

# Overcoming IT Limitations with Storage Federation

## Analysis of HP 3PAR Storage Federation

By Russ Fellows

August 2015



**Evaluator Group**

*Enabling you to make the best technology decisions*



Evaluator Group



## Executive Overview

Information is increasingly becoming the primary asset that a company controls. The efficiency with which a company manages, processes and retains information is what differentiates one well-run firm from another. Executives and IT professionals understand that, as information becomes more valuable, it is critical to retain and safeguard corporate data.

With the use of server virtualization technologies and an increasing number of Cloud deployments, organizations have a need for data movement to occur on an ongoing basis. Complicating this is the requirement that data movement be transparent to applications and may occur over geographic boundaries. Workloads must be balanced between multiple elements within a data-center and between data-centers spanning a metropolitan area. Transforming IT data-centers will require coordination of each aspect of technology—including applications, servers, networking and storage.

As the boundaries between internal and external IT diminish, Cloud Computing will become a requirement for how IT is delivered and consumed. This trend is continuing as organizations shift IT resources to common hardware and software elements. Delivering IT as a Service (ITaaS) requires standard components, extreme flexibility at scale and the capability to move applications and data without boundaries.

There have been several approaches and technologies used to enable transparent data movement; however, these solutions are not equivalent, with some creating additional issues. In this paper, we analyze several techniques used to enable transparent data movement and investigate the challenges and relative usefulness of these solutions.

As a part of this research, Evaluator Group took part in a hands-on lab evaluation of several specific capabilities of HP 3PAR StoreServ Storage, which are presented in this paper.

## Consolidated Workload Requirements

Consolidating workloads onto enterprise storage is done to increase efficiency, but this efficiency must not come at the cost of application availability. For this reason, enterprise storage systems need to maintain the ability to transparently rebalance resources without disruption while delivering high performance.

IT organizations understand that they can improve their efficiency by consolidating multiple applications onto fewer systems and platforms. However, the tradeoff has traditionally meant a reduction in agility and potentially greater risk to operations resulting from an outage. For these reasons, companies and IT organizations must carefully weigh the benefits of each architectural approach before making changes.

Consolidating storage workloads requires systems not only with greater capacity and performance, but also storage systems with higher levels of availability and manageability. With the recent emergence of solid-state storage arrays, organizations have found that they can attain high performance and capacity needs using multiple all-flash storage arrays. However, consolidation makes it imperative that the

chosen platform delivers high availability, enterprise data protection features and the performance necessary to support thousands of applications on a single platform.

## Scaling Enterprise Storage

There are a number of architectural approaches that have been used to enable storage systems to scale. Several examples from the HP storage portfolio are listed, in order of their availability:

- Distributed Shared-Nothing (i.e. Hadoop filesystem)
- Multiple Redundant Systems (Multiple mirrored controller systems – i.e. HP MSA 2040)
- Loosely Coupled Scale-Out (i.e. HP StoreVirtual)
- N + 1 Scale-Up (i.e. a single 3PAR StoreServ 20800 system )
- Federated Scale-Out (i.e. a federation of 3PAR StoreServ systems)

As listed, each approach has increasingly greater system-wide availability than the previous one. For example, the storage used by a Hadoop filesystem has low inherent availability, which is why the Hadoop filesystem creates multiple copies of data. Traditional, entry and midrange arrays use dual-controller architectures with no single point of failure, providing a good balance between cost and system availability. By loosely coupling multiple systems into a pooled resource, higher availability can be achieved. Tier-1 class storage systems tightly couple multiple nodes, any one of which can fail without loss of data access and limited system-wide impact.

With the movement to continual operations, companies and their customers now demand non-stop application availability, requiring systems that provide availability beyond “six nines”. Going beyond Tier-1 platforms is challenging, requiring new architectures and approaches to provide even higher levels of performance and availability features. The move to non-stop operations and internal IT as a Service also requires the ability to deliver storage access without any disruptions while still allowing companies to non-disruptively move data.

## Delivering Transparent Data Mobility

There are multiple occasions when data must be migrated from one storage system to another. The challenge is moving data transparently, thereby providing non-disruptive operations for applications. The high visibility of having thousands of applications offline is one of the primary reasons companies resist storage workload consolidation.

There are a variety of reasons for wanting to take a storage system offline, and typically include the following use cases:

- Storage system reconfiguration
- Application performance balancing
- Storage system upgrades or technology refresh

Technologies that can provide continual, transparent access to data from hosts are available. One solution to this problem has been to add external storage virtualization to enable transparent, continual access to data. However, some external storage virtualization products have their own limitations.

Another alternative is to federate multiple systems together, providing a pool of independent, highly available storage resources. The following sections will compare these two very different approaches.

## Analyzing External Storage Virtualization

Generally, virtualization is the process of creating a layer of abstraction that removes the direct visibility or access of underlying resources. For block storage virtualization specifically, storage virtualization hides underlying block storage resources, and creates a mapping between underlying storage devices and volumes that are presented to an application or higher layer. External storage virtualization products have been used to minimize storage disruptions due to data migrations, or system outages.

Although external storage virtualization can help address the shortcomings of a single system, it may also introduce additional complexity. One issue is the requirement to access storage and data availability services at the point of virtualization. In many cases, this requires the use of new data availability products and tools for managing storage provisioning and data protection capabilities.

External virtualization often adds cost and complexity, specifically issues including:

- Additional systems to configure and manage
- Additional points of failure requiring mitigation
- Ongoing hardware and software maintenance
- Limited scalability (lack of N+1 scale capability)
- Requires new data protection software

Perhaps the biggest issue with external storage virtualization is the fact that it adds another layer of complexity with new tools and management to existing Tier-1 storage platforms. Although the concept is powerful, storage virtualization has not gained widespread acceptance for these reasons. Clearly, alternative approaches should be analyzed before deploying external storage virtualization products.

## The Storage Federation Alternative

The concept of federation evolved from the need to increase agility and efficiency—including solving scalability and management issues—without adding overhead or complexity. The idea behind storage federation is that, by building native capabilities into the storage system architecture, IT administrators can increase efficiency and agility through enabling transparent, non-disruptive data movement while continuing to access the features inherent within each system. Thus, storage federation is the ability to natively join multiple systems together in peer-to-peer relationships rather than add additional layers to mask underlying storage systems.

---

*Evaluator Group comments: Storage federation provides multiple fault domains, while providing a logically consolidated storage platform for un-interrupted application availability, without adding the cost or complexity of external appliances.*

---

Federation works by allowing independent nodes to work together as a group of peers, forming a federation of systems. Storage federation enables scalability without additional cost or complexity and, with transparent data movement tools, can provide levels of availability beyond enterprise Tier-1 platforms, enabling continual availability required by cloud scale computing.

## HP Storage Federation

The HP storage portfolio is designed to leverage federation capabilities, both between peers of the same platform and between different architectures. HP 3PAR systems include all-flash 7450 and 20850 models, with the entire product line using the same controller and software architecture, enabling data protection compatibility and replication between models. Additionally, 3PAR systems may be federated, thereby providing another level of features and availability beyond the limitations of single Tier-1 systems.

As a consolidation platform, 3PAR StoreServ supports block and file protocols along with a RESTful interface along with features that ensure performance and efficiency for consolidated environments including:

- **Flash-Optimized Architecture** – Specifically designed for the unique behavior and performance of solid-state technology.
- **Multi-tenant Support** – The 3PAR StoreServ architecture includes N+1 controllers, with data distributed across the entire system, and supports role-based management.
- **Tier-1 Storage Architecture** – 3PAR systems are designed for virtual and cloud infrastructures using an active mesh architecture to enable simultaneous access to every volume via every controller in the system using the 3PAR StoreServ Gen5 ASIC.
- **Enterprise Data Protection** – The ability to support data-center availability with synchronous, asynchronous, three-site replication and new Remote Copy Async Streaming replication, all without external appliances or additional software.
- **Storage federation for data mobility** – HP 3PAR Storage supports federating systems to allow data movement between arrays and availability of services within a Metropolitan area. The availability of the entire federation is greater than that of any one system, while retaining the same data protection and management tools.

## Storage Federation with HP 3PAR Storage

Storage federation enables non-disruptive and transparent movement of application volumes between 3PAR systems. The critical aspect is the complete transparency of these activities, which enables storage federation to solve issues that have limited storage systems, regardless of their high-availability features or scale.

External virtualization technologies require their own management tools and typically require external data protection and replication technologies. In some cases, data protection is not included with

storage virtualization and requires a third set of products. In contrast, storage federation with HP 3PAR StoreServ enables these features by using a single platform and a single pane of glass for management.

In Table 1, we compare the principal use cases where storage virtualization and federation are deployed.

Use Case	Storage Virtualization	Storage Federation
<b>Transparent Workload Balancing</b>	<b>Yes:</b> Primary virtualization use case supports host transparent data movement between arrays	<b>Yes:</b> Supports movement of data between federated storage system peers using Peer Motion
<b>Local / Metro BC/DR</b>	<b>Yes:</b> Metro stretch clusters are supported with most storage virtualization products	<b>Yes:</b> Metro distances supported across members within a Federation when using Peer Persistence
<b>Remote BC/DR</b>	<b>Minimal:</b> Requires the use of additional technologies and data protection software. Data loss RPO tradeoffs between Sync and Async.	<b>Yes*:</b> Federation plus HP Peer Persistence, and Async Streaming Replication minimize potential loss, along with RTO and RPO metrics.
<b>Technology Upgrades</b>	<b>Limited:</b> Non-disruptive upgrades for storage; however, virtualization upgrades may not be supported for rolling upgrades.	<b>Yes:</b> Non-disruptive upgrades supported, old equipment may continue to be utilized

**Table 1: Comparing Workload Consolidation: Virtualization vs. Federation**

In Table 1 above, storage federation technology is shown to provide similar benefits as external virtualization. For remote BC/DR capabilities, storage virtualization requires additional tools and technologies. HP's 3PAR systems with storage federation can address the BC/DR needs when combined with 3PAR Peer Persistence and Async Streaming replication. Thus, storage federation alone does not meet BC/DR needs; however, federation plus HP replication does so with lower cost and complexity than the alternative.

By using HP 3PAR StoreServ federation technologies such as Peer Motion, Peer Persistence and Online Import, organizations are able to overcome many of the challenges outlined. Specifically, this combination of 3PAR StoreServ technologies can improve operational availability, flexibility and efficiency as systems are scaled from midrange to large configurations. Compared to external storage virtualization, storage federation delivers the required features without added cost or complexity.

## HP 3PAR Federation Features

The HP 3PAR platform includes several key capabilities, which enable 3PAR systems to solve several common challenges. These foundational technologies include: HP Online Import, Peer Motion and Peer Persistence.

### HP 3PAR Online Import

HP 3PAR Online Import software is a built-in native data service in the HP 3PAR Operating System that provides minimal to non-disruptive data migration capabilities from non-3PAR storage systems without

an added appliance in the data path or additional impact to host resources, all without the need for complex planning or professional services engagements.

### HP 3PAR Peer Motion

HP 3PAR Peer Motion software provides non-disruptive data mobility and load-balancing capabilities for enterprise storage without the need for complex planning or professional service engagements. HP 3PAR Peer Motion enables IT and storage administrators to balance I/O workloads across systems, or perform technology refreshes while eliminating application issues from these tasks. Peer Motion operations are managed through the 3PAR user interface, may be automated using CLI scripts, and supports throttling of data movement to ensure consistent application performance.

### HP 3PAR Peer Persistence

HP 3PAR Peer Persistence software enables 3PAR StoreServ systems located at metropolitan distances to act as peers to each other, presenting a single storage system image to attached hosts. This capability allows IT administrators to provide a high-availability solution between two sites or data centers where failover and failback is transparent to hosts and the associated applications. This solution is certified with VMware for metro storage clusters, known as vMSC and is also supported in Hyper-V and RedHat Linux environments.

## Workload Consolidation Platform Requirements

Enterprise IT organizations often cycle between having many, smaller systems and fewer, larger systems. Each approach has benefits, within limits. Using multiple systems reduces the impact of system failures. However, as the number of systems being managed increases, the sheer scale becomes an issue and companies naturally find it more efficient to consolidate their IT environments, both logically and physically. By federating systems, IT users are able to balance the centralized vs. distributed alternatives and achieve the benefits of both.

As described, transparent data access beyond “six nines” of availability (99.9999% +) is a requirement for consolidation, along with other Tier-1 features such as quality of service and secure multi-tenancy. However, in order to effectively consolidate storage, additional components must be considered including:

- Scalable performance and capacity
- Storage connectivity options

### Storage Performance Requirements

As workloads have evolved and the demand for lower latencies has increased, traditional enterprise platforms have not adapted to incorporate flash performance at a cost-effective price. Customers require both high aggregate I/O rates and low latencies for critical applications that require minimal storage delays.



Traditionally, large Tier-1 enterprise storage platforms have been utilized as a means of consolidating storage workloads. There are several reasons for this, including their high reliability and data availability capabilities, along with high aggregate performance and other features discussed.

Performance remains a primary consideration for companies of all sizes and is an even bigger consideration when consolidating applications. It has been a common practice to segregate mission-critical workloads onto different storage systems, both as a way to reduce the potential impact of an outage and as a method for ensuring consistent performance.

As workloads are consolidated, it becomes increasingly important to have tools available to manage workload prioritization. Implementing Quality of Service (or QoS) for a specific application requires the ability to prioritize one application ahead of another, in some instances limiting the ability of less important applications to consume too many resources.

### Storage Connectivity Requirements

With consistent high-performance available from new all-flash storage systems, organizations are using this ability to consolidate multiple workloads. As storage systems are able to aggregate a greater number of applications, the connectivity to these systems must be correspondingly faster.

Thus, I/O contention can occur as consolidation drives more I/O, which requires higher speed SAN connectivity. Due to queuing and I/O contention, it is possible to have connectivity delays while still not utilizing the full bandwidth potential of the storage connection.

Due to higher speed and technology enhancements, 16 Gb/s Fibre Channel connectivity offers several advantages over alternative connectivity options, including:

- Higher data transmission efficiency than 8 Gb/s FC
  - 17% less overhead using 16Gb/s, 64b/66b encoding
- Approximately 2.5X higher throughput than 8 Gb/s FC or FCoE
  - Combined greater efficiency plus lower I/O queuing
- Consistently lower latency than alternatives
- Backwards compatibility, provides transparent upgrades and consolidation
  - 16 Gb/s FC switches support 8 and 4 Gb/s connectivity to existing servers

It is also possible to achieve benefits from higher performance storage connectivity without upgrading an entire SAN. Highly subscribed inter-SAN-links (ISL connections) and connectivity between SAN switches and storage systems are often the first candidates for an upgrade. For these reasons, many companies now are beginning to transition from older 8 or 4 Gb/s Fibre Channel connectivity to new, 16 Gb/s FC SANs. In previous lab validation testing, Evaluator Group has tested the differences between 16 Gb/s FC and other storage connectivity options and found several benefits.<sup>1</sup>

---

<sup>1</sup> Evaluator Group Comparison of 8 Gb and 16 Gb SAN connectivity

## Analysis of HP Federated Storage Capabilities

In order to determine the practical usefulness of HP's federated storage capabilities using the HP 3PAR StoreServ platform specifically, Evaluator Group evaluated HP 3PAR Peer Motion and Peer Persistence in a laboratory environment.

---

*Evaluator Group comments: Evaluator Group observed several uses of federated storage, using both Peer Persistence and Peer Motion in a lab environment. The ability to transparently move data between storage systems was observed, both locally and at distance. The ability to perform these operations transparently is what sets 3PAR StoreServ Storage apart from other options.*

---

### Evaluator Group Evaluation

In particular, Evaluator Group observed how HP 3PAR StoreServ federated storage systems were able to support specific uses that are common in enterprise environments. The operations included the following:

- **Workload Balancing** – This involved moving a single storage volume from one peer member of a federation to another peer member.
- **Application Relocation** – This involved moving multiple storage volumes from one peer member of a federation to another peer member.

Both of the above features were analyzed in use as performed by HP lab personnel. Evaluator Group was able to direct activities and analyze the results of the system as HP and Evaluator Group interacted with the system. In both instances, the system performed the activities as claimed, validating the applicability of the capabilities.

### Use Case Analysis

#### Use Case 1 - Workload Balancing

The need to balance or adjust workloads is an on-going requirement for most enterprise IT organizations. Server virtualization technology has enabled these operations at the VM level, which allows IT administrators to move workloads to better utilize resources. Similarly, IT administrators need the ability to transparently move application and user data between systems to balance performance, or to assist with other IT tasks. The critical aspect is operational transparency, ensuring that applications continue to run without interruption.

The capability to migrate workloads without application downtime also enables IT administrators to offload data from a storage system in order to perform system maintenance. In this case, HP 3PAR Peer Motion software allows an application volume or set of volumes to be moved from one 3PAR system to another.

Currently, this solution supports the manual movement of workloads; however, in the future it is feasible these operations could be automated to allow unattended data migration for fine-grained workload balancing. With a single HP 3PAR StoreServ 20000 system scaling to eight controllers and multiple Petabytes of capacity, it is unlikely that continual workload balancing would be a common occurrence. More commonly, only a few applications would be migrated between federated 3PAR systems as new workloads are added to individual systems. Movement is supported between any current 3PAR systems, enabling data movement between a midrange 3PAR StoreServ and enterprise systems or other combinations as required.

Just as server virtualization has enabled IT administrators to evacuate a single system in order to perform maintenance, so too does 3PAR Peer Motion. 3PAR federated storage and the 3PAR Peer Motion software allow the storage infrastructure to continue serving applications without interruption while moving data to accommodate for changing conditions.

## Use Case 2 - Metropolitan High Availability

Extending the concept of local high-availability and failover to a second datacenter is a technique used by many enterprises to ensure continual operation for their most critical applications. Traditionally, this required complex and costly technologies such as Geographically Dispersed Parallel Sysplex (GDPS) with Mainframes. Many fault tolerant capabilities of Parallel Sysplex are now available at the application layer, or with assistance from hypervisor or operating system clustering. As a result, highly available clusters across data centers are possible at the server layer, but applications still require highly available, consistent storage at distance.

Enabling storage consistency at distance requires enterprise-class storage replication technologies. Although these capabilities exist in a growing number of arrays, the ability to do so without interruption to hosts is significantly more difficult. The three currently used methods of delivering this capability are:

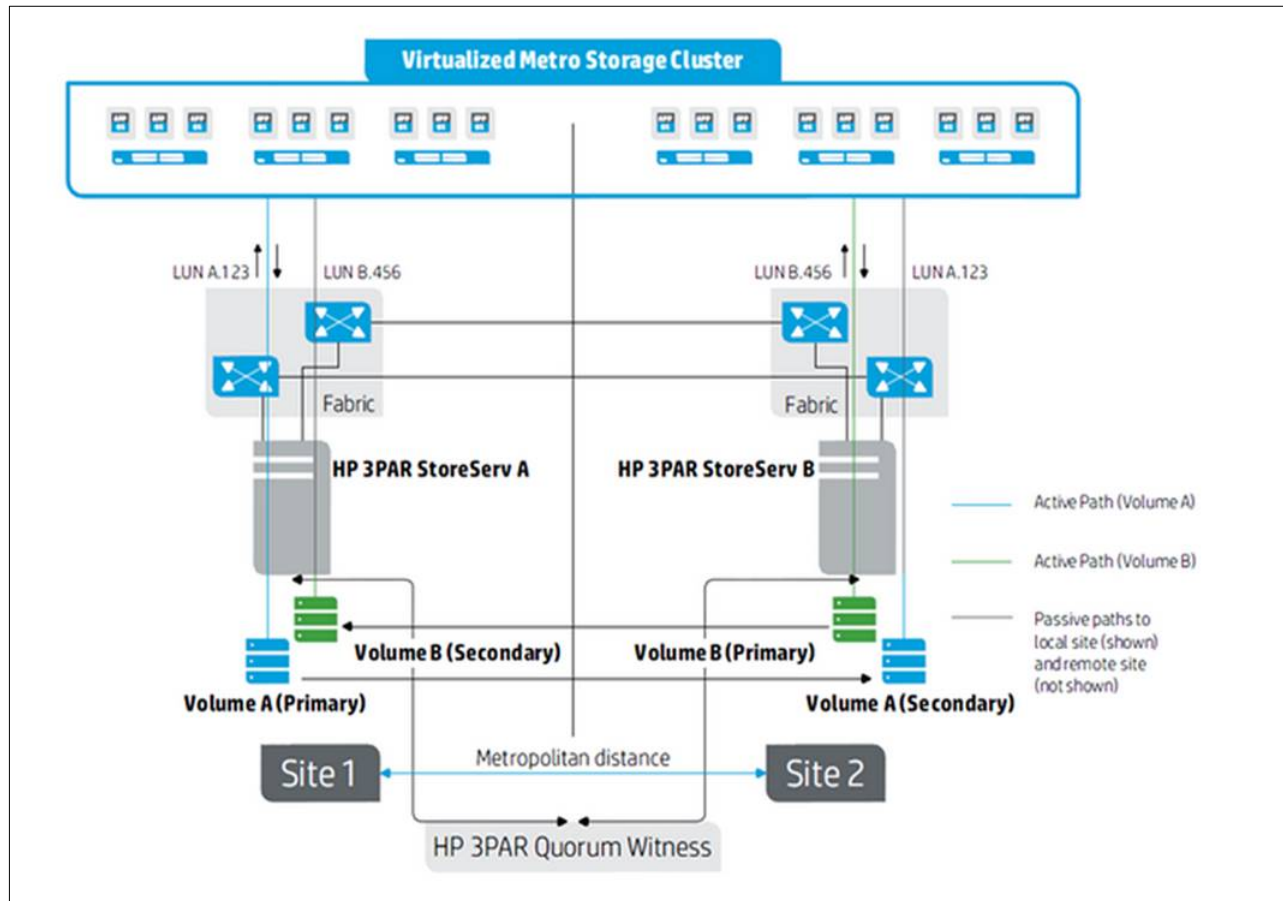
1. External storage virtualization with a split cluster
2. External storage business continuity and replication solutions
3. Persistent volume presentation within a distributed storage system

Although each of these approaches can address the need for high availability at metro distances, the first two options both require additional products that reside in the data path between the server and storage systems. Using a persistent volume presented across a distributed system does not require additional products, beyond the native replication technology of the storage system. Therefore, a federation of systems using HP 3PAR Peer Persistence provides less complexity and in turn reduces the cost and risk associated with managing additional data path elements.

### *Metro Storage Cluster*

One specific capability enabled with 3PAR Peer Persistence is support for VMware vSphere Metro Storage Cluster (vMSC), a certified configuration for stretched cluster storage. A vMSC configuration is designed to maintain data availability beyond a single physical or logical site. HP 3PAR StoreServ systems, including federated 3PAR configurations, are supported storage devices listed in VMware

compatibility guides for vMSC. Shown below in Figure 1 is a VMware Metro Storage Cluster configuration.



**Figure 1: HP 3PAR Peer Persistence HA with VMware (Source: HP)**

The concept of a Metro Storage Cluster is not limited to VMware environments. HP 3PAR Peer Persistence supports remote consistency of volumes for other environments, including MS Windows alone or with the MS Hyper-V hypervisor. In all instances, the addition of a third component is used to increase overall system availability.

The third component, known as a “Quorum Witness”, is an external application that monitors both sites for availability. The combination of these three elements ensures that the entire system can mark the correct site as being unavailable. Using a quorum provides for failover to the correct site; in some cases, an outage may appear to be a remote failure, when in fact it is the local site that is offline. These capabilities are all enabled by HP 3PAR federated storage with native Remote Copy and Peer Persistence software to transparently manage host access.

In order to provide transparent remote access, HP 3PAR Peer Persistence software uses underlying 3PAR Remote Copy software to maintain consistency at metropolitan distances. This capability allows a storage volume to be presented to hosts at two sites as the same volume. Consistency is maintained

using synchronous replication, with volume transparency accomplished within the storage SCSI layer to present the same volume as “active” to a local host, and the remote copy as a standby mirror to remote hosts. HP 3PAR Peer Persistence software manages host access to volumes, enabling highly available data to exist at two sites where switchover and switchback of access to the volumes across arrays remains transparent to the hosts and applications running on those hosts.

The effect of these technologies is that the failure of any single element has no impact on the running applications, including an OS or hypervisor at either location. Only a complete site failure would result in an outage, which would be limited in duration to the time necessary to restart virtual machines at the surviving site. Site failover would occur automatically with the use of 3PAR Quorum Witness using automated transparent failover feature.

### **Use Case 3 - Remote Business Continuity and Disaster Recovery**

The support for business continuity and disaster recovery (BC/DR) at remote distances presents challenges to IT organizations of all sizes. Beyond local metropolitan distances, replication speeds and delays necessitate the use of different storage replication technologies.

Transparent failover across distances greater than 100 ms of latency are typically not feasible due to the inherent delays created, which are beyond the tolerance for most applications. A competing product to HP 3PAR, EMC VPLEX uses external appliances to provide distributed storage access. In the case of VPLEX Geo, distance limitations remain at less than 100 ms latency to maintain consistency between sites; beyond these distances additional technologies are required.<sup>2</sup>

The solution to this challenge has been to adopt a “3 Site Replication” solution, named due to the use of three sites; a local site with synchronous replication and a 3<sup>rd</sup>, remote site for disaster scenarios with asynchronous replication. Only a small number of storage platforms support 3 Site Replication, typically only the largest Tier-1 storage systems. HP 3PAR supports this Tier-1 feature across its portfolio, including between midrange 3PAR 8000-series and large enterprise-class 20000-Series systems.

An important component to the HP solution is the use of 3PAR Streaming Asynchronous Replication. With this technology, the recovery time objective (RTO) and recovery point objective (RPO) windows can be decreased to seconds when recovery is required at the 3<sup>rd</sup>, remote site. With failover to a metro location, the HP 3PAR is able to failover without any loss of data, and without application disruption.

Most significant, these capabilities are available without the use of additional replication or virtualization products, as is required with some competing solutions.

### **Use Case 4 - Technology Refresh**

Although large-scale technology refreshes may occur infrequently, refreshes and upgrades must occur and are often cited as one of the most disruptive activities. It is not uncommon for data migration projects between storage platform to require months of planning and execution.

---

<sup>2</sup> <http://www.emc.com/collateral/hardware/white-papers/h8214-application-mobility-vplex-geo-wp.pdf>

Many enterprise class storage systems permits system upgrades and migration between systems of the same type. However, these operations often require application downtime unless additional software or external appliances are introduced. Since most storage systems do not support dissimilar hardware or software, these systems are limited in their ability to be upgraded without service disruption.

A comparison between storage federation and storage virtualization a technology refresh is useful. Although external storage virtualization enables the transparent movement of data between underlying storage systems, these systems typically do not permit a refresh of the virtualization platform without disruption. Thus, although data may be moved between underlying systems transparently, the access to storage through the virtualization product creates additional migration and upgrade concerns.

In contrast, federated HP 3PAR StoreServ systems are able to provide high availability and transparent access to data during a technology refresh. With HP 3PAR Online Import software, HP 3PAR StoreServ systems are able to migrate data between any generation 3PAR StoreServ Storage systems as well as between 3PAR and 3<sup>rd</sup> party systems without application interruptions. The list of systems supported for on-line import includes other HP systems, HDS and EMC storage products.

---

*Evaluator Group comments: The importance of the capability to transparently migrate data between old and new systems within a 3PAR federation cannot be overstated. It can reduce expensive, costly and operationally risky data migration projects from months to hours. With 3PAR Online Import, existing data residing on 3<sup>rd</sup> party platforms can be migrated to 3PAR systems, providing transparent migration for heterogeneous environments as well.*

---

HP 3PAR storage federation technology extends the clustering concept, providing a federation of the underlying Tier-1 systems. This additional layer supports dissimilar hardware and software versions. By combining this capability with transparent data migration already explained, it is now possible to perform rolling technology upgrades without interruption to host access or critical applications.

## Summary

In this paper, the concept of federation was examined in several typical use cases in IT environments, in particular for organizations that require highly available applications, irrespective of the underlying hardware or geographic location of individual data-centers. The stated strategy of HP is to enable federation between systems within the same family and interoperability between product lines. HP currently supports federation between HP 3PAR StoreServ systems as well as federating HP StoreOnce backup systems for data protection.

Additionally, interoperability between 3PAR primary and StoreOnce backup systems is available through the Express Protect feature of HP StoreOnce Recovery Manager Central software. This enables copying 3PAR snapshots directly to an HP StoreOnce system to be used as a data protection points, without additional backup software. Although currently supported only in conjunction with VMware vSphere, HP has stated plans to extend this capabilities to other environments.

Storage federation between HP 3PAR StoreServ systems enables IT departments to significantly improve application availability while decreasing operational risk and costs. Specifically on the 3PAR StoreServ platform, 3PAR federated storage capabilities are enabled with Peer Motion and Peer Persistence, providing higher levels of availability through transparent data access both locally and between data-centers.

---

*Evaluator Group comments: The HP approach to providing always-on applications is to use federated storage. HP 3PAR federated storage systems, together with native 3PAR data protection and availability tools, enable organizations to overcome the limitations of typical enterprise storage without the cost and complexity associated with storage virtualization.*

---

The concept of storage virtualization is powerful, yet the products have not gained widespread adoption for several reasons; for one, additional layers and data protection software add expense and complications. These solutions typically use an external storage virtualization appliance, additional replication products or in some instances, both. HP's solution using federation is simpler than these alternatives, in that no additional products are required. Placing new products into the data path has significant risk. By eliminating these additional elements, HP 3PAR federated storage capabilities offer a solution to common enterprise requirements without unnecessary risk, cost or complexity.

The ability to transparently move data between 3PAR systems to perform workload migration, disaster avoidance or online migration are all important considerations for IT departments. HP enables these use cases using HP 3PAR Peer Persistence, Peer Motion and Online Import, which provide high availability during system migrations, upgrades or between locations **without interruption** to applications. For companies that need to move beyond Tier-1 storage limitations, HP 3PAR federation technologies can deliver practical results.



HP Publication : 4AA3-6483ENW

## About Evaluator Group

*Evaluator Group Inc. is dedicated to helping **IT professionals** and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver **in-depth, unbiased analysis** on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and **education**. [www.evaluatorgroup.com](http://www.evaluatorgroup.com) Follow us on Twitter @evaluator\_group*

## Copyright 2015 Evaluator Group, Inc. All rights reserved.

*No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, consequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.*

This document was developed with HP funding. Although the document may utilize publicly available material from various vendors, including HP, it does not necessarily reflect the positions of such vendors on the issues addressed in this document.