HPE GreenLake

Technical white paper

Transitioning from HPE Nimble Storage and HPE Alletra 5000/6000 to HPE Alletra Storage MP B10000

Considerations when migrating data to HPE Alletra Storage MP B10000

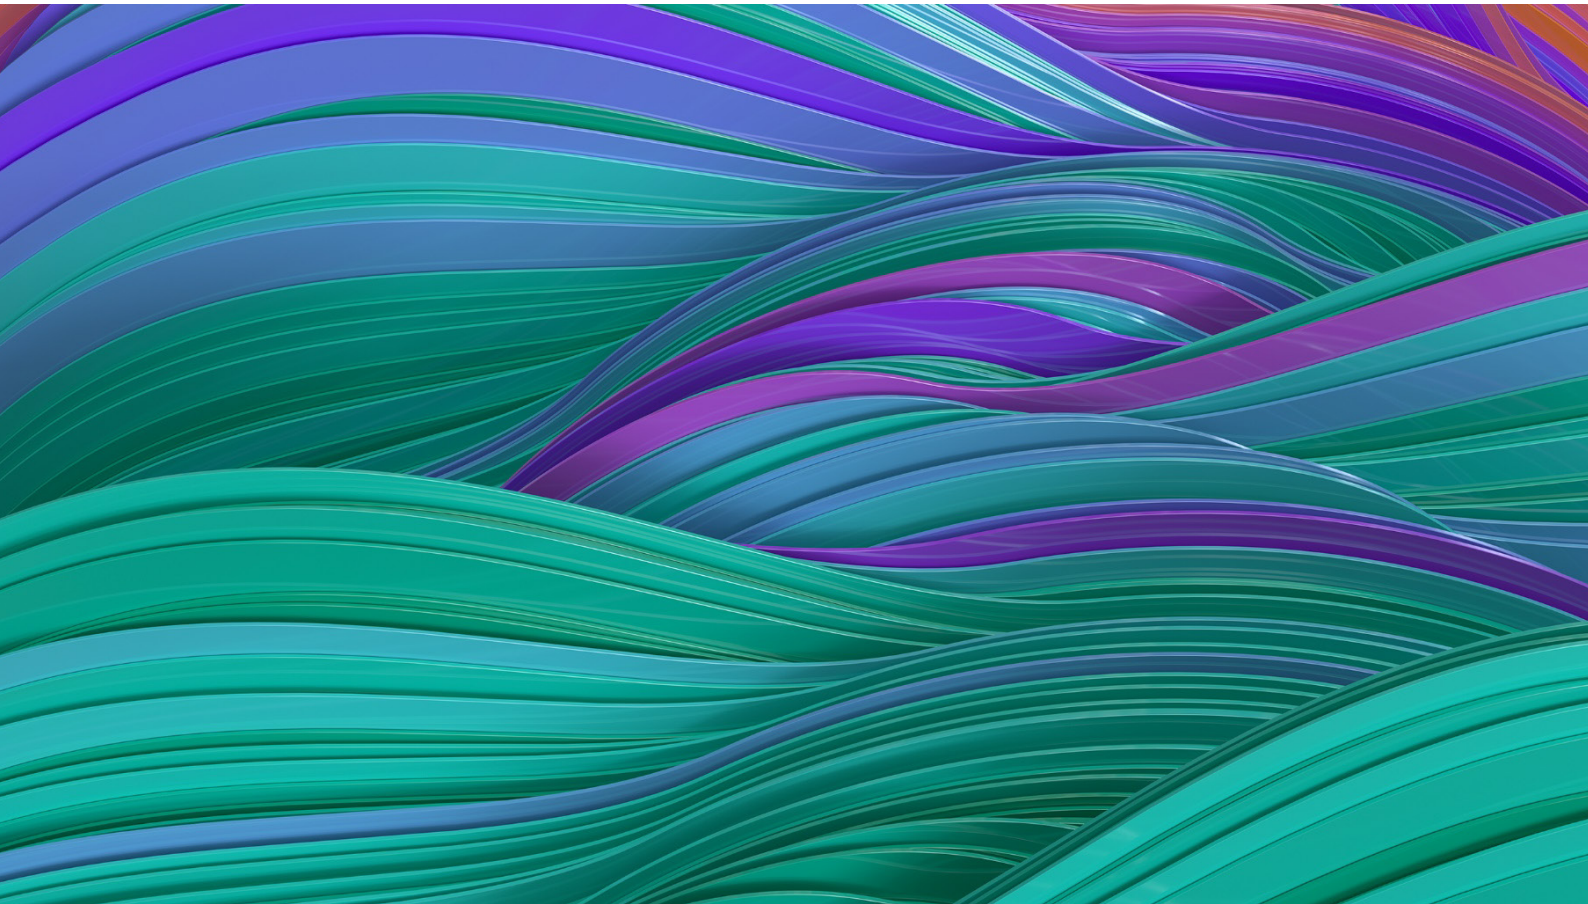

Contents

Executive summary

In 2023, Hewlett Packard Enterprise introduced HPE Alletra Storage MP B10000: intelligent, disaggregated, scale-out block storage that provides mission-critical performance and availability at a midrange price point. Built on the new HPE Alletra Storage MP hardware platform and managed through the HPE GreenLake cloud, the all-active NVMe architecture delivers extreme resiliency and ultra-low latency performance for mission-critical applications with an intuitive cloud management experience. As enterprises age of out of their existing HPE Nimble Storage and eventually out of HPE Alletra 5000/6000 storage solutions, HPE Alletra Storage MP B10000 is the perfect upgrade—offering higher performance, unmatched resiliency backed by a 100% availability guarantee[1], and a competitive, midrange price point.

This white paper is intended for solution architects, field consultants, and IT specialists, as they look to refresh their existing HPE Nimble Storage and HPE Alletra 5000/6000 systems with HPE Alletra Storage MP B10000. The objective of this document is to lay out the data migration options available to the storage administrator and guide them in choosing a data migration solution between storage platforms.

Although this paper covers some key platform differences and explains how they affect the planning and execution of data migration, it does not cover every difference and is not intended to replace the administrator guides or user guides for either platform.

Learning HPE Alletra Storage MP B10000 for HPE Nimble Storage users

The HPE Nimble Storage and HPE Alletra 5000/6000 platforms are based on different software and hardware architecture from that of HPE Alletra Storage MP B10000. The differences in the platforms mean some important considerations must be accounted for when preparing to migrate from one to the other. This section breaks down certain key considerations by topic:

- Architecture
- Performance and sizing
- Volumes
- Replication

Architecture

HPE Nimble Storage and HPE Alletra 5000/6000 platforms have an active-standby design: the active controller handles all I/O and management operations with the standby controller ready to take over in the event of an active controller failure. This architecture guarantees 100% controller headroom since, by design, array resources are never more than 50% busy—a node failure has no impact on performance. However, this design results in unused array resources, and individual arrays cannot be scaled out.

The all-active controller architecture of HPE Alletra Storage MP B10000 offers superior performance and resiliency. All nodes in the cluster are active and handling workload. Volume ownership is load balanced across the cluster, any node can process I/O, and all nodes see all drives. The scalable design of the switched architecture allows you to add controller nodes and capacity independently as your requirements increase. HPE Alletra Storage MP B10000 can withstand an entire controller node failure without impacting data availability. Volumes owned by a failed node are automatically redistributed evenly across the remaining nodes with no interruption to workload.

Table 1. Architecture comparison between HPE Nimble Storage and HPE Alletra Storage MP B10000

Feature	HPE Nimble Storage	HPE Alletra Storage MP B10000
Clustering	Active-standby	All-active
Controller nodes	2[1]	2–4
Capacity (raw)	• All-flash: 6 TB–1104 TB • Hybrid: 11 TB–1470 TB	15.36 TB–5.24 PB
Media type	• NVMe SSD (HPE Alletra 6000) • SAS SSD (AF and HF models) • HDD (HPE Alletra 5000 and HF models)	NVMe SSD (TLC and QLC)
Availability	6-Nines Availability Guarantee[2]	100% Availability Guarantee
Encryption	On a per-volume basis	On a system basis (using self-encrypting drives)

[1] HPE Nimble Storage allows you to group up to four arrays in a scale-out configuration. Eight controller nodes total can be managed in a group. This is not an all-active architecture with shared back end like HPE Alletra Storage MP B10000.
[2] Brochure: HPE Get 6-Nines Guarantee, September 2021.

Performance and capacity

Dramatic performance improvements can be expected when upgrading from HPE Nimble Storage to HPE Alletra Storage MP B10000 as a result of the all-active architecture, latest generation hardware, and faster drives.

The scalable architecture of HPE Alletra Storage MP B10000 also supports a larger amount of raw storage capacity than HPE Nimble Storage—over 5 petabytes of raw capacity in a single system with the current release (10.4.5). The B10000 supports adding capacity in two drive increments, enabling cost-effective capacity expansion.

Advances in data reduction extend that capacity with greater cost efficiency. HPE Alletra Storage MP B10000 includes several data reduction and data compaction technologies to reduce storage capacity utilization, backed up by a 4:1 data compaction ratio guarantee. These features include advanced deduplication with Express Indexing, Efficient Compression, Data Packing, Thin Provisioning, Adaptive Sparing, and Virtual Copy. For more information on the B10000's data reduction efficiency, refer to the HPE Alletra Storage MP B10000 data reduction efficiency technical white paper.

Volumes

Key differences between HPE Nimble Storage platforms and HPE Alletra Storage MP B10000 include supported maximum volume counts, sizes, and provisioning policy options. When preparing to migrate data to HPE Alletra Storage MP B10000 from HPE Nimble Storage, it is important to adjust volume and policy creation on the target array to account for the factors highlighted in Table 2.

Table 2. Volume limitations based on the latest OS release as of publication

Features	HPE Nimble Storage	HPE Alletra Storage MP B10000
Volume provisioning options	• Compression • Deduplication • Both	• Thin • Data reduced (both compression and deduplication enabled)
Max number of snapshots	Hybrid models: 190,000 All-flash models: 300,000	64,000
Volume Access Protocol	• iSCSI • Fibre Channel	• Fibre Channel • iSCSI • NVMeoF—FC • NVMeoF—TCP
Virtual Lock[3]	Yes	Yes
Max number of volumes	Hybrid models: 1024 All-flash models: 10,000	64,000
Max volume size	127 TiB	64 TiB

[3] Virtual Lock is a data protection technology that prevents volumes and snapshots from being deleted until a specified retention period has passed.

When preparing to migrate to HPE Alletra Storage MP B10000, some adjustments might be required on the target array before migration:

- **Volume size:** Volumes larger than 64 TiB must be broken into smaller volumes when they are created on the target array before data migration.
- **Volume name length:** When target volumes and snapshot policies are created on the new array, volume names cannot exceed 31 characters. If a replication or snapshot policy is applied using DSCC, the limit is 14 characters.
- **Provisioning options:** Capacity planning and performance sizing on the new array must account for whether workloads will be deployed on thin volumes or on data reduction volumes with both deduplication and compression enabled.
- **Snapshot scheduling limitations:** HPE Alletra Storage MP B10000 can support only one replication and one snapshot policy per volume/volume set, and snapshot retention policies should be adjusted to account for the upper limit of snapshots when reconfiguring snapshot schedules on the target array.

Note

For a comprehensive list of supported features for HPE Alletra Storage MP B10000 refer to the HPE Alletra Storage MP B10000 support matrix.

Replication

Both HPE Alletra Storage MP B10000 and the HPE Nimble Storage array family offer the ability to configure and remotely protect volumes with two main options:

- Synchronous replication
- Periodic asynchronous replication

For both platforms, replication can be applied as a policy to sets of volumes. On HPE Nimble Storage, these volume sets are referred to as **volume collections**, and **upstream/downstream** is used to refer to the direction of replication. On HPE Alletra Storage MP B10000, they are referred to as **volume sets**, or sometimes as **Remote Copy groups** (when managed in the CLI). **Primary** and **secondary** are used to refer to which array owns source volumes and the replication direction (for asynchronous replication and non-HA synchronous replication).

Table 3 provides more information about specific replication setup considerations and limits.

Table 3. Replication comparison based on the latest OS release as of publication

Features	HPE Nimble Storage	HPE Alletra Storage MP B10000
Replication link support	IP links (can be dedicated or shared with management and iSCSI links)	Dedicated IP links (must be on a different subnet from management and iSCSI) Dedicated FC links
Synchronous and asynchronous replication between the same two systems	No	Yes
Asynchronous replication of the same volume to multiple targets	Yes (up to 2)	No
Synchronous and asynchronous replication of the same volume to multiple targets	No	Yes (with Synchronous Long Distance and 3-Datacenter Active Peer Persistence)
Asynchronous minimum RPO	5 minutes	30 seconds
Maximum volumes replicated asynchronously	1024	2-node: 2400 4-node: 9000
Maximum volumes replicated synchronously	1024	2-node: 800 4-node: 9000
Maximum volumes in a protected set	50	2-node: 800 4-node: 1800

With HPE Nimble Storage and HPE Alletra 5000/6000, synchronous replication is possible only by grouping two arrays together as part of a Peer Persistence configuration. When two arrays have been grouped, they can exclusively perform synchronous replication between each other—they can no longer asynchronously replicate between each other. For more information, see the Peer Persistence and Active Peer Persistence section in this document.

For a comprehensive guide to the prerequisites and administration of replication, refer to the HPE Alletra Storage MP B10000 UI user guide or to HPE Alletra Storage MP B10000: Getting started with data replication using Remote Copy and the CLI user guide.

For limitations and supported features regarding replication on HPE Alletra Storage MP B10000, refer to the HPE Alletra Storage MP B10000 support matrix.

Peer Persistence and Active Peer Persistence

HPE Nimble Storage and HPE Alletra Storage MP B10000 offer high-availability synchronous replication solutions called, respectively, **Peer Persistence** and **Active Peer Persistence**. Both solutions enable zero-RPO and zero-RTO protection of volumes between two arrays with the help of a third-site quorum witness. However, there are considerable differences in setup and functionality.

With HPE Nimble Storage, synchronous replication with Peer Persistence requires the arrays to be grouped together. Grouping arrays for Peer Persistence has strict requirements, including the following:

- Both arrays must be of the same model.
- They must have identically configured subnets for management and data traffic.
- They must be on the same OS version.

After being grouped, the arrays are managed as a single entity with one array (the Group Leader) owning the primary management interface and maintaining group configuration data, and the second array (the backup Group Leader) taking over in the event of a Group Leader failure. Automatic switchover (ASO) applies to both the array group management and any replicated volume collections. Both arrays can own upstream or downstream volumes during regular operation; however, the Group Leader array takes over all volume collections in the event of failure. Although it is uncommon, this grouping of systems can result in losing access to the backup Group Leader array management and volumes if a quorum witness is not configured or is down and replication links are lost.

With Peer Persistence on HPE Nimble Storage, both the source and target volumes are presented to the hosts. The upstream (source) volumes are presented with active paths and the downstream (target) volumes are presented with standby paths. Only the array that owns the upstream volume can handle I/O to that volume.

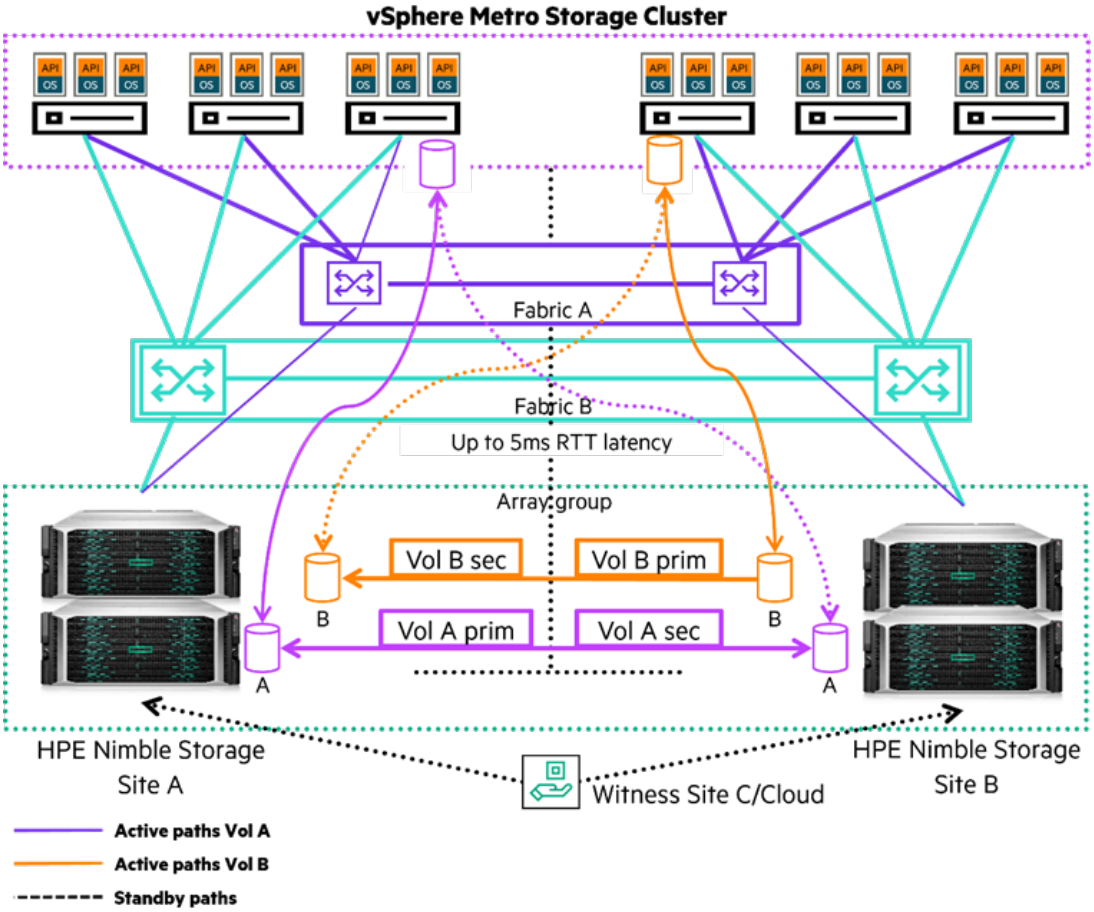

Figure 1. VMware vSphere® Metro Storage Cluster with HPE Nimble Storage Peer Persistence

With HPE Alletra Storage MP B10000, arrays performing synchronous replication are managed and treated as separate entities. Replication can be configured easily with fewer limitations and a simpler failure domain. There is no concept of a Group Leader array and Backup Group Leader array with HPE Alletra Storage MP B10000 and Active Peer Persistence. The labels **primary** and **secondary** are used to refer to which array owns a replicated volume set—arrays can be primary for some replicated volume sets and secondary for others. Like the primary/secondary role, the Active Peer Persistence and automatic failover policies are applied strictly on a per-volume set (sometimes referred to as Remote Copy group) basis. This enables more granularity in managing replicated volumes and allows the co-existence of Active Peer Persistence and plain synchronous replication. Arrays are referred to as primary or secondary exclusively in the context of replicated volume sets—not the entire replication configuration.

Because the arrays are not tightly coupled the way HPE Nimble Storage groups are, you do not run the risk of losing management capability for an array. There are no concerns about non-replicated volumes becoming unavailable on either array if a quorum witness is not configured or in the event of simultaneous quorum witness and replication link failure. The failure of all replication links in a configuration would not result in one array suddenly having to take over all I/O for protected volumes.

Important

With HPE Alletra Storage MP B10000, arrays are referred to as primary or secondary exclusively in the context of replicated volume sets—not in relation to the entire replication configuration. An array can be primary for some replicated volume sets and secondary for others.

Active Peer Persistence allows both arrays to actively service I/O to a replicated volume. Paths can be either active/optimized or active/non-optimized between a host and both arrays, and they are configured through host proximity settings. This active-active volume access grants important benefits over the active-standby implementation of Peer Persistence on the HPE Nimble Storage and HPE Alletra 5000/6000 platforms:

- **Localized I/O:** For stretched clusters—where some hosts are co-located with one array and other hosts are co-located with the second array in a different datacenter—a localized host proximity can be configured which enables I/O to be serviced locally for the same volume. The array local to the host services the I/O, regardless of whether that array is the primary or secondary. This grants substantial performance improvements for remote hosts, particularly for read I/O.
- **Symmetric I/O:** For hosts and arrays that are all located in the same data center, a symmetric host proximity can be configured. Having a symmetric host proximity enables I/O for a replicated volume to be serviced in a round-robin fashion across the active paths on both arrays rather than just the primary (upstream) array. This enables load balancing without the need to manually create separate replicated volume sets for each array, or to manually hand over volume sets between the primary and the secondary array.

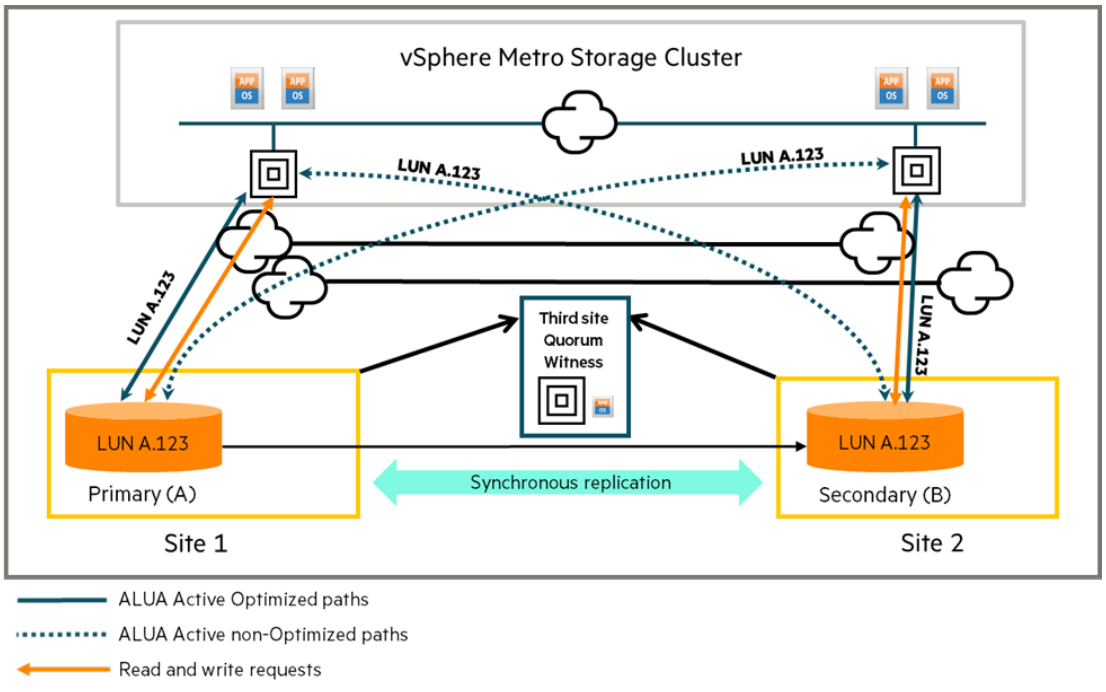

Figure 2. vSphere Metro Storage Cluster with Active Peer Persistence showing ALUA pathing for a single LUN

Active Peer Persistence also offers faster failover times, limiting the impact of failover on sensitive applications, and offers automatic failure recovery for a wider range of scenarios. Active Peer Persistence can transparently recover from the following failure scenarios:

- **Array failure:** If the primary array in the Active Peer Persistence configuration fails, the surviving array initiates an automatic transparent failover (ATF) of any groups owned by the failed array and services all I/Os.
- **Replication link failure:** The primary array for the Active Peer Persistence group continues to service I/O. No data is replicated to/from the secondary array. The ALUA paths to the secondary array for the replicated volumes transition to unavailable, and all host I/O is directed to the primary array for the Active Peer Persistence group.
- **Host disconnect** (host loses all connection to its optimized array): The host that loses connectivity on its active optimized paths uses its active non-optimized paths. If the primary array for the Active Peer Persistence group loses all paths to the hosts, it facilitates an ATF of group ownership to the secondary array. Replication continues as long as the replication links are still available.
- **Array reboot/shutdown:** If an array is manually rebooted or shut down, an ATF occurs for any Active Peer Persistence groups in which the rebooting/shutdown array is primary.
- **Primary RAID failure:** If the primary array undergoes a RAID set failure, an ATF to the secondary array will occur.
- **Metadata corruption:** If primary volumes in an Active Peer Persistence group experience unrecoverable metadata corruption, an ATF to the secondary array will occur.

Table 4. Peer Persistence comparison based on the latest OS release as of publication

Features	HPE Nimble Storage (6.1.2.400)	HPE Alletra Storage MP B10000 (10.3.0)
Peer Persistence volume access	Active-standby	Active-active
Peer Persistence subnet requirements for replication partners	Partner systems must be on the same subnet(s) to be grouped for Peer Persistence	Partner systems replication links can be on different subnets
Peer Persistence partner system requirements	Partner systems must be identical models to be grouped for Peer Persistence	Partner systems do not need to be identical models. HPE Alletra Storage MP B10000 can also do Active Peer Persistence with HPE Primera and HPE Alletra 9000. The best practice for Active Peer Persistence is to use similarly capable systems.
Peer Persistence maximum RTT latency between arrays	5 ms	10 ms
Failover time	~45 seconds	~10 seconds
Support for 3-Datacenter Peer Persistence	No	Yes

For comprehensive technical details about HPE Active Peer Persistence with HPE Alletra Storage MP B10000, refer to the HPE Active Peer Persistence technical white paper. For more information about HPE Peer Persistence with HPE Nimble Storage and HPE Alletra 5000/6000, refer to the HPE Nimble Storage and HPE Alletra 6000 Peer Persistence technical white paper.

Overview of data migration methods

From the standpoint of impact to the host, data migration methods can be classified as **online**, **minimally disruptive**, and **offline**:

- **Online migration:** Data from the source array to the target array is migrated without causing any disruption to the host I/O. Throughout the migration, applications and the operating systems are not disrupted during the data transfer.
- **Minimally disruptive migration:** With this approach, data from source array to the target array is migrated without causing any impact to the application. However, the host I/O is disrupted for a short period of time during the cutover process while you configure the host multipath solution on the destination storage system.
- **Offline migration:** Data is migrated from the source to the target array completely offline, suspending the host I/O throughout the migration process.

Table 5. High level comparison between migration methods

	HPE Peer Motion migration	Host-based migration	Migration with Zerto	HPE Services Direct Block Migration	HPE Services Appliance Migration
Migration type	Online, minimally disruptive, offline	Online	Minimally disruptive	Online, minimally disruptive, offline	Minimally disruptive
Protocol	iSCSI	FC, iSCSI	FC, iSCSI	FC	FC, iSCSI
Host/app disruption	No to minimal impact	No to minimal impact	Minimal disruption on VM cutover	No to minimal impact	No to minimal impact
Performance impact	Minimal impact	Method dependent	Minimal impact	Minimal impact	Minimal impact
Cost	Free (all-inclusive software license)	Method dependent	$ (Zerto license required)	$$ (Service engagement—no license required)	$$$ (Service engagement—license required)

HPE Peer Motion data migration

HPE Peer Motion software enables flexible, online, native block-level data migration to HPE Alletra Storage MP B10000. This software feature is a time-tested, reliable method for performing data migrations to the latest generation of storage array starting with the HPE 3PAR product line. Historically, it has only been an option for performing data migrations from HPE 3PAR-based source arrays (HPE 3PAR, HPE Primera, and HPE Alletra 9000) over Fibre Channel SAN protocol.

With the 10.4 release of HPE Alletra Storage MP B10000 OS, the ability to perform online iSCSI-based migrations from Gen3 and Gen5 HPE Nimble Storage arrays, HPE Alletra 5000, and HPE Alletra 6000 has been introduced.

Peer Motion data migration from HPE Nimble Storage can be orchestrated using the onboard HPE Alletra Storage MP B10000 UI or in Data Ops Manager on the HPE GreenLake.

The following entities can be migrated using Peer Motion:

- Volume-Scoped volumes
- Group-scoped volumes
- Hosts (initiator groups)

See the Migrating data from HPE Nimble Storage to HPE Alletra Storage MP B10000 using Peer Motion Technical White paper for comprehensive details on how Peer Motion works and step-by-step procedures to migrate data with it.

Host-based data migration

Host/application-based data migration refers to migrations that are performed at the operating system/application level. Examples of application-based data migration include Oracle® Automatic Storage Management (ASM), Oracle Data Guard, and Microsoft SQL Server migrations to Microsoft Azure. Commonly used host-based data migration methods include Logical Volume Manager (LVM) mirroring and VMware vSphere® Storage vMotion®.

Benefits of host or application-based migrations are that they are largely storage agnostic, protocol agnostic, and can typically be done completely online or with minimal disruption to the host or application. However, they can be complex to manage and methods like disk mirroring tend to cause performance impacts.

It is important to validate that the host type is supported by HPE Alletra Storage MP B10000 prior to performing data migration. Some data migrations, particularly for older deployments, require host OS upgrades in addition to storage migration.

Performing data migration using VMware vSphere Storage vMotion

vSphere Storage vMotion allows you to migrate your virtual machines (VMs) to different compute and storage resources within your vSphere environment with zero downtime. Using this technique, you can migrate VMs currently running on older storage to newly created datastores on the target HPE Alletra Storage MP B10000 array.

Using vSphere Storage vMotion to move data involves the following tasks:

1. **Prepare**: Create and export the new volumes from the HPE Alletra Storage MP B10000 system to the VMware ESXi™ hosts and configure the new VMFS datastores.

   Verify that vSphere Storage vMotion is enabled on the network and that the new datastores are present and accessible to all hosts if migrating a clustered environment.
2. **Initiate vSphere Storage vMotion**: Identify the VM you want to migrate in the vSphere Client, select **Migrate**, then choose **Change storage only**.
3. **Select destination**: Choose the destination datastore where you want to move the VM.

   This can be done as a batch configuration or on a per-disk basis.
4. **Review**: Verify that the migration settings are correct, including the selected datastore and any advanced options.
5. **Begin migration**: Start the migration process.

   vSphere Storage vMotion transfers the VM to the destination datastore without downtime.
6. **Monitor progress**: Monitor the migration progress in the vSphere Client to validate that it completes successfully.

   I/O begins to appear on the target HPE Alletra Storage MP B10000 array as the VM datastore and workloads running on the VM are migrated.
7. **Verify completion**: After the migration is complete, verify that the VM is running on the new datastore and that there are no issues with its operation.

Refer to the HPE Alletra Storage MP B10000: VMware ESXi Implementation guide for information about connecting the storage array to the host, protocol details, and required host settings.

Performing data migration using LVM disk mirroring

LVM mirroring, or disk mirroring, is a procedure that enables you to migrate data on a host by presenting the target block volumes to the host over a supported protocol such as iSCSI or Fibre Channel. Using this technique, data is mirrored between the original storage device presented from the old storage and the new device presented from the new HPE Alletra Storage MP B10000.

Using **lvconvert** for LVM mirroring involves the following tasks:

1. **Prepare**: Create and export the new volume from the HPE Alletra Storage MP B10000 system to the host.

   Validate the new raw device is present through the appropriate multipathing.
2. **Create the physical volume**: Use **pvcreate** to initialize partitions on the target disk for use with LVM.
3. **Add the new volume**: Use **vgextend** to add the new volume to the existing volume group (VG) containing the old volume.
4. **Mirror the logical volume**: Use the **lvconvert** command to begin mirroring data to the new volume.
5. **Monitor progress**: Use **lvs** or **lvdisplay** to monitor the mirroring process on the host.

   I/O begins to appear on the target HPE Alletra Storage MP B10000 array as the data is mirrored to the target disk.
6. **Remove the source disk**: After the mirroring has completed, the original source disk can be removed as a mirror, removed from the VG, and unpresented from the old storage array.

LVM is just one option for mirroring disks on Linux® hosts. Other host OSes and applications have similar options for disk mirroring. Refer to the documentation for the host OS as well as the appropriate HPE Alletra Storage MP B10000 implementation guide.

Zerto for data migration

Zerto is a powerful and flexible tool available in the HPE Storage portfolio that enables storage-agnostic data mobility across virtualized environments. The journal-based replication enables a near-zero recovery point objective on protected resources, and it is primarily used as a disaster and ransomware recovery tool.

In addition to local and remote protection, Zerto allows the configuration of Virtual Protection Groups (VPGs) specifically for data migration and the execution of zero-data loss move operations. The move operation with Zerto is minimally disruptive because the initial sync and subsequent journal syncs of the VPG are done online prior to the move operation. Because Zerto uses asynchronous journal-based replication, there is no performance impact to the VMs being migrated. When the move operation is executed, the VM is powered off, any outstanding data is synced, and then the VM is powered back on at the new datastore.

Performing data migration using the Zerto Move operation

Using Zerto Move to migrate data involves the following tasks:

1. Install the latest Zerto Virtual Manager (ZVM) and link your VMware vCenter® or Microsoft Hyper-V server.
2. Deploy virtual replication appliances (VRAs) on the hosts with migrating VMs.
3. Create and export volumes from the new HPE Alletra Storage MP B10000 to the host environment and configure the datastores.
4. Create data migration VPGs pointing to the new storage datastore as the recovery point.
5. Wait for the initial sync to kick off and complete.
6. Perform the move operation to cutover to the new datastore.

Zerto Move provides various other options, including the ability to manually revert or commit a move, begin reverse protection back to the original source after the move to the new storage has completed, or keep the original source VMs on the original storage for test or restoration.

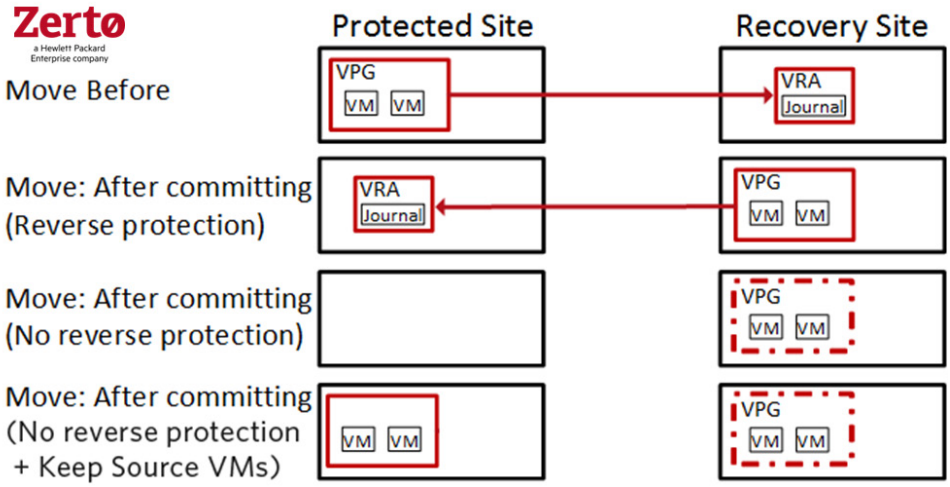

Figure 3. Zerto Move operation and options including reverse protection and keep source VM

The Zerto Move operation can be leveraged in existing Zerto deployments or can be licensed and installed specifically for data migration. For more information about using Zerto to migrate data, review the move process in the Zerto Self-Service Portal User Guide.

HPE Storage Data Migration Service

Data migration is a lengthy and often complex task. HPE offers the HPE Storage Data Migration Service, which provides managed end-to-end data migration services. These services leverage the depth of HPE technical capabilities combined with proven methodologies to help accelerate your transition to HPE Alletra Storage MP B10000. This data migration service helps you overcome data migration challenges such as the following:

- Qualifying and understanding data
- Transitioning from legacy storage solutions
- Reducing cost, risk, and time during data migration
- Integrating migrated data into established security, data protection, and governance frameworks

There are many ways to perform data migrations, and not all apply to every environment. It usually means finding the balance between what is supported based on your environment, what is available and cost-effective, and what provides the fastest cutover with the least downtime.

For migrations to HPE Alletra Storage MP B10000, HPE Storage Data Migration Service can employ entirely online array-based block migration methods or appliance-based migrations with minimal cutover. It can also coordinate host-based migrations based on the environment and requirements.

For more information, refer to the HPE Storage Data Migration Service data sheet.

Summary

Data migration is a major undertaking that requires careful analysis, planning, and execution. Without proper foresight and guidance, data migration programs can disrupt operations, diminish employee productivity, and impair revenue.

With a well-defined approach, an understanding of the new storage platform, and adherence to best practices, refreshing to HPE Alletra Storage MP B10000 can be straightforward and cost-effective, while helping to achieve seamless continuity of operations and maximizing the benefits of modern storage solutions for organizations of all sizes.

Resources

HPE Alletra Storage MP B10000 support matrix

HPE Alletra Storage MP B10000 UI user guide

HPE Alletra Storage MP B10000: Getting started with data replication using Remote Copy and the CLI user guide

HPE Alletra Storage MP B10000 data reduction efficiency

HPE Nimble Storage and HPE Alletra 6000 Peer Persistence technical white paper

HPE Active Peer Persistence technical white paper

HPE Storage Data Migration Service

HPE Alletra Storage MP B10000 migration experience

Zerto Self-Service Portal User Guide

Learn more at

HPE.com/storage

Visit **HPE.com**

Chat now (sales)

a00138873ENW, Rev. 2

Hewlett Packard Enterprise