

Migrating to HPE MSA Gen6 Storage

Strategy and technical reference guide



Contents

Executive summary.....	3
Target audience.....	3
Document purpose.....	3
Assumptions.....	4
Considerations for data migration.....	4
The discovery phase.....	5
Infrastructure.....	5
Data networks.....	5
Other systems.....	6
The planning phase.....	10
Creating a migration strategy.....	11
Defining source data.....	11
Choosing a general data movement strategy.....	11
Identifying data copy windows.....	13
Defining a running order of tasks.....	14
Setting up user acceptance testing.....	15
Planning for post-migration monitoring.....	15
Sizing the destination HPE MSA Gen6 storage system.....	15
Planning for cutover.....	15
Cutover examples.....	15
Disconnecting decommissioned storage systems.....	21
Conducting dry runs.....	21
Choosing candidates.....	22
Preparing for rollback.....	22
Stop triggers.....	22
Data migration examples.....	22
VMware Storage vMotion.....	23
Microsoft Hyper-V live migration.....	26
HPE MSA Remote Snap Replication.....	32
Summary.....	37



Executive summary

This paper investigates strategies and methodologies for migrating data from existing block-based storage systems to sixth-generation HPE MSA 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 fourth-generation and fifth-generation HPE MSA storage as the source system.

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 sixth-generation HPE MSA 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
- By using HPE Remote Snap replication (RSR) from fourth-generation or fifth-generation HPE MSA storage arrays configured with virtual storage

Important

Advice for migrating boot-from-SAN volumes is provided only in relation to RSR. 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 legacy fourth-generation or fifth-generation HPE MSA storage or other array families, whether within the HPE portfolio or within 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 to 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 sixth-generation 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 Gen6 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™ 6.7, Windows Server 2019, and HPE MSA firmware versions that were current as of October 2021.



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 will 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 missing 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 network 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.

Although offline migrations have their advantages, this paper describes methods that either keep applications fully online or involve very minimal downtime. 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.

Important

Because of new drive enclosures and drive carriers, data-in-place upgrades to Gen6 HPE MSA storage arrays are not possible.



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 interest 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, Hewlett Packard Enterprise 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 a backup-and-restore approach is used, the destination and source systems must be hosted in tandem for the duration of the migration. Therefore, 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 data paths, 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, array-based migrations to sixth-generation HPE MSA storage require the data source to be either fourth-generation or fifth-generation HPE MSA storage arrays. Where relevant to the source system, Hewlett Packard Enterprise 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 (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 third-generation HPE P2000 G3 arrays up to current Gen6 systems. [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, mouse-based approach to displaying pertinent information, as shown in Figure 1. Although some information is displayed in a manner that does not allow the user to select text for copy and pasting to an external application, the SMU is still the most accessible method.



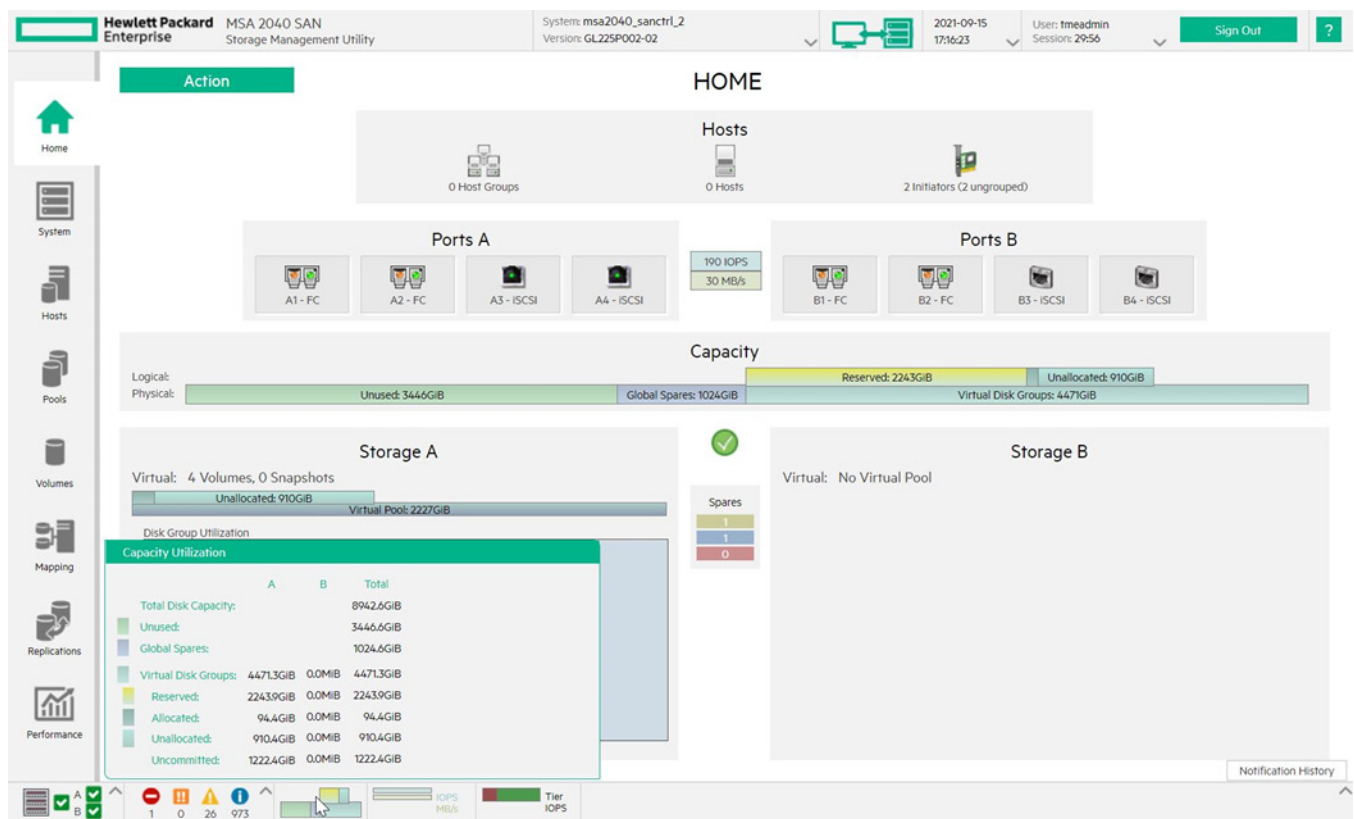


Figure 1. HPE MSA 2040 V3 SMU displaying information about capacity utilization

- **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 CLI Reference Guide for the relevant [fourth-generation](#)- or [fifth-generation](#) HPE MSA storage. Useful commands include the following:
 - show configuration
 - show pool-statistics
 - show disk-group-statistics
 - show maps

Note

For sixth-generation HPE MSA, 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

Sixth-generation HPE MSA storage arrays are tested and qualified for connection to HPE servers running specific support OSs through HPE Fibre Channel HBAs 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 [SPOCK](#) (Single Point of Connectivity Knowledge) 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 executed 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 will 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 executed within a given timeframe. Rushing a migration or cutting corners is likely to yield unexpected results.



Important

No two migrations will 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 Gen6 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 will 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 a sixth-generation HPE MSA 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 also 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 three potential approaches to copying data, each offering its own benefits and challenges:

- Synchronous replication
- Asynchronous replication
- Host-based replication

Synchronous replication helps ensure that partner volumes located on all participating systems are always 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 eliminated from consideration.

Asynchronous replication is a method of copying data from one location to another, either according to a schedule or when manually invoked. **Remote Snap Replication (RSR)** 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, RSR replicates snapshots from one



array to another. Although it was designed primarily as a disaster recovery solution, RSR is also well suited to the task of migrating data between compatible HPE MSA storage arrays.

Note

The RSR feature is not supported on SAS models of HPE MSA storage arrays of any generation. In addition, RSR requires both arrays to be connected to switches, and the arrays cannot be directly connected to each other.

Important

RSR requires both source and destination HPE MSA storage systems to be appropriately licensed.

The HPE MSA RSR 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
- Eliminates OS compatibility issues
- Works without host-based tools or features

Table 1 lists the multiple versions of RSR and the different features and compatibility it offers depending on which version each array is running.

Table 1. HPE MSA RSR compatibility and limitations across generations of HPE MSA storage arrays

Detail	Gen4	Gen4	Gen4	Gen5	Gen6
Data structure	Linear	Virtual	Virtual	Virtual	Virtual
Compatible with HPE MSA Gen6	No	N/A	Yes	Yes	Yes
RSR version	N/A	N/A	1	2	2
As of firmware version	GL100	GL100	GL220R005	VL100 / VE100	I100
Protocol support	<ul style="list-style-type: none"> • Fibre Channel • iSCSI 	N/A	iSCSI	<ul style="list-style-type: none"> • Fibre Channel • iSCSI 	<ul style="list-style-type: none"> • Fibre Channel • iSCSI
Failback support	Yes	N/A	No	No	Yes
Minimum scheduled replication interval	30 minutes	N/A	60 minutes	30 minutes	30 minutes
Replication queuing	No	N/A	No	Yes	Yes
Number of queued replications	N/A	N/A	N/A	1	1
Number of peer connections	N/A	N/A	1	<ul style="list-style-type: none"> • HPE MSA 1050: 1 • HPE MSA 2050/2: 4 	<ul style="list-style-type: none"> • HPE MSA 1060: 1 • HPE MSA 2060/2: 4
Peer connection authentication support	N/A	N/A	No	Yes	Yes
Ability to change peer connection protocol	N/A	N/A	No	Yes	No
Maximum replicated volumes per system	16	N/A	32	32	32
Maximum volumes per volume group	N/A	N/A	16	16	16
Maximum replication sets per volume	N/A	N/A	1	1	1
Notes	Virtual storage optional and not available until firmware GL200.	RSR not supported on virtual storage until firmware GL220R005.	Peer connections to Gen5 and Gen6 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 server. Host-based replication offers several advantages over array-based replication:

- It supports replication from arrays that do not support RSR or that are incompatible with RSR version 2.0 (for example, migrating from linear storage of HPE MSA arrays or third-party systems).
- It provides finer controller 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 (although 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 careful management.
- Human error is far more likely if appropriate tools are not used.

Important

It is outside 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](#).

Table 2. Capabilities by migration type

Detail	HPE MSA RSR	VMware vSphere Storage vMotion	Microsoft Hyper-V live migration
Array-based	Yes	No	No
Requires OS certification with both arrays	No	Yes	Yes
Can easily copy boot volumes	Yes	No	No
File system dependent	No	Yes	Yes
Can migrate VMs online	No	Yes	Yes
Can move data between any storage arrays supported by the OS	No	Yes	Yes
Risk of impact to application performance	Low	Medium	Medium

Identifying data copy windows

A copy window is an optimal period during which 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, bottlenecks 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 timeframe 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.



For example, RSR version 2 supports replication queuing, whereas previous versions do not. Therefore, if replications do not complete in time, it is possible that subsequent replications will fail.

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	1,649 GB	(1.5 TiB)
Transfer rate	0.125 GB/s ¹	(1 Gb/s)
Example	1,649 / 0.125 = 13,192 seconds / 3,600 = 3.67 hours = 3 hours 40 minutes / 2 = 2 × copy windows	

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) * 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	7,200 seconds	(2 hours)
Transfer rate	0.125 GB/s ¹	(1 Gb/s)
Example	7,200 × 0.125 = 900 GB	

Defining a running order of tasks

It is critical to execute 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 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, Hewlett Packard Enterprise 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 revolve around 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.

¹ Calculation in Base-10.



Setting up 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.

Hewlett Packard Enterprise 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 maximize participation, give the group a defined testing start time and duration in advance. A simple approach is to define specific workflows and then test 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, 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.

You should 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.

Sizing the destination HPE MSA Gen6 storage system

When you migrate 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 will need to offer greater performance and capacity than any individual source system. HPE makes it easy to choose an optimal HPE MSA Gen6 solution by using the [HPE Ninja Online for MSA](#) tool, which is a cloud-hosted application available to anyone with an HPE Passport. Hewlett Packard Enterprise also strongly recommends that you become familiar with the [HPE MSA 1060/2060/2062 Storage Arrays Best Practices](#) white paper, which offers guidance on how to achieve optimal configurations.

Caution

Array performance, availability, and the resulting suitability of the solution are closely related to appropriate array sizing. Therefore, Hewlett Packard Enterprise strongly recommends that you consult the [HPE Ninja Online for MSA](#) tool and the [HPE MSA 1060/2060/2062 Storage Arrays Best Practices](#) white paper before committing to the purchase and configuration of a new sixth-generation HPE MSA storage system.

Planning for cutover

Cutover is the process of redirecting applications from one system to another, specifically from a source storage system to a sixth-generation HPE MSA storage system. Planning this phase correctly is critical because there is the potential for applications to access different copies of the same data. During cutover, the principal goal is to prevent clients from connecting to the source system while enabling them to connect successfully and reliably to the destination.

Cutover examples

The cutover process can vary depending on whether a migration was carried out in a single step or in discrete stages. It also depends on which strategies were employed, 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 in the user interface.

Note

It is advisable to take an array-based snapshot of volumes located on the source storage system before attempting a migration or cutover.



vSphere Storage vMotion

Unlike mechanisms such as RSR, which copies data from one location to another, vSphere Storage vMotion moves VM data between datastores, which eliminates the possibility that older versions of data might exist and dramatically simplifies 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, if multiple VMs are configured to use a common .ISO image, 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 at the datastore level. Therefore, a task is created for each VM whose files are either partially or fully stored on 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.

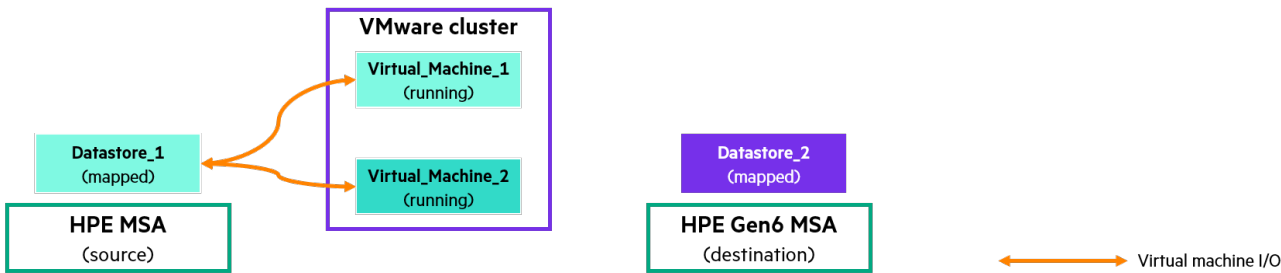


Figure 2. All VMs are running and located on the existing storage system

- 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 the destination datastores.

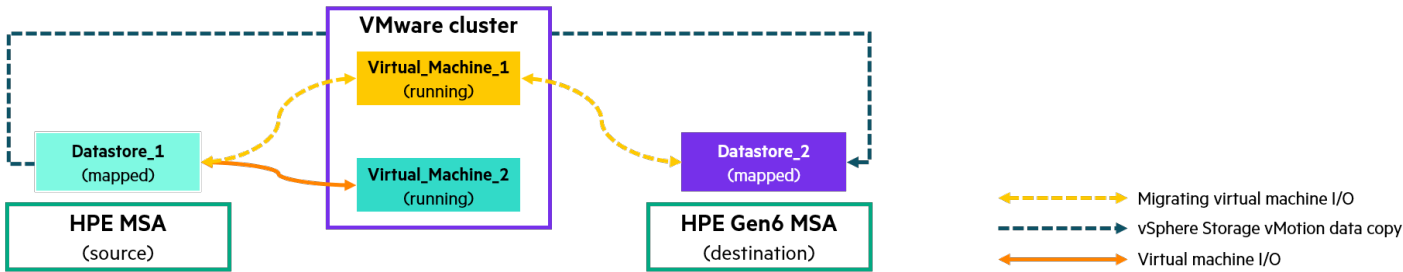


Figure 3. A vSphere Storage vMotion task is initiated for a single VM

- 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.

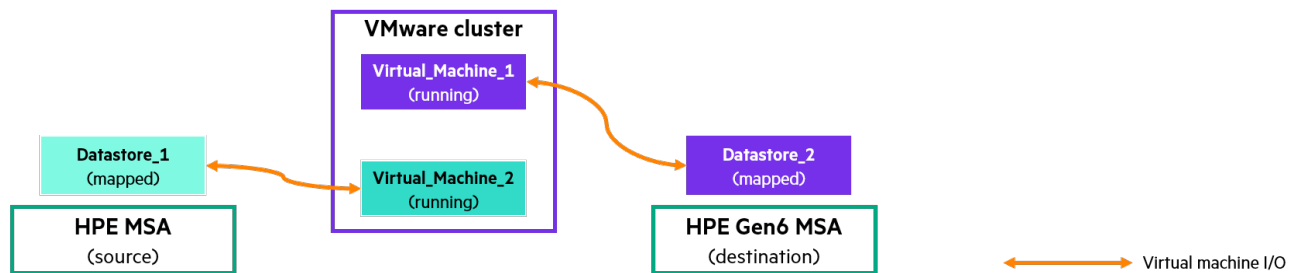


Figure 4. The migrated VM is now located on the new HPE MSA Gen6 storage system

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 Storage I/O Control are enabled for the datastore, disable them.
3. Make sure the datastore is not used for vSphere HA heartbeating.
4. In the 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 system²

1. Navigate to the Volume menu and select the volume hosting the datastore that is to be removed from the VMware host or hosts.
2. Click the **Maps** tab beneath the volume list.
3. Highlight all mappings by selecting the first entry and then, while holding the shift key, selecting the last one.
4. Right-click anywhere in the highlighted list or click **Actions** and select **Remove Mappings**.
5. 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 Gen6 array.

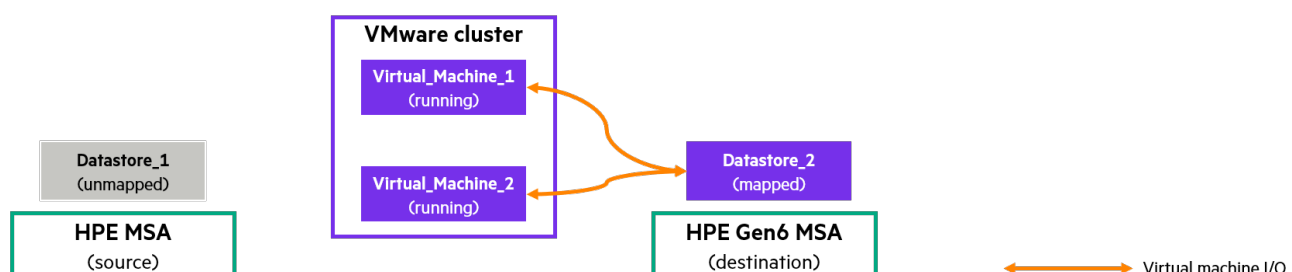


Figure 5. After all VMs and files are migrated, the original datastore can be unmounted, removed, and unmapped

² Example given for the V3 HPE MSA SMU. Consult the relevant documentation if the datastore is located on a different storage system.



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 eliminates 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, if multiple VMs are configured to use a common .ISO image, 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 that are located on a legacy HPE MSA array.

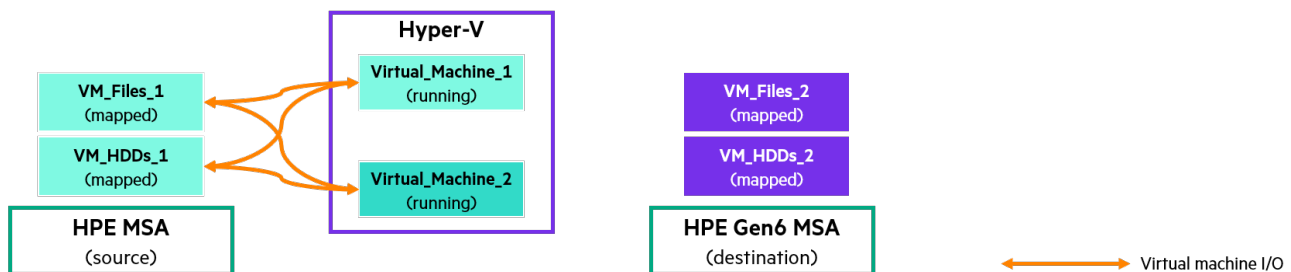


Figure 6. All VMs are running and are located on the existing storage system

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.

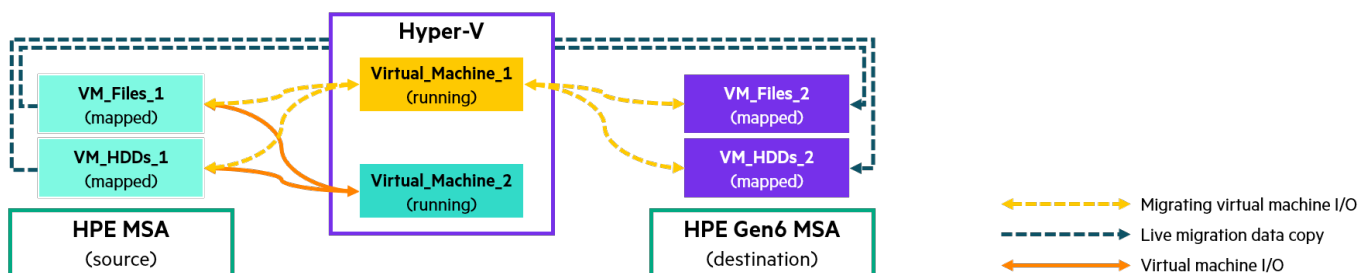


Figure 7. A live migration task is initiated for a single VM



- 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.

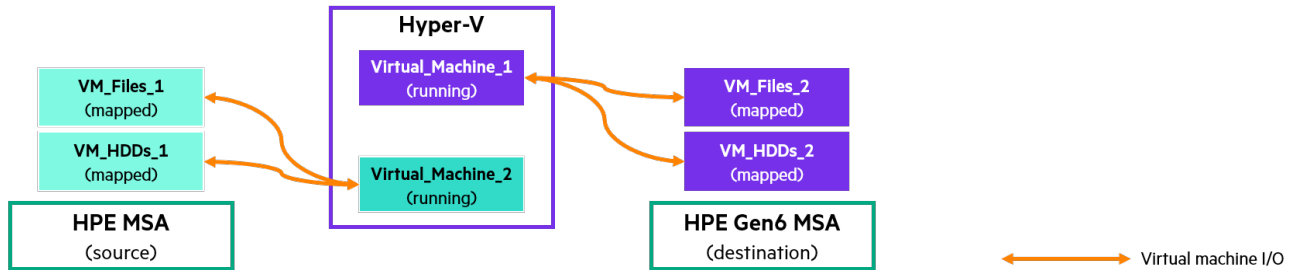


Figure 8. Replicate a source volume to the destination HPE MSA Gen6 storage system

After all 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

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

In the HPE MSA storage system³

- Navigate to the Volume menu and select the volume hosting the file system to be removed from the Windows host or hosts.
- Click the **Maps** tab beneath the volume list.
- Highlight all mappings by selecting the first entry, and while holding the shift key, selecting the last.
- Right-click anywhere in the highlighted list or click **Actions** and select **Remove Mappings**.
- 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 Gen6 array.

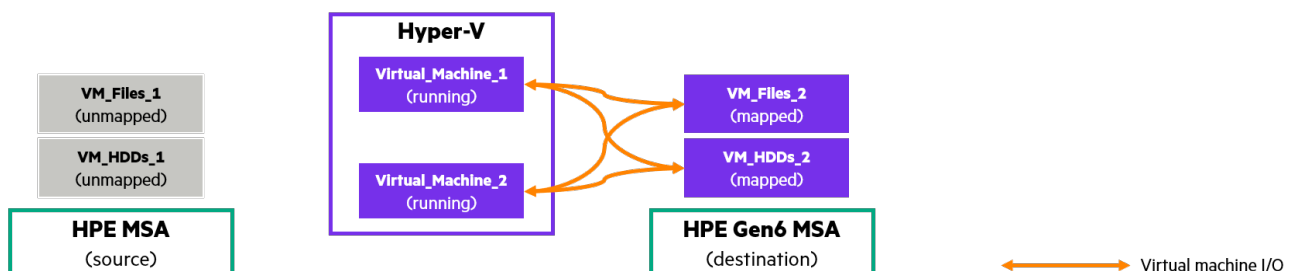


Figure 9. After all VMs and files are migrated, the original file system can be unmounted, removed, and unmapped

³ Example given for the V3 HPE MSA SMU. Consult the relevant documentation if the datastore is located on a different storage system.



HPE MSA Remote Snap Replication

Unlike vSphere Storage vMotion or Hyper-V live migration, as a snapshot-based technology, HPE MSA RSR 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, RSR 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, RSR is not the preferred migration technology for virtualized environments, especially when both convenience and uptime are accounted for. RSR is therefore best suited to bare-metal server installations, where it is more likely that applications are directly related to 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.

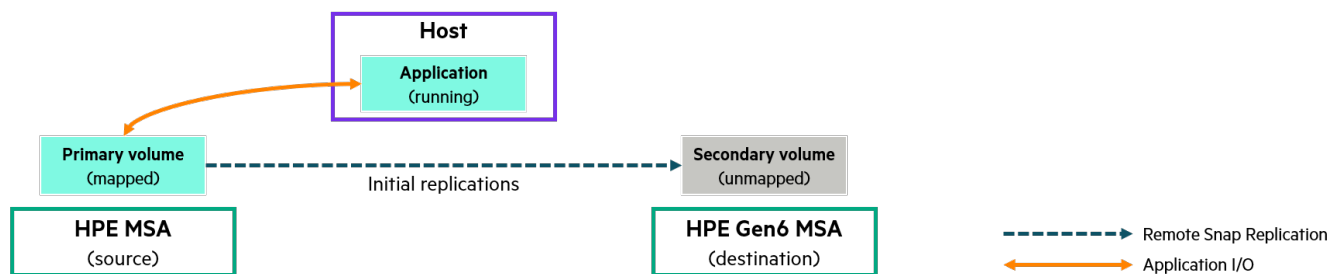


Figure 10. Applications continue to use their current volumes while snapshots are replicated to the destination array

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.

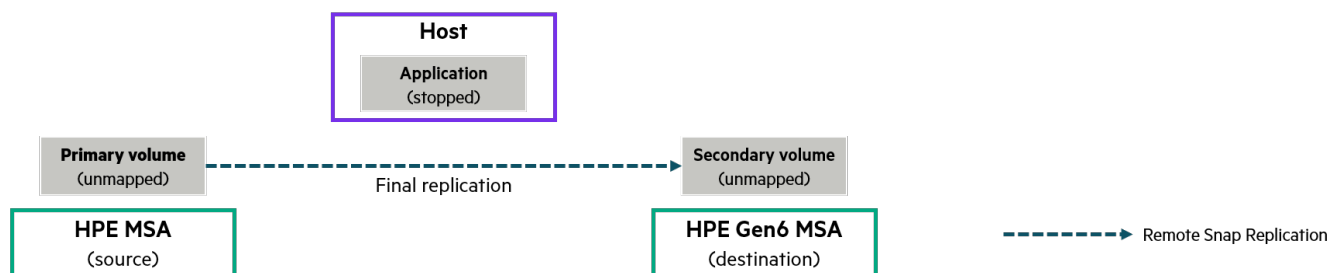


Figure 11. Applications are stopped and the volumes unmapped while a final replication is carried out to the destination array



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.

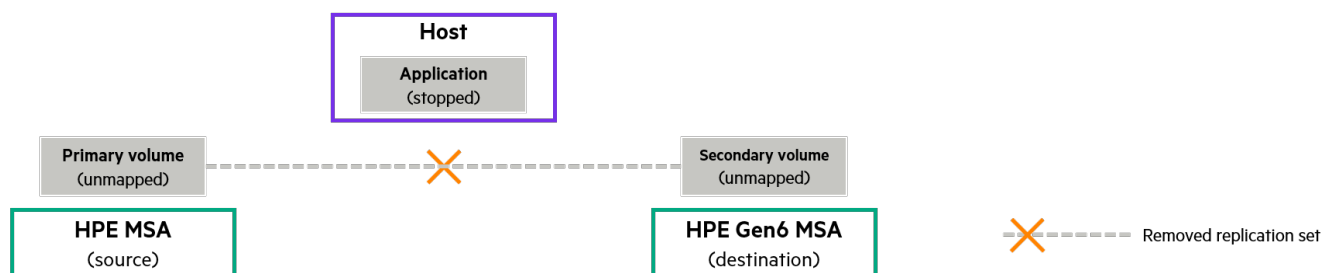


Figure 12. The replication set is removed, which in turn makes the destination volume a primary volume that can be mapped to a host

Note

After the replication set has been removed, Hewlett Packard Enterprise 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 mountpoint in the same location as the source, no application reconfiguration should be necessary.

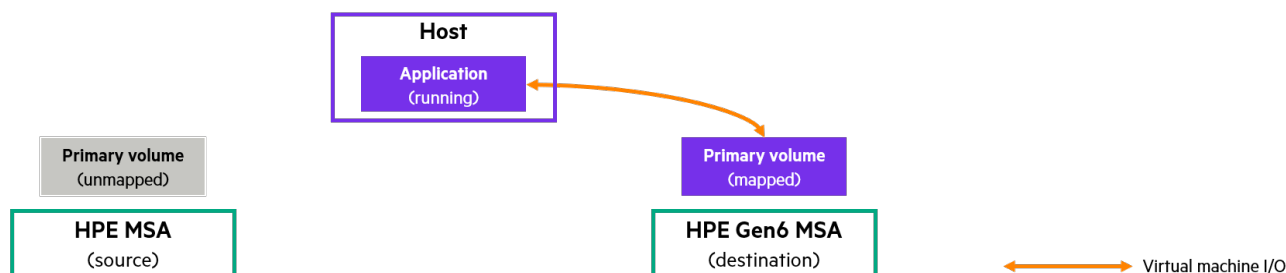


Figure 13. The destination volume is mapped to the host, and applications resume

Disconnecting decommissioned storage systems

Depending on which SAN protocol is in use, Hewlett Packard Enterprise 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 eliminates 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).
- **For iSCSI-attached SAN arrays**, remove persistent connections to the legacy targets (for example, Static Discovery entries for VMware and Favorite Targets for Windows).
- **For SAS-attached systems**, remove zones if using 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 eliminate 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.

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 will 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 applications workloads within the VM, using any preferred synthetic workload tool.

For application data that will be replicated by using RSR, Hewlett Packard Enterprise recommends taking a snapshot of one or more volumes and forming a replication set between the source array and the destination HPE MSA array.

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.

Hewlett Packard Enterprise recommends that dry runs include a simulated rollback.

Data migration examples

This paper explores three 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 Gen6 storage system by using vSphere Storage vMotion
- Migrating from any storage system to an HPE MSA Gen6 storage system by using Microsoft Hyper-V live migration
- Migrating from an HPE MSA Gen4⁴ or later system by using RSR

Both vSphere Storage vMotion and Microsoft Hyper-V live migration simplify the task of migrating the underlying storage of a VM and makes the process faster. Both hypervisors are designed in a manner that in most cases eliminates the need to bring VMs offline and

⁴ Gen4 HPE MSA arrays must be running firmware GL200 or later, and volumes must be located within a virtual pool. Linear storage is not supported for replicating to Gen5 or Gen6 systems.



requires no lengthy post-migration steps. In contrast, HPE MSA RSR 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_Gen6**. Each volume has been mapped to all hosts in the VMware cluster and formatted as a VFMS file system.




Figure 14. Visualization of an example storage migration of a VMware VM

Note
Hewlett Packard Enterprise recommends that you consult the most recent recommendations in the [Deployment Considerations for HPE MSA Gen6 with VMware vSphere 7](#) white paper before proceeding.

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**.



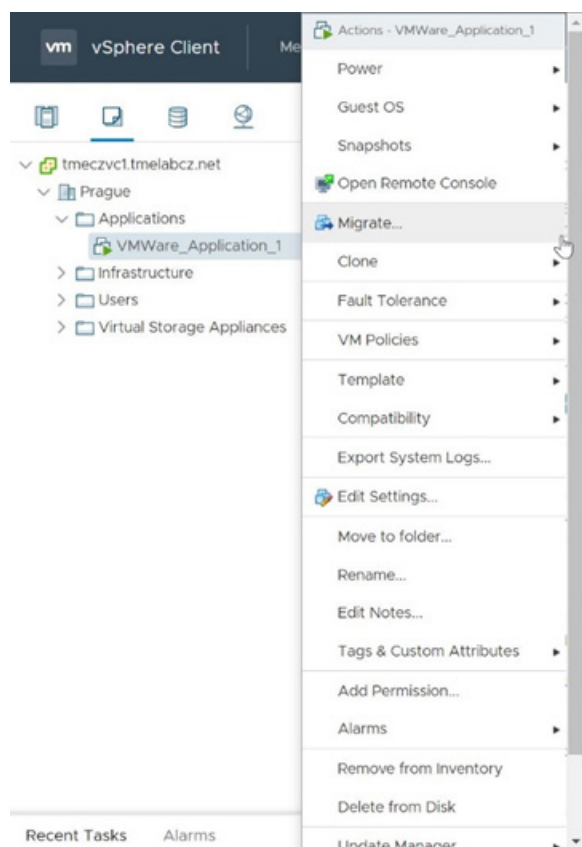


Figure 15. How to begin migrating a VM within VMware vCenter

2. Select **Change storage only**.

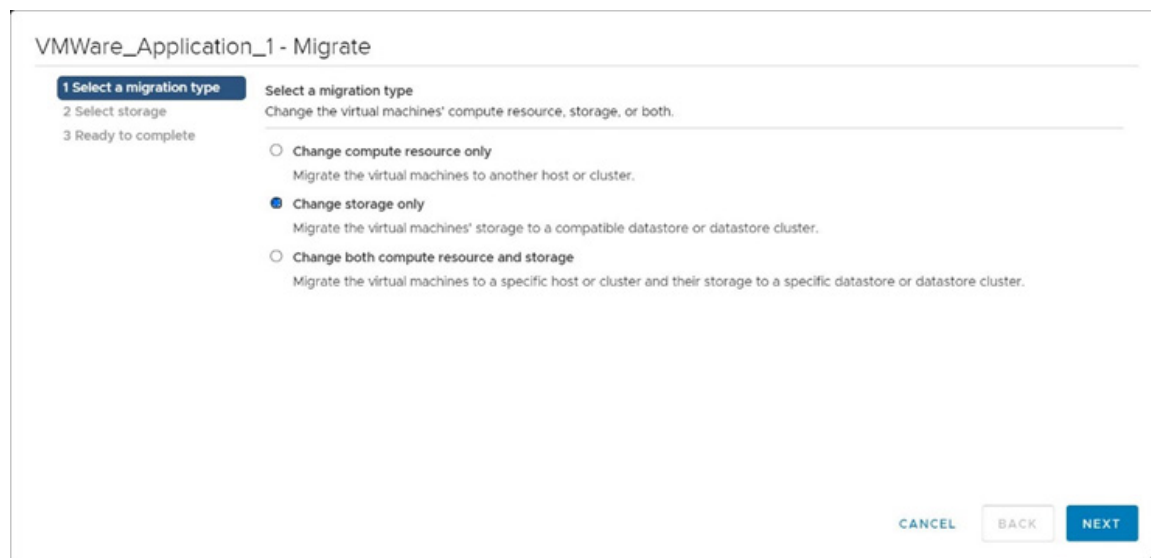


Figure 16. Selecting storage-only migration for selected VMs



3. Select a destination datastore that is located on the new HPE MSA Gen6 storage system.

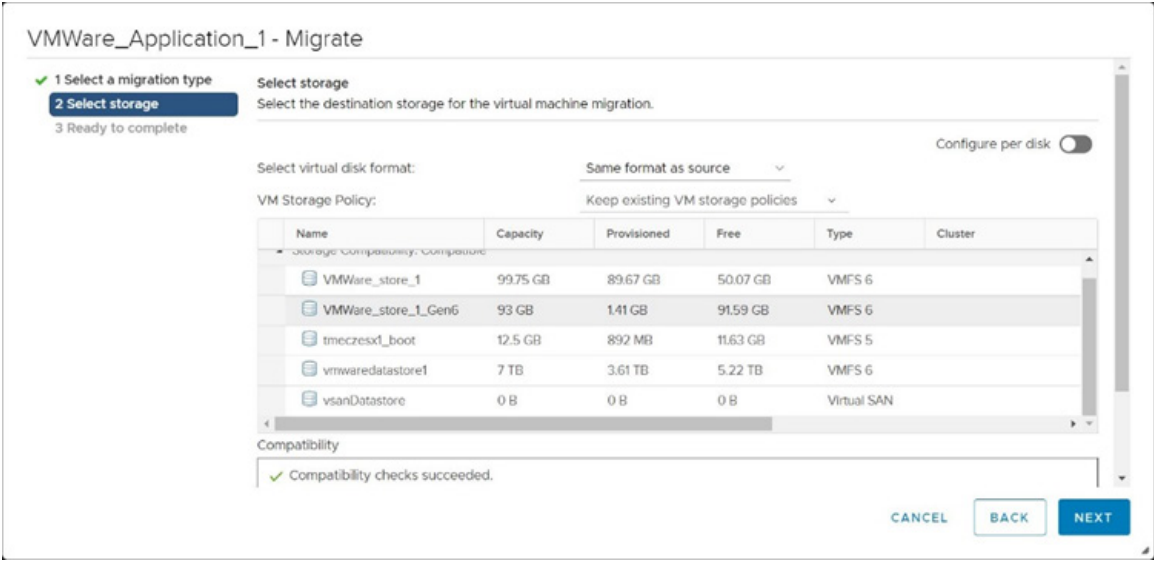


Figure 17. Select a destination datastore located on an HPE MSA Gen6 storage system

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

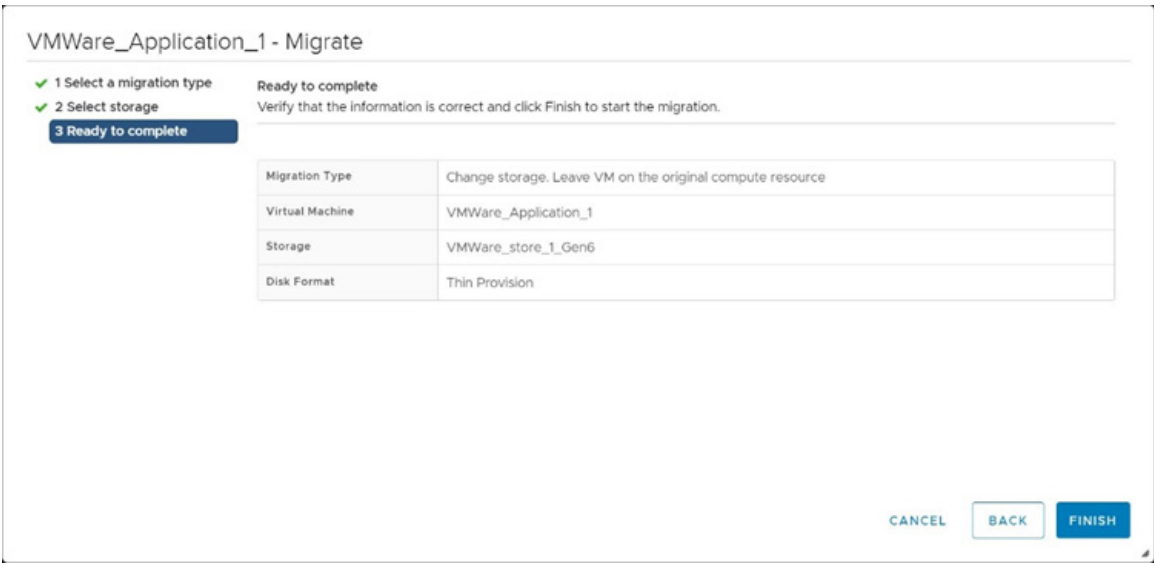


Figure 18. Select a destination datastore located on an HPE MSA Gen6 storage system



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.

Recent Tasks		Alarms						
Task Name	Target	Status	Initiator	Queued For	Start Time	Completion Time	Server	
Relocate virtual machine	VMware_Application_1	Completed	TMEIABCZ\imeadmin	3 ms	10/19/2021, 7:29:28 PM	10/19/2021, 7:30:53 PM	tmezcvc1tmeiabcz.net	

Figure 19. 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 standalone or in a failover cluster. However, the management console and subsequent steps differ slightly. The migration examples listed in Table 5 contain four volumes each.

Table 5. Example mountpoints and drive letters

Detail	Clustered	Standalone
Source VM files	C:\ClusterStorage\HyperV_VMFiles	E:\
Source VM virtual disks	C:\ClusterStorage\HyperV_HDDs	F:\
Destination VM files	C:\ClusterStorage\HyperV_VMFiles_Gen6	G:\
Destination VM virtual disks	C:\ClusterStorage\HyperV_HDDs_Gen6	H:\

Note

Hewlett Packard Enterprise recommends that you consult the most recent recommendations in the [HPE MSA Gen6 Storage with Microsoft Hyper-V Implementation Guide](#) before proceeding.

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 20, four volumes are mapped to hosts participating in the cluster and mounted in the folder C:\ClusterStorage in CSV format.




Figure 20. 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.



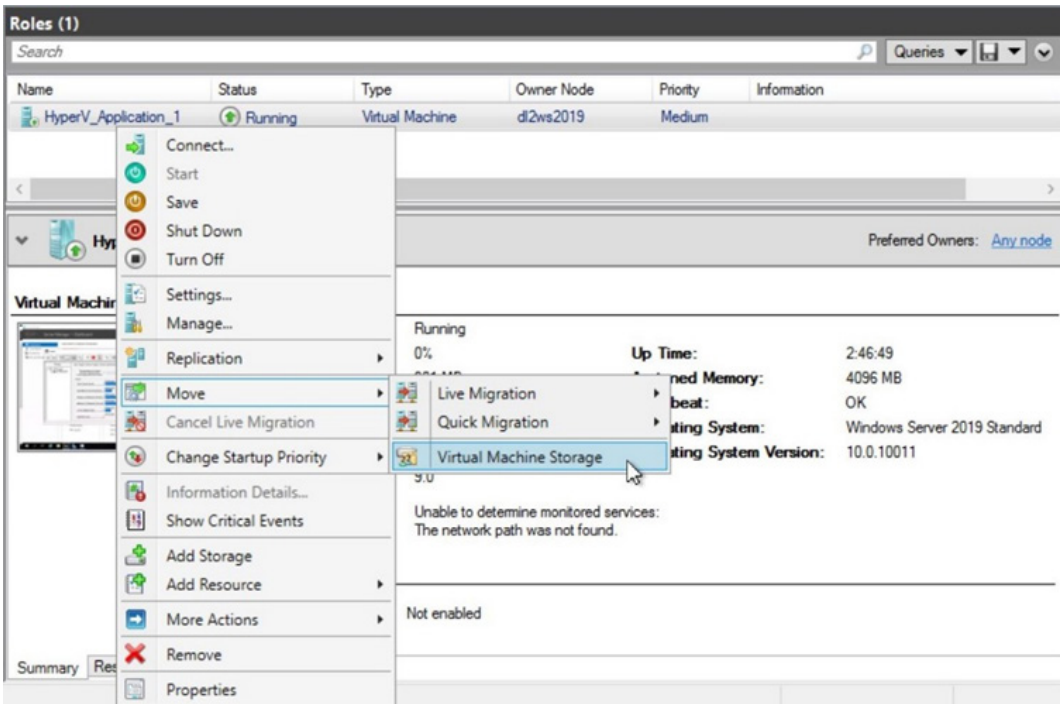


Figure 21. Selecting storage migration for a highlighted VM within the Microsoft Failover Cluster Manager console

2. Drag each item down to the relevant target volume in the Cluster Storage tree⁵ and click **Start**.

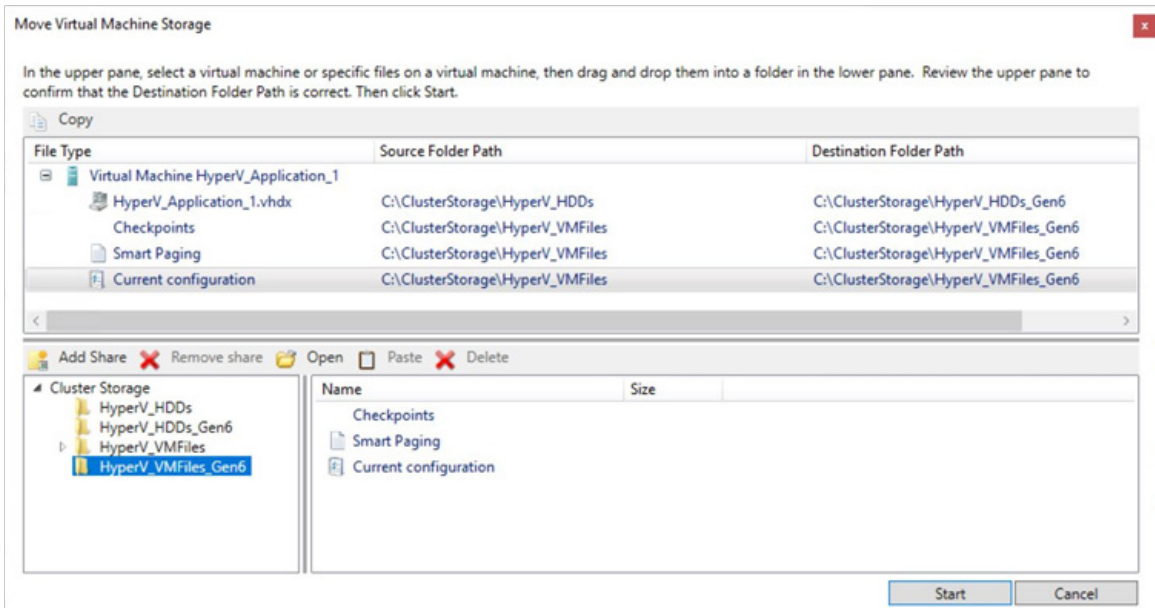


Figure 22. Defining a new destination folder path for the VM's related files

⁵ 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.

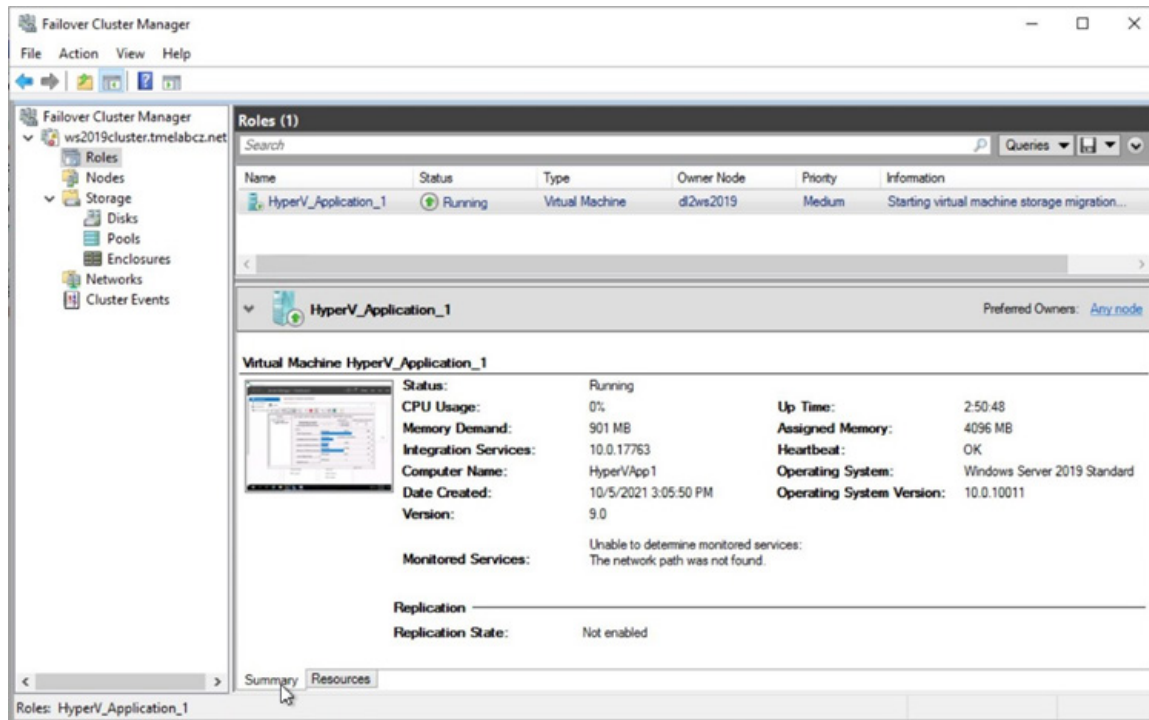


Figure 23. Select a destination datastore located on an HPE MSA Gen6 storage system

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.

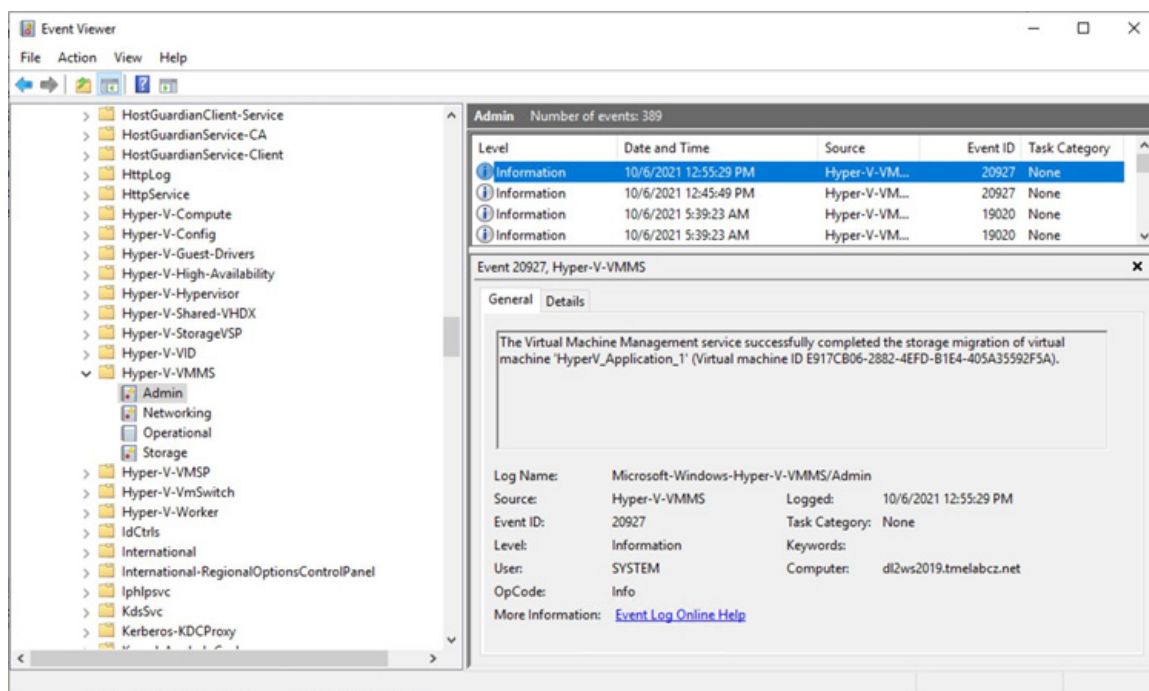


Figure 24. Event viewer entry confirming a successful storage migration



For VMs on a standalone host

The process of migrating storage for VMs on a standalone 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 mountpoint in an existing file system, in a way similar to the process for CSVs.)




Figure 25. Visualization of the example storage migration of a Hyper-V VM located on a standalone host

To migrate storage for VMs on a standalone 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**.

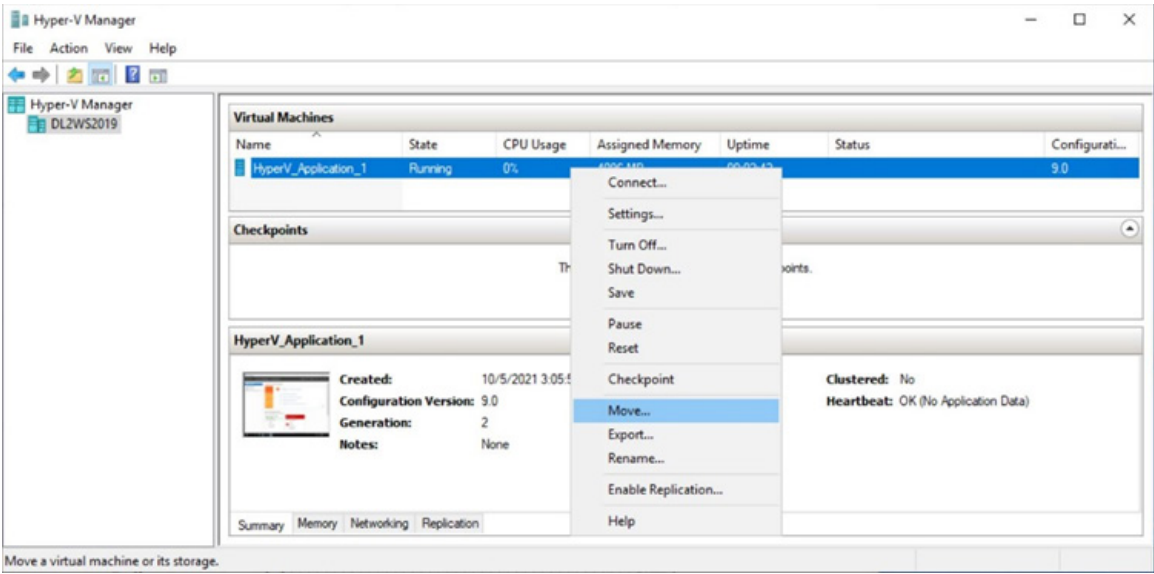


Figure 26. Starting a migration for a highlighted VM in the Microsoft Hyper-V Manager console



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

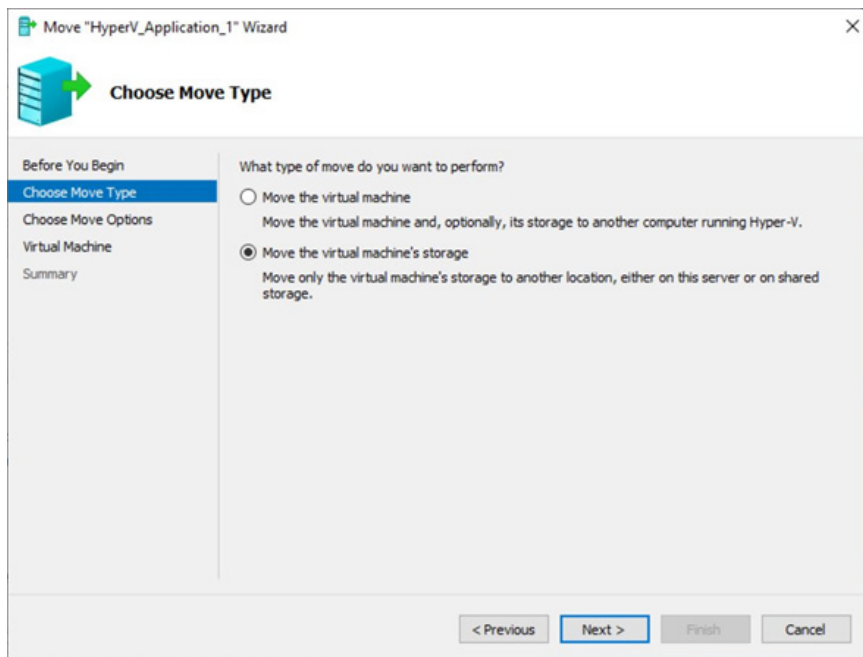


Figure 27. Selecting to move a VM's storage from the Hyper-V migration wizard

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

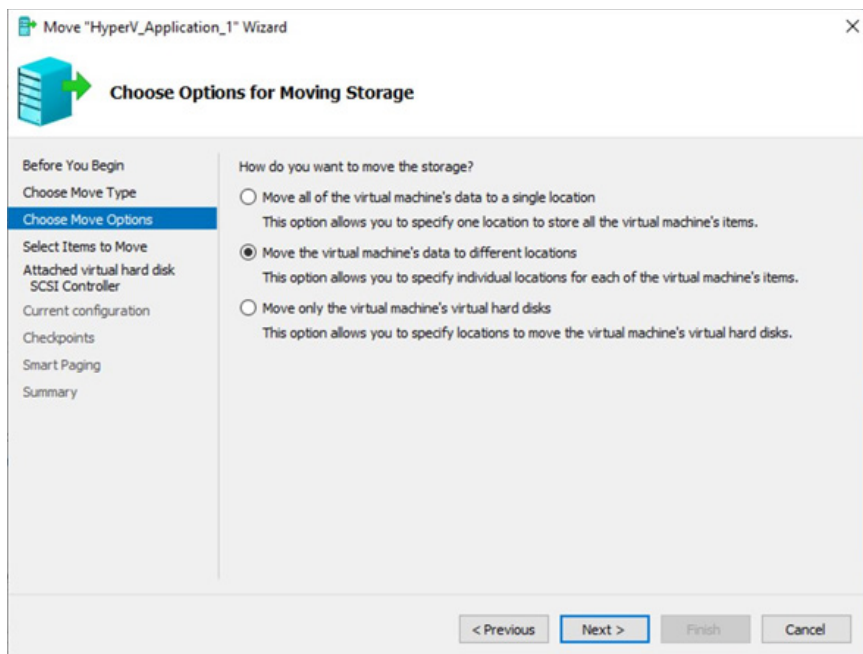


Figure 28. Selecting the option to move a VM's files by category from within the Hyper-V migration wizard



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

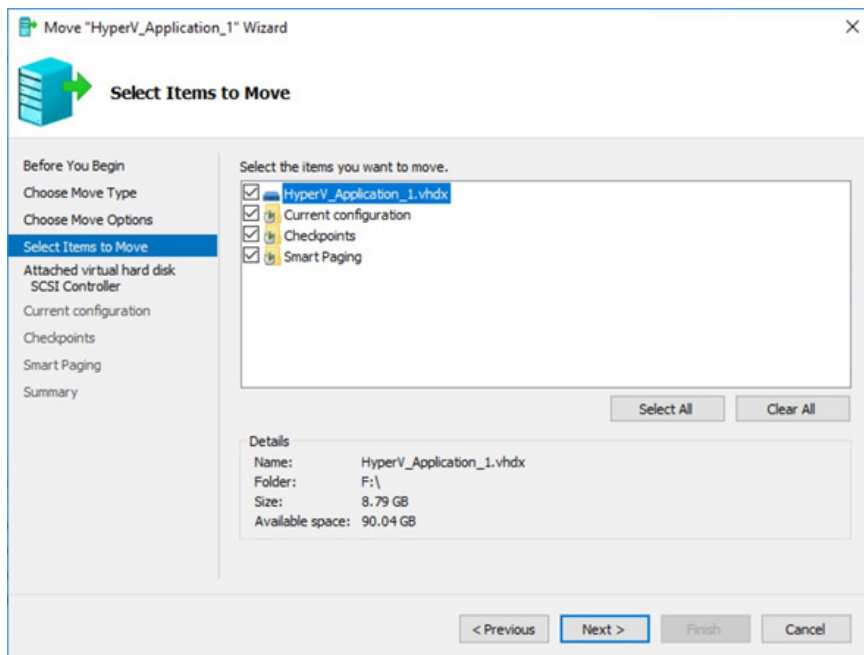


Figure 29. Selecting VM files to be migrated from within the Hyper-V migration wizard

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

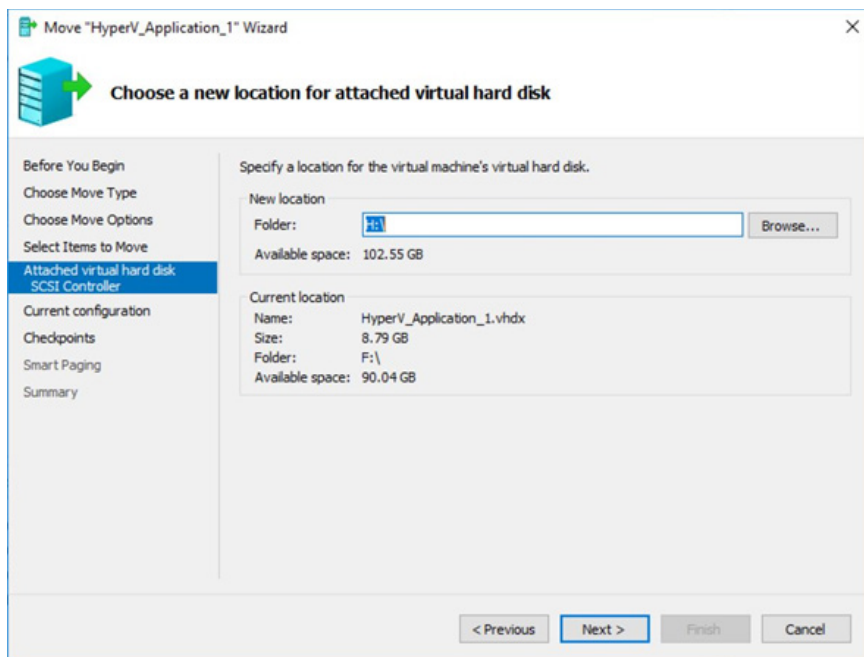


Figure 30. Specifying new location for VM data type



- 6. Click **Next** and repeat Step 5 for all data types.
- 7. Review the changes and click **Finish** to start the migration.

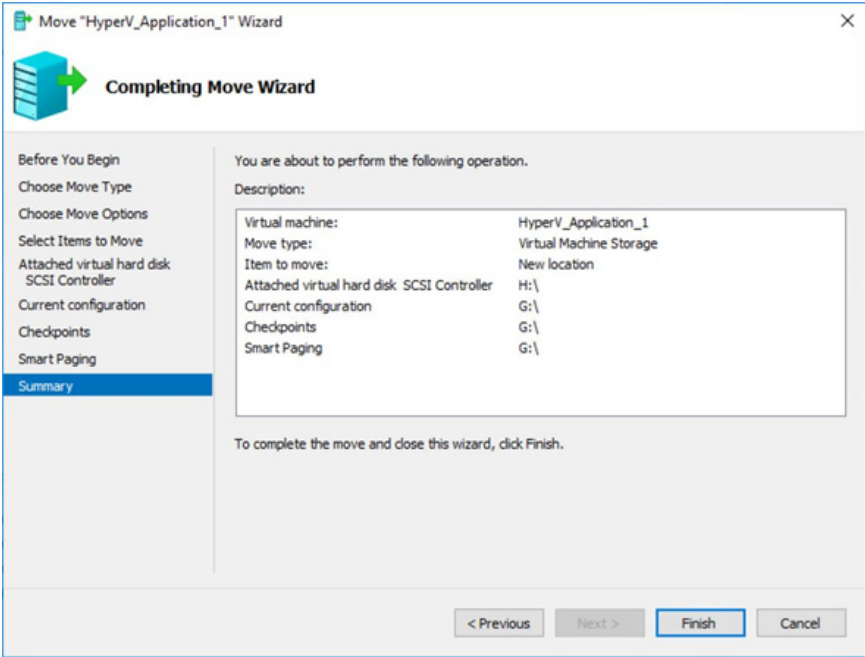


Figure 31. Confirmation screen of the Hyper-V migration wizard

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

HPE MSA Remote Snap Replication

The following procedure demonstrates how to use RSR to replicate volumes located on fourth-generation and fifth-generation HPE MSA arrays where volumes are located in virtual pools. To offer a potential alternative to using hypervisor-based migration technologies, the volumes chosen are those that hold Hyper-V file systems. However, the included steps are limited to the actions that must be carried out on the HPE MSA arrays, and not on the host OSs. RSR is best suited to the migration of bare-metal servers whose application data or boot volumes are located on HPE MSA arrays. Describing how to configure all supported OSs and HBAs would require more steps than can be included in this paper. Always consult the relevant documentation for all aspects of the specific environment before beginning a migration, and document all the necessary steps.

To replication volumes by using RSR, complete the following steps:

- 1. Create a peer connection.

Because fourth-generation HPE MSA arrays do not have an authentication mechanism, a fourth-generation array cannot create the peer connection (which is the fundamental relationship between the source and target arrays) for replication to either fifth-generation or sixth-generation HPE MSA arrays. In addition, for replication to a sixth-generation HPE MSA array, the peer connection can be configured only from the CLI. For guidance on using the sixth-generation HPE MSA array CLI, see the [HPE MSA1060/2060/2062 CLI Reference Guide](#).

For simplicity, this paper offers guidance only on how to use the CLI to create the peer connection. When creating a peer connection to a fourth-generation array, use **none** as both the user name and the password. When forming a peer connection to a fifth-generation system or later, use credentials that have either the **standard** or **manage** role.

Important

Refer to Table 1 to ensure that the migration plan does not exceed the maximum number of peer-connections supported by either HPE MSA array.



To create the peer connection, use the following command and insert the relevant system-specific information where **hostportaddress** is a host port WWPN or iSCSI IP of the source system and **peerconnectionname** is a nickname for the peer connection:

Command syntax: `create peer-connection remote-port-address hostportaddress remote-username username remote-password password peerconnectionname`

Example: `create peer-connection remote-port-address 207000c0ff1beb1f remote-username none remote-password none 2040_1060`

Only one host port address is necessary, and the arrays negotiate which other ports are available for use.⁶

2. Create a replication set for each volume.

A replication set is specific to an individual volume and contains the replication conditions for that volume. These conditions include the peer connection that should be used for replication, the destination pool, the volume name as it will appear on the remote system, and any schedules. Replication sets must be created by the source system.

Important

Refer to Table 1 to ensure that the migration plan does not exceed the maximum number of replication sets supported by either HPE MSA array.

To create a replication set, complete the following steps:

- From the source array, select the volume that is to be replicated from the Volumes menu on the left of the SMU's main screen, click **Action**, and select **Create Replication Set**.

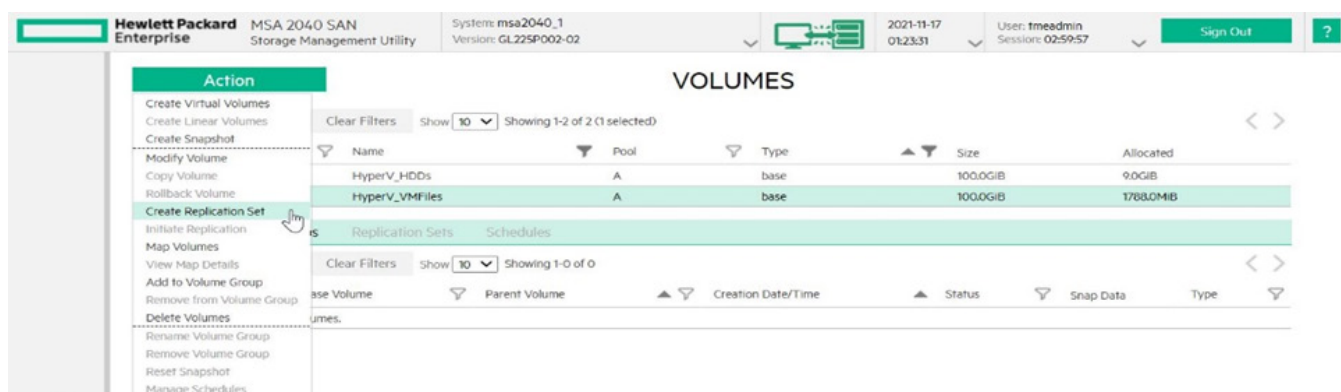


Figure 32. Create a replication set from within the V3 SMU

- Specify a replication set name as it will appear on both systems.

⁶ All connected ports between the two arrays will be used for replication. It is not possible to separate host and replication traffic.

Create Replication Set

Volume:

HyperV_HDDs

Replication Set Name:

HyperV_HDDs_RSR

Select Peer Connection:

2040_1060

✓

Secondary Pool:

☒ A ☐ B

Secondary Volume Name:

HyperV_HDDs

☐ Scheduled (configure recurring replications)

OK

Cancel

Figure 33. Entering replication set details within the V3 SMU

- c. Select the peer connection that was created in Step 1.
- d. Select the pool on the remote system where the replicated volume should be stored.
- e. Provide a secondary volume name for the volume as it will appear on the remote system.
For consistency of management after the migration is complete, Hewlett Packard Enterprise recommends that you use the same volume name that was used on the source array, spelled exactly as it appears there.
- f. Consider implementing schedules.
- g. If the migration plan supports it, click **Yes** to begin the initial replication, or click **No** to initiate it manually later.

Success

✓

The replication set was created. The first replication may take awhile. Do you want to begin a replication now?

Yes

No

Figure 34. Click Yes to begin the initial replication

Note

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 has been calculated to take many days or weeks, then it might simplify the management overhead to have the array automatically replicate volumes to a specific schedule. However, in the case of replicating from fourth-generation HPE MSA arrays, it is important to factor in the lack of replication queuing, which might lead to failures if too much demand is placed on the weakest links.

- 3. Repeat steps 2.a through 2.g for all volumes that are to be replicated.
- 4. If not started as part of step 2.g, begin replications for volumes that are to be migrated during a defined migration window. To initiate a replication, navigate to the Replication menu from the SMU’s main screen, select a volume, click **Action**, and choose **Initiate Replication**.



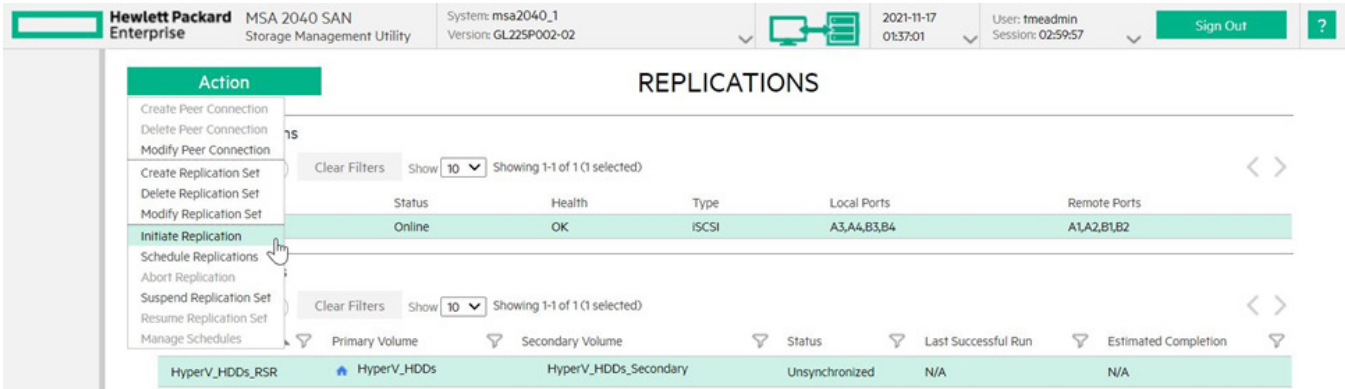


Figure 35. Manually initiating a replication for a volume from within the V3 SMU

5. Monitor the event log on the source array for event ID A98784 with the event message “A replication completed successfully” followed by the replication set name and serial number.

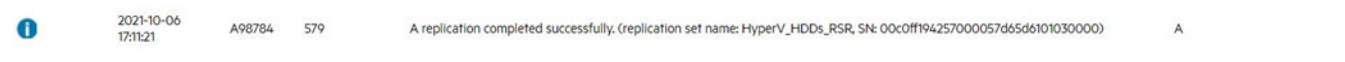


Figure 36. Event log in the V3 SMU indicating successful completion of a replication

6. After all volumes have been replicated, cease I/O to the volume from all hosts mapped to that volume and unmap it.
- There are multiple ways to unmap a volume, but one method is to select the volume from the Mappings menu of the SMU, click **Action**, and select **Map**. The existing mappings are displayed in the lower half of the dialog box. For each map in place, select **Delete** from the Action drop-down menu and click **OK**.

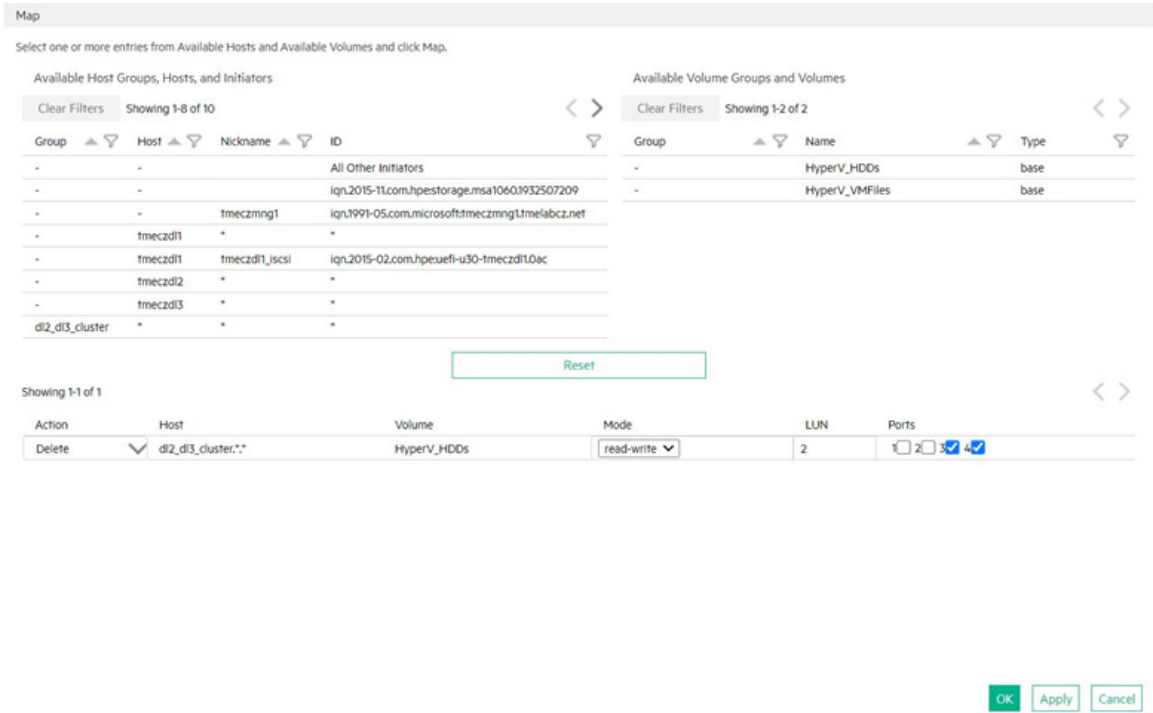


Figure 37. Unmapping a volume from within the V3 SMU

- 7. Repeat steps 4 and 5 to replicate the volume for the final time, which will include all changes since it was unmapped.
- 8. Move to the HPE MSA Gen6 array's SMU, select the secondary volume, and click the chevron on the right.

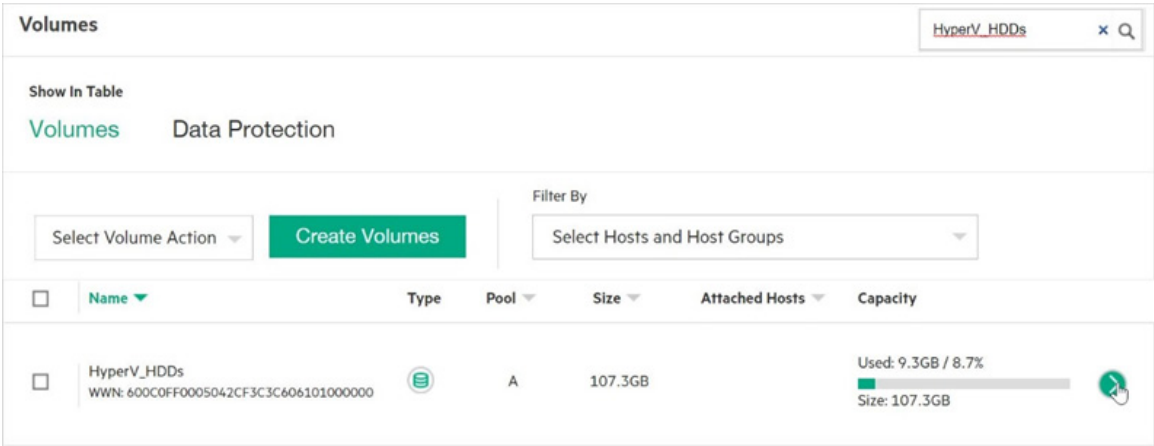


Figure 38. Selecting the secondary volume from within the HPE MSA Gen6 V4 SMU

- 9. Move to the Replications tab and click **Remove Replication**.
This action breaks that volume's relationship with the source array and allows it to be mapped to a host.

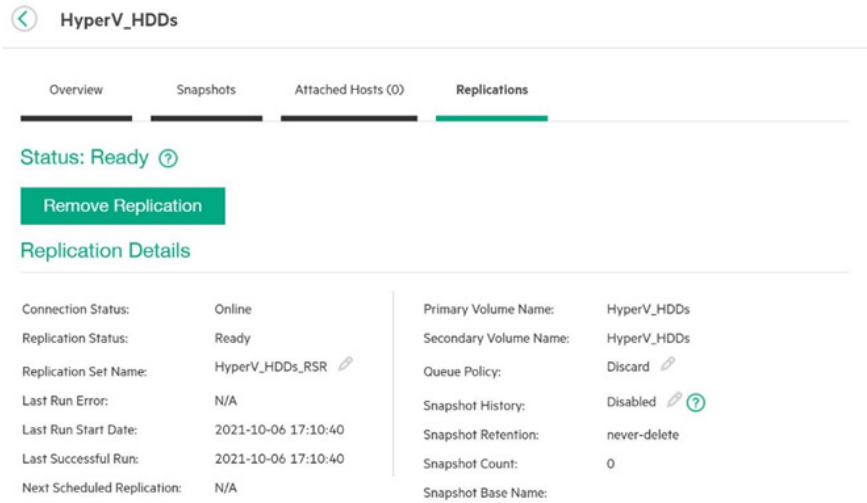


Figure 39. Replication information for a selected volume within the V4 SMU

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.

- 10. Follow the steps in the [HPE MSA 1060/2060/2062 Storage Management Guide](#) to configure hosts with the HPE MSA Gen6 array.
- 11. In 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 1060/2060/2062 Storage Management Guide](#) to complete the task.





Figure 40. Attached hosts panel from within the V4 SMU

12. Configure the host OS and any HBAs to connect to the HPE MSA Gen6 storage system.

13. 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. Although differences in the Gen6 array architecture exclude the data-in-place upgrades used by earlier generations, migration to the newest generation is simplified by a combination of ubiquitous host-based data movement technologies and seasoned in-array replication capabilities. With careful planning, any organization that owns aging systems — whether from HPE or another vendor — can move away from them easily and with minimal risk.

Resources

HPE MSA Gen6 Virtual Storage Technical Reference Guide

hpe.com/psnow/doc/a00103247enw?from=app§ion=search&isFutureVersion=true

HPE MSA 1060/2060/2062 Storage Best Practices

hpe.com/psnow/doc/a00105260enw?from=app§ion=search&isFutureVersion=true

Deployment considerations for HPE MSA Gen6 with VMware vSphere 7

hpe.com/psnow/doc/a00109648enw?from=app§ion=search&isFutureVersion=true

HPE MSA Gen6 storage with Microsoft Hyper-V implementation guide

hpe.com/psnow/doc/a00119216enw?from=app§ion=search&isFutureVersion=true

Migrating from any vendor's NAS to HPE StoreEasy 1000 storage in Active Directory integrated environments via SMB

hpe.com/psnow/doc/a00069469enw?from=app§ion=search&isFutureVersion=true

HPE MSA Health Check

msa.ext.hpe.com/

HPE Ninja Online for MSA

ninjaonline.ext.hpe.com/

HPE MSA 1060 Storage QuickSpecs

hpe.com/support/MSA1060QuickSpecs

HPE MSA 2060 Storage QuickSpecs

hpe.com/support/MSA2060QuickSpecs

HPE MSA 2062 Storage QuickSpecs

hpe.com/support/MSA2062QuickSpecs

HPE Support

support.hpe.com

Sign up for HPE updates


h41360.www4.hpe.com/alerts-signup.php



Learn more at

[HPE.com/us/en/storage/msa-shared-storage.html](https://hpe.com/us/en/storage/msa-shared-storage.html)

Explore **HPE GreenLake** 

 **Chat now (sales)**