HPE Storage

Technical white paper

Migrating to HPE MSA Gen7 Storage

Strategy and technical reference guide

HPE GreenLake

Contents

Executive summary

This paper investigates strategies and methodologies for migrating data from existing block-based storage systems to HPE MSA Gen7 Storage arrays. Although it is also possible to migrate data from third-party vendor arrays and other array families in the Hewlett Packard Enterprise portfolio, this paper emphasizes HPE MSA Storage systems as the source.

Countless combinations of data sources, storage arrays, operating systems (OSs), applications, and infrastructure make it impossible to cover every configuration and situation. For example, file system–based data copying methodologies are complex and highly dependent on the OS. Therefore, this paper provides guidance and suggestions specifically for migrating data to HPE MSA Gen7 Storage arrays in the following scenarios:

- From any storage system by using one of the following mechanisms:
  - VMware vSphere® Storage vMotion®
  - Microsoft Hyper-V live migration
- From an HPE MSA Gen6 Storage system by one of the following mechanisms:
- HPE MSA remote snap
- Data-in-place upgrade

Note
HPE recommends that end users take advantage HPE Services in-depth knowledge and extensive experience of delivering data migration services to our customers, rather than attempting a data migration themselves. HPE Services offer a broad set of expertise in performing what can be a complex task and can employ alternative hardware-based methods to move data, thereby accelerating the process while minimizing risk. To obtain further information or order services, contact a local HPE sales representative.

Important
Advice for migrating boot-from-SAN volumes is provided only in relation to data-in-place upgrades and remote snap. Migration from other storage arrays requires the use of third-party applications and is outside the scope of this paper.

Target audience

The target audience for this paper includes those who are tasked with designing and implementing a successful data migration from HPE MSA Storage or other array families, whether within the HPE portfolio or a third-party block-storage system. Readers should understand and have experience in administering relevant technologies, including the following:

- Basic networking concepts
- Fibre Channel or iSCSI SAN fabrics
- RAID and similar data protection technologies
- SAN storage systems
- OSs that are supported with HPE MSA Storage systems

Document purpose

In addition to exploring the practical differences among multiple data movement solutions, this paper offers guidance on devising a migration plan. A successful and robust migration plan should include an evaluation of the current systems, an identification of target array requirements, a listing of practical tasks, and a plan for cutover and rollbacks. Guidance is provided to assist in migrating data within complex environments, yet it is likely that smaller data centers would require a less exhaustive approach. Although thorough research reduces risk, striking a balance with efficiency is often necessary. This paper should therefore be considered a guide from which customers can take a subset of recommendations that suit a given scenario.

Although this paper aims to help in forming a successful plan and provides examples of how to configure systems for data migration, it is not a substitute for formal user guides, and it does not list all features or explain how to configure them. For detailed information about the features of HPE MSA Gen7 arrays, use the links on the last page of this document to access core documentation.

Note
To minimize inaccuracies that might be caused by new array firmware, OS updates, or other unpredictable future changes, this paper does not include guidance on how to configure an HPE MSA array with a given OS. The OS vendor usually has the most up-to-date documentation on how its OS should be configured. In some cases, HPE provides documents that explain how to deploy a Gen7 HPE MSA array or offer recommendations for a particular OS. For HPE documentation, see HPE Product & Solutions Now.

Variations in UI might occur over time and across supported OS versions. The example screenshots in this paper are from VMware ESXi™ 8.0 U1, Windows Server 2022, and HPE MSA firmware IN300R004.

Assumptions
This paper makes the following assumptions:

- Adequate consideration has been given to validate that the HPE MSA family of storage arrays are a suitable solution to meet performance and capacity requirements. However, some guidance is provided on how to size and configure optimal drive configurations.
- All participating systems are fully operational, and no hardware or software components can interfere with the expected movement or access to data. This assumption includes a networking topology that can route traffic to and from the source and destination systems as well as maintain a consistent link speed and latency.
- The destination system has the same data structure as the source (for example, the same number of volumes).
- Because backup and recovery from a failed migration is outside the scope of this paper, it is assumed that a functional and tested disaster recovery solution is in place.
- If systems are in multiple physical locations, the assumption is that they form a single logical site or that they can be treated as one.

Considerations for data migration
Migrating data from one system to another, especially online, presents many opportunities for unexpected events that might cause data loss or unavailability and impact business continuity. Therefore, a migration begins not with the physical movement of data, but with the planning phase.

Minimizing unknown factors is a cornerstone of risk management. Some unknowns require little investigation, but others are not easily discovered because such insights might require undocumented information such as application behavior and trends. Therefore, an exhaustive discovery and planning process is critical.

Risk is directly related to the topics of data availability and data integrity. To help ensure that the data retrieved from a storage system is identical to the data that was written to it, consider all data migration methods and evaluate their costs.

For example, one approach to tackling data integrity is to carry out a migration offline. If no additional data is accessed during the migration phase, data is less likely to be unintentionally altered. Copying data offline can also significantly reduce the time required to complete a migration because it might not be necessary to take an incremental approach. In addition, the storage systems and data networks are less burdened with competing traffic. However, the issue of availability becomes far more significant because the time required to copy all data might grow to days or longer. During this period, data would be unavailable, and because data is crucial for a business to function, few workloads are open to this approach.

This paper describes methods that keep applications online or involve very minimal downtime, as well as offline using data in-place upgrades. Whether implemented online or offline, migrations present an opportunity to consider the layout of data and the mechanisms in place to coordinate access to it, which might include the following:

- Changes to the underlying drive architecture, such as RAID and the distribution of volumes across drive groups, as well as the resulting capacity and performance
- Protocol migration (for example, from Fibre Channel to iSCSI)
- Encryption
- Host clustering

Attempts to tackle any of these transformations, especially online, significantly increase risk. However, post-migration changes are often more difficult to implement and might carry even greater risks, such as disk and volume manipulation.

The discovery phase

Although no single aspect of a migration can be considered the most important, the ability to choose the best strategy depends heavily on a successful discovery phase. During the discovery phase, a complete view of the current environment is documented, including the infrastructure, the data footprint, and how the data is accessed, as well as how often and how intensively it is used.

Discovery is a hands-on process that requires several approaches to gather all necessary information. In a well-documented environment, you can record the physical and logical system configuration with a minimum investment of time and effort. This task alone, however, is unlikely to yield a deep understanding of application behavior and data movement trends. A fuller understanding of application behavior—and to some extent of data movement insights—requires discussing the trends with those who own or have a vested interested in the data. These discussions have the potential to produce undocumented information and reveal differences of opinion within an organization about what is configured and how, as well as the fundamental application requirements. This outcome is not unusual and is a healthy approach to information consolidation that can lower risk during a migration and beyond.

During discovery, several pieces of information must be collected. Hewlett Packard Enterprise advises you to create a worksheet to record this data. Software for creating and editing spreadsheets would be perfect for this task.

Infrastructure

As the first step, HPE recommends that you evaluate the current, transitional, and future footprints of hardware locations, including these categories:

- Rack or floor space
- Power and cooling requirements
- Available network switch ports

Unless using a backup-and-restore or data-in-place approach, the destination and source systems must be hosted in tandem for the duration of the migration. In this case, you must have enough physical capacity to accommodate these systems and provide adequate cooling as well as connectivity to power and networking infrastructure.

Data networks

To migrate data online through the same network infrastructure as existing traffic, the network must be adequate. A migration plan often tries to target the smallest window of time to copy the data, which in turn depends heavily on the quality and bandwidth of the interconnecting datapaths, as well as the quantity of data that must be copied.

The following list includes examples of data network information that should be collected:

- Network topology
- Switch models
- Port:
  - Utilization
  - Supported link rates
  - Availability
  - Security
- Network segmentation:
  - Fibre Channel zoning
  - VLANs
  - Subnetting
- Naming conventions for components:
  - Ports
  - Devices
  - Fibre Channel zones

- Link quality:
  - Bandwidth
  - Latency
  - Traffic patterns
- Switches and management software:
  - FQDNs and IP addresses
  - Administrator credentials

Other systems

It is essential to document existing systems adequately, including the following:

- Storage systems
- Servers

Storage systems

Any storage system from any vendor can potentially serve as a source array for a host-based migration. However, data-in-place and array-based migrations to HPE MSA Gen7 Storage require the source to be an HPE MSA Gen6 Storage array. Where relevant to the source system, HPE recommends that you collect the following information to clarify the present configuration and simplify troubleshooting:

- Vendor and model
- Serial number
- Component firmware versions, which for HPE MSA Storage systems include these components:
  - Controller firmware
  - I/O module firmware of any expansion enclosures
  - Drive models and firmware
- Management IP addresses
- Administrator credentials
- Target port IP addresses (iSCSI systems only)
- Cluster IP addresses (If applicable)
- WWPNs (World Wide Port Name, for Fibre Channel and SAS systems only)
- Storage configuration, which for MSA systems include these components:
  - Installed drives:
    - Drive type (for example, SSD, 10K SAS, MDL SAS)
    - Drive capacities
    - Assignments (for example, disk group name, global spare, available)
  - Disk group and pool configurations:
    - RAID types
    - Drive counts per disk group
    - Tier capacity
    - Available tier capacity
    - Pool capacity
    - Available pool capacity

- Volumes:
  - Volume names
  - Snapshots
  - Snapshot schedules
  - Snapshot rate of change (unique data per snapshot)
- Capacity
- Remaining capacity
- Host mappings
- Host configurations:
  - Initiators IDs (IQN or WWNs)
  - Host names
  - Host group participation
  - Performance per pool:
- Potential performance:
  - IOPS = The highest tier within the pool
  - GB/s = The highest HDD tier within the pool
- Current utilization:
  - Minimum and maximum performance consumed
  - Trends (for example, day and time of peak loads, sustained average performance)

Note

HPE provides a free-to-use tool for collecting and reporting on HPE MSA Storage array health from HPE P2000 G3 onward. HPE MSA Health Check is a cloud-hosted, AI-driven tool that provides a convenient method of collecting some of the recommended information with minimal effort. Results can be saved to a PDF.

HPE MSA Health Check provides multiple insights that inform the user of configurations that do not align with current availability-related best practices, such as outdated firmware. Although it does not represent all recommended data for migration purposes, it can simplify some data collection, and it provides useful warnings that can help to optimize an HPE MSA array.

There are two ways to collect information directly from an HPE MSA array:

- **The Storage Management Utility (SMU),** the web-based HPE MSA interface, provides a simple-to-use approach to displaying pertinent information, as shown in Figure 1.

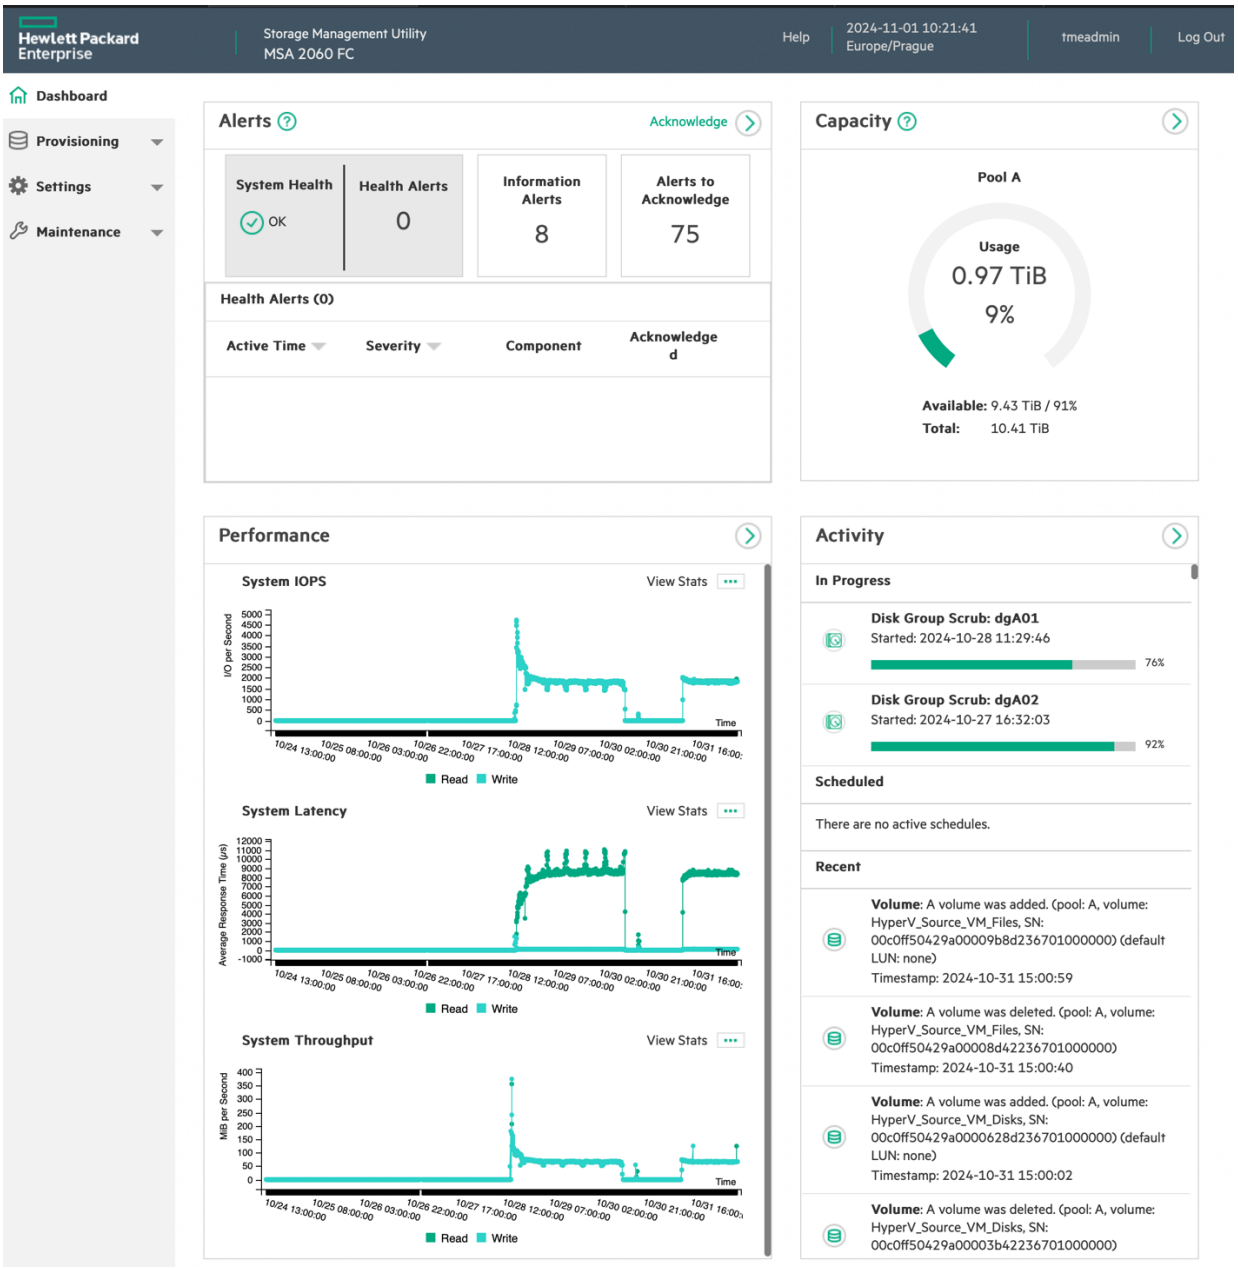

Figure 1. HPE MSA 2060 SMU dashboard

- **The CLI** provides the simplest method of collecting system-related data, but it requires a knowledge of the array commands. For a complete list of commands, see the HPE MSA 2070/2072 CLI Reference Guide for the HPE MSA Storage system:
  - show configuration
  - show pool-statistics
  - show disk-group-statistics
  - show initiators
  - show host-groups
  - show maps
  - show disks
  - show pools

  - show disk-groups
  - show volumes
  - show volume-groups
  - show schedules
  - show peer-connections
  - show replication-sets

Note
Since HPE Gen6 MSA Storage, the SMU Reference Guide has been renamed to the "Storage Management Guide."

Servers

Server configurations can have a varying impact on a migration, depending on whether data movement is host based, or array based. For example, host port bandwidth and available server performance might affect the time it takes to move data from source to target. Regardless of how data is moved, however, applications servers must be reconfigured to use the new data location. For that reason, it is advantageous to have a clear record of the current configuration.

Important
HPE MSA Storage arrays are tested and qualified for connection to HPE servers running specific support OSs through Fibre Channel HBAs from HPE or in-box iSCSI initiators. HPE offers support for third-party servers and a limited set of Linux® distributions under a best-effort policy. Before committing to a migration plan, check Single Point of Connectivity Knowledge (SPOCK) to confirm that support is available.

Important server-related information includes the following:

- Server brand and model
- OS:
  - Type and version
  - Administrator credentials
- Out-of-band management interface (for example, HPE iLO):
  - IP address and FQDN
  - Administrator credentials
  - Licensed features
- Application and data network configuration:
  - Cabling scheme
  - Adapter:
    - Model
    - Driver version
    - Supported link rate
    - Negotiated link rate
    - Physical presentation (for example, 10GBASE-T)
    - Port count
    - Port utilization
    - Port function (for example, iSCSI access, application traffic, hybrid)
    - Port VLAN assignments

  - Software switches:
    - Port participation
    - VLAN assignment
    - VM or application affinity
    - Enabled features (for example, vSphere Storage vMotion)
  - IP addresses and FQDNs
  - Link aggregation
  - VLAN assignments
  - iSCSI initiator:
    - IQN
    - Target portals
    - Favorite targets
- Fibre Channel:
  - HBA:
    - Model
    - Driver version
    - Supported link rate
    - Negotiated link rate
    - Port utilization
    - WWPNs
  - SAN disks
    - Drive letter or mount point
    - Capacity
    - Remaining capacity
    - Multipathing software (MPIO)
    - MPIO configuration for each volume (for example, round robin)
  - Applications:
    - Type
    - Performance requirements
    - Activity trends
    - Data storage locations
    - Availability requirements
    - Trusted users (for user acceptance testing)

System management information
All systems are accessed and managed differently and might offer both in-band and out-of-band options. Therefore, it is important to document in advance how to access these interfaces and what credentials are required to successfully authenticate users to the management interface. This information should be collected for every connected system. IP addresses should already have been collected during previous discovery steps. Common management interfaces include the following:

- **A web-based interface (WBI)** enables an internet browser to be used for management. WBIs are available in two varieties:
  - Target-based WBIs are hosted directly on the system to be managed. Some interfaces enable management of remote systems through this same target portal
  - A central appliance, which is an application hosted on a traditional OS that connects to and manages subordinate external systems (for example, VMware vCenter®)
- **The CLI** typically uses protocols such as the following:
  - SSH
  - Telnet
- **Host-based applications** are also installed and hosted by using traditional OSs but are interacted with as point-and-click programs. Therefore, they require access to a graphical desktop.

The planning phase

With the data collected during the discovery phase, it should be possible to start planning the migration. The goal of this phase is to build a clear order of events that results in a smooth migration and minimizes interruptions, as well as to establish a recovery plan in case something goes wrong.

The planning phase begins with an initial outline detailing how a migration would take place based on all information collected. The understanding is that this outline shall require changes. In creating the outline, consider what the business needs and expectations are and how they correspond to the capabilities of old and new components. Expect to revise this plan multiple times; such revisions are aided by rigorous trials and testing. Take care not to overexert the infrastructure or overestimate what can be carried out within a given time frame. Rushing a migration or cutting corners is likely to yield unexpected results.

Important
No two migrations are likely to have the same migration plan. The suggestions in this paper and the related examples are intended only to help form a suitable solution.

A typical planning phase covers the following tasks:

- Creating an overall migration strategy:
  - Defining the source data
  - Choosing a general data movement strategy
  - Identifying data copy windows
- Sizing the destination HPE MSA Gen7 Storage system
- Defining a running order of tasks
- Setting up user acceptance testing:
  - Ensuring access to data
  - Confirming accuracy of data
- Planning for post-migration monitoring
- Planning for cutover
- Planning for possible rollback
  - Defining stop triggers
- Setting up dry runs and testing:
  - Determining what data to migrate and how much

Creating a migration strategy

Creating a migration strategy involves deciding what data to migrate and how, as well as what approaches to use. The strategy defines the source data, describes the mechanisms that are to be used to copy it, and covers when and how to tackle any caveats and limitations. It normally evolves from an initial draft into a final version that is very detailed, especially after dry runs are completed.

Defining source data

The usual approach toward migrating data to an HPE MSA Gen7 array is to copy all source volumes 1:1 to the new storage system. In some cases, however, this approach might be unnecessary. For example, perhaps some volumes located on the source array were created for a development project that was completed, but those volumes have not yet been removed. Reducing the amount of data to be copied decreases the time required to complete the migration and reduces risk. The discovery phase should provide the insights necessary to define which volumes should be migrated.

Choosing a general data movement strategy

There are four common approaches to relocating data, each offering its benefits and challenges:

- Synchronous replication
- Asynchronous replication
- Host-based replication
- Media exchange

Table 1 outlines some of the advantages of each migration strategy demonstrated within this white paper and how they align with the common strategies listed.

Table 1. Highlight capabilities by migration type

Detail	Remote snap	Storage vMotion	Hyper-V live migration	Data-in-place upgrade
Type	Array-based asynchronous	Host based	Host based	Media exchange
Requires OS certification	No	Yes	Yes	No
Can easily relocate boot volumes	Yes	No	No	Yes
File system dependent	No	Yes	Yes	No
Can migrate VMs online	No	Yes	Yes	No
Array agnostic	No	Yes	Yes	No
Cost advantage	None	None	None	Significant
Complexity	Complex	Simple	Simple	Simple
Requires additional drive media	Yes	Yes	Yes	No
Application downtime per array	Minutes to an hour	None	None	Minutes to an hour
Performance impact	Low	Medium	Medium	n/a

Array-based synchronous replication helps ensure that partner volumes located on participating systems are identical. It usually operates at the block level. However, because it requires array-based features that are not found on any model of HPE MSA Storage arrays, synchronous replication is removed from consideration.

Note

Near-synchronous replication can be achieved using Zerto solutions and migrations can be carried out with offerings such as Zerto Migration Software. However, these solutions are beyond the scope of this paper.

Array-based asynchronous replication is a method of copying data from one location to another, either according to a schedule or when manually invoked. Remote snap is a licensed feature of HPE MSA Storage arrays that has been available since the HPE P2000 G3 to provide asynchronous replication over Fibre Channel or iSCSI fabrics. Unlike synchronous replication, remote snap replicates snapshots from one array to another. Although the feature was designed primarily as a disaster recovery solution, remote snap is also well suited to the task of migrating data between compatible HPE MSA Storage arrays.

Note
The remote snap feature is not supported on SAS models of HPE MSA Storage arrays of any generation. In addition, remote snap requires both arrays to be connected to switches, and they cannot be directly connected to each other.

Important
Remote snap requires both source and destination HPE MSA Storage systems to be appropriately licensed.

The HPE MSA remote snap feature offers significant advantages:

- Enables the offloading of data movement from hosts
- Simplifies task management
- Makes it easier to implement an incremental approach to copying data, which can mitigate link saturation and reduce cutover times
- Greatly simplifies migrating boot volumes
- Removes OS compatibility issues
- Works without host-based tools or features

Table 2 lists the multiple versions of remote snap and the different features and compatibility it offers depending on which version each array is running.

Table 2. HPE MSA remote snap compatibility and limitations across generations of HPE MSA Storage arrays

Detail	Gen4	Gen4	Gen5	Gen6	Gen7
Data structure	Linear	Virtual	Virtual	Virtual	Virtual
Remote snap version	1.0	1.5	2.0	3.0	3.0
SMU version	v2	v3	v3	v4	v4
Supported replication partner generation	• Gen4 • P2000 G3	• Gen5 • Gen4	• Gen6 • Gen5	• Gen7 • Gen6 • Gen5	• Gen7 • Gen6
As of firmware version	GL100	GL220R005	VL100/VE100	IN100	IN300
Protocol support	• Fibre Channel • iSCSI	iSCSI	• Fibre Channel • iSCSI	• Fibre Channel • iSCSI	• Fibre Channel • iSCSI
Failback support	Yes	No	No	Yes	Yes
Minimum scheduled replication interval	30 minutes	60 minutes	30 minutes	30 minutes	30 minutes
Replication queuing	No	No	Yes	Yes	Yes
Number of queued replications	n/a	n/a	1	1	1
Number of peer connections	n/a	1	• HPE MSA 1050:1 • HPE MSA 2050/2:4	• HPE MSA 1060:1 • HPE MSA 2060/2:4	4
Peer connection authentication support	n/a	No	Yes	Yes	Yes
Ability to change peer connection protocol	n/a	No	Yes	No	No

Table 2. HPE MSA remote snap compatibility and limitations across generations of HPE MSA Storage arrays (continued)

Maximum replicated volumes per system	16	32	32	32	32
Maximum volumes per volume group	n/a	16	16	16	16
Maximum replication sets per volume	n/a	1	1	1	1
Notes		Peer connections to Gen5 systems supported but cannot be created by a Gen4 system			

Host-based replication is the copying of data from one storage system to another through a commonly connected server. Host-based replication offers several advantages over array-based replication:

- It supports replication from arrays that do not support remote snap or that are incompatible with remote snap version 3.0 (for example, migrating from linear storage of obsolete HPE MSA arrays or third-party systems).
- It provides finer control over the data that is to be migrated.
- It can potentially enable seamless online data movement. For example, vSphere Storage vMotion enables virtual machine (VM) files to be moved from one datastore to another while retaining read-write access throughout, thus eliminating the need to bring applications offline.

Despite these advantages, host-based replication is not suitable for all environments, and it might increase complexity. Some OSs include more capable tools than others, and third-party applications might mitigate some challenges, often at additional cost. In general, however, the following considerations apply:

- A file system can increase data copy times. For example, many small files take longer to copy than a few large files.
- Volumes that are used for boot from SAN are difficult to copy.
- An OS might not be supported on both the source and the destination arrays.
- Because a single OS might not recognize the file systems used for all volumes that are to be copied, it might be necessary to use multiple hosts or applications that access the volume in RAW mode.
- Copying file system permissions can be complicated.
- File systems and shared folder configurations cannot easily be copied and must be prepared in advance.
- Encryption that is implemented at the file system level might require additional care.
- Human error is far more likely if appropriate tools and methods are not used.

Important

It is beyond the scope of this paper to provide guidance on host-based migrations except when using either vSphere Storage vMotion or Microsoft Hyper-V live migration. However, a similar paper on how to migrate to HPE StoreEasy 1000 Storage offers an in-depth investigation into and methodology for copying Windows NTFS file systems. For more information, see Migrating from any vendor's NAS to HPE StoreEasy 1000 Storage in Active Directory integrated environments via SMB.

Data-in-place upgrade is the process of physically moving drive media, expansion enclosures, and other supported hardware from one HPE MSA Storage system to another while retaining all data and disk group structures. Some of the most significant advantages of data-in-place upgrades include the following:

- The lowest cost method as no additional drives or expansion enclosures are required. When migrating between systems of the same protocol, it is probable that the same switching infrastructure including ports, licenses, transceivers, and cables can be reused.
- No other mechanism provides a comparable or faster method of moving large quantities of data from one location to another. It is theoretically possible to migrate petabytes of data in minutes.
- Provides a convenient method to migrate from one protocol to another. For example, an iSCSI HPE MSA Gen6 to a Fibre Channel HPE MSA Gen7 system.

Note

It is not supported to migrate from HPE MSA Gen7 to Gen6 Storage using the data-in-place method.

Though convenient, there are some caveats and disadvantages that may prevent a migration plan going forward using a data-in-place method. For example:

- Data migration is an ideal time to transform disk group layout and to improve an existing disk configuration. However, changing RAID type and optimizing drive counts and tiering, cannot easily be achieved unless there are additional and unused drives.
- There is no inherent way to maintain availability of volume data residing on drives that are to be moved. Some downtime is always required.
- As they age, mechanical drive media is susceptible to failure when power cycled. Moving drives that have not powered down in a long time and are outside of warranty presents an increasingly unpredictable risk the higher the drive count. The risk to data is also higher if source disk groups are configured in unsuitable RAID schemes. For example, 16 drives in RAID 5 are far less resistant to failure than those configured in HPE MSA-DP+.
- SSD media has a finite lifespan dictated by how much data has been written to it. Unlike an expected rate of failure, wear out is a definite value after which no further writing of data to that drive can occur. HPE provides a warranty for drive wear out for up to three years from purchase. It may therefore be inconvenient in the longer term and present unacceptable risk to reuse SSD media.
- Though a working backup plan is always recommended, it is essential when performing data-in-place upgrades due to elevated risk of data loss or unexpected downtime compared to other methods.

Identifying data copy windows

A copy window is an optimal period to migrate data, typically business quiet times when application activity is low. Many factors must be considered when choosing how much data can and should be replicated within a given copy window. For example, array-based replication might take the least amount of time to copy data but might lead to more time being required for disruptive host reconfiguration compared to using a host-based method such as vSphere Storage vMotion. In addition, bottlenecking caused by insufficient system performance or link saturation can come into effect if there is an attempt to copy too much data within too small a time frame or if there is competition from application traffic. Hewlett Packard Enterprise recommends that you carry out trials before committing to a migration strategy to ascertain the relative end-to-end performance characteristics and to measure and understand limitations.

Depending on the outcome of testing, the quantity of data that must be migrated, and the supported methods available, it might be possible to migrate all required data within a single data copy window. However, it might also prove impossible to both copy data and complete all post-migration tasks such as cutover and user acceptance testing in time or without elevating risk. A typical strategy might be to migrate data incrementally on a per-application basis, by individual host, or by volume.

After the available bandwidth and copy windows are known, it is a matter of basic mathematics to determine how much data can be copied during a given window. A simple formula is **total data / average transfer rate = time to transfer (in seconds)**.

In the following example, data can move uninhibited at the line rate of a single gigabit connection (125 MB/s). The resulting copy time of three hours and 40 minutes equates to two copy windows. This means that if a copy window presented itself once per day, it would take two days to copy 1.5 TiB of data.

Table 3. Sample calculation of how many data copy windows are required for a specified amount of data, using a known and consistent transfer rate

Type	Value	
Data copy window	2 hours	
Total data	1649 GB	(1.5 TiB)
Transfer rate	0.125 GB/s[1]	(1 Gb/s)
Example	1649 / 0.125 = 13,192 seconds / 3600 ≈ 3.67 hours = 3 hours 40 minutes / 2 ≈ 2 × copy windows	

[1] Calculation in base-10

Another approach is to calculate the maximum amount of data that can be copied within a defined window at a consistent average transfer rate. A simple formula is **data copy window (in seconds) x average transfer rate = total data**.

Table 4. Sample calculation for how much data can be moved within a specified data copy window by using a known and consistent transfer rate

Type	Value	
Data copy window	7200 seconds	(2 hours)
Transfer rate	0.125 GB/s[2]	(1 Gb/s)
Example	7200 × 0.125 = 900 GB	

Defining a running order of tasks

It is critical to initiate a data migration in a specific predefined order that was established during migration planning. Failure to do so might result in dependency issues and the interruption of all subsequent tasks. Principle factors affecting the running order should have been discovered in the discovery phase and for might include limitations such as these:

- Architecture
- Quantity of data
- Maintaining access to data to support-specific service levels
- Strict time allotment to complete the migration
- Data retention requirements

When you define a running order, HPE recommends differentiating between preparation and the eventual movement of data. Preparation tasks include actions such as physically installing and cabling hardware, preparing IP addresses and network settings, bringing the new HPE MSA Storage system online, and other similar groundwork. Data movement tasks are the steps needed to initiate the migration of data and the redirecting of applications to that data. There are also cleanup tasks that include the actions necessary to remove or potentially repurpose legacy hardware.

User acceptance testing

A failure to discover data access and integrity errors before migration is marked as complete can have severe implications for business continuity. User acceptance testing helps reduce the risk of undiscovered errors, and when carried out before a migration as part of a dry run, it can expose ineffective strategies before problems arise. User acceptance testing aims to ensure that applications can access data in the same way as before a migration began and that the data has been copied correctly. A test plan must be established that includes a representative subset of users and, where applicable, administrators of affected applications.

HPE recommends that you nominate a group of trusted users to perform a limited range of tests after each stage of migration is completed and that you meet with them in advance to confirm that they understand their roles. To minimize risk and improve participation, give the group a defined start time and duration for testing in advance. A simple approach is to define specific workflows and then confirm that they can be completed and that the data is accurate.

If user acceptance testing fails during a dry run, amend the migration plan and reattempt it until it is successful. If user acceptance testing fails post migration, consider a rollback.

Planning for post-migration monitoring

After user acceptance testing is successfully completed and where a data-in-place upgrade was not performed, both the source and the destination systems should be kept online and monitored regularly. In relation to a migration, monitoring is required to confirm that two goals have been achieved:

- Applications are no longer addressing the source system.
- Performance and availability of the destination system is adequate.

It is typically prudent to allocate a period measured in days to weeks during which the source system is kept online and regularly monitored for host access. Allowing this much time helps in discovering any undocumented applications that might have infrequently scheduled tasks.

[2] Calculation in base-10

Data-in-place upgrade considerations

Before commencing with a data-in-place upgrade, familiarize yourself with the following.

Supported components

HPE MSA Gen7 arrays support all HPE MSA Gen6 hardware options including those not offered for sale with Gen7. However, some combinations of options may require additional consideration.

Table 5. Supported components between HPE MSA generations

Item	Supported with HPE MSA 2070/2072	Notes
HPE Gen6 storage controllers	No	Moving controllers is not supported on any Gen6 or Gen7 system
HPE Gen6 array enclosures	No	Array enclosures cannot be migrated between generations
HPE Gen6 expansion enclosures	Yes	
HPE Gen6 TAA expansion enclosures	Yes	Destination array configuration must meet TAA requirements
Drive media	Yes	Including 15K Enterprise SAS drives and drive capacities of any drive type offered on Gen6 and not Gen7
Self-encrypting drive media	HPE MSA 2070 only	HPE MSA 2072 arrays ship with non-encrypting SSDs. All installed drives must be SEDs to support encryption. Drives can be moved only if encryption is disabled and there is no expectation to enable it.
Transceivers and DACs	Yes	
HPE MSA ADS Suite License	No	Licenses cannot be transferred between systems
Any component from a generation of HPE MSA prior to Gen6	No	

Firmware

HPE recommends that both source and destination systems have their firmware brought up to date prior to attempting a data-in-place migration. HPE MSA Storage Health Check can assess overall system health, check on the availability of newer firmware, and simplify the process of obtaining the relevant downloads.

Protocol migration

It is supported to migrate between both array generations and host protocols. For example, it is possible to remove drives from an iSCSI Gen6 HPE MSA Storage system and place them in a SAS array of the same or next generation. Table 6 includes relevant considerations and Table 7 provides explanations for each item.

Table 6. Considerations when migrating between protocols

Protocol	Fibre Channel (16 Gb)	Fibre Channel (32 Gb)	iSCSI (10/25 Gb)	iSCSI 10GBASE-T	SAS
Fibre Channel (16 Gb)	• WWPN	• Transceivers • WWPN	• Transceivers • Cables • Switches • WWPN/IQN • Mapping	• Cables • Switches • Mapping • WWPN/IQN	• Cables • WWPN • Mapping • No RS
Fibre Channel (32 Gb)	• Transceivers • WWPN	• WWPN	• Transceivers • Cables • Switches • WWPN/IQN • Mapping	• Cables • Switches • WWPN/IQN • Mapping	• Cables • WWPN • Mapping • No RS
iSCSI (10/25 Gb)	• Transceivers • Cables • Switches • WWPN/IQN • Mapping	• Transceivers • Cables • Switches • WWPN/IQN • Mapping	• IQN	• IQN • Cables • Switches	• Cables • WWPN/IQN • Mapping • No RS

Table 6. Considerations when migrating between protocols (continued)

iSCSI 10GBASE-T	• Transceivers • Cables • Switches • WWPN/IQN • Mapping	• Transceivers • Cables • Switches • WWPN/IQN • Mapping	• Transceivers • Cables • Switches • IQN	• IQN	• Cables • WWPN/IQN • Mapping • No RS
SAS	• Transceivers • Cables • Switches • WWPN • Mapping	• Transceivers • Cables • Switches • WWPN • Mapping	• Transceivers • Cables • Switches • WWPN/IQN • Mapping	• Cables • Switches • WWPN/IQN • Mapping	• WWPN • No RS

Table 7. Expanded detail for Table 6

Item	Detail
Transceivers	New transceivers or DACs may be required (array, switch/host)
Cables	New host port cabling required
Switches	New switch ports may be required if attaching to a fabric
WWPN	Host port WWPNs will change and require new zoning and initiator configuration
IQN	The array IQN will change and will require initiator reconfiguration
Mapping	Initiator nicknames will be lost, and volumes will need to be remapped to hosts
No RS	Remote snap not supported

Configuration data

Some configuration data is stored on drive media, and some is stored on the controllers. When drives are migrated to another system, data that is stored on array is not carried over. Although friendly names for initiators and hosts are not retained, volume mapping remains in place. If the destination array is the same protocol, then volumes continue to be accessible to the host once its initiators and network have been configured to access the new system.

Data that is stored on drives and is carried over include:

- Pool configuration and contents
- All disk groups and configuration data
- Volumes and their user and configuration data
- Volume mapping
- Snapshots and their user and configuration data

Information and configuration data that is not carried over include:

- Schedules
- Hostnames
- Initiator nicknames
- Host groups
- Volume groups

Encryption

All HPE MSA Gen6 and Gen7 Storage arrays can support encryption provided all installed drives are self-encrypting drives (SEDs). It is essential to follow the prescribed method in this white paper to ensure that data is accessible once drives are installed within the destination system.

Important

HPE MSA 2062 and 2072 Storage systems ship with two nonencrypting 1.92 TB SSDs. As encryption requires all installed drives to support full disk encryption (FDE), then unless the included drives are removed it is not possible to unlock and access data on encrypted drives.

TAA configurations

Table 8 outlines differences in TAA compliance between HPE MSA Gen6 and Gen7 Storage systems. To conform to TAA requirements, destination TAA-compliant HPE MSA Gen7 arrays must include a minimum of six factory-integrated drives.

It is supported to connect non-TAA compliant expansion enclosures and drives to an HPE MSA Gen7 system that is TAA compliant to achieve overall compliance. For example, when connecting a Gen6 expansion enclosure with non-TAA compliant drives to a TAA-compliant HPE MSA Gen7 array, the overall solution remains compliant.

Table 8. Differences in TAA-compliant configurations across HPE MSA generations

Item	Gen6	Gen7
Array enclosure	• TAA-specific models • Includes six TAA-compliant drives • Fixed starting capacity points	• TAA-specific models • Does not include drives • No restriction on drive type or capacity • Must be configured with at least six drives • Must be configured with factory integration services
Expansion enclosure	• Minimum of six TAA-compliant drives per enclosure • Does not include drives	• No minimum number of drives per enclosure
Drives	• Limited choice	• Can use any drive

Sizing the destination HPE MSA Gen7 Storage system

When migrating data to a new storage system, it is essential to plan how to configure disks for optimal performance and capacity. For example, if multiple source arrays are to be consolidated into a single system, it is likely that the destination array shall need to offer greater performance and capacity than any individual source system. HPE makes it easy to choose an optimal HPE MSA Gen7 solution by using the Ninja Online for HPE MSA tool, which is a cloud-hosted application available to anyone with an HPE Passport. HPE also strongly recommends that you become familiar with the HPE MSA 2070 and HPE MSA 2072 Storage arrays best practices white paper, which offers guidance on how to achieve optimal configurations and consider engaging with an HPE sales representative for expert guidance.

Planning for cutover

Cutover is the process of redirecting applications from one system to another, specifically from a source storage system to an HPE MSA Gen7 Storage system. Planning this phase correctly is critical because there is the potential for applications to access different copies of the same data when not performing a data-in-place upgrade. During cutover, the principal goal is to prevent hosts from attempting or successfully connecting to a source system while enabling them to connect successfully and reliably to the destination.

Cutover examples

The cutover process can vary depending on which migration strategy was employed, and whether a migration was carried out in a single step or in discrete stages. It also depends on how data is accessed, and for what purpose and by what OS. VMware vSphere Storage vMotion and Microsoft Hyper-V live migration share similar steps, and variation is generally limited to differences in terminology and the user interface.

Note

If migrating data rather than drives, it is advised to take an array-based snapshot of volumes located on the source storage system before attempting a migration or cutover.

vSphere Storage vMotion

Unlike remote snap, which copies data from one location to another, vSphere Storage vMotion moves VM data between datastores, removing the possibility that older versions of data might exist and dramatically simplifying the cutover process. If vSphere Storage vMotion is used, cutover begins when a VM is migrated between datastores, but the process cannot be considered complete until all other files have been moved and the datastore is removed from the host.

Important

Files that are not directly part of a VM are not moved by vSphere Storage vMotion migration. For example, suppose multiple VMs are configured to use a common ISO image. In that case, it must first be copied to the destination datastore and the affected VMs reconfigured to use the file in the new location.

Note

For current advice on how to migrate VMs that use raw device mapping (RDM) disks through vSphere Storage vMotion, see Migrating virtual machines with Raw Device Mappings (RDMs) (1005241).

vSphere Storage vMotion operates at the VM level, not the datastore level. Therefore, a task is created for each VM whose files are either partially or fully stored in datastores located on the source storage system. For each VM, there are three states:

1. The VM files are in their original locations.

   Figure 2 shows a representation of multiple VMs that share a single datastore located on a legacy HPE MSA array.

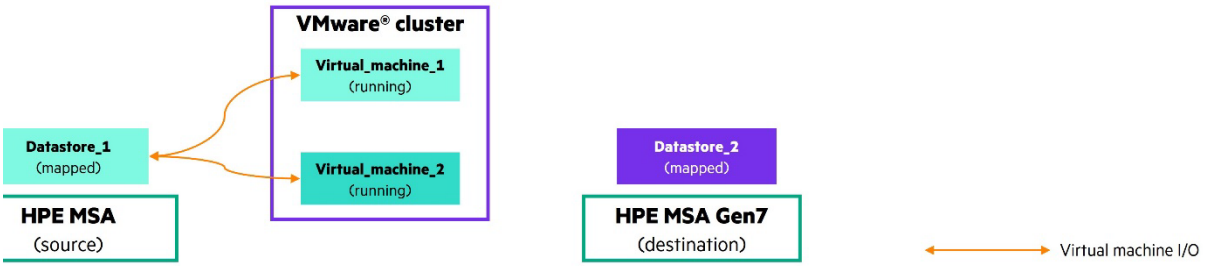

Figure 2. All VMs are running and located on the existing storage system

2. The vSphere Storage vMotion migration task starts.

   As Figure 3 shows, during this time, the VM remains in a running state and disk I/O is directed to the storage subsystems of both the source and destination datastores.

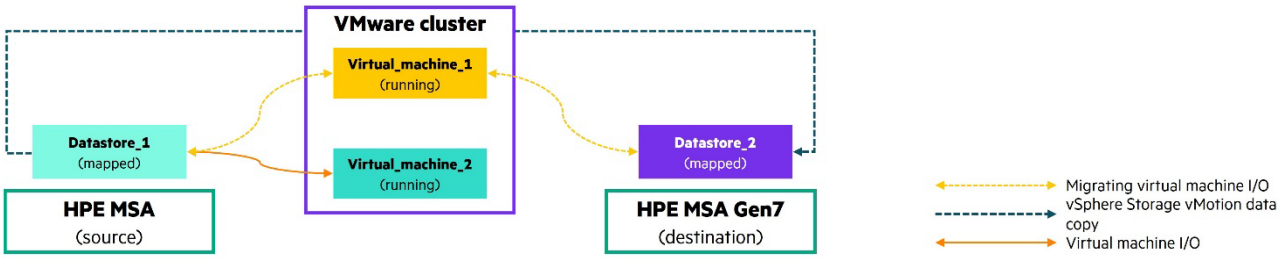

Figure 3. A vSphere Storage vMotion task is initiated for a single VM

3. After the vSphere Storage vMotion migration task is complete, the VM is located on the destination datastore, and all disk I/O for that VM is serviced by the storage subsystem that hosts it.

   As Figure 4 shows, the per-VM tasking of VMware means that it is possible that after the migration task is complete there might still be VMs located on the source datastore.

Diagram depicting

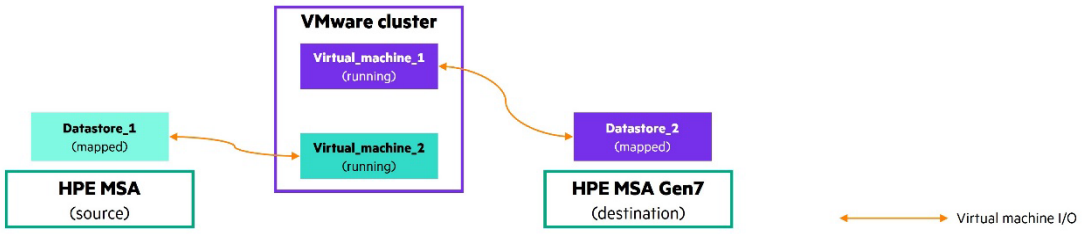

Figure 4. The migrated VM is now located on the new HPE MSA Gen7 Storage system

After all the VM files have been migrated and copied, the next task is to take the source volume that hosts the datastore offline. The steps vary depending on which system is involved:

In VMware vCenter or VMware ESXi

1. Unmount the datastore from all hosts.
2. If VMware vSphere® Storage DRS™ and VMware vSphere® Storage I/O Control are enabled for the datastore, disable them.
3. Make sure the datastore is not used for the VMware vSphere® High Availability heartbeat.
4. In the VMware vSphere® client, navigate to the datastore.
5. Right-click the datastore to be removed and select **Delete Datastore**.
6. Confirm that you want to remove the datastore.
7. **Optional:** Rename the datastore to give it the same name as the one that was removed.

In the HPE MSA Storage system

Version 4 SMU

1. Navigate to the **Volumes** view and select the volume hosting the datastore that is to be disconnected from the VMware hosts.
2. From the **Actions** menu, select **Detach from hosts**, check the relevant box for the hosts and proceed.

Version 3 SMU

1. Navigate to the **Volumes** view and select the volume hosting the datastore that is to be disconnected from the VMware hosts.
2. Highlight all mappings by selecting the first entry and then, while holding the shift key, selecting the last one.
3. Right-click anywhere in the highlighted list or click **Actions** and select **Remove Mappings**.
4. Click **OK** to confirm the action.

Figure 5 shows the result after all actions are completed. The legacy storage system is no longer used by VMware, and all VMs are located on the new HPE MSA Gen7 array.

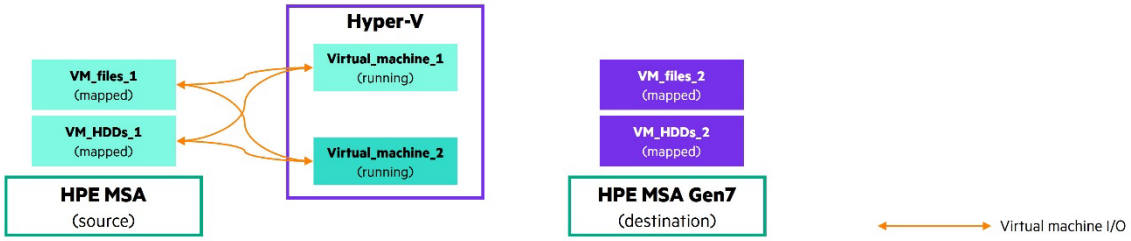

Figure 5. After all the VMs and files are migrated, the original datastore can be unmounted, removed, and unmapped

Microsoft Hyper-V live migration

Similarly to vSphere Storage vMotion, Microsoft Hyper-V live migration moves VM data between file systems. Therefore, live migration also simplifies the cutover process and removes the possibility that older versions of data might exist. When live migration is used, cutover begins when a VM is migrated between file systems, but the process cannot be considered complete until all files have been moved to their destination file systems and the physical disk is removed from the host.

Important

Hyper-V live migration does not move files that are not directly part of a VM. For example, suppose multiple VMs are configured to use a common ISO image. In that case, it must first be moved to the destination datastore, and the affected VMs must be reconfigured to use the file in the new location.

Just as with vSphere Storage vMotion, Hyper-V live migration operates at the VM level, and it must be configured for each VM whose files are either partially or fully stored within a file system located on the source storage system. The VM passes through three states during storage live migration:

1. The VM files are in their original locations.

   Figure 6 shows a representation of multiple VMs that share multiple file systems, which are located on a legacy HPE MSA array.

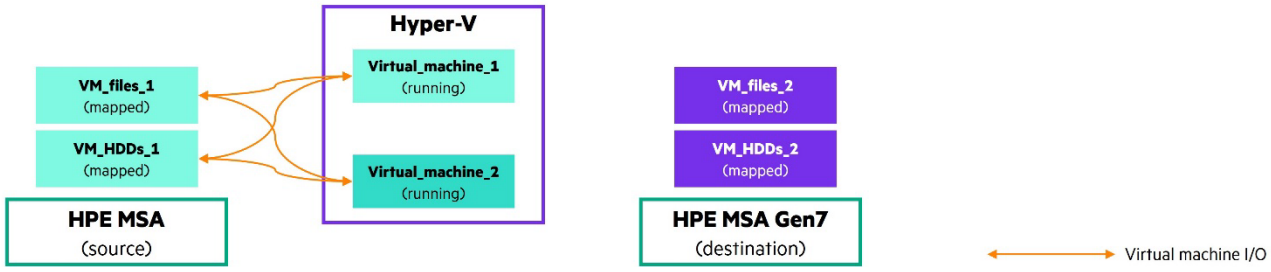

Figure 6. All VMs are running and are located on the existing storage system

2. The live migration task starts.

   As Figure 7 shows, during this time, the VM remains in a running state and disk I/O is directed to the storage subsystem of both the source and the destination file systems.

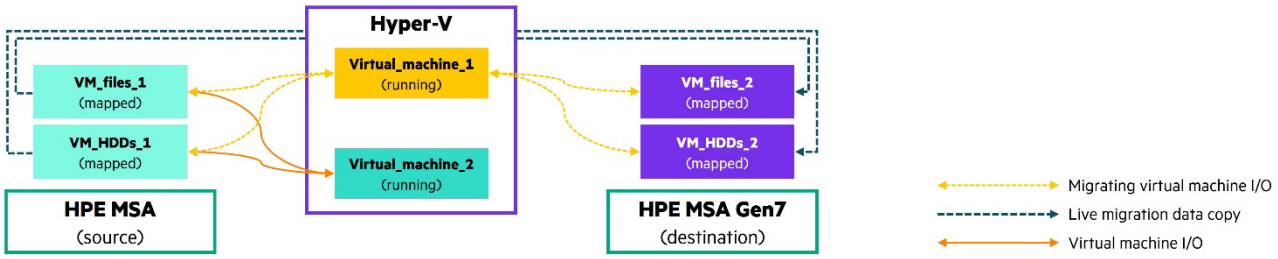

Figure 7. A live migration task is initiated for a single VM

3. After the live migration task is complete, the VM is located on the destination file systems, and all disk I/O for that VM is serviced by the storage subsystem that hosts it.

   As Figure 8 shows, because of the per-VM tasking of live migration, some VMs might still be located on the source file systems after the migration process is complete.

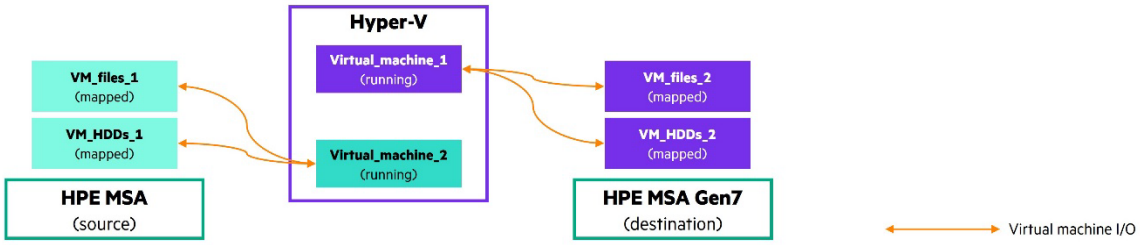

Figure 8. Replicate a source volume to the destination HPE MSA Gen7 Storage system

After all the VM files are migrated and copied, the next task is to take the source volumes that host the VM file systems offline. The steps vary depending on which system is involved:

In Windows Server

1. In the Microsoft Failover Cluster Manager console, navigate to **Disks**.
2. Right-click the now unused disk and select **Remove from Cluster Shared Volumes**.
3. Right-click the disk again and select **Remove**.

In the HPE MSA Storage system

Version 4 SMU

4. Navigate to the **Volumes** view and select the volume hosting the datastore that is to be disconnected from the VMware hosts.
5. From the **Actions** menu, select **Detach from hosts**, check the relevant box for the hosts and proceed.

Version 3 SMU

6. Navigate to the **Volumes** view and select the volume hosting the datastore that is to be disconnected from the VMware hosts.
7. Highlight all mappings by selecting the first entry and then, while holding the shift key, selecting the last one.
8. Right-click anywhere in the highlighted list or click **Actions** and select **Remove Mappings**.
9. Click **OK** to confirm the action.

Figure 9 shows the result after all actions are completed. The legacy storage system is no longer used by Windows, and all VMs are located on the new HPE MSA Gen7 array.

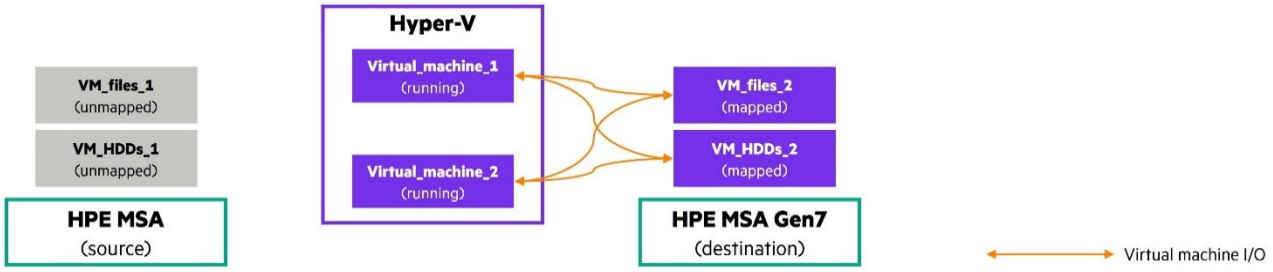

Figure 9. After all the VMs and files are migrated, the original file system can be unmounted, removed, and unmapped

HPE MSA remote snap

Unlike vSphere Storage vMotion or Hyper-V live migration, as a snapshot-based technology, HPE MSA remote snap does not support transparent VM migration. Instead of being moved, data is copied from an array-based point-in-time snapshot to a remote HPE MSA Storage system, and changes that occur to the volume after the snapshot is taken are not copied.

In addition, whereas hypervisor-level migration mechanisms are carried out on a per-VM basis, remote snap operates at the volume level. The most common deployment strategy for HPE MSA is to use large volumes that have a 1:1 relationship with a VMware datastore or a Windows file system and that contain multiple VM files or virtual HDDs. For that reason, remote snap is not the preferred migration technology for virtualized environments, especially when both convenience and uptime are accounted for. Remote snap is therefore well suited to bare-metal server installations, where it is more likely that applications are directly related to whole volumes.

To copy the data in full and with the most recent changes, it is necessary to stop I/O to a volume before the last replication. However, to minimize downtime, Hewlett Packard Enterprise recommends that you replicate data in stages:

1. As shown in Figure 10, initial replications of volumes can occur while the applications remain online. Because changes to a volume typically represent a small percentage of the overall allocated capacity a volume consumes, this approach takes the most time, but it reduces the time needed for subsequent replications.

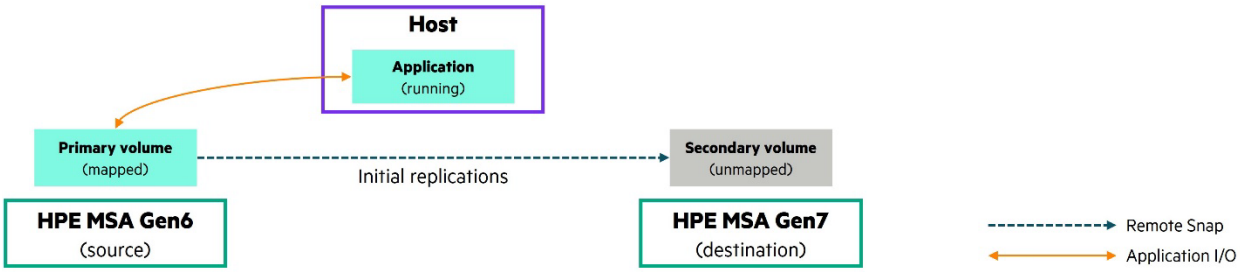

Figure 10. Applications continue to use their current volumes while snapshots are replicated to the destination array

2. After enough replications have occurred that change, deltas are reduced to their minimum, an application is brought offline, and a final replication is performed. Figure 11 shows that the destination volume remains, and the source volume should have its mapping removed. Unmapping the source volume prevents applications from unknowingly writing new data to a volume that is considered migrated.

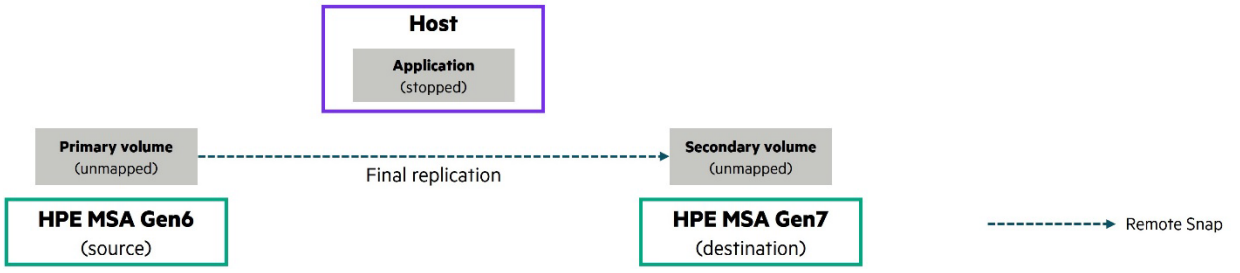

Figure 11. Applications are stopped and the volumes unmapped while a final replication is carried out to the destination array

3. As Figure 12 shows, after the final replication is complete, the replication set should be removed to make the destination volume available for mapping to a host.

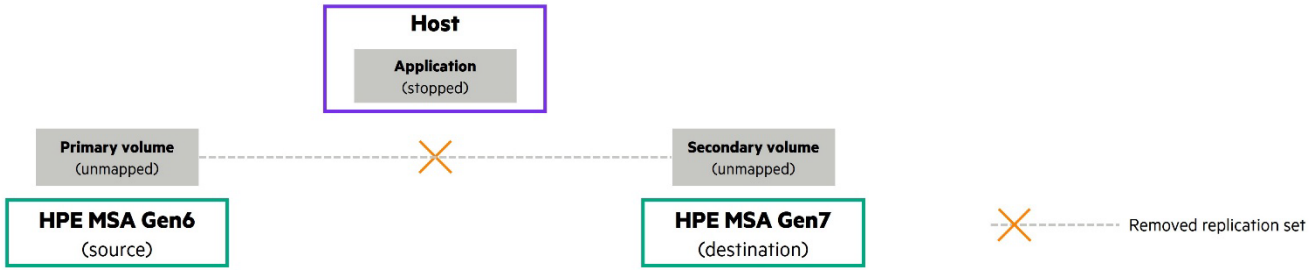

Figure 12. The replication set is removed, which in turn makes the destination volume a primary volume that can be mapped to a host

Note
After the replication set has been removed, HPE recommends that you take a snapshot of the destination volume before mapping it to a host.

4. Figure 13 shows the destination volume mapped to the host and the application resumed. The steps required to bring the destination volume online depend mainly on the host OS and fabric. However, if the destination volume is assigned the same drive letter or given a mount point in the same location as the source, no application reconfiguration should be necessary.

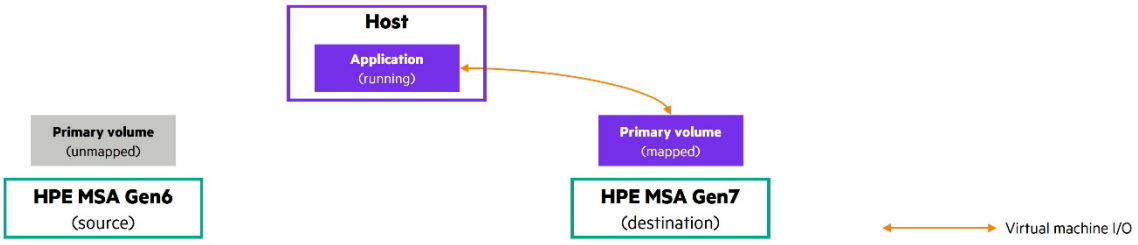

Figure 13. The destination volume is mapped to the host, and applications resume

Important
HPE supports remote snap between HPE MSA Gen7 Storage and HPE MSA Gen6 Storage only. As shown in Table 2, replication can work from older generations of HPE MSA arrays that support virtual storage and remote snap but is not supported.

Data-in-place upgrades

Data-in-place upgrades are theoretically simple yet carry extensive risks without sufficient planning. Figure 14 provides a high-level overview of the physical steps required.

1. Power down both source and destination systems.
2. Reinstall drives from the source to the destination system, making sure to place the drives in the same order.
3. Recable any attached expansion enclosures to the new system and in the same order.
4. Power on the destination system.

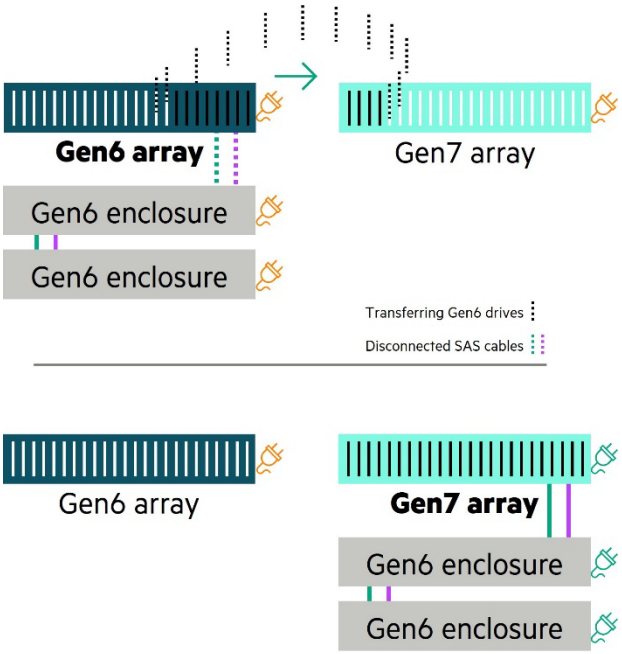

Figure 14. Example of the physical changes necessary for a data-in-place upgrade of an HPE MSA Storage system

The following running order of steps is arranged in blocks and should be carried out linearly. Careful consideration should be given to understand which are applicable and the actions required are both understood and feasible.

Destination HPE MSA system

1. Physically install and cable the array.
2. Power the array and proceed through the out-of-the-box experience (OOBE), configuring all standard management settings.
3. Confirm and remedy system health using HPE MSA Storage Health Check.
4. Upgrade system firmware to the latest version, downloadable from HPE Support.
5. Follow steps for configuring the broader infrastructure.
6. Optionally reassign or define friendly nicknames to host initiators and recreate hostnames and groups.
7. Shut down the controllers either from the SMU or through the CLI command: **shutdown both**
8. When the white LEDs on both controllers are illuminated, power down the system.

Important

Some models of HPE MSA Gen7 Storage ship with preconfigured disk groups, HPE MSA 2070 flash bundles, for example. It is highly important that all existing disk groups are removed prior to performing a data-in-place migration. Failure to carry out this procedure has a high probability of causing a conflict requiring a support case, extending application unavailability and risking data loss.

Broader infrastructure

1. If using a SAN, reconfigure so that hosts can reach the target ports of the new system. For example, create Fibre Channel zones that include connected array host ports and host HBA ports.
2. Reconfigure hosts initiators to target the host ports of the new HPE MSA system so that they are discovered by the HPE MSA.

Source HPE MSA system

1. Confirm and remedy system health using HPE MSA Storage Health Check.
2. Upgrade system firmware to the latest version, downloadable from HPE Support.
3. Remove any remote snap configuration including schedules and peer connections.
4. Cease all application I/O.
5. If using full disk encryption, secure the drives using the **Lock for transport** option within the SMU or through the CLI command: **clear fde-keys currentpassphrase**

Important
If the source system is secured with full disk encryption and the passphrase is not known, it will not be possible to perform a data-in-place upgrade. In such a scenario, all data must be copied to another location through a host-based technique and the drives erased using the **Repurpose secured disks** function.

6. Shut down the controllers either from the SMU or through the CLI command: **shutdown both**
7. When the white LEDs on both controllers are illuminated, power down the system.
8. Power down any expansion enclosures.
9. Remove drives and reinstall in the destination system in the same slot order.
10. Disconnect any expansion enclosures and recable to the destination system, maintaining the same enclosure order.

Destination HPE MSA system

1. Power on expansion enclosures and wait two minutes.
2. Power on the array.
3. If using full disk encryption, unlock the drives with the passphrase using the Secure system option within the SMU, or through the CLI using the commands: **set fde-lock-key passphrase currentpassphrase** and **set fde-state secure passphrase currentpassphrase**
4. Perform a rescan either through the SMU or using the CLI command: rescan
5. Verify host mappings or create new ones as needed

Disconnecting decommissioned storage systems

Depending on which SAN protocol is in use, HPE recommends that you remove the logical and physical connections between the host and the legacy storage system. Removing the connections limits the chance of any accidental use of a decommissioned storage system and stops timeouts to inactive systems or systems that are powered-off. The actions depend on the protocol:

- **For fabric-attached Fibre Channel SAN arrays**, remove the zones between the array and hosts from the active configuration; or for direct-attached configurations, remove the cables.

Caution
If a zone contains devices other than the HPE MSA Storage system and its hosts, make sure that alternative zoning is in place before removing it from the active configuration.

- **For iSCSI-attached SAN arrays**, remove persistent connections to the legacy targets. For example, **static discovery** entries for VMware and favorite targets for Microsoft Windows.
- **For SAS-attached systems**, remove zones if using HPE BladeSystem SAS switches; or for direct-attached arrays, remove the cables.

Conducting dry runs

A dry run is a simulated migration that is performed to uncover unexpected problems in a migration plan. A dry run can help to remove negative outcomes and address challenges by limiting the scope of a migration to a small yet representative subset of data. If problems are discovered during a dry run, the migration should be stopped and the problems solved. After the problems have been addressed, the dry run should be repeated. No migration should proceed without the completion of at least one successful dry run.

A dry run must include user acceptance testing by a subset of the users who will perform post-migration user acceptance testing.

Note
Dry runs cannot be carried out for migration strategies that lack a mechanism to return to the previous running configuration without risk of data loss or significant disruption. Dry runs are therefore unfeasible for data-in-place migrations outside of a lab environment.

Choosing candidates

A dry run for a block-level migration involves either VMs or array volumes. Because hypervisor-based migration technologies move data instead of copying it, it is not possible to perform a dry run on the same VMs that are to be migrated later. Instead, either create new VMs that have an equal footprint on disk or clone existing VMs. Regardless of which approach is chosen, be sure to simulate application workloads within the VM, using any preferred synthetic workload tool.

Preparing for rollback

Rollback is the process of reversing the decision to proceed with an active migration, or reverting after a migration is complete. A rollback can have significant implications, and complexity can increase quickly as more data is modified. The decision on how to proceed might vary considerably from one situation to another, and contingency planning must include the option of a full restore from backup.

It is beyond the scope of this paper to describe how to recover after a cutover has occurred, but with a careful approach, it might be possible to avoid the need for a rollback. If detailed dry runs are carried out and user acceptance testing is thorough, such a scenario is avoidable and highly unlikely.

Stop triggers

During dry runs, as well as during the final migration, several events can occur that indicate a failing or failed migration or indicate that the migration is having a negative impact on the production environment. Examples of reasons to halt a migration include the following:

- A high number of failed reads/retries from the source
- A high number of failed read/writes/retries to the destination
- Severe performance impact to the source system
- Severe performance impact to the destination, where previously migrated data is now located
- Severe performance impact to the data network

The triggers are usually a result of these conditions:

- A mistimed migration
- Too many concurrent operations
- Inadequate performance capability of participating components

In the absence of an anomaly, it is likely that the legacy source system is incapable of consistently matching the potential ingest rate of the destination and is saturated. It is also possible that the destination system was not properly sized. Both conditions are indicators of either an inadequate discovery phase or an insufficient number of dry runs.

HPE recommends that dry runs include a simulated rollback.

Data migration examples

This paper explores four typical scenarios and describes a potential running order for each. Each scenario is encapsulated by common tasks that are likely to form part of any migration. However, only those related directly to migrating data are defined in detail in this paper.

- Migrating from any storage system to an HPE MSA Gen7 Storage system by using vSphere Storage vMotion
- Migrating from any storage system to an HPE MSA Gen7 Storage system by using Microsoft Hyper-V live migration
- Migrating from an HPE MSA Gen6 Storage system through remote snap

Both vSphere Storage vMotion and Microsoft Hyper-V live migration simplify the task of migrating the underlying storage of a VM and make the process faster. Both hypervisors are designed so that in most cases there is no need to bring VMs offline and requires no lengthy post-migration steps. In contrast, HPE MSA remote snap offers the ability to migrate both application and boot volume data for any OS that is supported by both source and destination HPE MSA arrays, but at the cost of additional steps and planned application downtime.

With many of the most significant considerations addressed, these scenarios aim to serve as a high-level suggestion of the steps required to migrate data in simple environments. These examples should not be directly translated into full migration plans, but they are workable and tested solutions when taken with careful consideration.

VMware Storage vMotion

In the example shown in Figure 14, VM **VMware_application_1** is moved from its existing datastore, **VMware_store_1**, to a new datastore named **VMware_store_1_Gen7**. Each volume has been mapped to all hosts in the VMware cluster and formatted with the VFMS file system.

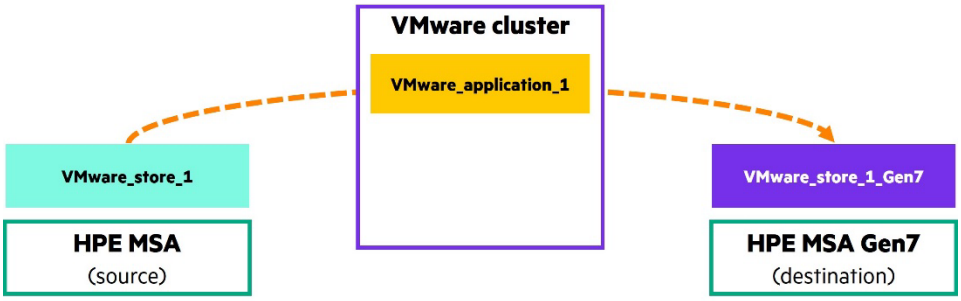

Figure 15. Visualization of an example storage migration of a VMware virtual machine

The following steps should in most cases enable the seamless migration of a VM's storage between datastores located on any storage system.

1. From VMware vCenter, either select multiple VMs in the same power state, or locate an individual VM, right-click, and select **Migrate...**.

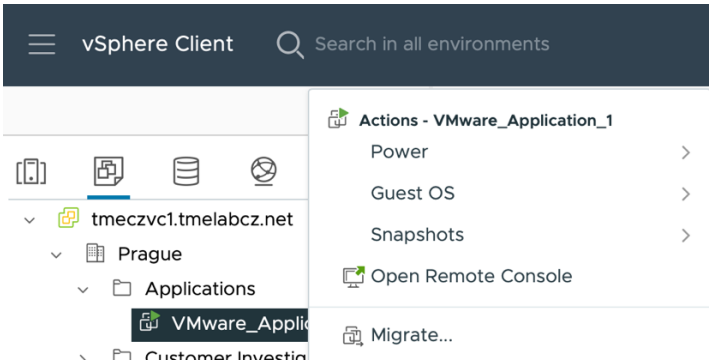

Figure 16. How to begin migrating a VM within VMware vCenter

2. Select **Change storage only**.

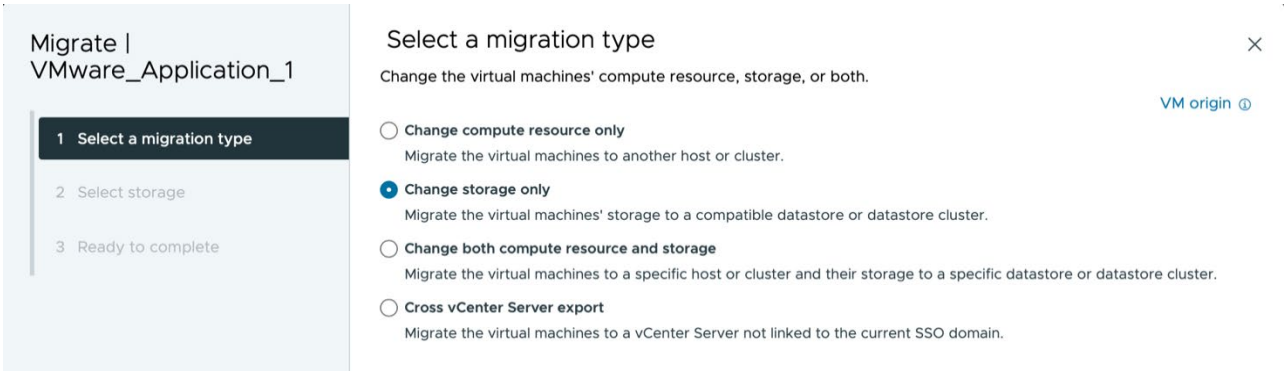

Figure 17. Selecting storage-only migration for selected VMs

3. Select a destination datastore that is located on the new HPE MSA Gen7 Storage system.

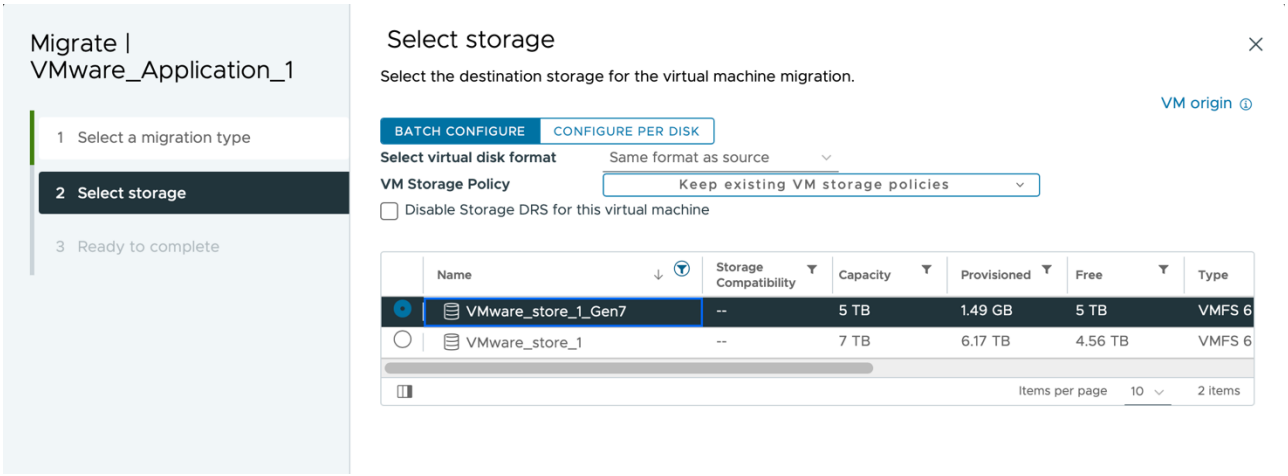

Figure 18. Select a destination datastore located on an HPE MSA Gen7 Storage system

4. Review the selected options and click **FINISH**.

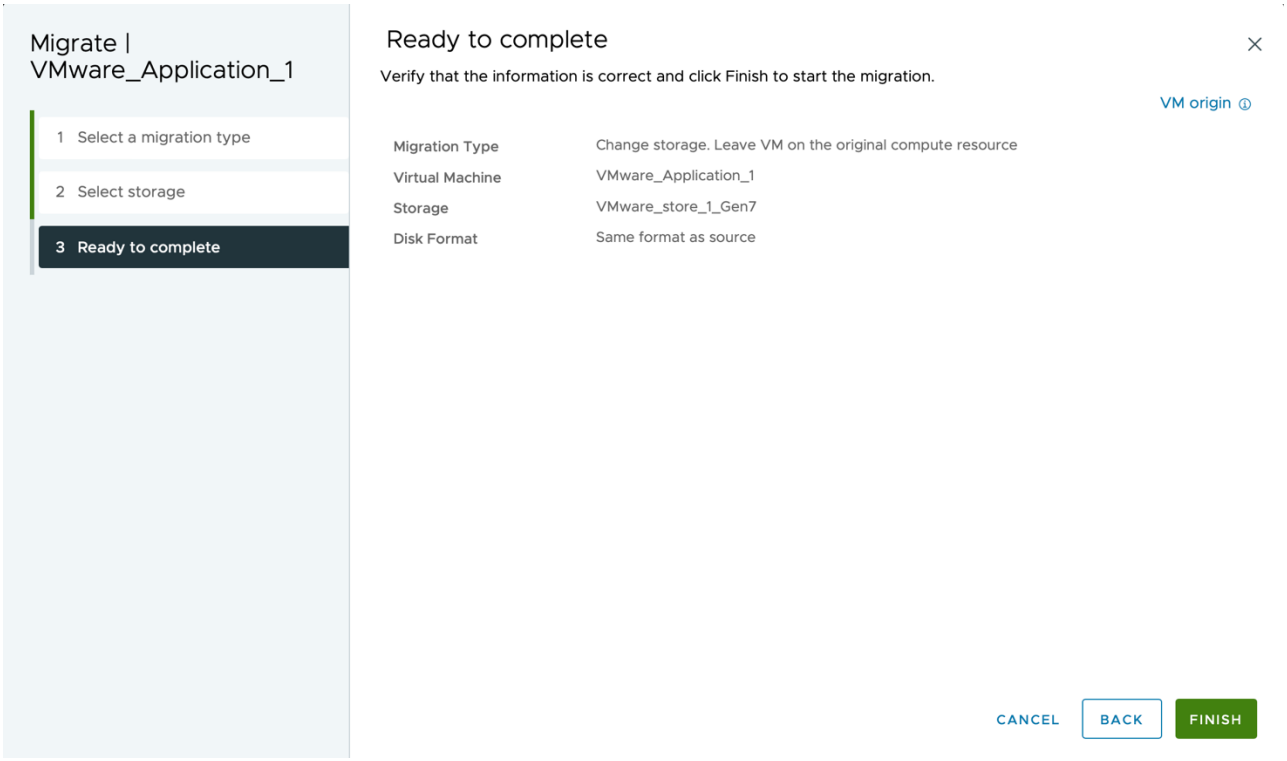

Figure 19. Confirm migration settings

5. In VMware vCenter, monitor the **Recent Tasks** status for the completion of the task.

   The time this task requires depends highly on the quantity and size of the virtual hard disk files as well as any competing I/O.

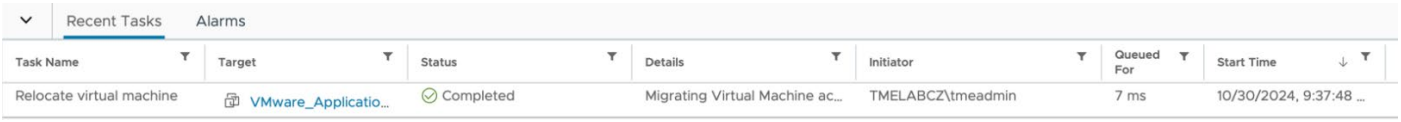

Figure 20. Task monitoring within vCenter showing the successful relocation of a VM to a new datastore

Microsoft Hyper-V live migration

Unlike vSphere Storage vMotion, Microsoft Hyper-V live migration can move VM storage between file systems, whether the hosts are stand-alone or in a failover cluster. However, the management console and subsequent steps differ slightly. The migration examples listed in Table 9 contain four volumes each.

Table 9. Example mount points and drive letters

Detail	Clustered	Stand-alone
Source VM files	C:\ClusterStorage\HyperV_VMFiles	E:\
Source VM virtual disks	C:\ClusterStorage\HyperV_HDDs	F:\
Destination VM files	C:\ClusterStorage\HyperV_VMFiles_Gen7	G:\
Destination VM virtual disks	C:\ClusterStorage\HyperV_HDDs_Gen7	H:\

For VMs in a failover cluster

When Microsoft Failover Clusters are used, physical disks presented to hosts from a SAN storage array are ordinarily mounted within an existing file system and configured as a Cluster Shared Volume (CSV). In the example shown in Figure 21, four volumes are mapped to hosts participating in the cluster and mounted in the folder **C:\ClusterStorage** in CSV format.

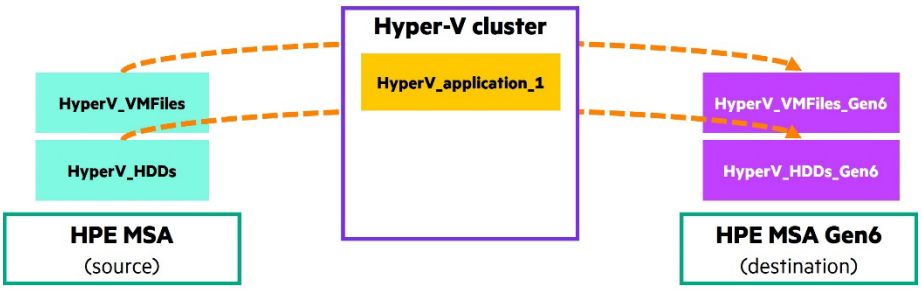

Figure 21. Visualization of the example storage migration of a Hyper-V VM located within a cluster

To perform a storage live migration, complete the following steps:

1. Open the Failover Cluster Manager console and, in the Roles view pane, highlight the VM you want to migrate, right-click, and select **Virtual Machine Storage** from the **Move** menu options.

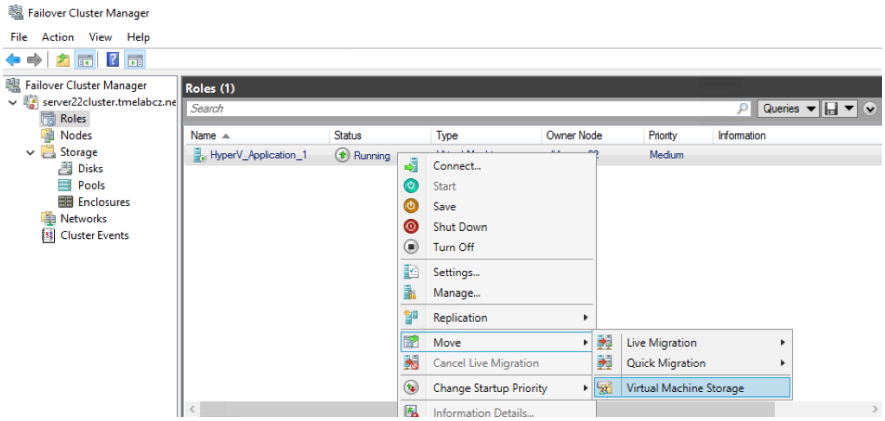

Figure 22. Selecting storage migration for a highlighted VM within the Microsoft Failover Cluster Manager console

2. Drag each item down to the relevant target volume in the Cluster Storage tree[3] and click **Start**.

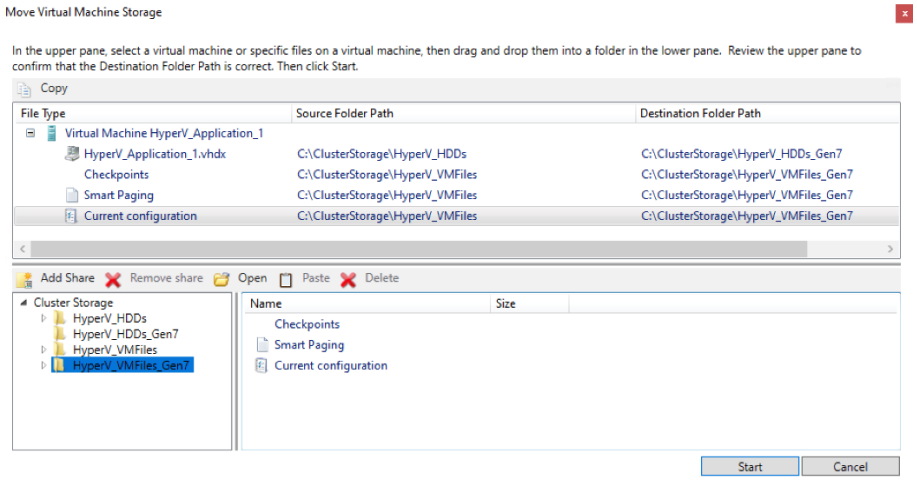

Figure 23. Defining a new destination folder path for the VM related files

[3] For convenience, each volume has been given a friendly name in advance. This can be done through Windows File Explorer and should match the name in the Disks panel of the Failover Cluster Manager console. It should also match the volume name in the HPE MSA SMU.

3. Check the **Information** column in the **Roles** view pane to confirm that migration has started.

   Hyper-V begins migrating the VM to the new file systems, but it provides minimum information related to its progress.

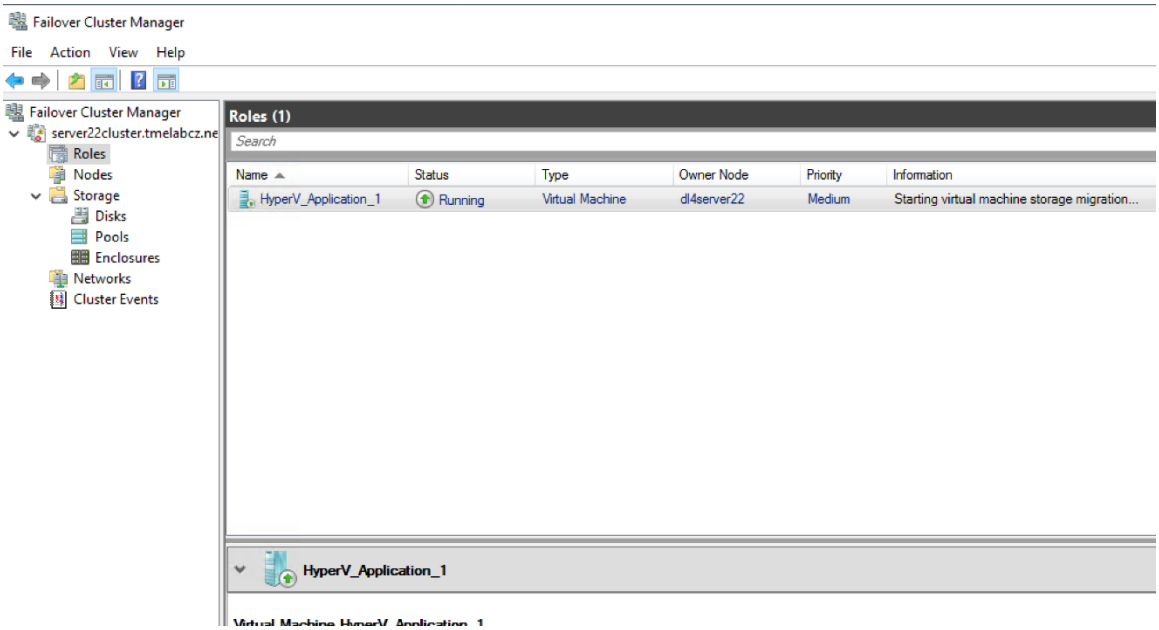

Figure 24. Hyper-V **Roles** view with **Information** showing a running migration

4. **Optional:** For more detailed evidence of a successful migration, open the **Event Viewer** for the cluster node on which the VM is running, browse to **Event Viewer → Applications and Services → Logs → Microsoft → Windows → Hyper-V-VMMS → Admin**, and look for Event ID 20927.

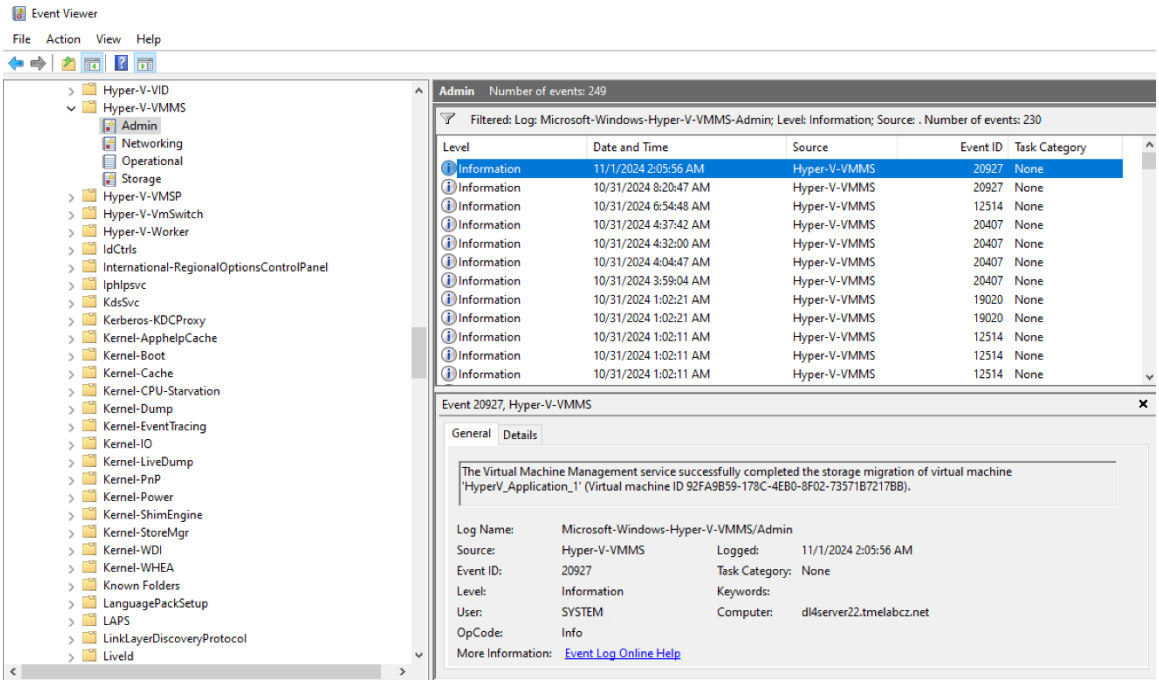

Figure 25. Event viewer entry confirming a successful storage migration

For VMs on a stand-alone host

The process of migrating storage for VMs on a stand-alone host is like that of migrating storage for VMs in a cluster, except that it is managed through the Hyper-V Manager console. In the example shown in Figure 25, data is moved between drive letters. (It is also possible to mount a volume as a mount point in an existing file system, in a way like the process for CSVs.)

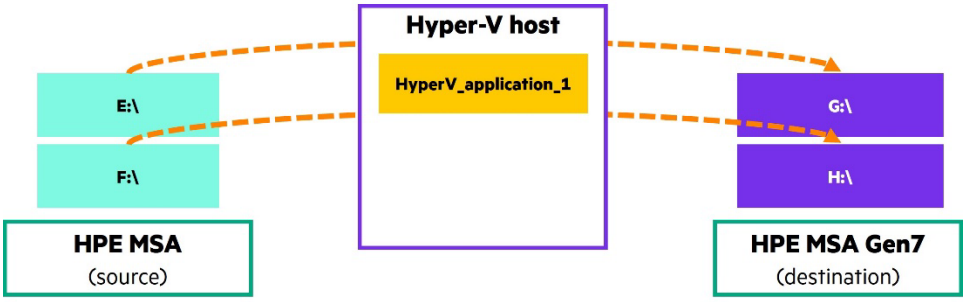

Figure 26. Visualization of the example storage migration of a Hyper-V VM located on a stand-alone host

To migrate storage for VMs on a stand-alone host, complete the following steps:

1. Open the Hyper-V Manager console, select the host that is running the VM, highlight the VM to migrate, right-click, and select **Move...**.

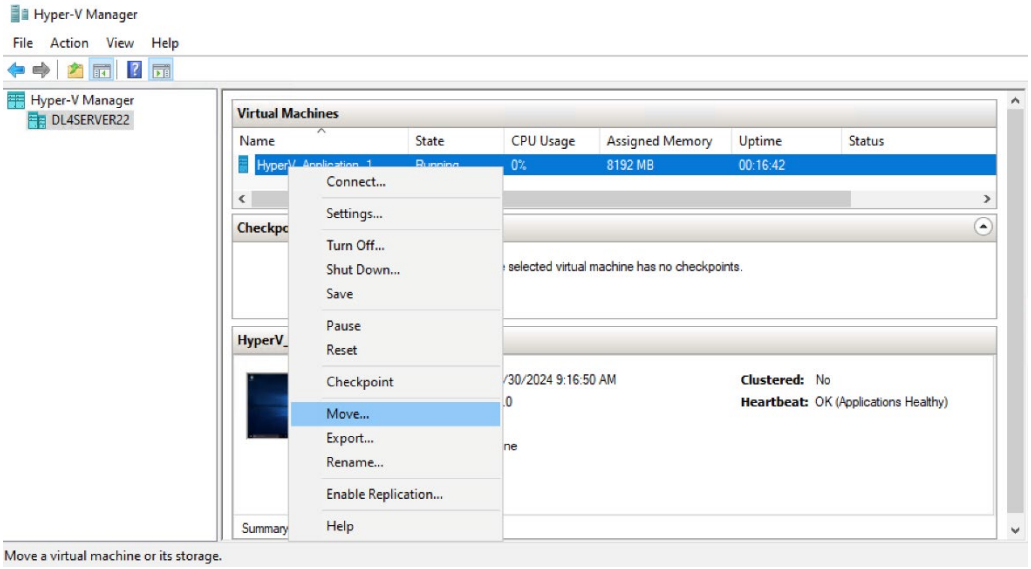

Figure 27. Starting a migration for a highlighted VM in the Microsoft Hyper-V Manager console

2. Select **Move the virtual machine's storage** and click **Next >**.

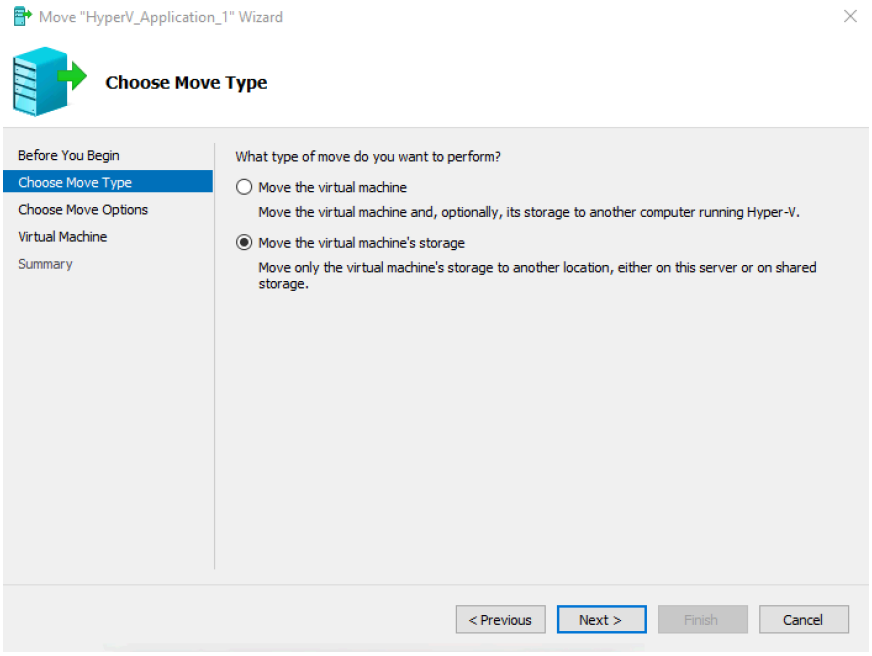

Figure 28. Selecting to move a VM's storage from the Hyper-V migration wizard

3. Select **Move the virtual machine's data to different locations** and click **Next >**.

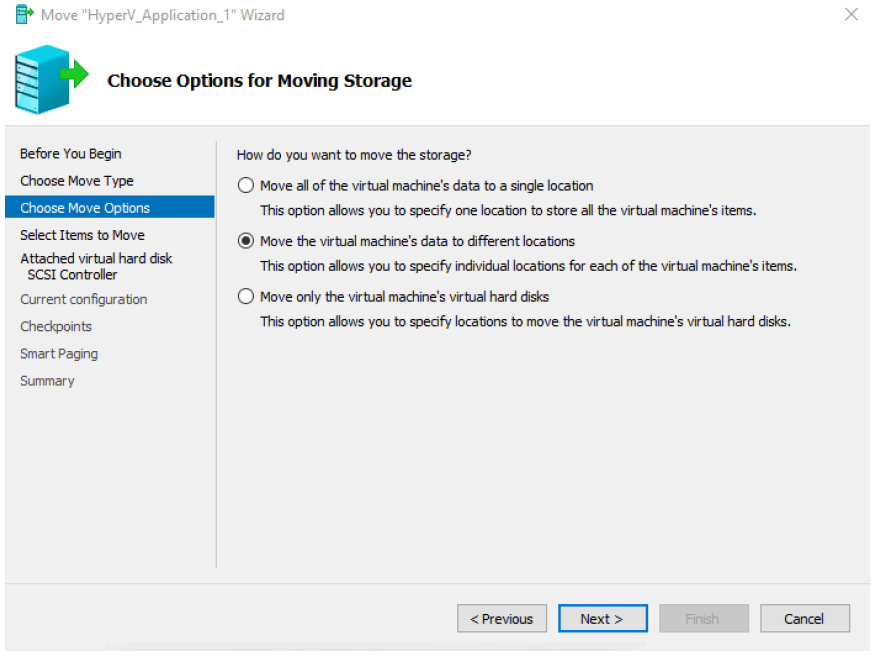

Figure 29. Selecting the option to move a VM's files by category from within the Hyper-V migration wizard

4. Leave all items for the VM selected and click **Next >**.

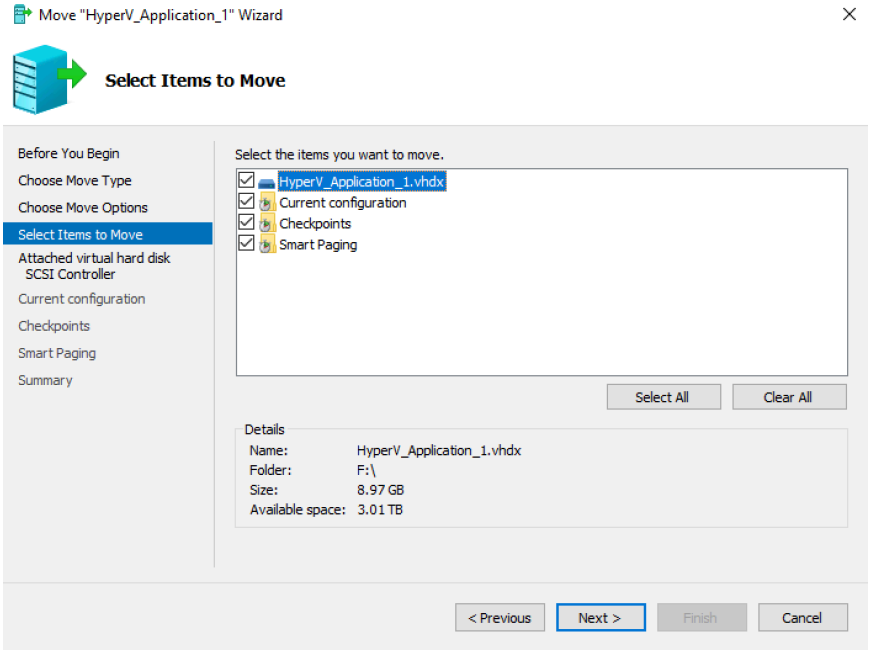

Figure 30. Selecting VM files to be migrated from within the Hyper-V migration wizard

5. Click **Browse...** and specify an appropriate new location for the relevant VM data type.

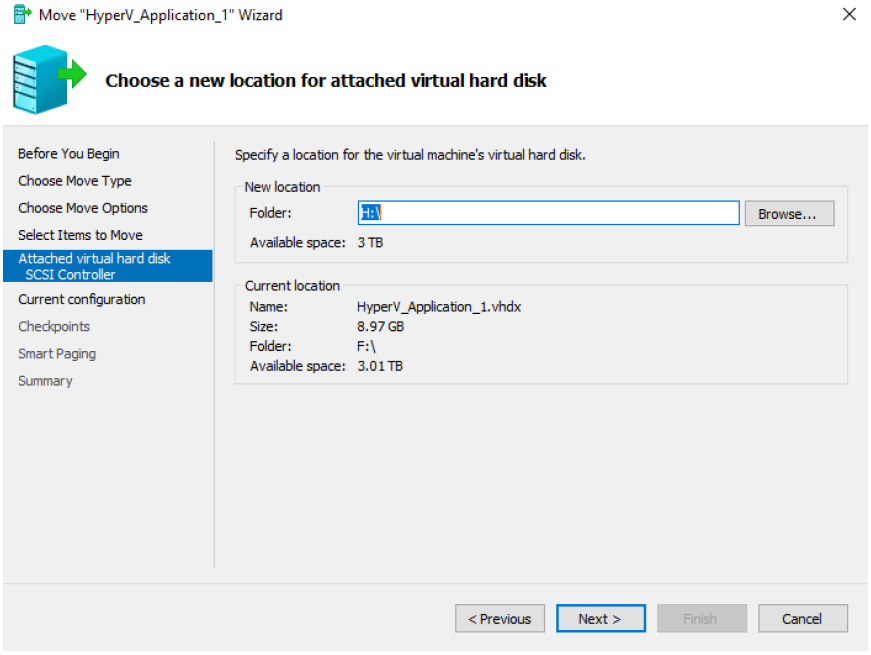

Figure 31. Specify new location for VM data type

6. Click **Next** and repeat **Step 5** for all data types.

7. Review the changes and click **Finish** to start the migration.

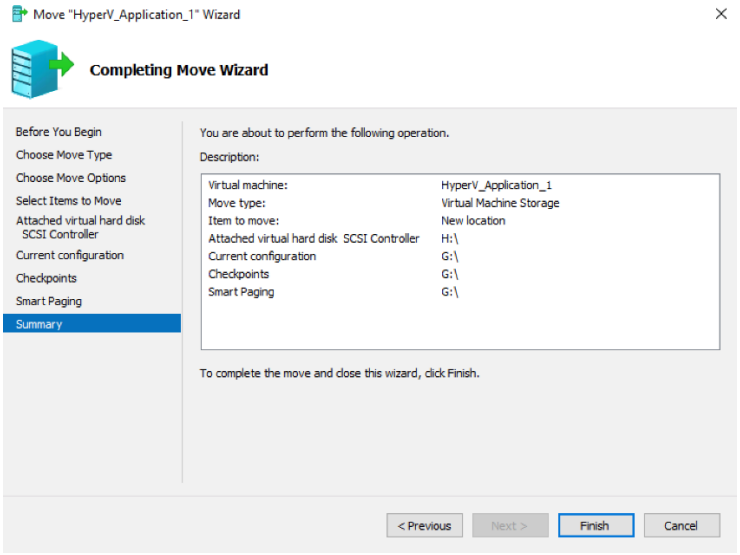

Figure 32. Confirmation screen of the Hyper-V migration wizard

Unlike what happens when you migrate a VM in a cluster, a progress bar blocks further actions until the migration is completed or fails.

Remote snap

The following procedure demonstrates how to use remote snap to replicate volumes from HPE MSA Gen6 to HPE MSA Gen7. The steps demonstrated are limited to the actions that must be carried out on the HPE MSA arrays. Remote snap is best suited to the migration of bare-metal servers whose application data or boot volume are located on HPE MSA arrays and where a data-in-place upgrade is unsuitable. Describing how to configure all supported OSs and HBAs would require more steps than can be included in this paper. Always refer the relevant documentation for all aspects of the specific environment before beginning a migration and document all the necessary steps.

Important

See Table 2 to ensure that the migration plan does not exceed the maximum number of peer connections supported by either HPE MSA array.

To replicate volumes by using remote snap for the first time, complete the following steps:

1. Add data protection to a volume.

   From within the SMU of the HPE MSA Gen6 Storage system, as shown in Figure 33, navigate to the **Volumes** menu and select the volumes you wish to replicate by selecting **Add Data Protection** from the drop-down menu.

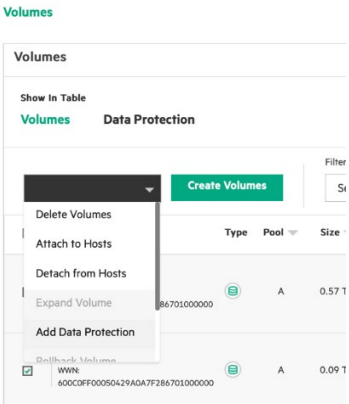

Figure 33. Adding data protection from an HPE MSA Gen6 to selected volumes

As shown in Figure 34, confirm that **Remote Replication** should be added to the selected volumes.

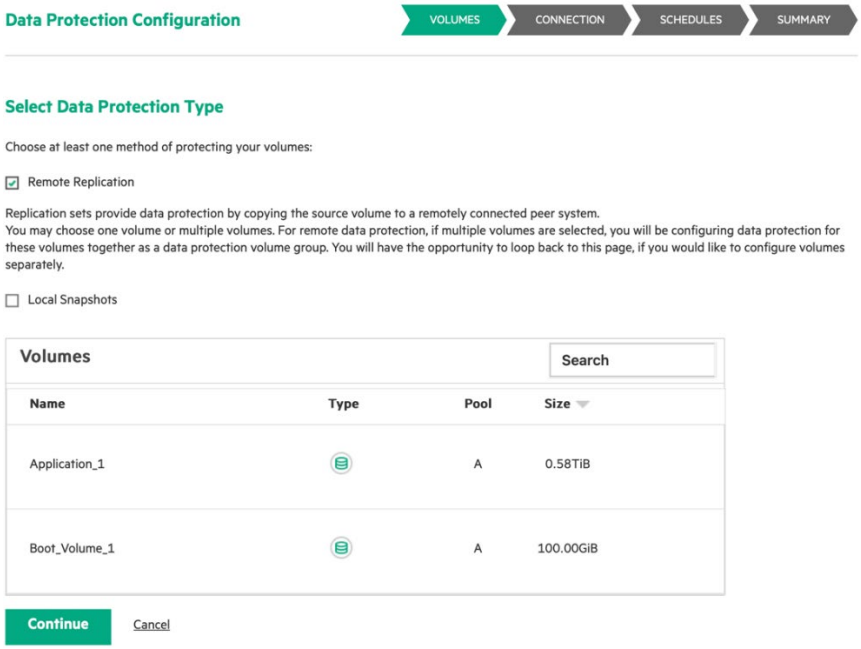

Figure 34. The first step of adding remote replication to selected volumes from an HPE MSA Gen6

2. Create a peer connection.

   A peer connection is the name given to a relationship between two arrays and contains one or more datapaths through which replication traffic is passed. To create a peer connection there must be an appropriately configured storage network between participating array host ports. iSCSI systems require that at least one host port of each HPE MSA array be configured within the same VLAN and be on the same subnet, or that routing is configured so as to allow traffic to pass between disconnected networks. Fibre Channel systems require that at least one array host port from each system is configured within the same zone. As per HPE best practices, it is not necessary or recommended that all host ports be in the same subnet or zone. Host ports with connectivity to each other are automatically utilized to provide additional bandwidth and tolerance to link failure.

   When creating a peer connection, enter a host port address of the remote system to be used for path discovery. Figures 35 and 36 show the defining of a remote iSCSI host port IPv4 address and a Fibre Channel WWPN respectively.

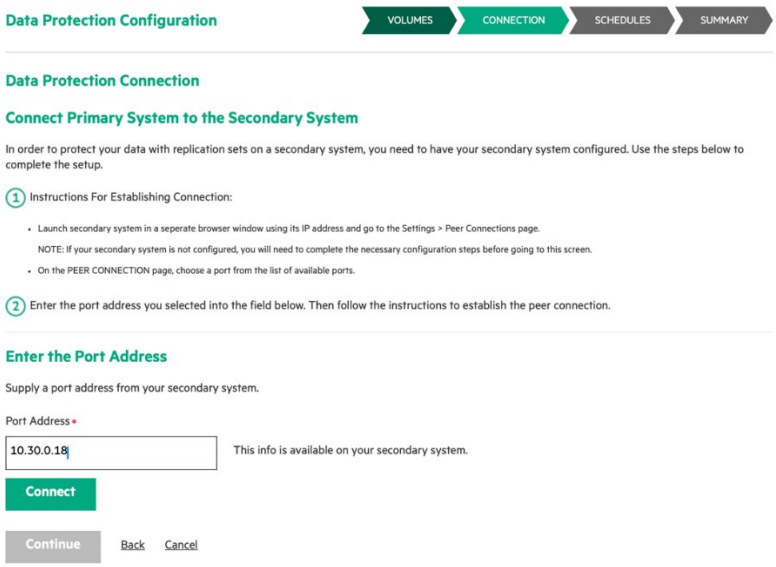

Figure 35. Creating a peer connection to the remote system over iSCSI

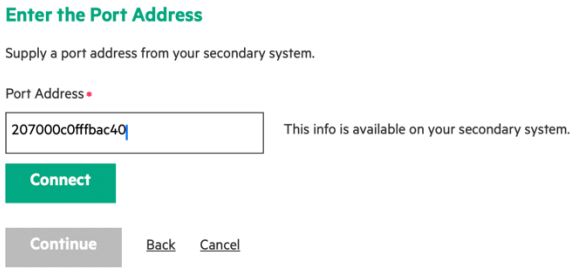

Figure 36. Creating a peer connection to the remote system over FC

3. Create a replication set.

   A replication set contains the replication conditions for volume or group of volumes. These conditions include the peer connection to be used, the destination pool, and the volume name as it will appear on the remote system as well as schedules if configured. Replication sets must be created by the source system.

Important

An HPE MSA array supports up to 32 replicated volumes per system. A volume can only be a member of one replication set and a replication set can support up to 16 volumes.

Figure 37 shows the defining of **Replication Set Name**, which should align to what is being replicated. In this example, multiple volumes are grouped into a single replication set called **data_migration**. Additionally, the destination pool and whether to queue replications can be set differently to their defaults.

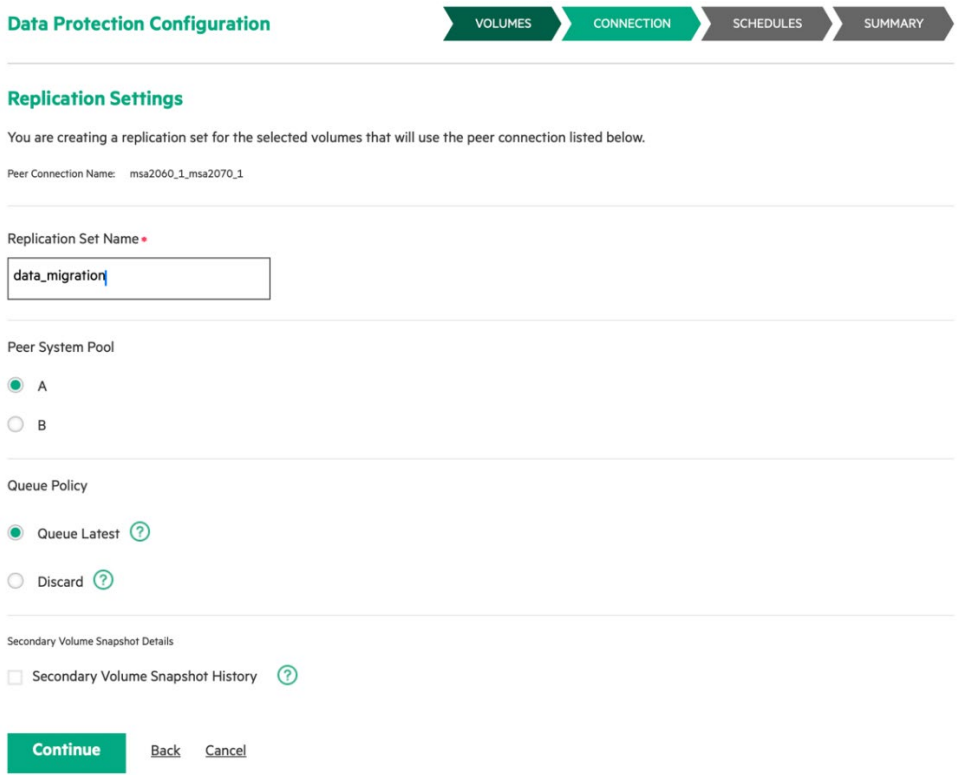

Figure 37. Defining a replication set and setting parameters

Next, confirm whether the replication should initiate for the replication set immediately after completing the wizard and if required, define a replication schedule, as shown in Figure 38. The decision to use schedules depends on what the discovery phase uncovered in terms of the total data that is to be migrated, and the constraints imposed. For example, if all volumes can be replicated within a relatively small period, schedules are likely not necessary and can be initiated manually as needed. Conversely, if the initial replication of all volumes is expected to take many days or weeks, then it might simplify the management overhead to have the array automatically replicate volumes to a specific schedule.

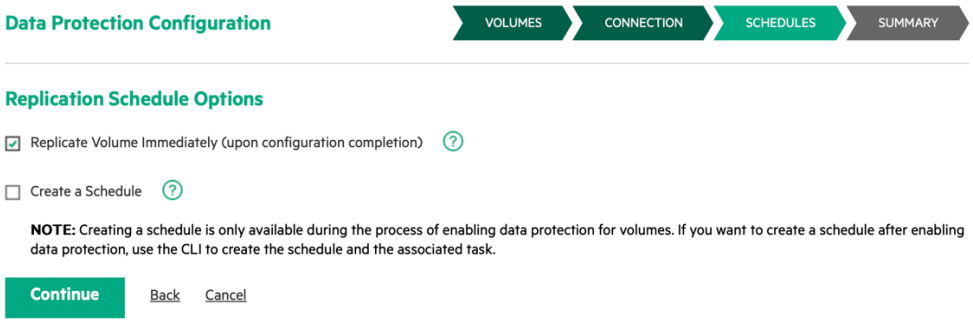

Figure 38. Confirming and defining a replication set's scheduling options

Finally, as seen in Figure 39, review all options before clicking **Apply** to finish.

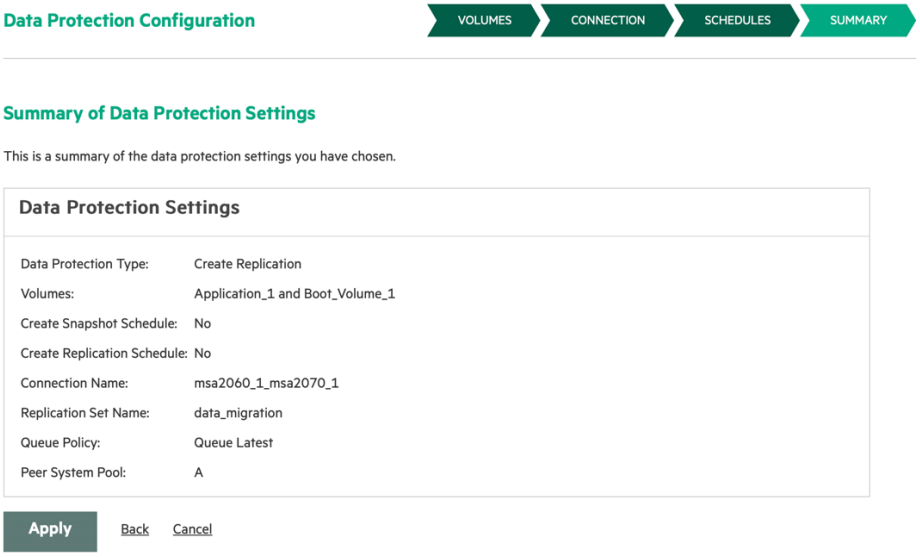

Figure 39. Data protection wizard settings review

4. Replicate all volumes.

   Whether adding volumes to individual replication sets or creating multiple sets with specific volume groups, additional replications are configured by the same process of selecting volumes by choosing **Add data protection**. However, if replicating to the same destination array, choose **Select from an existing peer connection** and select the relevant peer connection from within the drop-down box, as shown in Figure 40.

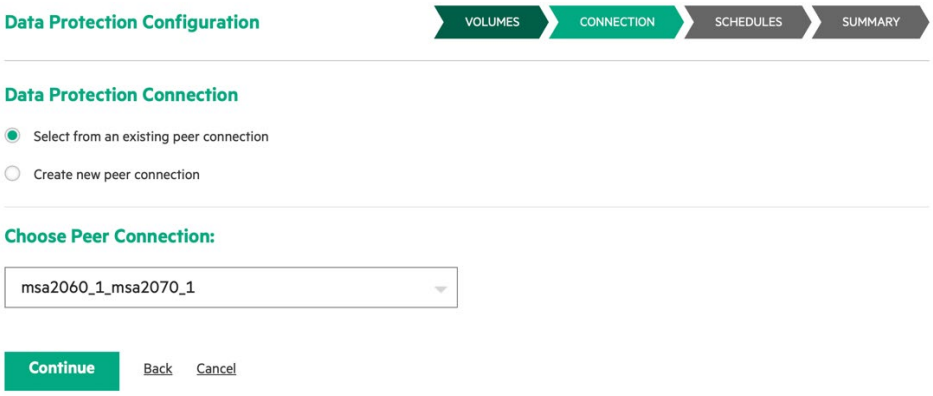

Figure 40. Selecting an existing peer connection for a new replication set

To manually start or restart replication for a volume or a volume group, navigate to the Volumes view, click the icon next to a volume that is configured for replication, select the **Replications** tab, and click **Start Replication**, as shown in Figure 41.

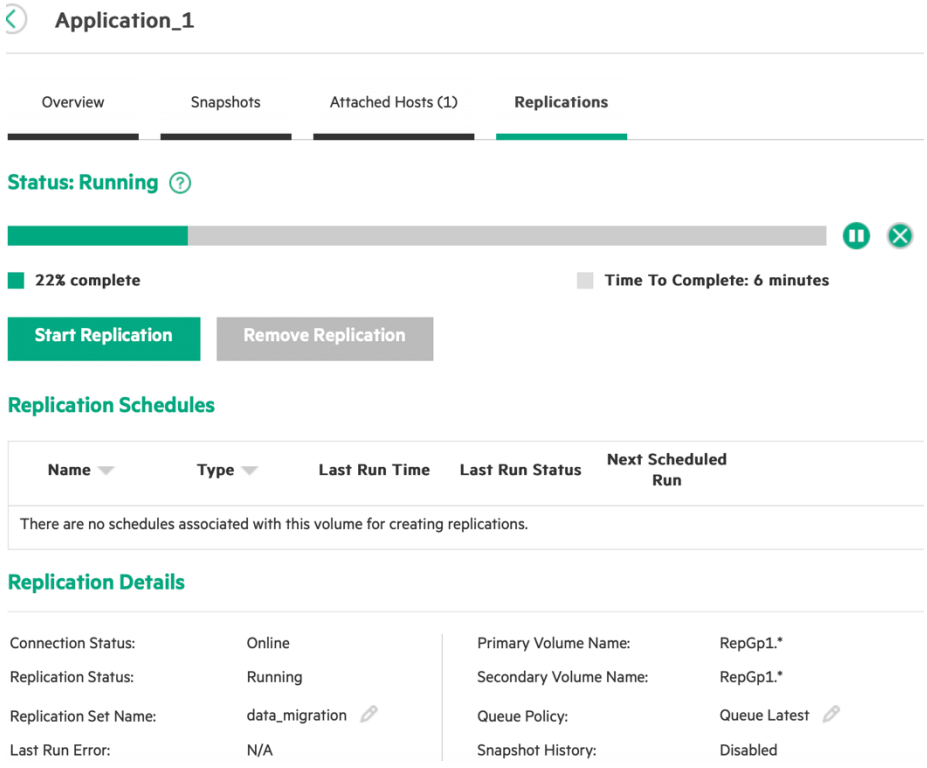

Figure 41. Manually initiating and monitoring a replication

5. Monitor replication progress and confirm within the **Activity** view on the source array of replication completion. As shown in Figure 42, filter **Recent** activity by Replication and look for descriptions including **A replication completed successfully** followed by the replication set name and serial number.

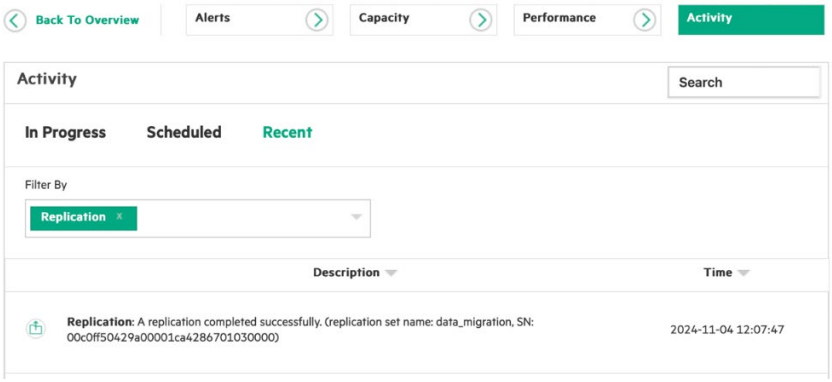

Figure 42. Confirming replication completion from the source array activity view

6. When all have replicated, successfully cease all I/O to the source volumes. The safest way to do this is to temporarily shut hosts down before performing an **unmap** from the source HPE MSA Storage system. There are multiple ways to unmap a volume but the most common method as shown in Figure 43 is to navigate to the volume properties, select the **Attached Hosts** tab, and click **Unmap** for all hosts.

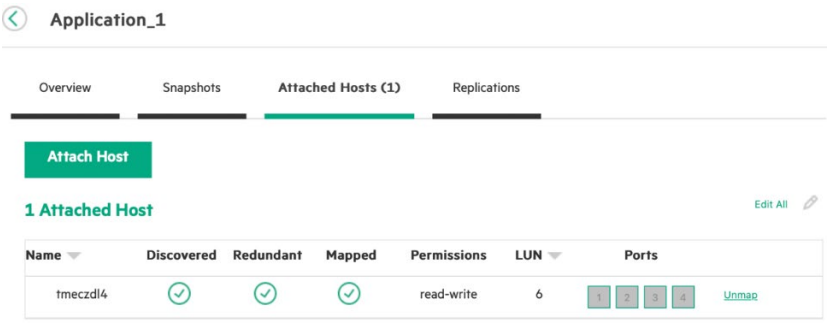

Figure 43. Unmapping a volume from the HPE MSA Gen6 array and a host

7. Repeat Steps 4 and 5 to replicate the volume for the final time to ensure any changes to the volume since it was unmapped and replicated.
8. Move to the SMU of HPE MSA Gen7 array and click the chevron to see the properties of a replicated volume from within the **Volumes** view, as shown in Figure 44.

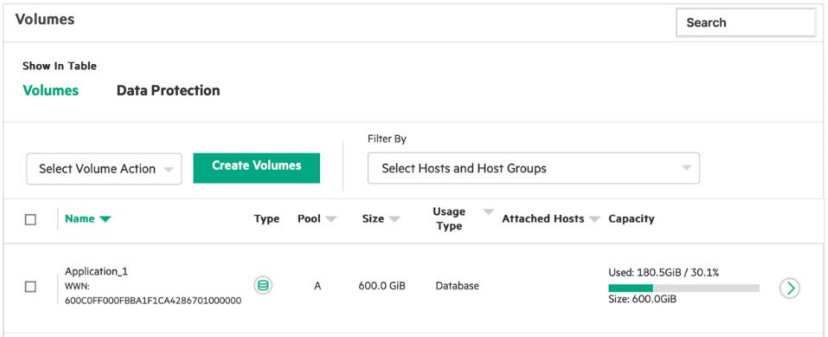

Figure 44. Viewing the secondary volume within the HPE MSA Gen7 SMU

As shown in Figure 45, move to the **Replications** tab and click **Remove Replication** to break the volume's relationship with the source array and allow it to be mapped to a host.

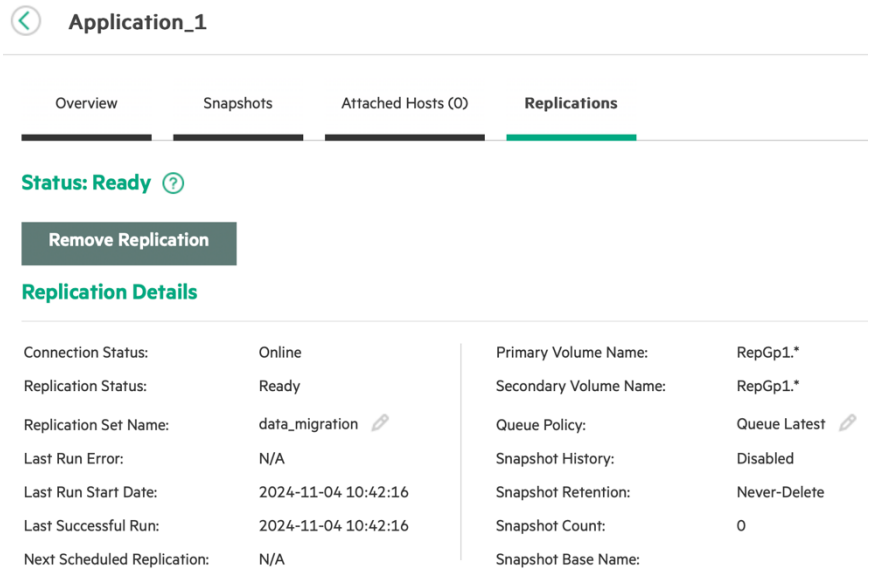

Figure 45. Replication information for a selected volume on a destination HPE MSA Gen7 array

Note
Hewlett Packard Enterprise recommends that you take array-based snapshots of both the source and the destination volumes. These snapshots can assist in recovering data if problems arise later in the migration.

9. Map migrated volumes to their original hosts through the **Attached Hosts** tab of the volume's properties, click **Attach Host** to map the replicated volume to hosts. Follow the steps in the HPE MSA 2070/2072 Storage Management Guide for more detailed information on how to complete this task.

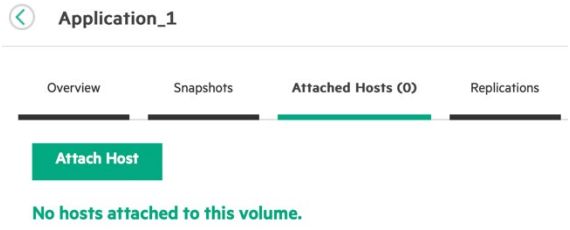

Figure 46. Attached hosts panel from within the HPE MSA Gen7 SMU

10. Configure the host OS and any HBAs to connect to the HPE MSA Gen7 Storage system.
11. Resume application I/O and repeat until all volumes are replicated and all application I/O has resumed.

Summary

HPE MSA Storage systems belong to an ever-evolving platform that has seen continuous development for over a decade. Useful new features, improved performance, and consistent pricing across each new generation have kept it a reliable choice for SMBs and large-scale customers alike. With careful planning and the help of this white paper or HPE, any organization that owns aging systems—whether from HPE or another vendor—can move toward HPE MSA Gen7 Storage easily and with minimal risk.

Resources

HPE MSA Gen7 virtual storage technical reference guide

HPE MSA 2070 and HPE MSA 2072 Storage arrays best practices

HPE MSA Storage Health Check

Ninja Online for HPE MSA

HPE MSA Gen7 Storage QuickSpecs

HPE Support Center

Sign up for HPE updates

Learn more at

HPE.com/us/en/storage/msa-shared-storage.html

a50011807ENW

Hewlett Packard Enterprise