more petabytes than ever to manage, storage administrators welcome any concept that cuts down on time spent to daily routine tasks and larger projects like migrations and asset refreshes. Here are a few proof points stating how storage federation saves time for storage administrators:

- Asset refresh of storage federation members is orchestrated from the well-known SSMC graphical interface by clicking of a few buttons. No CLI commands are to master or an appliance to install saving days or weeks of time.
- The SAN zoning for storage federation is set up at the time the federation is created or when a member is added or removed from the federation. No zoning changes are required for N:N data mobility between any federation members. This means the time to prepare a data migration between federation members from a storage array and SAN zoning perspective is zero.

- Thanks to the pre-set SAN zoning, coordination with SAN support team for determining a time window when they can make zoning changes is eliminated. This saves communication and frees up time for the SAN team.
- The SSMC takes care, at the click of a button, of synchronizing the configuration parameters between federation members—a frequent chore that is time-consuming and error-prone.
- Data mobility between federation members is non-disruptive and transparent for applications. This eliminates the requirement to negotiate with the application owners on the right time for the cutover from one storage system to another one.

The above features of storage federation return hours or days per month to storage administrators.

Save money

With the ever-increasing amount of data to keep online, the cost of storage goes up. Advances in thin provisioning, compression, and deduplication have optimized the consumption of physical disk space. On a data center level, we still frequently see storage systems being overdimensioned at the time of purchase. These systems may never get near using their full capacity while other systems need a capacity upgrade by the same disk type that the first systems have on surplus. The obvious suggestion to relocate disks physically may not be an option on disparate architectural grounds or because of its complexity. Migrating the application to a more suitable storage system can be such a formidable task that the application owners decide on spending capital to acquire new disk hardware, rather than undertaking a risky and complex application migration maneuver. This is the environment where storage federation provides savings. The following points demonstrate the value of HPE 3PAR Storage Federation when it comes to reducing capital expenses:

- Storage federation has been designed with the customer in mind: with just one click in the SSMC graphical interface, HPE leads the industry in multi-array bidirectional data mobility between systems.
- Storage systems for applications that show a high variation in seasonal storage resource consumption can be sized for the steady state load. An application can be moved temporarily to a more powerful system during the high season by a single mouse click. This temporary migration and its return to the original system becomes a routine operation for storage system administrators. Capital expenditure is reduced because storage systems must not all be sized for their peak resource consumption.
- No expensive professional services team needs to be engaged: even moderately skilled storage administrators can migrate data between federated systems.
- Storage federation offers data mobility without an external appliance that competing vendors require. This saves on space, power, and cooling in the data center, on security investigation and administration time, and on risk of downtime when bringing the new device into the data path.
- The cost of data retention is lowered by moving data non-disruptively with a single mouse click to Nearline drives on a system that is part of the storage federation.
- Unused storage capacity on overdimensioned systems can be put in use by relocating applications on overloaded or overprovisioned systems by a single mouse click.
- End-of-support or end-of-lease systems are decommissioned and replaced by new systems in time and without complex planning with a few clicks of a mouse. This eliminates extended support or lease costs because of delayed migrations bringing in substantial savings.

Technical prerequisites

The following are prerequisites for an HPE 3PAR StoreServ system to enter an HPE 3PAR Storage Federation:

- Model is HPE 3PAR StoreServ 7000, 8000, 10000, or 20000
- HPE 3PAR OS version is 3.2.2 or later
- A free FC port on one controller node and one on its partner node, both ports to be configured in Peer connection mode
- An FC port on one controller node and one on its partner node, configured in Host Connection Mode as Target; can be shared with other hosts
- Valid Peer Motion license

The following are prerequisites to create a HPE 3PAR Storage Federation for bidirectional data mobility:

- An instance of the SSMC 2.2 or later to create and manage the federation and initiate mobility
- An account with the Super role on each system intended to join the federation
- An FC SAN fabric (dual redundant ones by preference) for Peer links and host interconnects
- Two SAN zones containing the Peer port WWNs and associated Host ports
- The host under migration is defined on all federation members the migration is planned for

The following are prerequisites for migration sources that attach to a federated system:

- Model is HPE 3PAR StoreServ 7000, 8000, 10000 or 20000 running HPE 3PAR OS 3.1.3 or later, supported EMC, HDS or IBM system satisfying minimum firmware/OS/microcode version, see SPOCK.⁴
- An FC port on one controller node and one on the partner node of the migration source, configured in Host mode; can be shared with other hosts
- Two SAN zones for FC interconnect to the federation member, single-initiator-single-target zoning preferred

The following are prerequisites for the SAN switches to which federated systems, their migration sources and hosts are attached:

- A supported model of FC switch running supported firmware; storage federation enforces no particular model or firmware beyond what is supported for HPE 3PAR StoreServ systems; see the SAN Design Guide⁵ for more information
- N-port ID Virtualization (NPIV) enabled on the ports of the SAN switch that the FC Peer ports connect to

Exactly two Peer ports are supported per system, regardless of the number of controller nodes or FC interfaces in a system. The HPE 3PAR Online Import Utility (OIU) and the Peer Motion Utility (PMU) can be operated concurrently with Storage Federation. Data transfer by OIU and PMU can take place simultaneously to the same federation member and concurrently with bidirectional migrations happening on that federation member. OIU, PMU and Storage Federation use the same Peer ports on the federated system. It is not supported nor needed to create a second set of Peer ports on the federation member for exclusive use by OIU and/or PMU. The controller node numbers and the speed of the FC ports for the Peer links do not need to be equal on the source and destination.

Supported federation configurations

An HPE 3PAR Storage Federation counts between one and four HPE 3PAR StoreServ member systems. Each system can be a member of only one federation at a time. All members of a federation must be managed by the same SSMC instance. The maximum number of federations that can be managed by an SSMC instance is determined by its maximum count of 32 systems: up to 16 two-member federations can be managed, up to 10 three-member ones, up to eight four-member ones and any mix of these that honors the maximum of 32 systems. The number of source systems for a federation member has been increased from four to eight in SSMC 3.0. Below are the current bounds of a storage federation in tabular form:

Table 1. Current bounds for HPE 3PAR Storage Federation

#	Item	Number
1	Number of member systems in a storage federation	1-4
2	Number of migration sources attached to a storage federation member	0-4
3	Number of member systems and migration sources in a storage federation	1-8
4	Number of migration destinations attached to a storage federation member	0-3
5	Number of migration sources attached to a storage federation	0-6

⁴ hpe.com/storage/spock

⁵ h20566.www2.hpe.com/hpsc/doc/public/display?docId=emr_na-c00403562&lang=en-us&cc=us

The value of four in item (2) requires a single-member storage federation. In this setup the system in the federation system has no destinations. The maximum value of 8 in item (3) for the member systems and migration sources in a federation imposes the bounds in items (4) and (5). The value of three in item (4) requires a four-member storage federation. The value of six in item (5) is valid when the federation is composed out of two members. In this configuration, three migration sources are added to each federation member. Based on table 1, up to five migration sources can be added to a three-member federation. Data from migration sources is migrated using Online Import Utility.

A migration source can be zoned for data transfer to up to all four federation members. Up to four migration sources can transfer data concurrently into one federation member. File Persona volumes cannot be migrated, this is a limitation of Peer Motion, not of the SSMC or storage federation. The use of Smart SAN⁶ is not supported with storage federation.

Orchestrating data mobility

Table 2 contains the overview of the orchestration tools for data mobility from a source to a destination system. The destination system is in each case an HPE 3PAR StoreServ running HPE 3PAR OS 3.2.2 and a member of a storage federation.

Table 2. The migration orchestration tools available for the different types of systems supported. The destination system is in each case a Federation member running HPE 3PAR OS 3.2.2

Tool	Source system type	Source system HPE 3PAR OS version	Capability
StoreServ Management Console (SSMC)	7000/8000/10000/20000	3.2.2	Bidirectional N:N
StoreServ Management Console (SSMC)	7000/10000	3.1.3, 3.2.1	Unidirectional 1:1, N:1, 1:M, N:M
Management Console (MC)	7000/10000	3.1.2, 3.1.3, 3.2.1	Unidirectional 1:1
Management Console (MC)	F- and T-class	3.1.2, 3.1.3	Unidirectional 1:1
Peer Motion Utility (PMU)	F- and T-class, 7000/8000/10000/20000	3.1.2, 3.1.3, 3.2.1, 3.2.2	Unidirectional 1:1, N:1, 1:M, N:M
Online Import Utility (OIU)	EMC, HDS, IBM XIV	See SPOCK ⁴	Unidirectional 1:1, N:1, 1:M,N:M
EVA P6000 Command View	HPE EVA P6000	See SPOCK ⁴	Unidirectional 1:1

The SSMC is the mandatory tool for concurrent, bidirectional data mobility between federation systems. The SSMC is the preferred tool for configuring HPE 3PAR StoreServ 7000 and 10000 systems on 3.1.3 or later as migration sources to a federation. Volumes of these systems are migrated unidirectionally into one or more federation systems. Reordering the pending migration tasks from different migrations sources on priority happens on a continuous basis on the destination system.

The Peer Motion wizard in the HPE 3PAR Management Console (MC) does not show destination HPE 3PAR StoreServ systems running HPE 3PAR OS 3.2.2 or later. Legacy F- and T-class systems and the HPE 3PAR StoreServ 7000 and 10000 can have their data relocated using the HPE 3PAR MC. The MC and Peer Motion Utility (PMU) can setup and initiate data relocation from migration sources attached to a storage federation member. Peer Motion Utility is a command line oriented, scriptable alternative to SSMC but can only set up a single, unidirectional migration, not a bidirectional one. Both are not federation-aware and hence cannot create or manage a federation. For the HPE EVA P6000 invoke the Online Import wizard inside HPE EVA P6000 Command View. For non-HPE systems, use the command line driven HPE 3PAR Online Import Utility.

Check SPOCK⁴ for the list of supported HPE and non-HPE storage systems that can act as a migration source and the prerequisites for them.

⁶ h20564.www2.hp.com/portal/site/hpsc/public/kb/docDisplay?docId=c04778666

Features

This section contains the description of the features of HPE 3PAR OS 3.2.2 and SSMC 3.0 that affect HPE 3PAR Storage Federation.

Federation auto-detection

SSMC 3.0 supports auto-detection of an existing federation. Assume you created a federation in an SSMC instance and for some reason set up a secondary SSMC instance. After adding the same systems into the new instance, the federation will be auto-detected from state information kept on the federation member systems. The Federations page in the new SSMC will display the existing federation(s) automatically. The user does not have to perform a file copy of the federation details from the Persistent Store on the primary SSMC instance, and restore the federation from the Federations page. The federation state for migration sources is stored on the federation member it is connected to. The Activity log on new SSMC instance is identical to the one on the "old" SSMC.

Improved orchestration

The Peer Motion dialog allows for a Basic and an Advanced mode in the SSMC. The majority of the users will be fine using the Basic mode which is the default. Clicking "Advanced options" on the "Actions" -> "Start Peer Motion" page offers these three options:

- Delete source virtual volumes

This is set to "Yes" by default, like in Basic mode. When set to "No", the source volumes will be kept on the source after a successful migration. Note that the migrated source volume becomes stale immediately after the migration completes.

- Minimally disruptive migration

Set to "No" by default. If Peer Motion requires a host migration via MDM method, the option can be changed to "Yes". This allows the user to shut down their host, make SAN zoning changes and configure MPIO, reboot and then click "Continue/Resume" on the Federations page. For online migrations, this option is not selectable.

- Pause Peer Motion before starting data migration activity

Set to "No" by default. When changed to "Yes", online Peer Motion will be suspended before importing virtual volumes, even if a host rescan or MPIO configuration is not needed. It is used when hosts or volumes are migrated that have Veritas software in the I/O stack to allow the source paths to be unzoned before importing the volume.

Filtering federations

SSMC 3.0 allows for selectively showing existing federations in the left pane of the Federations screen. Filtering is on the name of the federation. This is convenient when an SSMC instance manages a large number of federations.

Unlimited data mobility

The SSMC is the orchestration tool for 1:1, 1:M, N:1 and N:M bidirectional data mobility between storage federation members. The SSMC imposes no maximum limit on the number of migrations executed in parallel between federation members. Volumes can be moved back and forth between the same or different federation members an unlimited number of times.

Migration resilience

The SSMC shows a "Retry" button to restart the migration for those volumes that did not completely transfer. The restart moves the entire volume from the start.

Concurrent Peer Motion tasks

The SSMC enforces no hard limit on the number of Peer Motion operations that can be started between federation members and between a federation member and a migration source. The import of a volume is controlled by an import task that executes within the destination HPE 3PAR StoreServ environment. Up to nine import tasks execute concurrently, the rest of them is queued and become active when an ongoing task completes. There is no hard limit for the number of import tasks that a destination system can queue. Queuing import tasks are ordered by priority (high, medium, low) and inside a particular priority by time of creation. This ordering is discussed in the paragraph titled Import task priority.

Import task priority

For destination systems on HPE 3PAR OS 3.2.2 or later, every volume import task created by the SSMC receives a priority tag. The default priority is medium; low and high can be selected as well. The legacy FIFO queue on the destination HPE 3PAR StoreServ that holds import tasks until they become active is replaced in HPE 3PAR OS 3.2.2 by the equivalent of three queues, one for each priority level. Inside each of these queues, import tasks are ordered by time of arrival on the queue with the oldest task in front. The migration subsystem will activate import tasks first from the high-priority queue until depleted, upon which the medium, and finally the low-priority queue tasks are scheduled. In this arrangement, lower priority tasks may starve if higher priority tasks are submitted to the volume import queues on a regular basis. To resolve this, the oldest task on a lower priority queue will become active once it has been skipped for activation by a higher priority task nine times. This mechanism ensures that lower priority tasks are not completely starved out. When made active, high-priority import tasks move up to six regions of 256 MB before the queue manager turns to the next task to assign migration resources to it, medium-priority tasks move up to four, and low-priority tasks move up to two regions. The legacy global limit of 36 region moves per time slice for volume import has not changed in HPE 3PAR OS 3.2.2. Consequently, when nine high-priority import tasks are active, each of them gets on average four regions moved per time slice. Once active, an import task does not get pre-empted.

A task's priority can be changed after its task became active from the SSMC or the HPE 3PAR CLI. Import task prioritization is also available for data ingesting from volumes on a migration source into a storage federation member.

Peer links over 16 Gbps FC

Although the impact of data migration upon hosts is negligible, most customers prefer to perform a data migration in the least impactful time. With storage federation, we support 16 Gbps FC ports for the Peer links in Peer Motion and Online Import. These ports are integrated in the controller nodes of the 8000 and the 20000 HPE 3PAR StoreServ, in the new units of the 7000 series and available on a plug-in card in the controllers of all mentioned systems. With 16 Gbps SAN switches, end-to-end 16 Gbps means twice the Peer link speed compared to 8 Gbps rendering migrating workloads the benefits from the destination platform in half the time. The 16 Gbps ports are compatible with 8 Gbps FC ports when creating the Peer links for a storage federation. The faster Peer links are compatible with 8 and 4 Gbps FC ports on migration sources.

Automatic host rescan

In HPE 3PAR OS 3.2.1 and prior, Peer Motion migration workflow requires a manual rescan on the host before starting the import task for the host to discover the new destination paths. Post migration, a manual rescan removes the source volume path to clean up stale source paths. This requires a coordination between host and storage administrators and defeats the "one-click" online migration simplification that was in plan for HPE 3PAR OS 3.2.2 and later releases.

For an online migration between HPE 3PAR StoreServ systems in a storage federation, the SSMC runs through the familiar operations of the admit phase to create the Peer volumes on the destination system and export them to the host on the source system. At the end of this phase, the SSMC issues a selective SCSI LIP command to the host's FC ports causing it to rediscover devices. This rescan completes in a few seconds. The volume import command next inquires if the host(s) under migration executes SCSI TUR, INQ, or MAINT_IN commands over every path from the host to the Peer volumes on the destination system followed by whether READ, WRITE, and PROUT activity take place over the same paths. If successful, the volume import will start automatically. If none of the checks returns success, the SSMC will prompt the user in the Peer Motion task details for a manual disk rescan on the host. Minimally disruptive migration always requires a manual rescan after the SAN zoning changes.

Simplified zoning

At larger customers, the storage teams and the SAN zoning teams are different entities operating by work orders coming in. Legacy Peer Motion requires SAN zones for the Peer links to be created during the setup and modifications to them when changing the source, the destination or the direction of data flow. With HPE 3PAR Storage Federation, the SAN zones between the intended federation members must be created and enabled at birth of the federation. Multiple initiators to multiple targets per zone are supported for federation members. This reduces the number of zones required in a federation topology to just two, dramatically simplifying the initial zoning setup and the changes needed when adding or removing arrays to and from the Federation. The SSMC wizard for creating a federation advantageously produces the list of WWNs that need to enter into each of the two zones. If adhering to the single-initiator-single-target, a two-member federation requires 36 zones, a three-member one 108 zones and a four-member one 216 zones.

Local reads during migration

During the import phase of Peer Motion, all I/O requests to the Peer volume on the destination array are mirrored to both the source and target volume involved. This is crucial for application write operations to maintain source volume integrity while the import is ongoing but induces

double writes in an increasing amount as the migration advances to completion. In HPE 3PAR OS 3.2.1 and earlier, all host read I/O requests are forwarded to the source array as well. HPE 3PAR Storage Federation makes use of the optimized read algorithm in HPE 3PAR OS 3.2.2 that serves read I/O requests from the destination array for already migrated data regions. This avoids an extra hop on the SAN from the HPE 3PAR destination array to the source array over the Peer links without compromising data integrity. This feature produces lower read latencies resulting in progressively improved response time to application I/O as volume migration evolves. It compensates for the potential rise in application service time due to the increasing double writes mentioned earlier in this paragraph.

LUN ID conflict resolution

Volumes migrated using Peer Motion or Online Import retain their name, size, WWN, export host, and LUN ID on the destination array. This can give rise to a LUN ID conflict if the LUN ID of the migrating volume on the source is already in use on the destination array for the particular host. The SSMC discovers this conflict and proposes the next available ID for the volume. One can cancel the migration at this point and visit the LUN on the source system to give it an ID that is unused on the host or visit the destination system to change the duplicate LUN ID. Peer Motion Utility discovers this LUN ID conflict by default and assigns a new LUN number to the incoming volume on the destination array. This may cause an issue to applications that require data volumes to be present on a particular LUN number. The option `-autoreresolve false` causes a Peer Motion Utility migration to fail in case of a LUN ID conflict.

Support for migration of a subset of VVset

With SSMC 3.0 a subset of volumes in a VVset can be migrated to a destination HPE 3PAR StoreServ system that is a member of a federation. Volumes in VVsets usually belong together, for example to form an LVM structure or because we want to take a consistent point in time snapshot of all of them. Migrating only a subset of the member volumes of a VVset is considered an advanced feature that should be used with care. A VVset with the same name as the source system is created on the destination system if it did not exist yet. Volumes of a VVset that are migrated are removed from the source VVset, the option to keep them on the source system is not available. When the last volume in the VVset is migrated, the VVset is removed from the source system. The subset of volumes can optionally be migrated inside a consistency group.

HPE 3PAR Storage federation advantages

The data migration technology behind HPE 3PAR Storage Federation is the next-generation Peer Motion software that permits two independent Peer links pairs to connect two HPE 3PAR StoreServ systems in a federation. The SSMC uses this enhanced capability to execute concurrent bidirectional migrations between members of the same federation with a single mouse click. The standalone HPE 3PAR Peer Motion Utility can migrate volumes between HPE 3PAR StoreServ systems on HPE 3PAR OS 3.2.2 but no concurrent or bidirectional migrations are possible.

With federated systems, the SAN zoning between the systems is created and enabled upfront at the genesis of the federation. No zoning changes are needed when changing the direction of the migration or when changing the source or destination system, empowering less-experienced HPE 3PAR administrators or application managers to execute migrations. With Peer Motion Utility, SAN zoning adaptations are needed when changing the direction of the migration or when changing the source or destination system involving the cooperation of the SAN zoning team.

Federated systems can maintain an identical configuration across them by the click of a button. Without automation, this activity can be time-consuming and error-prone. This, together with the very easy setup and management of a federation and the effortless bidirectional data mobility, is a strong incentive for building storage federations in the data center.

Comparing HPE 3PAR Storage Federation with appliance-based storage migration tools

Storage federation's multi-array bidirectional data mobility offers unique advantages compared to appliance-based storage migration tools from other vendors.

Ease of installation

The HPE 3PAR Storage Federation software subsystem comes built-in into HPE 3PAR OS 3.2.2 and later, no post-installation is needed. A valid Peer Motion license on the system is all that is presumed to switch on the federation concept. Installing and configuring an external migration appliance potentially complemented by virtualization of the storage systems involves expensive professional services due to its convoluted composition. Storage federation's appliance-less data mobility concept returns weeks per year in setup and maintenance time to system administrators.

Ease of management

The SSMC adds or removes an HPE 3PAR StoreServ system to or from a federation in just a few mouse clicks and non-disruptively to the other federation members. Appliances add an extra layer of management and complexity. Adding or removing a storage system from a virtualized storage environment is complicated because of the tight coupling between the arrays and the appliance executing the virtualization. Migrations in HPE 3PAR Storage Federation initiate with one mouse click in the familiar SSMC interface already used for volume and host management.

Appliances have a user interface that is different from array management tools incurring a time-consuming learning curve for the storage administration staff. In this context, executing migrations using an appliance is not a routine affair reserving it for professional services staff only.

Performance

Migration in a storage federation is point-to-point between the members systems keeping latency low compared to an external appliance that operates inside the data path.

Retain data services

Appliances and storage system virtualization may render certain advanced or unique functions on the array inaccessible to hosts. Being a member of a storage federation, an HPE 3PAR StoreServ system retains all features of snapshots, provisioning, replication, tiering, priority optimization (quality of service), and more. In fact, by joining a federation a system's data services are expanded by the N:N concurrent, bidirectional migration option.

Reduction of cost

Appliance-based solutions use specialized hardware devices, sometimes mixed with industry-standard servers. Due to the low-volume sales of appliances, a substantial acquisition cost for hardware, software, and licenses is not uncommon. Not always accounted during the planning phase is the considerable expense for floor space, power, and cooling an appliance imposes in the data center. Some appliances are closed, mandating a vendor's intervention to make modifications decreasing agility and raising cost of ownership.

HPE 3PAR Storage Federation requires no installation, is managed from within the familiar SSMC, sports native array-to-array performance, and offers unrestricted access to all array features such as hardware-assisted thin provisioning and space reclamation, consistency groups, quality-of-service, unique snapshot capabilities, and replication protection during the migration. This renders HPE 3PAR Storage Federation clearly superior in features, functionality, and cost compared to migration appliances.

Best practices

The SSMC is a server-based management console accessed via a web browser using HTML5. Multiple federations can co-exist and be managed within a single SSMC instance. This means even large workgroups do not need more than one instance of the SSMC if the number of HPE 3PAR StoreServ systems does not exceed 32. Thanks to the auto-discovery feature in SSMC 3.0, a federation created on a primary SSMC 3.0 instance becomes visible on any secondary SSMC instance. Hewlett Packard Enterprise discourages and does not support managing a federation and a Peer Motion migration from another SSMC instance than the one the object was created on. This is because once a Peer Motion workflow is initiated, it will be internally tagged with an ownership that is connected to its SSMC instance. Any action required for this workflow (like Resume or Retry) can only be done by the owner. Only in the case of disaster recovery when the primary SSMC instance is no longer available, a secondary SSMC instance should be commissioned. In the case of disaster, the user can click the "Takeover" link on the federation's View page in the new SSMC to change the ownership for a particular workflow.

Merging federations is not possible, you need to remove the federation members from one federation to add them to a second one.

Migrating volumes of an application between systems is extremely simple with HPE 3PAR Storage Federation. While data relocation now becomes a routine operation, it should not preclude storage administrators from studying in depth the impact of the migration. Every component of the destination HPE 3PAR StoreServ will influence the performance of the application migrated to it. The following parameters should be compared between the source and the destination system to estimate the approximate performance of an application:

- System type (7000, 8000, 10000, or 20000)
- Number of controller nodes and associated on-board cache
- Disk technology (SSD, 10,000/15,000 rpm FC/SAS, Nearline)
- Number of disk drives in the source and destination CPG

- RAID level of the source and destination CPG
- Use of volume tiering by Adaptive Optimization
- The presence or absence of Adaptive Flash Cache and its size
- The I/O throughput to hosts (4, 8 or 16 Gbps for FC, 1 or 10 Gbps for iSCSI/FCoE)

While the destination system by itself may be fit for substantially improving the performance of the migrating application, its current workload may eventually preclude the migration of a demanding application. For a first time migration, the footprint of the application regarding CPU and cache load, number of IOPS to disk and their size, I/O to the SAN and/or the network, service time, wait time, and queue length has to be collected for a substantial amount of time (minimum a week) using HPE 3PAR System Reporter. The same exercise should be completed during the same timeframe on the destination system. With HPE Services estimating the frontend and backend IOPS and throughput capacity, the spare cycles on the destination system can be determined. Using this information, the storage administrator can make an informed decision based on facts whether migrating an application to another federation member is beneficial to the application and viable for the receiving system.

The SSMC creates the physical and virtual Peer ports from unused Host ports on the participating systems as part of the formation of a new federation or when adding a new member to it. Alternatively they can be created upfront using the SSMC or HPE 3PAR StoreServ CLI. If the SAN zoning for the Peer links already exists, the SSMC discovers the Peer and Host ports in use. If the SAN zones do not yet exist, verify the proposed FC ports for the Peer links are the ones you discussed with the customer. Changes to the proposed FC ports can be made in the SSMC in Actions -> Edit and clicking the pencil at the right of the lines listing the system name and the ports. Each change immediately adapts the list of WWNs in both fabrics in the section "Fiber Channel Zoning" on the Edit screen. Clicking the "OK" button will add the new Peer ports and remove the old ones. Once the zones are created for the Peer links, verify the Health section in the Federations screen for the "Normal" state.

When deleting a federation, the SSMC does not convert the Peer ports of its members back into Host ports. This is because the federation members may be engaged in setups for Online Import Utility or Peer Motion Utility that require the Peer ports to be present. All migration sources connected to a federation need to be removed first before the federation can be deleted.

The Host ports that are part of the Peer links on the source system can be shared with other hosts for I/O. This may adversely impact the speed of the migration and the I/O latency for these other hosts when a migration is ongoing. Starting concurrent migrations using the same Host ports will aggravate the situation. Careful planning to perform the migration(s) during periods of low I/O minimizes the impact to hosts and maximizes the efficiency of the Peer links shortening the migration time. In general, it is recommended to select dedicated Host ports on the source system as starting point for Peer links, and definitely so if concurrent migrations are planned over the same source Host ports.

The physical and virtual Peer ports and the Host ports they connect can be assembled in two SAN zones simplifying zoning significantly. HPE recommends taking advantage of this simplified zoning to mitigate the complexity of interconnects. Application hosts should be zoned to every member in the federation for seamless, one-click migration of their data across all federation members. The zones for hosts connected to the federation systems can be of the "single-initiator-multiple-target" type, reducing their number to the number of HBAs on the host although HPE prefers single-initiator-single-target zoning.

Systems can enter and leave a federation with a few clicks of the mouse. When a system enters the federation, the SAN zoning topology between the systems must be updated manually on the switches. Thanks to the simplified zoning, WWNs only have to be added to the two federation zones, no new zones have to be created. To eliminate zoning errors, navigate the SSMC to retrieve the WWNs per zone during or after the system was added to the federation. When a system leaves the federation, the SAN zones need to be updated as well. Look up the list of WWNs per zone in the SSMC after the system has left the federation to update the SAN zones.

An HPE 3PAR StoreServ on HPE 3PAR OS 3.2.2 without Peer Motion license cannot enter a storage federation but it can send its data to federation members attached as a migration source. A migration source can be connected to more than one federated system. Depending on the model and the HPE 3PAR OS installed on it, the SSMC, Peer Motion Utility, or the HPE 3PAR MC can migrate data unidirectionally from the migration source to federation members.

Federated HPE 3PAR StoreServ systems are installed and managed in the same way as non-federated ones. This means that operational activities like host, CPG, volume and snapshot creation, tuning of CPGs with Dynamic Optimization, tiering with Adaptive Optimization, I/O balancing with Priority Optimization and Adaptive Flash Cache, space reclamation for Thin provisioning, replication with Remote Copy, and more can be executed in the same way as on non-federated systems. The same holds for hardware maintenance activities like adding, removing, and replacing disks and disk enclosures, controller nodes and interface cards. All best practices guidance for managing a HPE 3PAR StoreServ system and the applications executing on it apply to federated systems.

If the HPE 3PAR StoreServ has a valid Peer Motion license, you can create a single system federation. Up to four legacy HPE 3PAR and non-HPE 3PAR systems can be added as migration sources to a single system federation allowing unidirectional migration to it. When adding an HPE 3PAR migration source, the SSMC shows the list of WWNs that need to go in the SAN zones for the Peer links. The Peer links for the migration sources share the Peer ports on the federated system.

Peer Motion tasks tabulated in the “Activity” screen of the SSMC receive a time stamp local to the destination federation system executing the import task. If the real time clock on the Federation members are not synchronized, the ordering of executing and executed Peer Motion tasks does not reflect reality. This may cause confusion for administrators when they analyze when a particular volume was migrated. Hewlett Packard Enterprise recommends to use a common Network Time Protocol (NTP) server for all members in the federation, all migration sources, all hosts under migration and the server executing SSMC.

In SSMC versions prior to 3.0, the federation state and the ongoing and past workflow per member system are kept in the SSMC Persistent Store located on the server executing the SSMC instance. With SSMC 3.0 and later, these details are stored on the member system. This obviates the need to back up the files in the Persistent Store regularly to cope with disaster.

The SSMC cannot migrate the same volume to more than one destination system, as this would present the volume to the same host from two arrays. A clone of that volume exported to a different host can be migrated to another federation member. Use cases for this include presenting a database for mining or a code repository for testing from a second array to a second host. An individual snapshot or a snapshot tree hanging off a base volume cannot be migrated between federation members. After migration of the base volume, its snapshots are deleted with the base volume on the source system unless selected otherwise. A snapshot can be migrated to the destination after promoting it to a base volume but that may inflate its size substantially.

A host can be connected to two federation members by a different number of FC ports. This is not recommended because application performance may degrade after migrating its volumes to a federation member with a lower FC port array connection count for the host.

The SSMC can synchronize objects like LDAP, SNMP, NTP, and syslog settings next to users, domains, and domain sets, and hosts and host sets between all federation members and from a migration source to its federation member. All items per object are passed on to all federation systems; no selective synchronization is possible.

The SSMC requires the host under migration to exist on the destination HPE 3PAR system; the host is not created automatically at migration time. This creation can be done by object synchronization with the source system using “Actions -> Sync federation” or by hand. Note that the host needs to be zoned to the source and destination with the same number of paths.

System Reporter in the SSMC contains no template showing the historical or real time IOPS, throughput and latency for the Peer ports on an array. For historical data, go to “Reports -> Create report” and select the template “Host Port—Performance Statistics”. Next click “Add Rules”, add the rule for the “Peer” port type, remove the “Host” type and finish by clicking the “Create” button. For real-time statistics on the Peer ports, select the template “Port (Data)—Real-Time Performance Statistics,” click “Add objects,” next click “Select objects,” make your selection for the Y axis and the “Type” and select both Peer ports for plotting an individual or combined series for each object. Next click “Add objects”, then remove the other nine predefined objects in the chart to increase clarity and finish by clicking the “Create” button.

When HPE 3PAR systems are within the physical distance their host and Peer ports can join the same SAN zone, a federation can be created for bidirectional data mobility. The same holds for migration sources and their federation member destination. The caveat is that, because the Peer Motion and Online Import do not migrate the server, the I/O latency for applications can grow because of the inter-site travel time. Relocating the server closer to the destination storage system is the solution.

One VVset on the source HPE 3PAR StoreServ system can be migrated at a time using the SSMC. The volumes of a VVset are automatically placed in a Consistency Group for migration; it is advised to keep this selection. To migrate more than one VVset, you need to select the volumes in them. The VVsets of the volumes are created on the destination system. If the default migration priority of medium is changed, it is recommended to set the same priority to all the volumes in the VVset. The migration priority of a volume has been added to the output of the HPE 3PAR CLI command `showtask -d` for an import task. HPE 3PAR Priority Optimization rules on volumes and VVsets are not created on the destination system after migration.

Interrupt coalescence is a widespread optimization technique in network and Fibre Channel HBA cards to moderate the number of processor interrupts when I/O takes place. Except on a high-volume single-stream I/O, HPE recommends to disable interrupt coalescence as the modern, fast CPUs in the controller nodes of the current HPE 3PAR StoreServ models handle thousands of interrupts easily with room to spare for other operations. Converting a Host port to a Peer port enables the interrupt coalescence feature automatically. It should be disabled using the SSMC or CLI.

The region mover subsystem in HPE 3PAR OS is common between tasks executing Dynamic Optimization, Adaptive Optimization, `tunevv`, `tunesys`, snapshot promotions, Peer Motion and Online Import. This means that every instance of these tasks executing during a Peer Motion operation will reduce the maximum number of nine simultaneous volume import tasks by one on the destination array lowering the migration throughput. The result of `tunesys` is compromised when running Peer Motion to a CPG subject to `tunesys`: regions dropped “late” by Peer Motion in the CPG under `tunesys` will stay in that “wrong” CPG forcing a second `tunesys` after Peer Motion ended.

Storage federation provides bidirectional data mobility between volumes of different provisioning types on different disk hardware tiers. HPE 3PAR OS 3.2.1 MU1 and later supports deduplicated volumes on the SSD tier. When a volume on SSD with deduplication enabled is migrated, the volume is rehydrated inline before traveling over the Peer links to the destination system. The deduplicated volume is not expanded upfront before the migration meaning the transfer does not incur space on the source system. If the destination volume is deduplicated as well, its final size may differ from the source volume because the “deduplication dictionary” on the destination can be different from the source one.

Do not make changes to the credentials of arrays with volumes under migration. The next authentication will fail and the migration will halt. You cannot edit a storage federation setup while a Peer Motion or Online Import migration is in progress, even when it is paused for a host rescan. You should add or remove an HPE 3PAR StoreServ system to or from a federation when no migrations are taking place. A system can be removed from a federation if no more migration sources are connected to it.

Logging

Migrations orchestrated by the SSMC are executed inside the embedded version of Peer Motion. The default location for the log directory for Peer Motion operations is the standard one for SSMC, located at `<Install drive/folder>\Hewlett Packard Enterprise\SSMC\ssmcbase\data\logs` for Windows and `/var/opt/hpe/ssmc/data/logs` for Linux. In this directory several separate logs files are kept for different types of activities.

Detailed information about past migrations by the SSMC can be filtered out of the HPE 3PAR Event Log by executing the HPE 3PAR CLI command `showeventlog -min <minutes> -msg import` on the destination system. The `-min` parameter, expressed in minutes, indicates how long to go back in time to find matching events. The command `showeventlog -min <minutes> -msg <VV name>` shows detailed information about the steps in the migration of a particular volume. Other detailed information can be obtained with `showtask -d <task number>` using the task number obtained in the `showeventlog` output.

Licensing

HPE 3PAR Storage Federation mandates a valid Peer Motion license on all federation members. The free one-year license for Online Import on HPE 3PAR StoreServ with HPE 3PAR OS 3.2.2 cannot be used for Peer Motion migration and federation. Migration sources do not require a Peer Motion license to transfer data to a federation member.

The Peer Motion license is spindle-based and can be acquired as an individual software title or as part of the HPE 3PAR Data Optimization Software Suite and its V2 version. The Peer Motion license has no restrictions on the amount of data moved, on the number of times a migration can be executed between federation members or on the number of imports started from a migration source towards a federation member.

Delivery model

Hewlett Packard Enterprise has designed HPE 3PAR Storage Federation with ease of use in mind. As a result, the preparation and SAN zoning, the actual migration and the cleanup can be executed by customers. Assistance with the setup and management of storage federation and the migrations is available from HPE Technical Services Consulting. HPE 3PAR Storage Federation can be part of a pre-packaged or a customized data migration service. Each type of service will bring in expertise, best practices and automation to deliver a successful end-to-end migration solution. Consult your HPE representative or HPE partner and this web page⁷ for more information about HPE Data Migration Services.

Troubleshooting HPE 3PAR Storage Federation

Migrations in storage federation can fail because of an issue internal to it or external. The reason of the import failure will be shown in the SSMC under Federation -> Peer Motion Summary -> Failed. External reasons can be degraded federation members on the hardware level or on the SAN or Peer link connectivity. If the reason is external, the SSMC will offer to retry the import task when the root cause was fixed. If the SSMC does not offer the retry option, recovery is outside the scope of the SSMC.

⁷ [h20195.www2.hpe.com/V2/GetPDF.aspx/4AA5-3759ENW.pdf](https://www2.hpe.com/V2/GetPDF.aspx/4AA5-3759ENW.pdf)

In general, all import is handled by the destination HPE 3PAR StoreServ alone; the SSMC just checks the prerequisites and monitors the migration progress. No specific tool for diagnosing federation issues exists, the files generated for standard SSMC debugging contain the necessary information. See the section titled “Logging” in this white paper for the location of the log files. The HPE 3PAR StoreServ Event log complements the SSMC logs with information from the viewpoint of the destination array.

The HPE 3PAR Peer Motion and HPE 3PAR Online Import User Guide³ contains an extensive chapter on how to troubleshoot the various problems that may arise during the preparation and the execution of a migration. Please consult this guide for more information.

Information collection when contacting support

When contacting HPE Support on a Storage Federation issue, you should proactively collection information and attach it to your request for help. The steps for collecting the information are outlined here. We have assumed the default installation directory for SSMC 3.0 at `C:\Program Files\Hewlett Packard Enterprise\SSMC`. If the SSMC was installed in a different directory, change the path in the steps accordingly.

1. On all HPE 3PAR StoreServ systems in the Storage Federation: gather the output of the following HPE 3PAR StoreServ CLI commands:
 - a. `showsys -d`
 - b. `shownode -d`
 - c. `showversion -a -b`
 - d. `showport -peer`
 - e. `showpeer`
 - f. `showportdev all n:s:p` with `n:s:p` the location of each of the Peer ports
 - g. `showtarget -rescan` followed by `showtarget`
 - h. `showtarget -lun`
 - i. `showeventlog -debug -online -min X` with `X` the number of minutes from now to look back for information in the event log. Make sure the value of `X` is such that it covers the entire migration operation
2. On the server executing the SSMC instance: run the file `support.bat` located in `C:\Program Files\Hewlett Packard Enterprise\SSMC` as Administrator on Windows. On Linux run `support.sh`, located at `/opt/hpe/ssmc`. This will generate a compressed file in the same directory.

Compress all output generated above and forward it to the HPE Support group when requested.

Learn more at

hpe.com/storage/3PAR



Sign up for updates



© Copyright 2016–2017 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.