



**Hewlett Packard  
Enterprise**

**Technical white paper**

Check if the document is available  
in the language of your choice.



# **DATA MIGRATION FROM HPE 3PAR TO HPE PRIMERA**

Data migration guide

---

# CONTENTS

|   |    |
|---|----|
| Executive summary.....                                  | 3  |
| Scope and target audience.....                          | 3  |
| Overview .....  | 3  |
| Array-based data migration.....                         | 3  |
| Appliance-based data migration.....                     | 4  |
| Third-party data migration.....                         | 4  |
| Host/application-based data migration.....              | 4  |
| HPE Storage Data Migration Service.....                 | 4  |
| HPE Peer Motion Overview.....                           | 5  |
| Workflow.....   | 6  |
| Migrating using HPE Peer Motion in the SSMC .....       | 7  |
| Prerequisites, planning, and design considerations..... | 7  |
| Step-by-step procedure.....                             | 8  |
| Migrating using the HPE Peer Motion Utility.....        | 15 |
| Prerequisites, planning, and design considerations..... | 15 |
| Multi-link support.....                                 | 16 |
| Performance considerations with HPE Peer Motion.....    | 16 |
| Summary.....  | 17 |



## EXECUTIVE SUMMARY

As enterprises go through the process of refreshing their existing storage technology, a solid data migration strategy is potentially the most important factor that influences the choice of the right technology. Data mobility for HPE Storage platforms is one of the key differentiators in the marketplace with complementary HPE tools such as Peer Motion, Online Import, and Peer Copy. This white paper outlines the available options for migrating data from HPE 3PAR to HPE Primera platforms and goes into depth regarding performing migrations using HPE Peer Motion in the HPE StoreServ Management Console (SSMC).

### Scope and target audience

This white paper is intended for solution architects, field consultants, and IT specialists who need to refresh their existing HPE 3PAR systems with HPE Primera systems. The objective of this document is to lay out the data migration options available to the storage administrator and provide guidance in choosing a data migration solution between storage platforms. This paper does not address the installation and setup of an HPE Primera array. It is assumed that these tasks have already been performed and the HPE Primera array is ready to be used in production. This white paper is also not intended to replace the administrator guide or user guide for either platform.

## OVERVIEW

Data migration is a major undertaking that requires careful analysis, planning, and execution. Many IT organizations lack the experience and resources to carry out a migration successfully. Without proper foresight and guidance, data migration programs can disrupt operations, diminish employee productivity, and impact revenues.

From a host impact standpoint, data migration can be classified as online, minimally disruptive, and offline:

- **Online migration:** Data from the source array to the target array is migrated without causing any disruption to the host I/O. Throughout the migration, the applications and the operating systems are not impacted.
- **Minimally disruptive migration:** With this approach, data from source array to the target array is migrated without causing any impact to the application. However, the host I/O is disrupted during the cutover process, which occurs after the host is successfully reconfigured and all I/O is then handled by the destination array, while the host multipath solution is configured on the destination storage system.
- **Offline migration:** Data is migrated from the source to the target array completely offline suspending the host I/O throughout the migration process.

**TABLE 1.** High-level comparison between available migration methods

| Approach                   | Array-based                    | Appliance-based                               | Third-party                                 | Host/app based                       |
|----------------------------|--------------------------------|---|---|--------------------------------------|
| <b>Migration type</b>      | Online                         | Minimally disruptive                          | Minimally disruptive                        | Online or minimally disruptive       |
| <b>Source array</b>        | HPE 3PAR, EMC, HDS, and IBM    | HPE Primera, HPE 3PAR, and HPE Nimble Storage | Any HPE or third-party array                | Any HPE or third-party array         |
| <b>Target array</b>        | HPE Primera, HPE 3PAR          | HPE Primera, HPE 3PAR, and HPE Nimble Storage | Any HPE or third-party array                | Any HPE or third-party array         |
| <b>Transport</b>           | Fibre Channel to Fibre Channel | Fibre Channel or iSCSI                        | Fibre Channel to Fibre Channel <sup>1</sup> | Fibre Channel or iSCSI               |
| <b>Host/app disruption</b> | No minimal impact              | Minimal disruption on host cutover            | Minimal disruption on host cutover          | No to minimal impact                 |
| <b>Performance impact</b>  | Minimal impact                 | Minimal impact <sup>2</sup>                   | Minimal impact                              | Performance impact <sup>3</sup>      |
| <b>Cost</b>                | Zero                           | Zero  | \$\$\$                                      | Appropriate licenses may be required |

### Array-based data migration

Array-based migrations are performed by leveraging built-in features to migrate from one array to another. HPE Peer Motion and HPE Online Import Utility are tools that can be used to perform array-based migrations.

<sup>1</sup> Check the appropriate third-party migration technology for supported transport protocol.

<sup>2</sup> Follow best practices to avoid performance impact during migration by isolating host connectivity to migration path.

<sup>3</sup> VMware vSphere® vMotion® handles host performance well when configured with VMware vSphere® High Availability as per best practices compared to logical volume manager (LVM) mirroring, for example.



- **HPE Peer Motion** allows you to nondisruptively migrate data from one storage system to another without impacting the application or the host operating system. HPE Peer Motion works between HPE 3PAR to HPE Primera systems and HPE 3PAR to HPE 3PAR systems. HPE Peer Motion migrations can be initiated and monitored through either the CLI-based Peer Motion Utility (PMU) or the SSMC.
- **HPE Online Import Utility** is built on HPE Peer Motion technology that allows you to nondisruptively migrate data from third-party storage systems to HPE Primera or HPE 3PAR systems.

### Appliance-based data migration

Appliance-based migrations are performed at the block level using an appliance to orchestrate and move data between source and target arrays. HPE Peer Copy technology is in this category. HPE Peer Copy allows you to migrate data between different HPE 3PAR, HPE Primera, and HPE Nimble Storage arrays.

---

#### NOTE

Refer to the [Migrating Data from HPE 3PAR to HPE Nimble Storage](#) white paper for an in-depth guide on migrating data from HPE 3PAR using HPE Recovery Manager Central (RMC) Peer Copy.

---

### Third-party data migration

Third-party data migrations referenced in this section refer to migrations performed at the block level using third-party software that is licensed separately. Examples of this type of software include Cirrus Data, Seven10 Storage Software, Carbonite, and Arsan. Migrations involving these solutions are typically done with a licensed agent or media server between the data path of the source array and the target array.

### Host/application-based data migration

Host/application-based data migrations are performed at the operating system/application level. Examples of host-based data migration include LVM mirroring and vSphere vMotion. Examples of application-based data migration include Oracle® ASM, data guard, and Microsoft SQL Server migrations to an Azure SQL database. Refer to the appropriate operating system distribution documentation for details.

**LVM or disk mirroring** can allow you to migrate data on a host by presenting the target block volumes to the host over iSCSI or Fibre Channel. Using this technique, data is mirrored to the target system while the data is written to the original storage device. After the mirroring is complete, the volume from the source system can be detached from the mirror. Follow the appropriate user documentation based on the host operating system.

**vSphere vMotion** allows you to migrate VMs to different compute and storage resources within your vSphere environment with zero downtime. Using this technique, you can migrate VM/application storage online to newly created datastores on the target HPE Primera array. Refer to the appropriate VMware® documentation to follow VMware best practices and step-by-step procedure. Application-based migrations are similar to host-based migrations and are specific to the application itself.

## HPE STORAGE DATA MIGRATION SERVICE

In addition to the HPE self-service migration tools, HPE offers the HPE Storage Data Migration Service, which provides managed end-to-end data migration services. These services leverage the HPE depth of technical capabilities combined with proven methodologies to help accelerate your transition to HPE Primera. This data migration service helps you overcome data migration challenges such as:

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

There are many methods for performing data migrations and all are not equal. It usually requires finding the balance between what is supported based on the environment, what is available and cost-effective, and what provides the fastest cutover with the least downtime.

HPE can help you make a decision tailored to your needs. Contact your local HPE account representative for more information and engagement.



**TABLE 2.** Comparison of available SKUs for HPE Data Migration Services

|   | SKU     | Service                                | Description  | Duration  |
|---|---------|--|--|---|
| 1 | HU7C5A1 | HPE Storage DM 40 Hours Phone Supp SVC | Provides telephone support to help assist customers who have adopted a do-it-yourself approach to data migration. HPE will determine the appropriate migration methodology to use.                         | Up to forty hours of telephone support <sup>4</sup> |
| 2 | HR2L6A1 | HPE Data Migration 5 Server Base SVC   | Provides an off-site delivery for data migration of up to five servers to HPE Nimble Storage, HPE Primera, or HPE 3PAR. Maximum of one source array. HPE Peer Motion is used as the data migration method. | Up to three days of off-site delivery               |
| 3 | HR2L7A1 | HPE Data Migration 10 Server Base SVC  | Provides an off-site delivery for data migration of up to 10 servers to HPE Nimble Storage, HPE Primera, or HPE 3PAR. Maximum of one source array.   | Up to four days of off-site delivery                |
| 4 | HR2L8A1 | HPE Data Migration 50 Server Base SVC  | Provides an off-site delivery for data migration of up to 50 servers to HPE Nimble Storage, HPE Primera, or HPE 3PAR. Maximum of two source arrays and up to two sites.                                    | Up to 18 days of off-site delivery                  |
| 5 | HR2L9A1 | HPE Data Migration 100 Server Base SVC | Provides an off-site delivery for data migration of up to 100 servers to HPE Nimble Storage, HPE Primera, or HPE 3PAR. Maximum of two source arrays and up to two sites.                                   | Up to 34 days of off-site delivery                  |
| 6 | HR2M0A1 | HPE Data Migration 250 Server Base SVC | Provides an off-site delivery for data migration of up to 250 servers to HPE Nimble Storage, HPE Primera, or HPE 3PAR. Maximum of four source arrays and up to four sites.                                 | Up to 64 days of off-site delivery                  |
| 7 | HL3P5A1 | HPE Data Migration Addon 10 Server SVC | Allows for an additional 10 servers to be migrated. A single Base SKU, HR2L7A1, HR2L8A1, HR2L9A1, or HR2M0A1 must be ordered as a prerequisite before ordering the Add-on SKU.                             | Up to three and a half days of off-site delivery    |
| 8 | HL3P6A1 | HPE Data Migration Addon 50 Server SVC | Allows for an additional 50 servers to be migrated. A single Base SKU, HR2L8A1, HR2L9A1, or HR2M0A1 must be ordered as a prerequisite before ordering the Add-on SKU.                                      | Up to 15 days of off-site delivery                  |

**NOTE**

Even though there are multiple data migration techniques available, this paper covers data migrations from HPE 3PAR to HPE Primera arrays using HPE Peer Motion in the SSMC.

**HPE PEER MOTION OVERVIEW**

Peer Motion is a complementary software feature available on HPE 3PAR and HPE Primera arrays that enables and facilitates the non-disruptive migration of block data and array configurations. Because Peer Motion software resides natively on the arrays, migrations can be configured, initiated, and monitored without an external appliance in the data path. Inline conversions allow users to migrate data from volumes without compression or deduplication to data reduction volumes with compression and deduplication on the target HPE Primera arrays.

Because of the common architecture shared between the HPE 3PAR and HPE Primera platforms, Peer Motion migrations can import configuration details from the source array and migrate objects including virtual volumes, virtual volume sets, hosts, and host sets, dramatically simplifying the administrative overhead traditionally involved in data migrations from one platform to another.

For most host configurations, users have the option of entirely online migrations with Peer Motion that involve zero host downtime and leverage native multipathing on the hosts. For hosts that do not support online migrations, Peer Motion software allows you to initiate a minimally disruptive migration (MDM) where host I/O is only interrupted during a planned cutover period.

<sup>4</sup> Telephone support only. This will use Peer Motion.






FIGURE 1. Peer Motion architecture

### Workflow

The migration procedure with HPE Peer Motion is very intuitive and can be done entirely through the SSMC or Peer Motion Utility. Figure illustrates the high-level workflow associated with Peer Motion data migration. The following sections cover the specifics about what happens at each of the steps and what is expected.

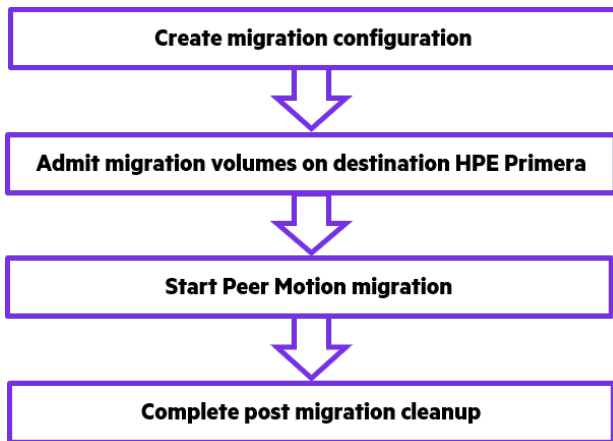


FIGURE 2. Peer Motion migration workflow



## MIGRATING USING HPE PEER MOTION IN THE SSMC

Starting with SSMC version 3.7, Peer Motion migrations from HPE 3PAR to HPE Primera arrays can be initiated and managed within the SSMC interface.

### Prerequisites, planning, and design considerations

You must meet the following prerequisites before configuring a Peer Motion migration using the SSMC:

- The source HPE 3PAR array must be running a supported version of HPE 3PAR OS and be reachable and managed by an SSMC instance. The SSMC instance should be updated to the latest supported version.
- The source HPE 3PAR array must have a Fibre Channel port pair configured as either host or free for the migration. A port pair has the same slot and port on two partner nodes.
- The target HPE Primera array must be properly installed, networked, zoned to hosts, and added into the same SSMC instance as the source HPE 3PAR array via the SSMC Administrator Console.
  - The destination HPE Primera array must be zoned to the same ports on the hosts being migrated as the source HPE 3PAR system.
- The target HPE Primera array must have a Fibre Channel port pair dedicated for peer traffic located on the same SAN fabrics as the specified HPE 3PAR host ports. These peer ports cannot be zoned to other targets or initiators and must not be configured as Remote Copy over Fibre Channel (RCFC) or host ports.
  - iSCSI is not supported for peer links.
  - Direct Fibre Channel or Fibre Channel over Ethernet (FCoE) are not supported for peer links.
- Redundant SAN fabrics for the peer links should be used for optimal fault tolerance and to prevent data unavailability during the data migration.

---

### NOTE

Do not perform any zoning between the source HPE 3PAR array and the target HPE Primera array until the HPE Primera peer ports have been configured. Configuring Fibre Channel ports as peer ports changes the WWN of the ports and creates additional virtual WWNs for a given port that will not function properly with existing WWN-based aliases and zoning.

---

### Key considerations

#### Migration preparation

HPE recommends recording a performance baseline for storage systems, SAN switches, and hosts involved during the week preceding the migration. This activity collects data points on application latency, IOPS, and throughput on the source HPE 3PAR array. HPE 3PAR System Reporter and SAN switch performance collectors can generate this baseline, enabling you to plan migrations during periods where applications are less busy. Increases in I/O, system, and SAN load over the recorded baseline can be attributed to the Peer Motion tasks.

#### Host considerations

Most migrations of volumes exported to an ALUA-capable host with native multipathing configured can be done entirely online from HPE 3PAR to HPE Primera arrays. You can specify a migration to be MDM if you need short, planned host downtime or are performing other tasks related to migration such as SAN fabric migrations. For Linux® and Microsoft Windows host clusters, online migrations can be performed on clusters with a maximum of four nodes. Volumes exported to clusters with more than four nodes must be MDM.

Offline migrations are also possible and can be coordinated with HPE Peer Motion. Source volumes need to stay unexported from the HPE 3PAR and HPE Primera arrays during the entire migration.

For a complete list of supported host operating system versions and their supported migration types, review the [HPE Peer Motion support matrix for HPE Primera and HPE 3PAR](#). Migrations of volumes exported to unlisted host operating systems and clusters can be performed offline using Peer Motion.

### Managing migration tasks

Up to nine Peer Motion import tasks can execute concurrently, and subsequent import tasks are queued until ongoing tasks complete. However, SSMC does not enforce a hard limit on the number of Peer Motion operations that can be started between a migration source and destination and there is no maximum number of tasks that can be queued. Queued tasks can be ordered by priority (high, medium, and low). Within each priority level, tasks are ordered by time of creation with the oldest task executing first. Ongoing import tasks are not preempted even if a higher priority task enters the queue.



### Step-by-step procedure

The following section provides step-by-step instructions for configuring a Peer Motion migration and migrating various objects from an HPE 3PAR array to an HPE Primera array by using the SSMC. For additional procedures and usage regarding Peer Motion using the SSMC, reference the [Migrating data from HPE 3PAR to HPE Primera using the HPE SSMC user guide](#).

#### Creating a migration configuration

1. Navigate to the **Federation & Migration Configuration** page in the SSMC and click **+Create configuration**.
2. Give the configuration a name, and then click **Add systems** to add the destination HPE Primera array.

The SSMC automatically selects an unused Fibre Channel port pair to designate as peer ports. If you want to change the peer ports, click the pencil icon next to the system details and select a different set of available ports.

**Create Configuration** | General ▾ | ?

**General**

Name:

Comments:  optional

**Systems**

Only Smart SAN enabled systems

1 systems

| System            | Model         | HPE Storage OS Version | Peer Ports (Fabric A, Fabric B) | Host Ports (Fabric A, Fabric B) |  |
|-------------------|---------------|------------------------|---------------------------------|---------------------------------|--|
| prdf-C650-r21a221 | HPE_3PAR C650 | 4.2.1                  | 2:3:1, 3:3:1                    | —, —                            |  |

**Add systems**

**FIGURE 3.** Creating a migration configuration

#### NOTE

If no unused Fibre Channel port pairs are found, the HPE Primera system will show as unavailable in the Add systems menu. The system cannot be used to create a migration configuration until suitable ports are available.

3. Verify that the name, model, operating system version, and peer ports of the destination HPE Primera array are correct. Then click **Create**.
4. Add the source HPE 3PAR array to the newly created configuration by clicking **Actions** and then **Add Migration source**.
5. Click **Select source** to add the source HPE 3PAR array.

The SSMC will automatically select a Fibre Channel port pair on the source HPE 3PAR array that is configured either as host mode or free. The ports must be cabled to the fabric switch and in a ready state. If you want to change the ports that will be used for migration traffic, click the pencil icon next to the system details and select a different set of available ports.





**Select** Source System
?

Total 2

| System            | Model         | Version     | Status      |
|-------------------|---------------|-------------|-------------|
| prdf-8450-r20a223 | HPE_3PAR 8450 | 3.3.1 (MU5) | Available   |
| prdf-C650-r21a221 | HPE_3PAR C650 | 4.2.1       | Unavailable |

Add systems

Select

Cancel

**FIGURE 4.** Adding a migration source system

- Click **Add destination** to specify the HPE Primera array as the destination and verify the details of both systems.

**Add Migration Source** vMSC-PM
General ▾
?

**General**

Federation name vMSC-PM

**Source System**

Only Smart SAN enabled systems

1 systems

| Source System     | Model         | Version     | Host Ports<br>(Fabric A, Fabric B) |     |
|-------------------|---------------|-------------|------------------------------------|-----|
| prdf-8450-r20a223 | HPE_3PAR 8450 | 3.3.1 (MU5) | 0:0:1, 1:0:1                       | ✎ ✕ |

Select source

**Destination Systems**

1 systems

| Destination System | Model         | HPE Storage OS<br>Version | Peer Ports<br>(Fabric A, Fabric B) |   |
|--------------------|---------------|---------------------------|------------------------------------|---|
| prdf-C650-r21a221  | HPE_3PAR C650 | 4.2.1                     | 2:3:1, 3:3:1                       | ✕ |

Add destinations

**FIGURE 5.** Specifying the migration destination system



- Review the Fibre Channel zoning recommendations that populate after adding both the source and destination arrays to the configuration. If you are not using Smart SAN to automatically zone the ports, note the zoning recommendations and create the appropriate zones. Zoning recommendations can be revisited later by navigating to the Recommended Zones menu on a selected configuration.

| Migration Sources  |   |       |                  |
|--|---|-------|------------------|
| prdf-8450-r20a223 to prdf-C650-r21a221 Zone 1 (Fabric A) |   |       |                  |
| System   | ▲ | Port  | Port WWN         |
| prdf-8450-r20a223  |   | 0:0:1 | 20010002AC01D7F0 |
| prdf-C650-r21a221  |   | 2:3:1 | 22310202AC07EAC9 |
| prdf-C650-r21a221  |   | 2:3:1 | 22310202AC17EAC9 |
| prdf-C650-r21a221  |   | 2:3:1 | 22310202AC27EAC9 |
| prdf-C650-r21a221  |   | 2:3:1 | 22310202AC37EAC9 |
| prdf-C650-r21a221  |   | 2:3:1 | 22310202AC47EAC9 |
| prdf-C650-r21a221  |   | 2:3:1 | 22310202AC57EAC9 |
| prdf-C650-r21a221  |   | 2:3:1 | 22310202AC67EAC9 |
| prdf-C650-r21a221  |   | 2:3:1 | 22310202AC77EAC9 |
| prdf-C650-r21a221  |   | 2:3:1 | 22310202AC87EAC9 |

**FIGURE 6.** Reviewing the Fibre Channel zoning recommendations

The SSMC automatically configures eight virtual peer ports for each physical peer port specified on the destination HPE Primera array. These virtual peer ports serve as paths for handling SCSI reservation traffic when migrating Windows and Linux clusters. Even if the volumes being migrated do not reside on a Windows or Linux cluster, the virtual ports are still created and should be zoned as recommended.

**NOTE**

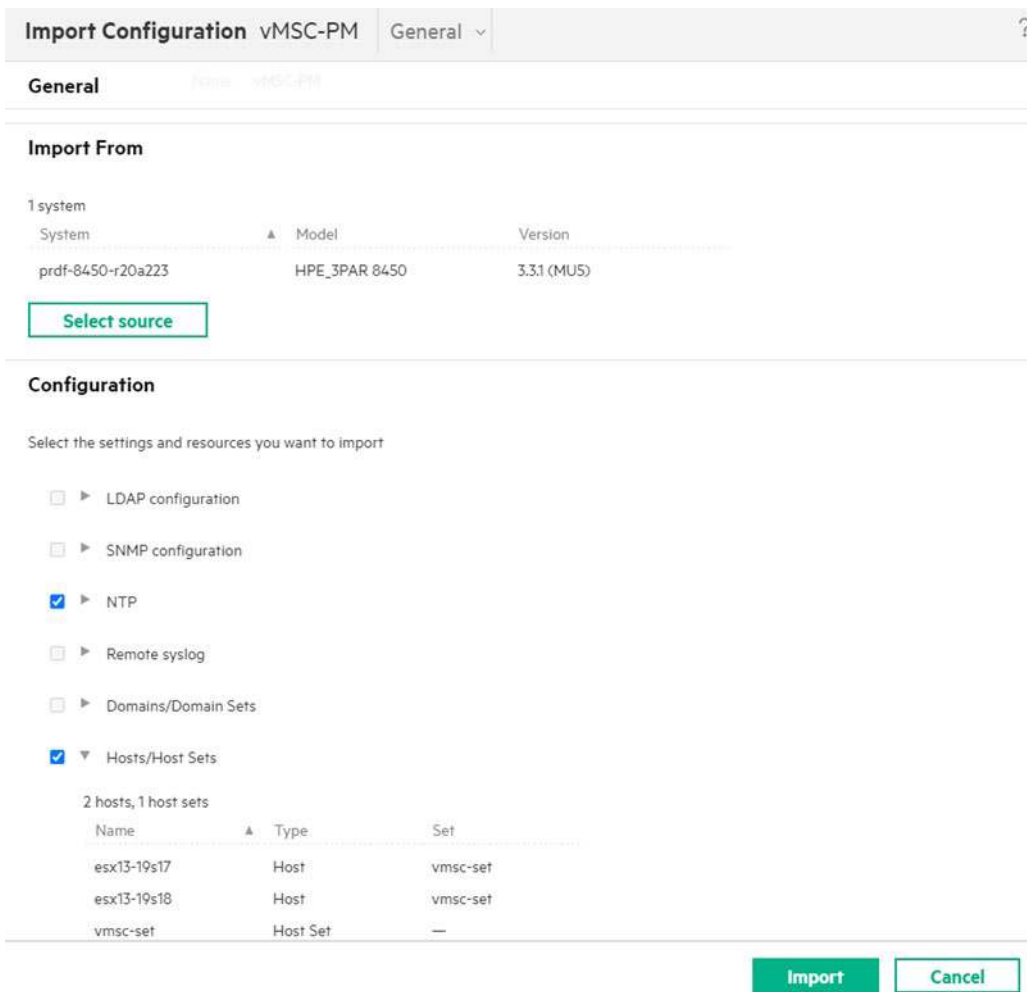
If you have checked the **Only Smart SAN enabled systems** box and the source array has Smart SAN capable host ports, zoning will happen automatically between the destination peer ports and the source host ports.

**Importing a configuration**

Peer Motion offers flexible migration granularity, enabling you to migrate various objects and configuration settings from a source HPE 3PAR array to a destination HPE Primera array. Before migrating block data, you can import existing settings on the source HPE 3PAR including LDAP configurations, SNMP configurations, NTP servers, remote syslog forwarding settings, domains, domain sets, host definitions, host sets, and users. To import configuration settings from a source array:

- From the Federation & Migration Configuration page, select the configuration. Click **Actions**, and then click **Import Configuration**.
- Select the source HPE 3PAR system and review the settings and resources that can be imported. Click the arrows next to a specific setting to expand it and see specific details about the import. Settings and resources that are suitable for import from the source array are automatically selected. Click **Import** after settings have been selected and reviewed.





**FIGURE 7.** Importing a configuration from the source system

**NOTE**

In order to import hosts and host sets from the source HPE 3PAR to the destination HPE Primera array, the hosts must be properly zoned to the destination array.

After the migration configuration with the source and destination arrays has been created, the configuration can be used immediately to start Peer Motion migrations or can remain in place until you are ready to initiate migration tasks. During this time, the peer links on the destination HPE Primera array will be consumed and unable to be assigned for other purposes such as Remote Copy or host ports but will not handle any traffic until a migration task has been created.

**Migrating virtual volumes and virtual volume sets**

Volumes can be migrated individually or as part of a consistency set using Peer Motion.

1. Navigate to the Virtual Volumes or Virtual Volume Sets tab of the SSMC menu. Select the virtual volume or volume set to migrate, click the **Actions** drop-down menu, and select **Start Peer Motion**.



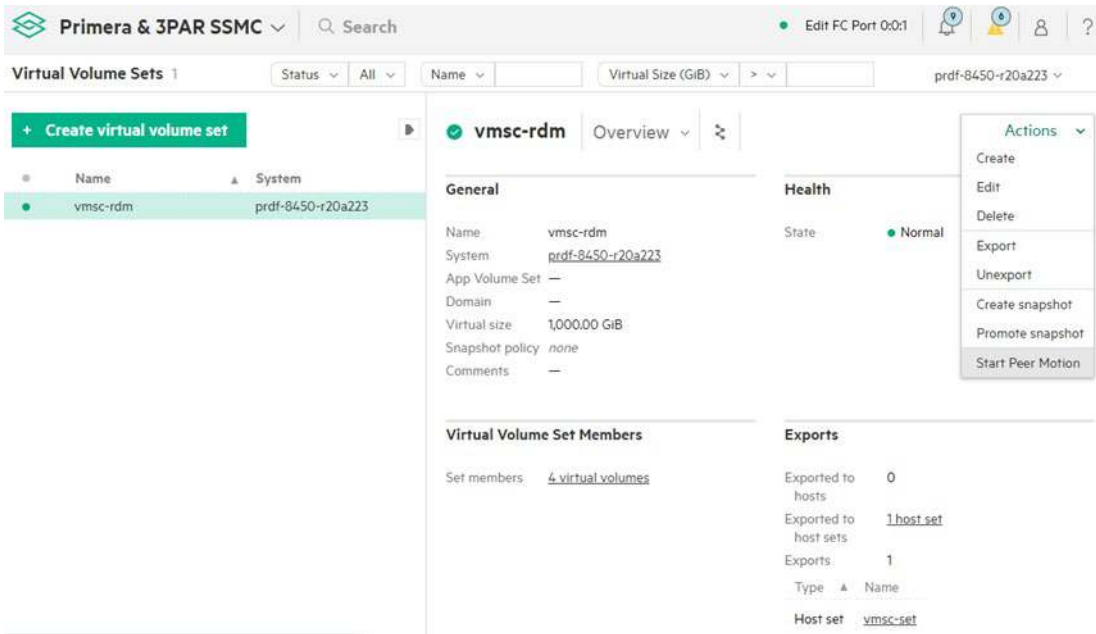


FIGURE 8. Starting the Peer Motion of a virtual volume set

2. Name the Peer Motion task or leave the automatically generated name. In this menu, you can also select whether to have email alerts sent after the migration task is complete if an SMTP server has been enabled.
3. In the Peer Motion Settings section of the menu, specify the destination HPE Primera system and the destination CPG on which the peer volumes will be created. Here, you can also specify if the migration should be an MDM, whether the source volumes should be automatically deleted upon completion, and if the Peer Motion should be suspended before importing volumes. You can also specify a future start time for the migration. Otherwise, the task starts immediately.

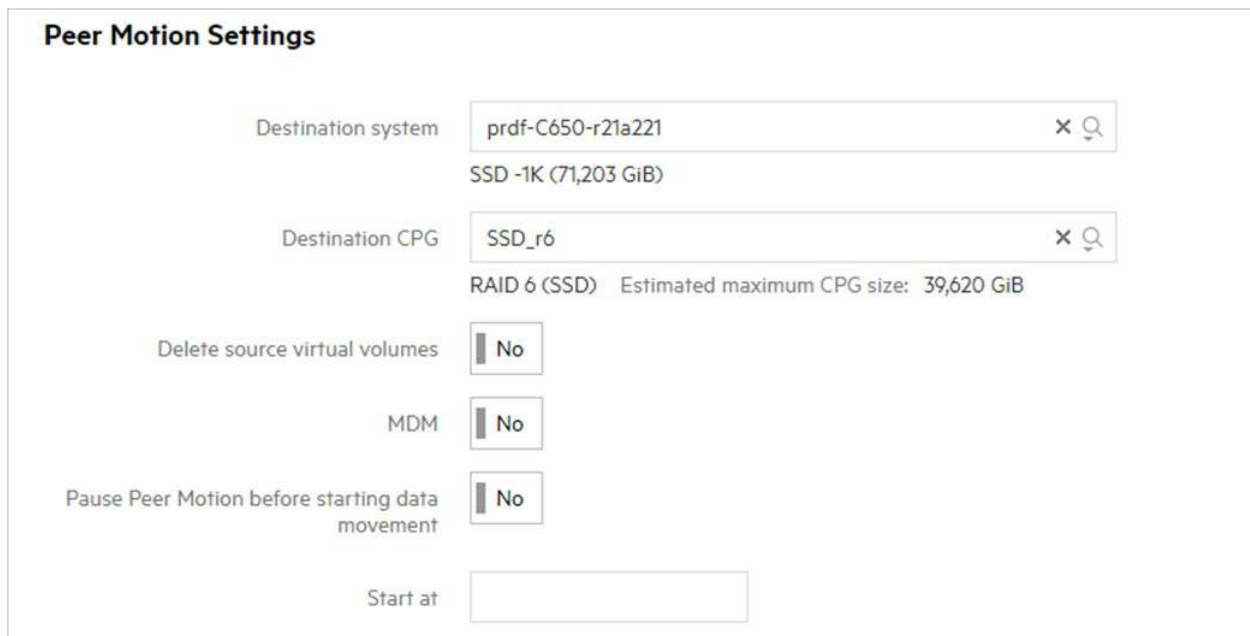


FIGURE 9. Specifying Peer Motion settings for a virtual volume set



- Under Virtual Volume settings, you can edit details about the target volumes before the volumes are created. The reserved size must be identical to the source volume, but the CPG type, provisioning type, and data reduction settings can be edited for the target volumes. You can also specify the priority of migration for each volume selected. Higher-priority migration tasks will take precedence in the queue if many tasks have been initiated.

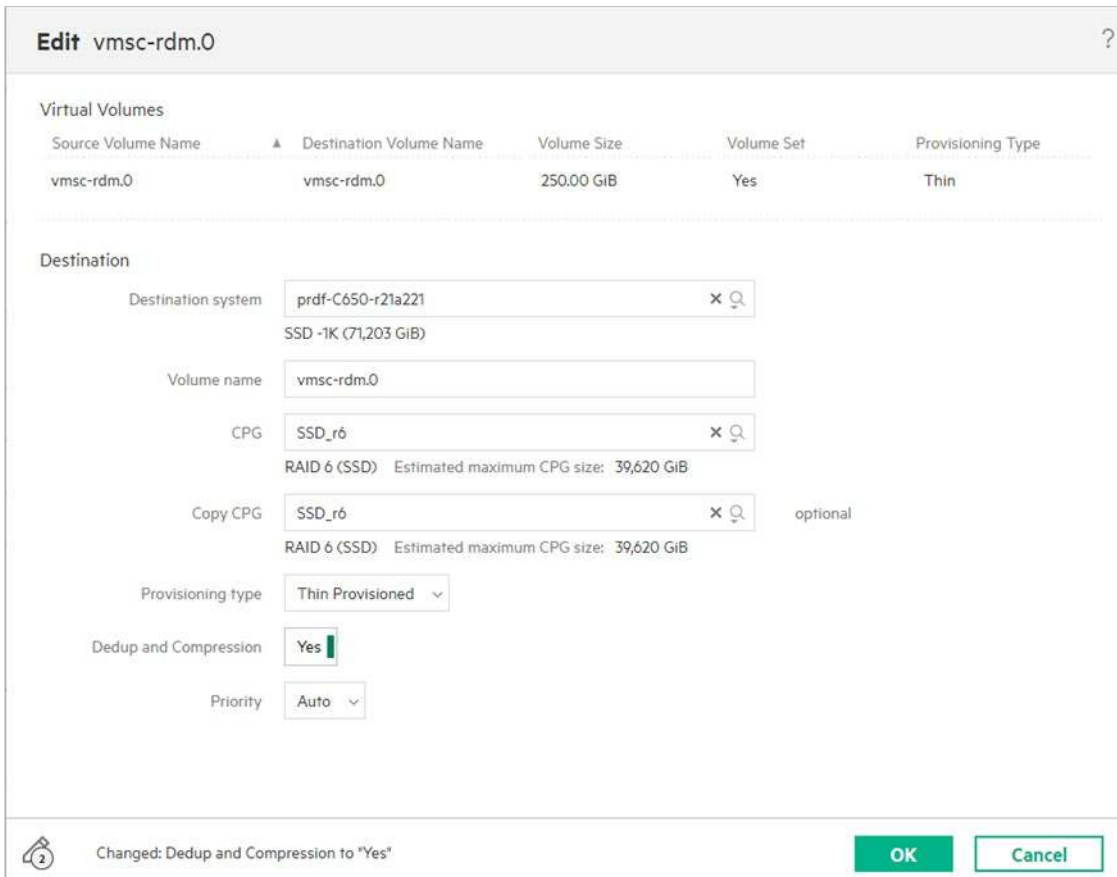


FIGURE 10. Specifying target volume settings for a Peer Motion volume

- If a virtual volume set was selected for a Peer Motion task or if an individual volume selected is part of a set, there will be an additional Virtual Volume Set settings section that allows you to edit set-specific settings including migration task priority and whether the set should be migrated as a consistency group. A consistency group specifies that all the volumes in the virtual volume set will be kept consistent on source and destination arrays until the migration of all volumes in the set are complete and the entire set can be cut over as a collective unit.
- After you have configured all appropriate settings, click **Start** to initiate the setup stage for the Peer Motion migration. Click the task to see the progress in the Peer Motion tab of the SSMC.

There are a few key distinctions between online migrations and minimally disruptive migrations to note at this point:

- For **Online Migrations**: The target peer volumes will be created on the destination array with the same WWNs as the source volumes and exported to the host or host set. The Peer Motion task will pause and instruct you to perform an HBA rescan on the host.

The host's multipath software will discover the new paths from the destination array and should display active paths from both the source HPE 3PAR array and the destination HPE Primera array for the volumes that are part of the current migration task. I/O can travel down all available paths during this transition period before the Peer Motion is resumed and import begins. Host I/Os that are sent down the source array's paths are handled as usual, and host I/Os that are sent down the destination array's paths are proxied over the peer links to the source array. During this time, no I/O is written to or read from the back-end disks of the destination array, and no data is being imported from the source volumes.

After you have verified that the new export paths between the destination array and the host are active, the Peer Motion task can be resumed from the Peer Motions screen. Resuming the Peer Motion will do two things:

- Transition the ALUA state of the paths on the source array from active to standby. As a result, all host I/O will pass through the active paths between the host and the destination array.



- Initiate the import of data over the peer links from the volumes on the source array to the target peer volumes on the destination array.

---

**NOTE**

After resuming the task, you should validate from the host that the host is only seeing standby paths from the source array and active paths from the destination array.

---

- For **Minimally Disruptive Migrations**: The target peer volumes will be created on the destination array with the same WWNs as the source volumes but will not be exported to the host or host set yet. The Peer Motion task will pause and instruct you to shut down all hosts that are part of the migration.

After the hosts are down, the Peer Motion can be resumed which will begin the import of data over the peer links from the source array's volumes to the destination array's target peer volumes. The target peer volumes will be exported to the host or host set. The hosts can be rebooted any time after this point and application I/O restarted with the following considerations:

- Hosts with a non-ALUA persona must be unzoned from the source HPE 3PAR array or the volumes must be unexported from the source array before rebooting to prevent data inconsistencies.
- Hosts configured with an ALUA-capable persona do not require any manual intervention with unexporting source volumes or un zoning from the source array. When the hosts come back online, the paths from the source array will be in standby mode and the paths from the destination array will be active.

When the migration of a virtual volume or volume set has completed, the volume or set is cut over. The target volumes on the destination HPE Primera will now show provisioning as thin rather than peer, and the volumes can now permanently reside on the destination HPE Primera array. The source volumes on the HPE 3PAR are unexported from the destination HPE Primera and are either removed automatically or can be removed manually later.

---

**NOTE**

After the migration of a volume or volume set has completed and the cutover occurs, the original volumes on the source array immediately become stale and should be unrepresented from hosts and removed to prevent data consistency problems if they are no longer required.

---

**Read I/O handling during an online or MDM migration**

Host reads initiated to a region that has not yet moved to the destination are proxied through the destination HPE Primera to the source HPE 3PAR array over the peer links. The read is then returned from the source to the host through the destination HPE Primera array. This means that earlier in the migration process, read latencies will be higher because reads to more regions require proxying over the peer links.

Host reads initiated to a region that has already moved to the destination are served locally by the destination array. As a migration progresses, read latencies will start to decrease as more regions can be read from locally without proxying. To help keep read latencies to frequently accessed blocks as low as possible throughout the migration, the destination virtual volume has read caching enabled.

**Write I/O handling during an online or MDM migration**

Host writes initiated to a region that has not yet moved to the destination are proxied through the destination HPE Primera to the source HPE 3PAR over the peer links. The write is committed to disk on the source HPE 3PAR array but not on the destination array. When that region is ready to be migrated, the destination HPE Primera will directly import the region from the source HPE 3PAR array and then commit it to disk.

Host writes for regions that have already been migrated are written to the destination array and mirrored over the peer links to the source array.

All writes require acknowledgement back from the source array to maintain volume consistency on both sites during the migration. This allows for an in-progress migration task to be safely aborted and rolled back to the source array without compromising data integrity.

When a consistency set has been defined for a migration, the import of the set is not complete until the imports have completed for all virtual volumes that are members of the set. In order to preserve data consistency at the set level, volumes that are done with their region moves early must continue to mirror their writes to the source array until all other volumes that are part of the set are also done with their region moves. All volumes in a set are cut over together at the end of the migration task.



**Cleaning up after migration**

After all Peer Motion tasks have successfully completed from the source HPE 3PAR array to the destination HPE Primera array and all objects have been migrated, the migration configuration should be cleaned up and the source HPE 3PAR array can be removed.

1. Remove the migrated volumes from the source HPE 3PAR array.
2. In the SSMC Federation & Migration Configuration page, select the configuration to be removed. In the Actions drop-down menu, select **Delete**. Deleting the migration configuration will remove all information about the migration from the participating storage systems.

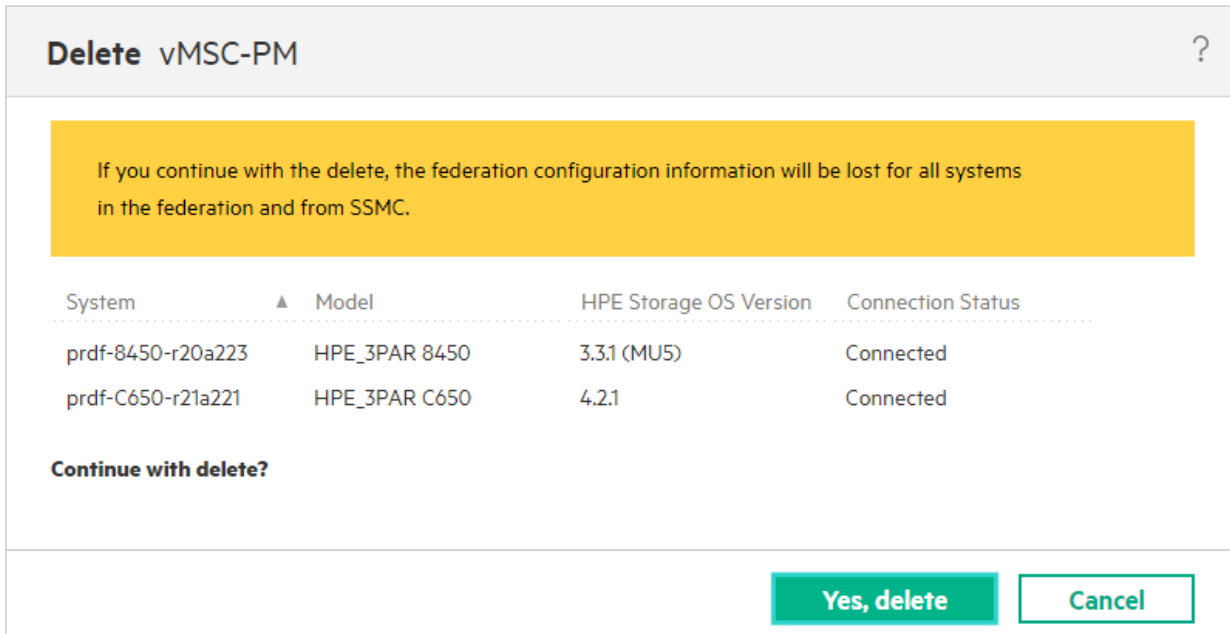


FIGURE 11. Deleting a migration configuration

**MIGRATING USING THE HPE PEER MOTION UTILITY**

For users that prefer a command-line experience, Peer Motion migrations from HPE 3PAR to HPE Primera can be controlled through the Peer Motion Utility (PMU). For a comprehensive guide on procedures and commands using the PMU, reference the [Migrating data from HPE 3PAR to HPE Primera](#) user guide.

**Prerequisites, planning, and design considerations**

You must meet the following prerequisites before configuring a Peer Motion migration using the HPE Peer Motion Utility:

- Have a networked Windows system to install the PMU client and server software.
  - TCP ports 2390 and 2388 must be open to facilitate Peer Motion migrations.
  - The Windows system must be able to reach both source HPE 3PAR array and the destination HPE Primera array management interfaces.
  - The system can be physical or virtual.
- The source HPE 3PAR array must be running a supported version of HPE 3PAR OS.
- The source HPE 3PAR array must have at least one Fibre Channel port pair configured as either host or free for the migration. A port pair has the same slot and port on two partner nodes.
- The target HPE Primera array must be properly installed, networked, and zoned to hosts.
  - Ensure the hosts that will be migrated are zoned to the destination HPE Primera array that uses the same host port WWNs as the source HPE 3PAR system.



- The target HPE Primera array must have at least one Fibre Channel port pair dedicated for peer traffic located on the same SAN fabrics as the HPE 3PAR host ports. These peer ports cannot be zoned to other targets or initiators and must not be configured as RCFC or host ports.
  - iSCSI is not supported for peer links.
- Direct Fibre Channel or FCoE is not supported for peer links.
- Redundant SAN fabrics for the peer links should be used for optimal fault tolerance and to prevent data unavailability during the data migration.

### Multi-link support

Starting with version 2.4 of the HPE Peer Motion Utility, multiple pairs of peer links are supported for migrations from HPE 3PAR to HPE Primera arrays. Configuring multiple peer link pairs helps reduce migration time by distributing host and migration load during Peer Motion across all available peer link pairs. Best practices for multi-link configurations are:

- Distribute peer links evenly across all available nodes on the destination storage system.
  - Destination peer link pairs must be from partner nodes.
  - Up to four peer link pairs (eight peer links) can be configured on one destination system.
- Use the same number of corresponding source host ports as peer ports.
  - Connect each source host port to one destination peer port.

## PERFORMANCE CONSIDERATIONS WITH HPE PEER MOTION

Migration in a Peer Motion configuration is point-to-point between the source HPE 3PAR array and the destination HPE Primera array, resulting in the shortest path possible between the participating storage systems. This keeps latency low compared to migrations that use an external appliance operating inside of the data path.

To maximize throughput and minimize host application latency and migration duration, HPE recommends executing online or minimally disruptive migrations during times of low application activity. Although the source HPE 3PAR array does require dedicated host ports for Peer Motion migration traffic, selecting host ports that are unused or carry minimal traffic from other hosts can help minimize host latency during active migration tasks.

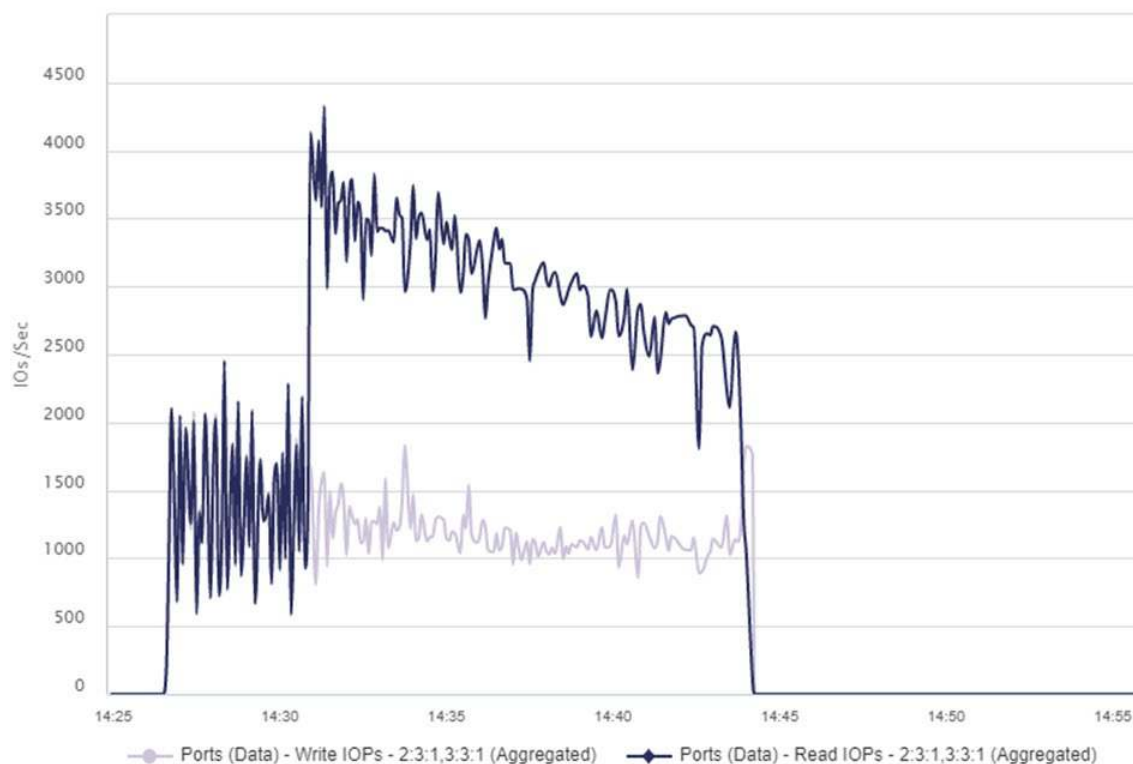
Peer Motion migrations operate best between co-located storage arrays. However, migrations can be initiated between arrays in separate data centers. If the source and destination systems are installed in different sites, ensure that the SAN intersite links (ISLs) have sufficient buffer-to-buffer credits to handle migration traffic as well as any other host or replication traffic traversing the ISLs. Because all writes are synchronously mirrored to the source array during active migrations, the application latency of the host might greatly increase during the migration if the distance between the source and destination arrays is significant.

The HPE SSMC can be used to graph real-time performance of the migration over the HPE Primera peer links. The figure below demonstrates an example destination HPE Primera array's aggregated peer link performance from start to finish for a Peer Motion online migration of a volume set.





Real Time Peer Link Performance



**FIGURE 12.** Example of peer link performance observed before, during, and after Peer Motion migration of a volume set

Before a Peer Motion task of a volume or volume set is started, the peer links are unused and will report no IOPs. When the Peer Motion task is created, the peer volumes are created and exported, and the host starts sending I/O down all available paths from both the source and destination array. During this time, all I/Os sent to the destination HPE Primera are proxied to the source HPE 3PAR array over the peer links and so some portion of host traffic appears on the peer links.

After the Peer Motion migration is resumed, the destination HPE Primera begins to handle all host traffic for the migrated volumes (mirroring all writes and proxying reads for regions that have not migrated yet) while simultaneously importing data from the source volumes. The peer link IOPs depicts both the traffic from proxied host I/Os as well as the ongoing import traffic.

As the migration progresses, reads to more regions can be serviced locally by the HPE Primera array. As a result, the amount of host read traffic sent over the peer links will gently decline. After the migration is complete, the volumes are cut over to the destination HPE Primera and traffic over the peer links is halted.

**SUMMARY**

Data migration from HPE 3PAR to HPE Primera can be performed seamlessly and nondisruptively using the complementary array native HPE Peer Motion software. Leveraging point-to-point connectivity, Peer Motion can ensure fast, easy migrations and there is no licensing per TiB moved.

Although software like Peer Motion removes much of complexity involving data migrations, HPE also offers industry-leading consulting services to make your move to HPE Primera with a managed migration service for customers that are hesitant to migrate their data themselves.



**Resources**

[Migrating data from HPE 3PAR to HPE Primera](#)

[Migrating data from HPE 3PAR to HPE Primera using the HPE SSMC 3.7](#)

[Migrating data from HPE 3PAR to HPE Nimble](#)

**LEARN MORE AT**

<https://www.hpe.com/us/en/storage/hpe-primera.html>

**Make the right purchase decision.  
Contact our presales specialists.**



**Chat**



**Email**



**Call**



**Get updates**

---

© Copyright 2021 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Azure, Microsoft, SQL Server, and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. VMware, VMware vSphere, VMware vSphere High Availability, and VMware vSphere vMotion are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. Oracle is a registered trademark of Oracle and/or its affiliates. All third-party marks are property of their respective owners. All third-party marks are property of their respective owners.

a00110464ENW, February 2021