

# EMC POWERPATH MIGRATION ENABLER HOST COPY

## A Detailed Review

### Abstract

This white paper focuses on EMC® PowerPath® Migration Enabler Host Copy, which is entirely host-based. As with other versions of PowerPath Migration Enabler, it moves data non-disruptively from a source to a target while cloning application writes between them. This paper covers installation, usage, and deployment scenarios, and provides recommendations.

March 2011

Copyright © 2010, 2011 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is”. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on [EMC.com](http://EMC.com).

All other trademarks used herein are the property of their respective owners.

Part Number h6927.1

## Table of Contents

<b>Executive summary</b> .....	<b>5</b>
Audience.....	5
<b>Supported platforms</b> .....	<b>6</b>
<b>Protocol support</b> .....	<b>6</b>
<b>Types of migrations</b> .....	<b>6</b>
<b>PowerPath framework</b> .....	<b>7</b>
PowerPath Migration Enabler and PowerPath Multipathing.....	8
Understanding PowerPath pseudo devices.....	8
Using native devices with PowerPath installed .....	10
<b>Overview of PowerPath Migration Enabler Host Copy steps</b> .....	<b>11</b>
<b>PowerPath Migration Enabler state diagram</b> .....	<b>12</b>
PowerPath Migration Enabler commands .....	12
Migration process.....	13
Pre-migration .....	13
Setup .....	14
Synchronizing .....	14
Source Selected .....	16
Target Selected .....	16
Commit .....	17
Cleanup.....	18
Post-migration.....	19
Migrating clusters.....	19
Graceful host reboot.....	19
Host crash .....	19
<b>Migrating to a virtually provisioned target</b> .....	<b>21</b>
<b>Migration speed and host impact</b> .....	<b>23</b>
Understanding throttle impact.....	24
AIX .....	25
Windows, Linux, Solaris, and HP-UX.....	27
<b>Thick to thin migration example</b> .....	<b>27</b>
<b>Logging</b> .....	<b>35</b>
Audit log.....	35
Error log.....	36
Log locations.....	36
<b>Measuring elapsed migration time</b> .....	<b>36</b>
<b>Encryption migration</b> .....	<b>37</b>

Throttle throughput example .....	38
Installing PowerPath Migration Enabler .....	39
Licensing PowerPath Migration Enabler .....	40
Obtaining a license .....	41
Licensing migration of encrypted devices .....	41
Applying a license .....	41
Restrictions .....	41
Conclusion .....	42
References .....	42

## Executive summary

Customers are facing an ever-increasing number of migration challenges. Information is growing faster than ever before. Infrastructure complexities and dependencies exist. Migration windows are variable and shortened. Disruptions need to be reduced. So reducing management, hardware, and personnel costs associated with operational overhead has become a critical factor in data center migrations. EMC® PowerPath® Migration Enabler Host Copy is a host-based migration tool that addresses these common customer concerns.

Storage administrators have numerous migration options from which to choose, including array-based migrations, virtualization products, and host-based tools. Choosing the appropriate solution depends on several factors. Ease of configuration and management, hardware capabilities, ability to throttle the rate of data movement, and determining application impact are critical when making a choice. The best solution in one migration may not necessarily be the best solution in the next migration. Also, migrations across the enterprise may differ among the functional business units or organizations. No one-size-fits-all migration tool exists. Each migration tool has its own set of advantages, disadvantages, challenges, and nuances. In the end, it is the data center administrator's responsibility to select a tool that has the necessary features and functions to solve these challenges.

PowerPath is known widely by data center administrators as a premier multipathing solution for SAN-based Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), and iSCSI storage. PowerPath Migration Enabler is an advanced PowerPath feature. It can use multiple EMC technologies to perform bulk data migration operations. Using virtual LUN abstraction, PowerPath creates pseudo devices that abstract the physical device by creating a virtual device. PowerPath Migration Enabler uses this feature to make migrations transparent to host applications.

## Audience

This white paper is intended for the technology professional who performs data migrations in an environment with EMC and third-party disk arrays. It is specifically targeted at EMC field technical staff and EMC customers who require non-disruptive data migration to a new array, to a thin LUN, or to an encrypted LUN using PowerPath Migration Enabler Host Copy.

## Supported platforms

PowerPath Migration Enabler (PPME) is supported on multiple operating systems including Windows, Linux, AIX, HP-UX, and Solaris. For specific OS and kernel versions, refer to the *PowerPath Migration Enabler Release Notes* on [Powerlink](#)<sup>®</sup> or the *PowerPath Migration Enabler Support Matrix* at [E-Lab](#)<sup>™</sup>. Access to both sites is required.

---

Note: PowerPath/VE for vSphere does not support PowerPath Migration Enabler.

---

## Protocol support

PowerPath Migration Enabler supports FC, iSCSI, and FCoE. Refer to the *PowerPath Protocol Support Table* on E-Lab Navigator for the latest support.

## Types of migrations

PowerPath Migration Enabler Host Copy supports multiple migration scenarios.

### Standard to standard devices

Host Copy can be used to migrate from a standard source device to a standard target device. Neither the RAID protection types nor the device sizes have to be identical.

---

Note: PowerPath Migration Enabler does not support migrating to a smaller target device.

---

### Virtually provisioned devices

Host Copy can be used to migrate to and from virtually provisioned devices (thin devices). As with standard devices (thick devices), the RAID protection types do not have to be identical when using virtually provisioned devices.

### Encrypted devices

PowerPath Encryption with RSA<sup>®</sup> is a software-based encryption tool that leverages RSA technology to encrypt user-specified PowerPath controlled devices. It is a LUN-based encryption tool; therefore, all file systems running on the LUN are also encrypted. Writes are encrypted to and reads are decrypted from the encrypted LUN. In order to encrypt existing data (plain- or clear-text), the data must be migrated to an encrypted target device. PowerPath Migration Enabler Host Copy is a non-disruptive method of performing this migration.

### Pseudo and native devices

PowerPath Migration Enabler supports migration using PowerPath devices (pseudo devices) or native OS devices. Migrating with pseudo devices has significant benefit. If the applications are configured to point at the pseudo devices, the migration will be non-disruptive. If the applications are configured to point at native devices, then the administrator must manually reconfigure the application after the migration so that

the host sees the new target device. For more information on PowerPath pseudo devices, refer to the next section. Table 1 lists all supported migration types.

---

Note: Pseudo to native migrations are not supported.

---

**Table 1. Supported PPME Host Copy migrations**

Migration type	Device attributes
Standard to standard devices	RAID <sub>x</sub> to RAID <sub>x</sub> RAID <sub>x</sub> to RAID <sub>y</sub> RAID <sub>x</sub> to RAID <sub>x</sub> with larger target RAID <sub>x</sub> to RAID <sub>y</sub> with larger target
Virtually provisioned devices	Thick to thick (same as standard to standard) Thick to thin Thin to thin Thin to thick
Encrypted devices	Plain-text to encrypted Encrypted to plain-text Encrypted to encrypted (rekeying operation)
Pseudo and native devices	Pseudo to pseudo (non-disruptive) Native to pseudo (disruptive) Native to native (disruptive)

## PowerPath framework

PowerPath is host-based software that resides below logical volume managers, file systems, and applications, but above the SCSI drivers. This location of PowerPath means that every I/O from the host to the storage system has to pass through PowerPath. This allows PowerPath to work with the storage system to provide intelligent I/O path management. This management includes dynamic load balancing of I/O requests and automatic detection and recovery from host-to-storage path failures.

PowerPath also has extensions. Extensions are platform-independent kernel modules that operate within the abstraction that the PowerPath framework provides.

Extensions are layered on top of each other. This means that an I/O request will flow through each extension, starting at the top, and may be redirected along the way if that extension isn't used for a particular I/O operation. Depending on the features in use (Multipathing, Migration Enabler, and Encryption), not every I/O will necessarily be redirected through each extension.

In Figure 1, the use of extensions depends on how an administrator uses PowerPath. The vast majority of SAN administrators using PowerPath use the Multipathing Extension.

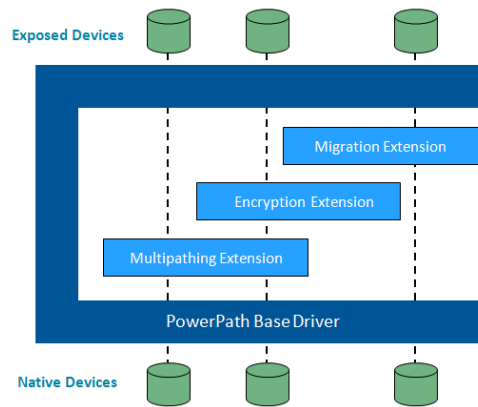


Figure 1. Basic components of the PowerPath framework

## PowerPath Migration Enabler and PowerPath Multipathing

As described in the previous section, PowerPath has multiple extensions that provide load balancing, failover, migration, and encryption functionality to the data center administrator.

To use PowerPath Migration Enabler, PowerPath multipathing is installed because the multipathing extension is part of the Migration Enabler base install. PowerPath *must* manage the devices. Another multipathing product like native MPIO cannot manage those devices. The multipathing license is not required. PowerPath Migration Enabler can support a single active path while redundant paths are in an “unlicensed” state. However, customers will achieve higher throughput and availability when multipathing is enabled.

## Understanding PowerPath pseudo devices

For each logical device that PowerPath discovers, it creates a PowerPath device. A one-to-one relationship exists between PowerPath devices and logical devices. PowerPath devices are presented differently depending on the platform. Much of this difference is due to the design of the host operating system. Depending on the operating system, PowerPath can present PowerPath devices as native devices or pseudo devices. A pseudo device is a special kind of device (which can be defined as an operating system object used to access devices) that PowerPath creates. Once PowerPath has been installed, it is path-independent, as are native devices. Only one pseudo device exists per path set. For example, PowerPath supports up to 32 paths per device. PowerPath creates one pseudo device to represent all 32 paths.

---

Note: PowerPath supports pseudo devices on all supported platforms except HP-UX.

---

One or more native paths to the same storage LUN are presented as a single pseudo device. A pseudo device provides load-balancing and failover support.

Figure 2 shows the output of the **fdisk -l** command on a RHEL5 server. One EMC CLARiiON® CX4 series 20 GB device has been assigned to this host. There are four paths to this CX4 — two paths to SP A and two paths to SP B.

The output shows four distinct native paths (/dev/sdb1, c1, d1, and e1). Also one PowerPath pseudo device represents all physical paths (/dev/emcpowera).

The use of pseudo devices is important when using PowerPath Migration Enabler. In Figure 2, the device name is **/dev/emcpowera**. This is simply a name that represents the device. If the device is configured, the application uses this name for it. During the migration process, Migration Enabler will continue to use the same name (/dev/emcpowera) with the application; however, the device itself will change from the source to the target.

```
Disk /dev/sdb: 21.4 GB, 21474836480 bytes
255 heads, 63 sectors/track, 2610 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

```
Device Boot  Start    End  Blocks  Id System
/dev/sdb1    1      2611  20968448  7 HPFS/NTFS
```

```
Disk /dev/sdc: 21.4 GB, 21474836480 bytes
255 heads, 63 sectors/track, 2610 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

```
Device Boot  Start    End  Blocks  Id System
/dev/sdc1    1      2611  20968448  7 HPFS/NTFS
```

```
Disk /dev/sdd: 21.4 GB, 21474836480 bytes
255 heads, 63 sectors/track, 2610 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

```
Device Boot  Start    End  Blocks  Id System
/dev/sdd1    1      2611  20968448  7 HPFS/NTFS
```

```
Disk /dev/sde: 21.4 GB, 21474836480 bytes
255 heads, 63 sectors/track, 2610 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

```
Device Boot  Start    End  Blocks  Id System
/dev/sde1    1      2611  20968448  7 HPFS/NTFS
```

```

Disk /dev/emcpowera: 21.4 GB, 21474836480 bytes
255 heads, 63 sectors/track, 2610 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

```

Device	Boot	Start	End	Blocks	Id	System
/dev/emcpowera1		1	2611	20968448	7	HPFS/NTFS

Figure 2. Output of the *fdisk-l* command on Red Hat 5

The PowerPath CLI is called **powermt**. It displays information about HBAs and devices configured for and managed by PowerPath. The **powermt display dev=all** command shows all PowerPath managed devices. Similar to the **fdisk-l** command, native I/O path information is shown as well as the EMC pseudo device. Figure 3 shows the output of the **powermt display dev=all** command on the same RHEL5 server. The pseudo device represents all four I/O paths.

```

[root@lppa020 ~]# powermt display dev=all
Pseudo name=emcpowera
CLARiiON ID=FNM00092000177 [lppa020]
Logical device ID=6006016038E02400E64F749233AADE11 [LUN 22]
state=alive; policy=CLAROpt; priority=0; queued-I/Os=0
Owner: default=SP A, current=SP A   Array failover mode: 4
=====
----- Host ----- - Stor -   - I/O Path -   -- Stats --
### HW Path   I/O Paths   Interf.   Mode   State   Q-I/Os   Errors
=====
1 qla2xxx   sdb           SP A1    active  alive   0   0
1 qla2xxx   sdc           SP B0    active  alive   0   0
3 qla2xxx   sdd           SP A0    active  alive   0   0
3 qla2xxx   sde           SP B1    active  alive   0   0

```

Figure 3. Output of the *powermt display* command

PowerPath Migration Enabler Host Copy will leverage these pseudo devices to non-disruptively migrate data without any application impact.

### Using native devices with PowerPath installed

PowerPath automatically creates pseudo devices in supported operating systems. Applications can be configured to use pseudo devices. However, administrators can choose to use native devices instead.

When using native devices, PowerPath is transparent to applications. PowerPath maintains the correspondence between an individual native device and the path to which it belongs. On Solaris and Linux systems, an administrator has the option of

either using native devices (with no conversion of applications) or converting the application to use pseudo devices.

For example, suppose three native devices exist in a path set (a pseudo device representing multiple paths). PowerPath maintains the association among these paths. When an application writes to any one of them, PowerPath redirects the I/O to whichever native device (in the path set) will provide the best throughput based on the load-balancing policy setting. (For more information on load-balancing settings, refer to the *PowerPath Family CLI and System Messages Reference Guide*.) Also, a problem with one native device does not disrupt data access. Instead, PowerPath shifts I/O processing to another native device in the path set, allowing applications to continue reading from and writing to native devices in the same path set.

Even though PowerPath Migration Enabler does not require the user to specify pseudo devices in the application configurations, it is a best practice to do so. The [Overview of PowerPath Migration Enabler Host Copy steps](#) section covers the migration steps for pseudo and native devices. Customers using native devices will require downtime to manually redirect the application to the target, whereas using PowerPath pseudo devices is non-disruptive because no manual redirection is needed.

## Overview of PowerPath Migration Enabler Host Copy steps

PowerPath Migration Enabler has a common set of processes and procedures for all supported technology types: Host Copy, Open Replicator, TimeFinder®/Clone, and Encapsulation. This paper focuses on Host Copy. For information about the other technology types, see the *PowerPath Family CLI and System Messages Reference Guide*. The following list shows the high-level steps for a migration when the source device is accessed with a PowerPath pseudo device name:

1. Install PowerPath.
2. Install the PowerPath Migration Enabler license.
3. Zone and mask the target LUN to the host.
4. Identify the source and target LUNs for the migration.
5. Set up the migration. Use the **powermig setup** command.
6. Start bulk data copying and cloning of application writes. Use the **powermig sync** command.
7. Query the migration status. Use the **powermig query** command.
8. Pause/resume the migration as needed. Use the **powermig pause/resume** command. (Pause and resume are not required. For a description of pause and resume, see the [Migration speed and host impact](#) section.)
9. After the bulk data copy completes, test the result by reading from the target using the **powermig selectTarget** command.
10. Commit the migration. Use the **powermig commit** command.

11. Clean up the migration. Use the **powermig cleanup** command.
12. Remove the source device.
13. Reconfigure any replication technologies to point at the target device if necessary (for example, SRDF<sup>®</sup> software).
14. Update any Storage Management Resource software if necessary (for example, Solutions Enabler or EMC Ionix<sup>™</sup> software).

## PowerPath Migration Enabler state diagram

Figure 4 depicts the PowerPath Migration Enabler state process. The shaded areas represent the state of the migration. The other text represents the **powermig** CLI commands.

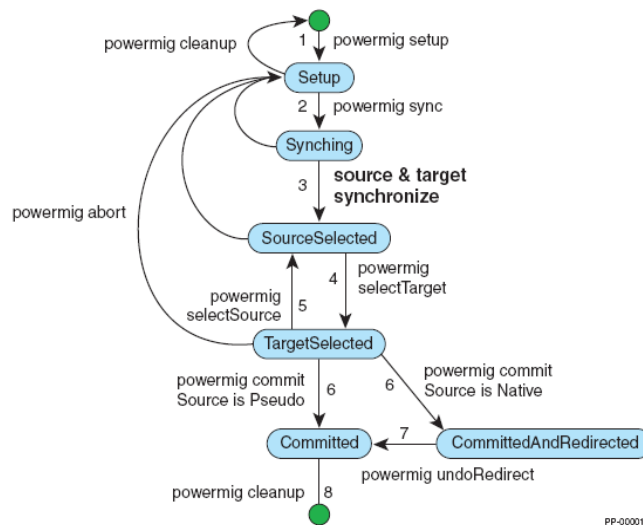


Figure 4. PowerPath Migration Enabler state diagram

## PowerPath Migration Enabler commands

PowerPath Migration Enabler uses a set of commands for all supported technology types, such as Host Copy, Open Replicator, TimeFinder/Clone, and Encapsulation. The use and meaning of these commands vary depending on which technology type is used. Table 2 provides a list and description of the PowerPath Migration Enabler commands when used with Host Copy.

Table 2. PowerPath Migration Enabler commands

<i>powermig</i> ⟨operation⟩	Description
<i>setup</i>	Specifies the source and target devices and technology type used in the migration
<i>sync</i>	Starts the bulk data copy and write cloning process while the application continues uninterrupted
<i>info</i>	Displays information about one or all migration sessions

<i>query</i>	Displays information about one migration session
<i>pause</i>	Suspends a migration in the Synching state
<i>resume</i>	Restarts a paused migration
<i>throttle</i>	Changes the bulk data rate for a migration
<i>abort</i>	Aborts a migration and returns the migration session to the Setup state
<i>selectSource</i>	Designates the source logical unit as the recipient of all I/O requests
<i>selectTarget</i>	Designates the target logical unit as the recipient of all I/O requests
<i>commit</i>	Commits the migration by permanently designating the target as the recipient of all I/O requests. The source device is no longer synchronized with the target
<i>undoRedirect</i>	Stops the redirection of I/O from the source to the target. Use this command only when the source logical unit is a native device and when application I/O is stopped
<i>cleanup</i>	Removes selected data on the source or target logical unit so that the abandoned device is left in a safe state
<i>getHandle</i>	Retrieves the handle for a migration session in progress
<i>version</i>	Displays the version of PowerPath Migration Enabler running on the host

The *PowerPath Family CLI and System Messages Reference Guide* provides an explanation of how these commands are used with the other technology types.

## Migration process

The two primary components of PowerPath Migration Enabler are the user space and the kernel space modules. PowerPath Migration Enabler executes CLI commands (**powermig**) in the user space. The movement of data occurs in the kernel space, but the migration is controlled by the user space. The kernel performs several key functions including write cloning, bulk data copying, access control to a LUN, and I/O redirection when necessary. The following sections cover the process from start to finish of a Host Copy migration.

### Pre-migration

The pre-migration process includes steps 1-4 from the [Overview of PowerPath Migration Enabler Host Copy steps](#) section. Prior to a migration, the administrator must provision one or more additional LUNs to the host. Those LUNs will be the target devices for the migrations. The new devices will be masked to the host like any other LUN. An existing array or a new array must be accessible by the host on the SAN. PowerPath Migration Enabler Host Copy can migrate data from a source to a target LUN on the same or different array whether it is an EMC Symmetrix<sup>®</sup>, VNX<sup>™</sup>, CLARiiON, or third-party array that PowerPath supports. However, the LUNs have to be devices that PowerPath manages. It is not necessary to format or have a file system present on the target device. When the migration completes, the target device will assume the attributes of the source device.

---

Note: There are certain cases in which it is necessary to label a Solaris target device before the Setup state. Refer to the *PowerPath Migration Enabler User Guide*.

---

## Setup

This is step 5 from the [Overview of PowerPath Migration Enabler Host Copy steps](#) section. To get to the Setup state, the administrator specifies the source and target pair used in the migration. PowerPath Migration Enabler issues a unique handle to the migration. The handle range is 1 to 30,000. Handles start at 1 and are incremented by 1 for each new migration session. Multiple migrations can be set up.

In this stage, the administrator:

- Specifies the source and target devices.
- Specifies the technology type (for example, **techType=hostcopy**).
- Sets throttle and suspend time (if applicable) values to configure the migration speed.
- Cleans up the migration. Before any data is moved, the migration can be aborted. The handle number is removed and the migration no longer exists. The same handle number will not be used again even if the same source and target are paired for another migration.

PowerPath Migration Enabler runs checks or identifies device attributes in this stage:

- It checks the devices for accessibility. For example, in Windows, it checks any device that is not in an offline state. The target device must be available to Host Copy and is available only in the online state.
- It makes sure that the devices are not part of any other migration. Only one source can be paired with one target.
- It makes sure that the target device is at least the same size as the source. It can be the same or larger, but it cannot be smaller. In the case of PowerPath Encryption with RSA, if the source device is plain-text and the target device is encrypted, then the target must be at least 65 KB larger to accommodate additional metadata.
- It checks devices to see if they are virtually provisioned devices.

---

Note: By default, the throttle value is set to 2 and the suspend time value is set to 5. These will be discussed in more detail in the [Migration speed and host impact](#) section.

---

Once the setup is completed, the migration is initialized and ready for the next stage.

## Synchronizing

This is step 6 from the [Overview of PowerPath Migration Enabler Host Copy steps](#) section. Bulk data copying for the migration starts when the administrator executes the **powermig sync** command. Migrations start on a per-handle basis. When moving to this state, the user space module tells the kernel space module to begin two processes:

- To initiate the cloning of application writes. Any new I/O write to the source device is automatically cloned and written to the target device.
- To initiate the bulk data copy. Data is read from the source device and then written to the target device. For Host Copy, a daemon in UNIX and a service in Windows manage the bulk data copy. They initiate periodic calls to the kernel to perform copy operations, keep its own log, and regularly checkpoint bulk data copy progress.

In this state, all I/O reads are directed to the source device while I/O writes are cloned to the source and target. The target devices are not accessible to any other application.

During the synchronization stage, the administrator can take the following actions to modify parameters, obtain status information, or cease migration operations:

- Query one or more migrations for percentage completion and throttle value.
- Pause the migration. Because Host Copy does not use an underlying technology like Open Replicator or TimeFinder/Clone, the bulk data copy is host-based. Administrators can temporarily pause the migration during peak I/O times. In this state, new application I/O writes are still cloned to the target. However, the bulk data copies are stopped.
- Resume the migration. Once restarted, the migration will continue from where it left off.
- Change the bulk data copy rate. The throttle and suspend time values can be updated at any time during the Synching stage to speed up or slow down the rate at which bulk data copies occur. This operation is performed on a per-migration basis based on the handle number.
- Abort the migration. The migration is returned to the Setup state. New application I/O writes are no longer cloned and bulk data copy stops. The bulk data copy must be restarted from the beginning.

### **Bulk data copy operation**

During the synchronization stage, PowerPath Migration Enabler Host Copy reads from the source device and writes to the target device in chunks of 256 KB.

### **Increasing or decreasing the number of synchronizing migration sessions**

An administrator can use Host Copy to migrate as many LUNs as needed. A user can put as many migrations into the Synching state as needed. At any given moment, eight migrations can perform bulk data copies simultaneously. If more than eight migrations are in the Synching state, copying will progress on all of them. Completion times will slightly vary.

Since copying uses host CPU cycles and I/O bandwidth, the host may need those resources for applications.

---

Note: Be careful not to impact host performance when synchronizing multiple migrations. If an application appears to be affected after a migration session starts, one or more migrations can be paused or the throttle values can be modified.

---

### Source Selected

Immediately following the completion of bulk data copy operations from the source to the target devices, the migration automatically enters the Source Selected state. The migration is not yet committed. This state has the following characteristics:

- Bulk data copy is complete.
- I/O reads are directed to the source.
- I/O writes are written to the source and cloned to the target.
- Source and target remain synchronized.

The administrator can take one of the following actions:

- Move to the Target Selected state (explained in the [Target Selected](#) section).
- Remain in this state for an extended period of time.
- Query one or more migrations.
- Abort the migration. The migration is returned to the Setup state. The target device is not accessible to applications.

### Target Selected

This is step 9 from the [Overview of PowerPath Migration Enabler Host Copy steps](#) section. The administrator can transition to the Target Selected state from the Source Selected state. This state has the following characteristics:

- Bulk data copy is complete.
- I/O reads are directed to the target.
- I/O writes are written to the target and cloned to the source.
- Source and target remain synchronized.

The administrator can take one of the following actions:

- Return to the Source Selected state. The administrator has the ability to cycle back and forth between the Source Selected and Target Selected states.

The administrator can allow “testing out” reads from the new LUNs before committing to the migration. This feature provides flexibility to the administrator to determine that data is intact and service levels are in line with application requirements on the new device.

- Remain in this state for an extended period of time.
- Query one or more migrations.

- Abort the migration. The migration is returned to the Setup state. The target device is not accessible to applications.
- Commit the migration.

## Commit

This is step 10 from the [Overview of PowerPath Migration Enabler Host Copy steps](#) section. The **commit** command can lead to two different states depending on whether the source is a pseudo or native device. Once the commit state is entered, Migration Enabler cannot return to the source device. Administrators must be sure that moving to the target device is desired.

When Migration Enabler uses pseudo devices as a source, the **commit** command puts the migration into the Committed state. This state has the following characteristics:

- I/O reads and writes are directed to the target.
- Writes are no longer sent to the source, which means that the source and target devices are no longer synchronized.
- The source device is not accessible to applications.
- Device personalities are swapped. When pseudo devices are used in the migration, the underlying hardware paths are changed; that is, the application continues to see the original PowerPath pseudo device name, but the underlying hardware path has changed to the target device. This is a key benefit to using PowerPath pseudo devices in the migration because no disruption to the application occurs.

When Migration Enabler uses native operating system devices as the source, the **commit** command puts the migration into the CommittedAndRedirected state. This state has the following characteristics:

- I/O reads and writes are redirected to the target.
- Writes are no longer sent to the source, which means that the source and target devices are no longer synchronized.
- Target device is not accessible to applications.
- To proceed to the next step, the administrator must manually stop the application and reconfigure it to point to the new target device name. If the source device is native, this is a necessary step. The administrator can proceed after issuing the **undoRedirect** command. The administrator will need to stop the application, execute the **undoRedirect** command, reconfigure the application to point to the new device, and restart the application. This command puts the migration session in the Commit state. It is identical to the pseudo device name migration except that the names are not swapped.

## Cleanup

This is step 11 from the [Overview of PowerPath Migration Enabler Host Copy steps](#) section. The **cleanup** command is the final part of the migration process. Executing this command produces the following:

- It removes the handle.
- After a migration has completed between two LUNs exposed to a host, they will have nearly identical data. This could confuse the host OS or applications. In order to avoid this, some part of the source device will be erased to avoid any confusion. The location and amount of the data erased is dependent on the host OS.

Table 3 shows how each operating system performs different actions during the cleanup process.

**Table 3. Cleanup process on source devices**

Operating system	Cleanup action on source device
Windows	Cleanup changes the disk signature and modifies the MBR partition table to make the physical disk appear as an empty disk.
Solaris	Cleanup writes zeros to blocks 1 through 34 to remove the EFI/GPT primary label. Zeros are written to the last track in the last cylinder to remove VTOC and EFI backup labels. A new label is installed using <b>powerformat</b> . When the source device is part of a VxVM volume, <b>vxdisk destroy</b> is run as well. The <b>-format</b> option on the <b>powermig cleanup</b> command does write zeros to the entire disk.
Linux	Cleanup erases the label on an LVM disk device so that the LVM will no longer recognize it as a physical volume. For VxVM, the disk is uninitialized. Sector 0 is zeroed out (writing zero filled buffers).
AIX	Cleanup zeros out the LVM and VxVM sectors. The PVID/VIID on block/sector zero are zeroed out and then transfer the PVID copy in the AIX ODM from the source to the target.
HP-UX	Cleanup zeros out the LVM and VxVM sectors.

---

Note: The **-format** option is not supported with encrypted devices.

---

---

Note: Do not consider the operations performed during the cleanup state a full data scrubbing. EMC Certified Data Erasure overwrites physical storage data with a pattern of random data in one or more iterations to render the underlying data unreadable. EMC uses tools that are custom-developed for the storage platform to be erased in the customer's facility. The erasure can require three to seven overwrites of the existing data with a combination of the random patterns in alignment with DoD 5220.22-M recommendations for clearing the rigid magnetic media.

---

## Post-migration

After migration is complete, the administrator can remove the source disk from the host.

Any replication technologies used to mirror the source device must be configured to replicate the new target device.

Storage Resource Management tools such as Solutions Enabler or Ionix should be updated.

---

Note: Array-based replication can be set up prior to configuring the target device to the host.

---

## Migrating clusters

In a cluster configuration, the administrator must select one node to remain active during migration. Only one node with access to the source and target can remain active. All failover groups must be moved to the active node.

Remember that PowerPath Migration Enabler Host Copy uses host resources to do bulk data copies from source to target and when cloning application writes. Administrators must determine how fast to run the migration based on throttle options within Host Copy. Since the standby nodes will not be available, the cluster is not as highly available. An appropriate time for performing the migration must be determined.

The *PowerPath Migration Enabler User Guide* contains detailed instructions on migrating devices in a cluster.

## Graceful host reboot

PowerPath Migration Enabler Host Copy uses a flag to indicate that a host was gracefully shut down. It also records a checkpoint in the Migration Enabler database where all records of the migration pair and current state are held. The migration progress checkpoint is recorded approximately once every 30 seconds on a local or SAN disk depending on where the boot disk is located. There is no Host Copy metadata on the source or target device. If the host is gracefully shut down, then the migration can resume from its last checkpoint when the host starts up. This behavior is for migrations that are synchronizing or paused. There is no impact to the migration when the host is gracefully shut down.

## Host crash

PowerPath Migration Enabler keeps the data at least as safe as when no migration is in process. Since the success of application writes cannot be guaranteed when a host crashes, PowerPath Migration Enabler cannot guarantee that the data on the source and target devices is the same after such a crash. The fault state (either TargetLUfault or SourceLUfault) appears in the **powermig info** output. PowerPath Migration Enabler reacts according to the state of the migration at the time of the

crash. Faults are recorded in the Synching, Source Selected, and Target Selected states.

Table 4 shows the three states when Logical Unit faults can be reported. Each state — Synching, Source Selected, and Target Selected — will cause a different result when the host is rebooted after the crash:

- In the **Synching** state, synchronization has stopped. PowerPath Migration Enabler will report a TargetLUfault. Migration Enabler returns the migration to the Setup state. Application reads and writes will still be directed to the source device. The migration can be restarted with the **powermig sync** command.
- In the **Source Selected** state, the migration has completed. Both source and target devices are fully synchronized. However, when the host crashes, PowerPath Migration Enabler will report a TargetLUfault. I/O writes may have been in flight during the time of the crash. Therefore, Host Copy can no longer guarantee that the data is the same on both the source and the target. The migration will return to the Setup state and can be restarted with the **powermig sync** command.
- In the **Target Selected** state, the migration has completed. Both source and target devices are fully synchronized. However, when the host crashes, PowerPath Migration Enabler will report a SourceLUfault. I/O writes may have been in flight during the time of the crash. Therefore, Host Copy can no longer guarantee that the data is the same on both the source and the target. The migration will proceed to the Commit state.

**Table 4. Reported faults after a host crash**

State at time of host crash	Host crash	Fault state reported	State after crash		Action ( <i>powermig</i> command)	Data loss
Synching	✓	TargetLUfault	Returns to Setup state		<i>sync or cleanup</i>	No
SourceSelected	✓	TargetLUfault	Returns to Setup state		<i>sync or cleanup</i>	No
TargetSelected	✓	SourceLUfault	pseudo	Moves to Committed state	<i>cleanup</i>	No
TargetSelected	✓	SourceLUfault	native	Moves to Committed and Redirected state	<i>undoRedirect</i>	No

When a host crashes, PowerPath Migration Enabler returns errors to both reads and writes for any device that is in a migration where the source and target are kept synchronized until PowerPath Migration Enabler startup recover is run. This protects data integrity until any device fault states that were recorded but not processed prior

to the crash are processed. Once startup recovery is run, it is safe for reads and writes to flow to these devices.

### Other reported faults

In addition to a host crash, it is possible that errors or faults can occur due to changes in the storage, host, or SAN. Table 5 shows the three states when logical unit faults are reported. If an I/O failure occurs in the course of a write I/O operation and if Migration Enabler is currently cloning write I/Os, then Migration Enabler records a logical unit fault, and the fault state (either TargetLUfault or SourceLUfault) appears in the **powermig info** output. Migration Enabler reports different types of faults depending on the state of the migration at the time the I/O failure occurs. Faults are recorded in the Synching, Source Selected, and Target Selected states.

**Table 5. Other reported faults**

Current state	Source device failure	Target device failure	Fault state reported	Cause	Action ( <i>powermig</i> command)	Data loss
Synching	✓		SourceLUfault	I/O write to source fails	<i>abort</i> <sup>1</sup>	No
		✓	TargetLUfault	I/O write to target fails	<i>abort</i>	No
SourceSelected	✓		SourceLUfault	I/O write to source fails	<i>selectTarget</i> <sup>1</sup>	No
		✓	TargetLUfault	I/O write to target fails	<i>abort</i>	No
TargetSelected	✓		SourceLUfault	I/O write to source fails	<i>commit</i>	No
		✓	TargetLUfault	I/O write to target fails	<i>abort</i> <sup>1</sup>	No

<sup>1</sup> When this fault occurs under these conditions, errors will be returned on all I/Os until the administrator performs this transition.

### Migrating to a virtually provisioned target

EMC Virtual Provisioning™ addresses storage administrator challenges with capacity utilization, management simplification, application disruption, and unused storage. It builds on the basic thin provisioning functionality, which is the ability to have a large “thin” device (that is, volume) configured and presented to the host while consuming physical storage from a shared pool only as needed. Symmetrix, VNX, and

CLARiiON Virtual Provisioning can improve storage capacity utilization and simplify storage management by providing sufficient capacity to an application for an extended period of time, which reduces the need to provision new storage frequently and avoids costly allocated but unused storage. This technology is appropriate for most storage tiers.

For Symmetrix, a thin device extent is 768 KB (or 12 tracks) in size and is the default unit of storage at which allocations occur. Therefore, aligned 768 KB chunks of zeros are needed to reclaim space. For VNX and CLARiiON, aligned 256 KB chunks of zeros are needed to reclaim space using Host Copy. (VNX and CLARiiON allocate 8 KB chunks.)

PowerPath Migration Enabler Host Copy facilitates the movement to virtually provisioned devices without disruption to the application. There are only a few differences between a thick to thin device migration and a standard migration.

In a Host Copy migration, the Setup state identifies target devices as either thick or thin. For thin target devices when the copy operation begins, PowerPath Migration Enabler copies only those blocks with data (Figure 6). It does not write blocks that are all zeros and part of an aligned chunk (where all constituent blocks are all zeros) as long as the target device is all zeros for that chunk as well. If one byte of any block in the chunk is non-zero, then Migration Enabler copies every block in the chunk. Administrators now have the ability to reclaim all-zero space including both host-unwritten extents/chunks and chunks that contain all zeros due to file system or database formatting methods. The following figure shows the basics of space reclamation on the target device.



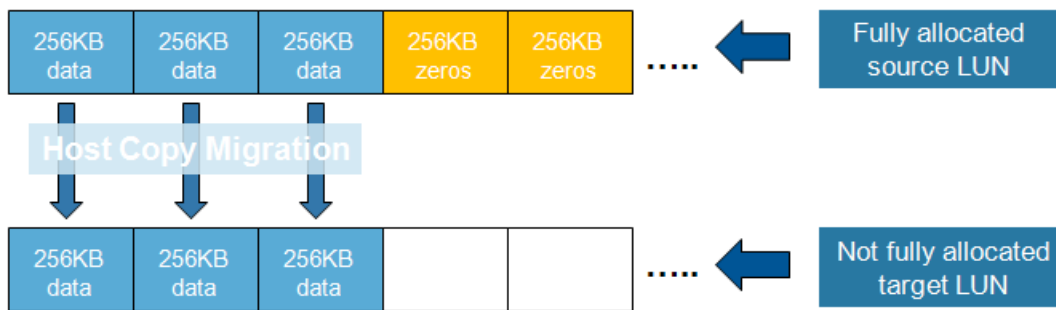
Figure 5. A logical representation of space reclamation

Host Copy will operate in this thin copying mode only if it recognizes the target device as a virtually provisioned device. It does this by checking an array-specific SCSI-mode page value. This value is not based on a standard. The array-specific SCSI-mode page value allows Host Copy to determine Symmetrix, VNX, or CLARiiON virtual LUNs. Host Copy cannot determine whether or not a LUN is thin on a third-party array.

---

Note: If the administrator wants to do a normal copy even though the target is thin, Host Copy has a “no thin option.” Choosing this option removes the overhead of checking for thin devices. The CLI command uses the **-nothin** operand.

---



**Figure 6. Migrating to a virtually provisioned target**

Note: Space efficiency of going thick to thin is strictly dependent on the amount of contiguous zero blocks in the thick LUN. This might require external tools to maximize the zero blocks. As a best practice, use a newly created target device. If the target is a repurposed device that already has allocated regions, less storage may be recovered when performing a thick to thin migration. With the latest Symmetrix Enginuity™ service release, administrators now have the ability to reclaim all-zero space, including both host-unwritten extents/chunks and chunks that contain all zeros due to file system or database formatting methods. The space reclamation process is non-disruptive and can be executed with the targeted thin device ready and able to be read and written to operating systems and applications.

## Migration speed and host impact

PowerPath Migration Enabler Host Copy is entirely host-based. When the migration enters the Syncing state, Migration Enabler Host Copy clones application writes to the source and target devices. At the same time, it reads the data from the source device and then writes it to the target. EMC has noted that the cloning of application writes has not produced any measureable overhead on the host. However, the bulk data copy can negatively impact the host. Multiple factors can contribute to host performance when using Host Copy:

- Host hardware resources
- Competing loads on host resources
- SAN configuration
- HBA speed
- Storage array configuration
- Number of simultaneous migrations
- Throttle setting
- Suspend time setting

Each host can expect to operate differently depending on all these factors. Host Copy has several options for correcting performance issues that can arise during the migration:

- Use **powermig pause** and **resume** capability.

After a migration starts, it is unlikely that application users will see an impact to performance. (See the [Understanding throttle impact](#) section for more information on the potential impact to an application.) However, if multiple Host Copy migrations are running in parallel on a server with minimal resources, it is possible that users may see a slowdown in host performance. The administrator can pause one or more of the migrations for an indefinite period of time. The Host Copy checkpoint feature allows the migration to resume from where it left off when the administrator restarts it.

---

Note: A checkpoint is always recorded at the time a pause is executed.

---

- Use the **powermig throttle –throttlevalue** command.

When a migration is in the Setup and Syncing states, the administrator can modify the throttle value to increase or decrease the rate at which the bulk data copy is progressing. The throttle value has different meanings depending on the PowerPath Migration Enabler technology being used. Here, only the Host Copy throttle is discussed.

- Use the **powermig options -hostcopy\_ceiling <value>** command.

When a migration is in the Setup and Syncing states, the administrator can modify the **powermig options –hostcopy\_ceiling** value to specify an upper limit for the aggregate rate of copying for all Host Copy migrations. This is a host-wide setting that overrides individual **powermig throttle** settings. It also allows more than eight migrations to perform bulk data copy operations simultaneously.

- Use the **powermig throttle –suspendtime** command.

When a migration is in the Setup and Syncing states, the administrator can modify the suspend time value to increase or decrease the amount of time that application write I/O is held while bulk data copy occurs.

---

Note: The suspend time setting will eventually be phased out on all platforms supporting Host Copy.

---

## Understanding throttle impact

PowerPath Migration Enabler Host Copy has different copying characteristics depending on how recently PowerPath was released on each OS platform. At this time, [PowerPath for AIX](#) uses the original method. The [PowerPath versions for Windows, Linux, Solaris, and HP-UX](#) use the new method.

## AIX

Host Copy copies data from the source to the target in cycles. Each migration pair has a suspend and a throttle time. These settings operate independent of each other. Modification to one or both will affect migration completion time.

### Suspend time

Suspend time is the length of time each unit of copying can last before resuming I/O writes. While I/O is suspended, Host Copy copies as much data as possible during that time. Copying for an individual migration occurs only while I/O writes are suspended. Changing the suspend time should be used with applications that are sensitive to held writes. EMC does not recommend setting the suspend time to more than 60 seconds.

---

Note: The suspend time value can be set above 60 seconds if the application is not running.

---

---

Note: Suspend time is scheduled to be phased out for AIX with PowerPath 5.5.

---

### Throttle

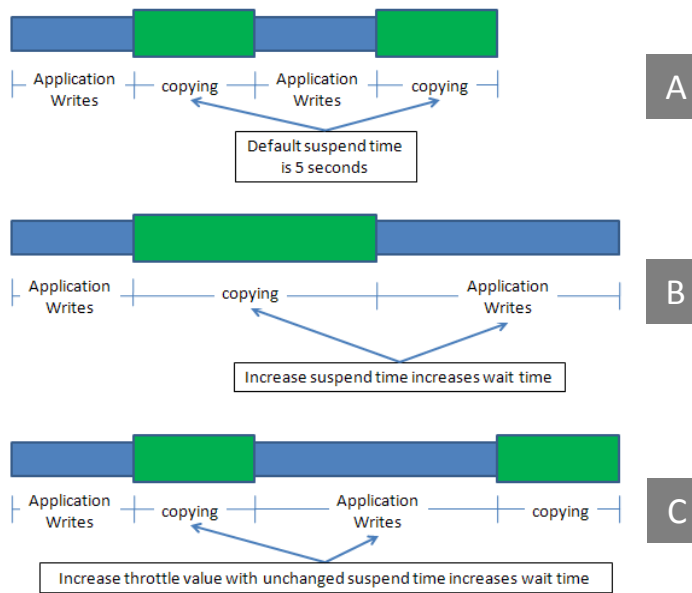
Throttle sets the time to wait between suspends. Once I/O writes resume (after the suspend time ends), copying stops and the migration waits for the amount of time determined by the throttle value. At the end of the throttle value time, I/O writes are again suspended and copying resumes for the duration of the suspend time. This cycle continually repeats itself until the bulk data copy completes. EMC recommends modifying the throttle value before making changes to the suspend time to maximize performance.

---

Note: If there are more than eight migrations in the Syncing state, it can take longer to resume copying.

---

Figure 7 shows the interaction between throttle and suspend.



**Figure 7. Throttle and suspend time interaction**

For example in Figure 7, **A** represents the default suspend time value of five seconds. **B** shows the impact of doubling the suspend time value. This has the effect of doubling the wait time between copy sessions even though the throttle value was not changed. In essence, roughly the same amount of data will be copied. Change the suspend time when the sensitivity of the application to write latency is important.

**C** shows a change only in the wait time. An increase to the throttle value increases the amount of time between the copy sessions which, in turn, increases the time to complete the migration. If there is a performance impact to the application when one or more migrations are synchronizing, begin by modifying the throttle value before changing the suspend time.

The copying rate with the throttle mechanism can be roughly translated using the values in Figure 8. 2 is the default throttle setting for Host Copy.

Throttle Value	0	1	2	3	4	5	6	7	8	9
Percentage of time the host spends copying data	100%	60%	36%	22%	13%	7.8%	4.7%	2.8%	1.7%	1.0%

**Figure 8. Host Copy throttle values per migration**

Ultimately, it is the combination of suspend time settings and throttle values along with host resources, competing workloads, and array characteristics that determine how long a migration will take to complete. Because every environment is different, PowerPath Migration Enabler Host Copy provides the tools for administrators to manage and monitor each migration.

## Windows, Linux, Solaris, and HP-UX

In the latest versions of PowerPath for Windows, Linux, Solaris, and HP-UX, PowerPath Migration Enabler Host Copy no longer uses suspend time; that is, the description provided in Figure 7 is not relevant to these platforms. However, the throttle values described in Figure 8 are unchanged.

In the latest versions, bulk data copy operations occur simultaneously with application I/O writes except in the occasional scenario where a write occurs within the 256 KB region being copied. In this rare case, the application write would be temporarily delayed until the copy of that 256 KB region completes.

In the latest versions, setting the throttle value is the only tool the administrator can use to modify the copy rate. The latest versions no longer use the wait time between copy cycles as seen in Figure 7. Increasing or decreasing the throttle value will simply increase or decrease the amount of time spent copying data within each Host Copy copying cycle, based on the copying rate shown in Figure 8. The new method is an improvement because the impact of the migration on I/O load on the host is smoother. Figure 9 shows the individual copy cycles running in parallel with application write I/O. The white stripes on the bottom (copying) show that the waiting periods to implement the throttle are finer-grained. As the throttle is increased, the width of those stripes increases to indicate the percentage of time spent copying.




Figure 9. Copy cycles without suspend time in the new design

## Thick to thin migration example

The following section details the steps and command outputs associated with a Host Copy migration from a CLARiiON CX4 device to a Symmetrix VMAX™ virtually provisioned device. In this example, a migration occurs from a thick source device to a thin target device on a Red Hat 5 server.

### Environment

Host	RHEL5 U3
PowerPath	5.3 SP1 (a.k.a. 5.3.1)
Source array	CLARiiON CX4 (FLARE® 29)
Source device	20 GB, thick LUN (LUN22)
Target array	Symmetrix VMAX (5774 ucode)
Target device	20 GB, thin LUN (01FA)

The PowerPath CLI provides a list of devices that PowerPath manages, as well as device-specific information PowerPath assigns like the pseudo devices, array serial number, and logical device ID. Two devices are visible to this host: a CLARiiON CX4 and a Symmetrix VMAX LUN. The **powermt display dev=all** command displays the list. The shaded area in Figure 10 provides the device-specific information for this migration. **emcpowerd** and **emcpowera** represent the PowerPath names for the source and target devices and not the devices themselves.

```
[root@lppa020 ~]# powermt display dev=all
```

Pseudo name=emcpowerd

←

Thin Target Device

```
Symmetrix ID=000194900335
Logical device ID=01FA
state=alive; policy=SymmOpt; priority=0; queued-I/Os=0
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path  I/O Paths  Interf.  Mode  State  Q-I/Os  Errors
=====
  1 qla2xxx   sdf    FA 7gA  active  alive  0    0
  3 qla2xxx   sdk    FA 8eA  active  alive  0    0

Pseudo name=emcpowera
```

Pseudo name=emcpowera

←

Thick Source Device

```
CLARiiON ID=FNM00092000177 [lppa020]
Logical device ID=6006016038E02400E64F749233AADE11 [LUN 22]
state=alive; policy=CLAROpt; priority=0; queued-I/Os=0
Owner: default=SP A, current=SP A   Array failover mode: 4
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path  I/O Paths  Interf.  Mode  State  Q-I/Os  Errors
=====
  1 qla2xxx   sdb    SP A1  active  alive  0    0
  1 qla2xxx   sdc    SP B0  active  alive  0    0
  3 qla2xxx   sdh    SP A0  active  alive  0    0
  3 qla2xxx   sdi    SP B1  active  alive  0    0
```

**Figure 10. Output of the *powermt display* command**

The output of **powermt display** monitors HBAs and devices. Path state, load-balancing policy, device priority, and more are shown. An administrator using an operating system management tool or PowerPath GUI and CLI cannot identify whether a device is virtually provisioned or not. (During the Setup state of a Host Copy migration, PowerPath does recognize a virtually provisioned device. However, it is not displayed to the user.) For this information, Symmetrix Management Console can be used to view device properties.

Figure 11 shows a thin pool called “lppa020\_thin” with one data device and two thin devices.

Note: Solutions Enabler CLI commands could also be used to view the same data. Refer to the *EMC Solutions Enabler Symmetrix Array Management CLI Product Guide*.

Note: For simplicity, only one data device with two thin devices is shown. Symmetrix best practices recommend having multiple data devices in the thin pool.

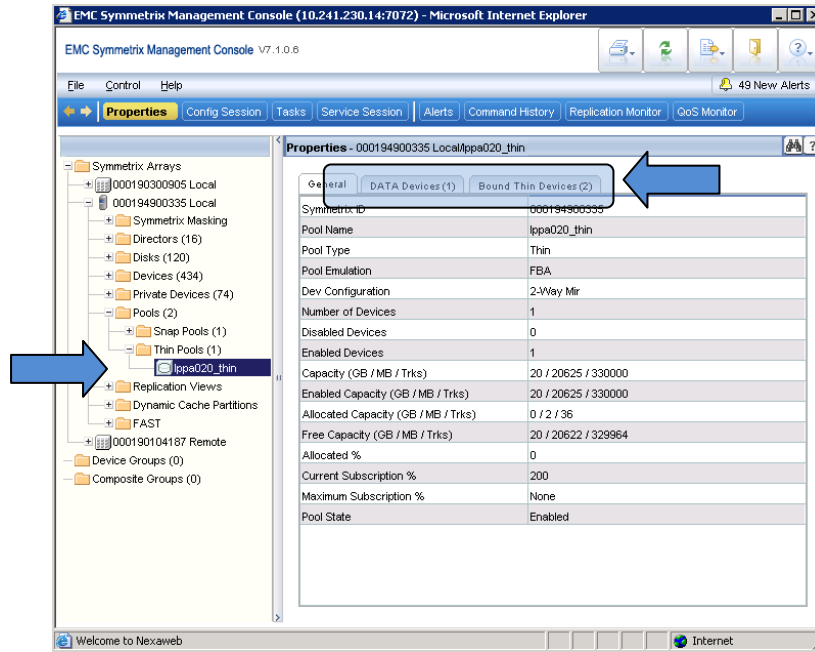


Figure 11. Symmetrix Management Console thin pool display

Figure 12 displays the properties for the Bound Thin Devices and DATA Devices. Two 20 GB bound thin devices exist. Prior to the migration with newly created devices, they are 0% allocated (except for metadata requirements on one thin device). Figure 12 also shows one 20 GB data device. With two 20 GB thin devices and one 20 GB data device, this thin pool is oversubscribed by 200%.

General		DATA Devices (1)		Bound Thin Devices (2)		
Dev	Pool	Cap (MB)	Subscribed %	Alloc Cap (MB)		
01FA	lppa020_thin	20625	100	1		
01FB	lppa020_thin	20625	100	0		

General		DATA Devices (1)		Bound Thin Devices (2)		
Dev	Emulation	Dev Config	State	Free (MB)	% Used	
01F4	FBA	2-Way Mir	Enabled	20623	0	

Figure 12. Properties of thin devices

As previously seen in Figure 10, both source and target devices were assigned EMC pseudo devices. The CX4 device is **emcpowera**, and the VMAX device is **emcpowerd**. As discussed earlier, these are merely the names for the devices and not the actual physical devices. Using the CX4 device, a file system has been created on this device and mounted to the /temp1 subdirectory. (This is our source device for the migration.) One 2 GB file exists in this subdirectory. As a result the file system is only 10% utilized. Executing **df -k** and **ls -l** commands on the RHEL5 server shows currently mounted file systems and directory files, respectively (Figure 13).

```
[root@lppa020 ~]# df -k
Filesystem                1K-blocks  Used   Available Use%   Mounted on
/dev/mapper/VolGroup00-LogVol 62785344 4478528 55066064  8%    /
/dev/sda1                  101086    16461   79406    18%   /boot
tmpfs                      1958372    0       1958372  0%    /dev/shm
/dev/emcpowera             20642428 1916704 17677148 10%   /temp1

[root@lppa020 temp1]# ls -l
total 1871728
-rw-r--r-- 1 root root 1914756751 Dec 18 11:22 testfile.zip
drwx----- 2 root root 16384 Dec 17 17:53 lost+found
```

**Figure 13. Source device file system information**

With the source and target devices identified, the migration can now be set up using the PowerPath Migration Enabler CLI commands. The **powermig setup** command specifies the source and target devices as well as the technology type. Handle **3** has been assigned to this migration.

```
[root@lppa020 ~]# powermig setup -src emcpowera -tgt emcpowerd -tectype hostcopy
Setup migration? [yes]/no: y
Migration Handle = 3
```

The **powermig info** command displays all migrations no matter what state they are in. In the output below, the handle is **3** and the state is Setup.

```
[root@lppa020 ~]# powermig info -all
=====
Hnd Source Target Tech State
=== =====
3 emcpowera emcpowerd HostCopy setup
```

The **powermig query** command displays similar information as the “**info**” command. However, **query** is specific to a particular migration.

```
[root@lppa020 ~]# powermig query -handle 3
Handle: 3
Source: emcpowera
Target: emcpowerd
Technology: HostCopy
Migration state: setup
```

The **powermig sync** command starts the bulk data copy and cloned application writes for one migration pair based on the handle.

```
[root@lppa020 ~]# powermig sync -handle 3
Start sync? [yes]/no: y
```

The **powermig query** command can be executed at any time during the migration to check on its status. (This is also true for **powermig info**.) The output below shows new information including the throttle and suspend time values. By default, throttle value is set to **2** and suspend time is set to **5**.

```
[root@lppa020 ~]# powermig query -handle 3
Handle: 3
Source: emcpowera
Target: emcpowerd
Technology: HostCopy
Migration state: syncing
Percent InSync: 10%
Throttle Value: 2
Suspend time: 5
```

The **powermig throttle** command allows the user to speed up or slow down the rate at which data is being copied. In the example below, the throttle value has been changed from **2** to **9**. Throttle value (and suspend time) is set on a per-migration basis and not on a per-host basis.

```
[root@lppa020 ~]# powermig throttle -throttlevale 9 -handle 3
Set throttle option(s)? [yes]/no: y
```

The **powermig pause** command allows the user to temporarily pause the migration. The **query** command displays the new status.

```
[root@lppa020 ~]# powermig pause -handle 3
```

```
Pause migration? [yes]/no: y
```

```
[root@lppa020 ~]# powermig query -handle 3
```

```
Handle: 3
```

```
Source: emcpowera
```

```
Target: emcpowerd
```

```
Technology: HostCopy
```

```
Migration state: syncing
```

```
Percent InSync: 21% [PAUSED]
```

```
Throttle Value: 9
```

```
Suspend time: 5
```

The **powermig resume** command allows the user to restart the migration from where it left off.

```
[root@lppa020 ~]# powermig resume -handle 3
```

```
Resume migration? [yes]/no: y
```

After the bulk data copy completes, the migration automatically moves to the Source Selected state.

```
[root@lppa020 ~]# powermig query -handle 3
```

```
Handle: 3
```

```
Source: emcpowera
```

```
Target: emcpowerd
```

```
Technology: HostCopy
```

```
Migration state: sourceSelected
```

The **powermig selectTarget** command allows the user to move from the Source Selected state to the Target Selected mode. The administrator can switch back and forth between Source Selected and Target Selected states. The **powermig query** displays the new state.

```
[root@lppa020 ~]# powermig selectTarget -handle 3
```

```
Transition to targetSelected state? [yes]/no: y
```

```
[root@lppa020 ~]# powermig query -handle 3
```

```
Handle: 3
```

```
Source: emcpowera
```

```
Target: emcpowerd
```

Technology: HostCopy  
Migration state: targetSelected

The **powermig commit** command ends the migration and stops cloning application writes to the source device. The **powermig query** displays the new state.

```
[root@lppa020 ~]# powermig commit -handle 3  
Commit migration? [yes]/no: y
```

```
[root@lppa020 ~]# powermig query -handle 3  
Handle: 3  
Source: emcpowera  
Target: emcpowerd  
Technology: HostCopy  
Migration state: committed
```

The **powermt display dev=all** command displays the pseudo device number change after executing the **commit** command. Note the changes in the following figure.

Before the migration:

**emcpowera** was the CX4 device.


**emcpowerd** was the virtually provisioned VMAX device.

After the migration:

**emcpowera** is now the virtually provisioned VMAX device.

**emcpowerd** is now the CX4 device.

```
[root@lppa020 ~]# powermt display dev=all
```



```
state=alive; policy=SymmOpt; priority=0; queued-IOs=0
```

```
=====
```

Host	Stor	I/O Path	Stats
### HW Path	I/O Paths	Interf. Mode	State Q-IOs Errors
1 q1a2xxx	sdf	FA 7gA active	alive 0 0

```
=====
```

```

3 qla2xxx    sdk      FA 8eA active alive 0 0
Pseudo name=emcpowerd
CLARiiON ID=FNM00092000177 [lppa020]
Logical device ID=6006016038E02400E64F749233AADE11 [LUN 22]
state=alive; policy=CLAROpt; priority=0; queued-IOs=0
Owner: default=SP A, current=SP A   Array failover mode: 4
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path      I/O Paths  Interf. Mode  State  Q-IOs Errors
=====
1 qla2xxx    sdb      SP A1  active  alive  0  0
1 qla2xxx    sdc      SP B0  active  alive  0  0
3 qla2xxx    sdh      SP A0  active  alive  0  0
3 qla2xxx    sdi      SP B1  active  alive  0  0

```

Figure 14. Output of the `powermt display dev=all` command

The **powermig cleanup** command deletes the migration handle and removes enough information from the former source device to ensure that the operating system or applications do not become confused. The **powermig query** shows that handle 3 no longer exists.

```

[root@lppa020 ~]# powermig cleanup -handle 3
Cleanup migration? [yes]/no: y

[root@lppa020 ~]# powermig query -handle 3
PPME error(6): Handle not found

```

From the point of view of the operating system, the pseudo device, the application, the mount point, or the file is the same. As expected, executing **df -k** and **ls -l** display the same information as before the migration.

```

[root@lppa020 ~]# df -k
Filesystem                1K-blocks  Used    Available Use%    Mounted on
/dev/mapper/VolGroup00-LogVol 62785344 4478528 55066064   8%    /
/dev/sda1                  101086    16461    79406    18%    /boot
tmpfs                      1958372    0        1958372   0%    /dev/shm
/dev/emcpowera             20642428 1916704 17677148  10%    /temp1

[root@lppa020 temp1]# ls -l
total 1871728

```

```

-rw-r--r-- 1 root root 1914756751 Dec 18 11:22 testfile.zip
drwx----- 2 root root 16384 Dec 17 17:53 lost+found

```

The following figure depicts Symmetrix Management Console displaying the new device usage and allocation of the thin pool. The new thin device has not been fully allocated. Only the used blocks were copied to the VMAX virtually provisioned device.

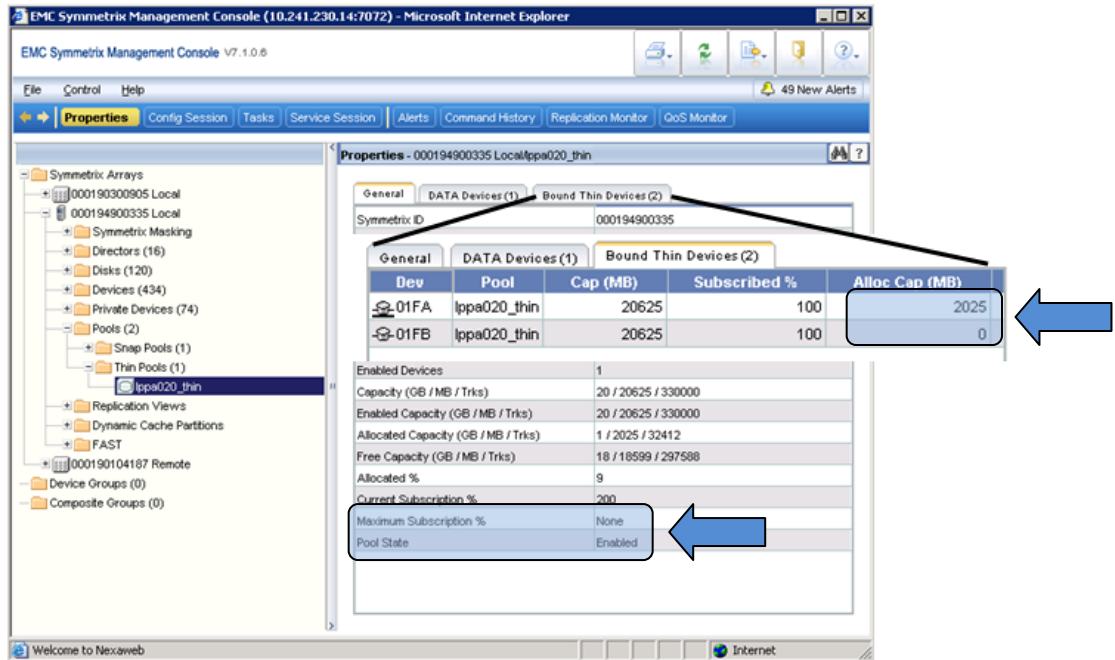


Figure 15. Symmetrix Management Console

## Logging

PowerPath Migration Enabler has two types of logs, [audit](#) and [error](#) logs. Audit messages show migration state transitions and logicalunit fault conditions, and generally record the migration workflow. Error and warning messages record unexpected events. Logging is enabled by default.

### Audit log

The audit log records all events that change the migration state, including:

- Command execution
- Transition to the Source Selected state when a bulk data copy operation completes
- Detection of a logicalunit fault

Here is an example of an audit log message:

*Dec 21 15:00:23 lppa020 emcpAudit: PPME: Info: handle=3, event=stateChanged, state=syncing, cmd=sync*

This log entry indicates that a **powermig sync** command was executed on migration with handle number 3 on December 21<sup>st</sup> at approximately 3 p.m. The state of the migration is now synchronizing, which means that bulk data copy and cloning of application writes have started.

## Error log

The error log messages capture unexpected events that occur during the course of a migration. Some error log messages convey information that also appears on the screen when a **powermig** command fails.

Here is an example of an error log message:

*Dec 21 14:49:16 lppa020 PPME: API: Error: Setup failed: PPME error(14): Devices already involved in a migration*

This log entry indicates that the **powermig setup** command failed because either the source or the target device was already part of another Migration Enabler migration session.

## Log locations

PowerPath reports any errors, diagnostic messages, and failover recovery messages through the syslog file or Event viewer (Windows) that the administrator has specified.

**Table 6. Default log locations**

Platform	Default location
Windows	Application Event Log
Linux	/var/log/messages
AIX	/usr/safe.log
Solaris	/var/adm/messages
HP-UX	/var/adm/syslog/syslog.log

For more details on logging, refer to the *PowerPath Family CLI and Systems Messages Reference Guide*.

## Measuring elapsed migration time

The audit log entries contain information about migration time. However, Host Copy has a separate log file that shows the progress of a migration. With UNIX, an administrator can look in the /etc/emc/ppme/emcpmigd.log file. The timestamps show when a migration begins and when it ends.

For example, to get all messages for handle **8** with source emcpower29a and target emcpower30a, the administrator can run the following command in the file:

```
"grep "8:emcpower29a_emcpower30a" /etc/emc/ppme/emcpmigd.log"
```

Locate the start message:

“2009/12/01 10:01:40 [8:emcpower29a\_emcpower30a] STARTING HostCopy at block 0”

Then look for the copy complete message:

“2009/12/01 10:13:39 [8:emcpower29a\_emcpower30a] copy complete.”

In Windows, open the hostcopy.log file. It is located at the <install directory>\PowerCommon\Log\Logs subdirectory. It will have the same start and complete messages.

## Encryption migration

PowerPath Migration Enabler Host Copy provides data migration capabilities for PowerPath Encryption with RSA. Because data cannot be encrypted in place, Host Copy migrates the plain-text data non-disruptively to an encrypted LUN while keeping the application online (Figure 16).

When migrating from a plain-text device to an encrypted device, the target must be at least 65 KB larger than the plain-text device to accommodate the PowerPath encryption metadata. Host Copy will read from the plain-text source device. PowerPath Encryption with RSA will then encrypt the data. Host Copy completes the operation by writing the data to the target device as cipher-text. Figure 1 on page 8 shows the order of PowerPath extensions.




Figure 16. Migrating to an encrypted device with Host Copy

## Throttle throughput example

Several factors contribute to the speed of a migration. Host Copy has no one-size-fits-all performance metric. As discussed earlier, it is the combination of suspend time settings and throttle values along with host resources, competing workloads, and array characteristics that determine how long a migration will take to complete. Adjusting the throttle is the primary tool for increasing or decreasing the migration rate.

The following configuration is a test environment. In this configuration, one server running PowerPath 5.3 has 20 simultaneous Host Copy sessions in the Syncing state. Each session is migrating a 4 GB LUN from a Symmetrix VMAX to a VMAX. The migrations are thick to thick devices. This test environment has no Virtual Provisioning or encryption. It includes two sets of tests with a total of 20 iterations. The only difference between each test set is the number of HBA ports. Figure 17 depicts the copy rate for each test set.

### Test Set #1:

- Two HBA ports and two Fibre Adapter (FA) ports
- 20 migration sessions with 4 GB VMAX LUNs
- 10 individual tests using throttles 0 through 9
- Figure 17 (square line)

### Test Set #2:

- Four HBA ports and four FA ports
- 20 migration sessions with 4 GB VMAX LUNs
- 10 individual tests using throttles 0 through 9
- Figure 17 (diamond line)

Every customer configuration will have different variables yielding different results. The following test results are only one example.

Operating System	Solaris 10 (SPARC)
PowerPath	5.3 for Solaris
Server	Sun Fire T200
HBA	Emulex LP11002
SAN	Brocade 5300
Source storage	Symmetrix VMAX, 4 GB LUNs
Target storage	Symmetrix VMAX, 4 GB LUNs

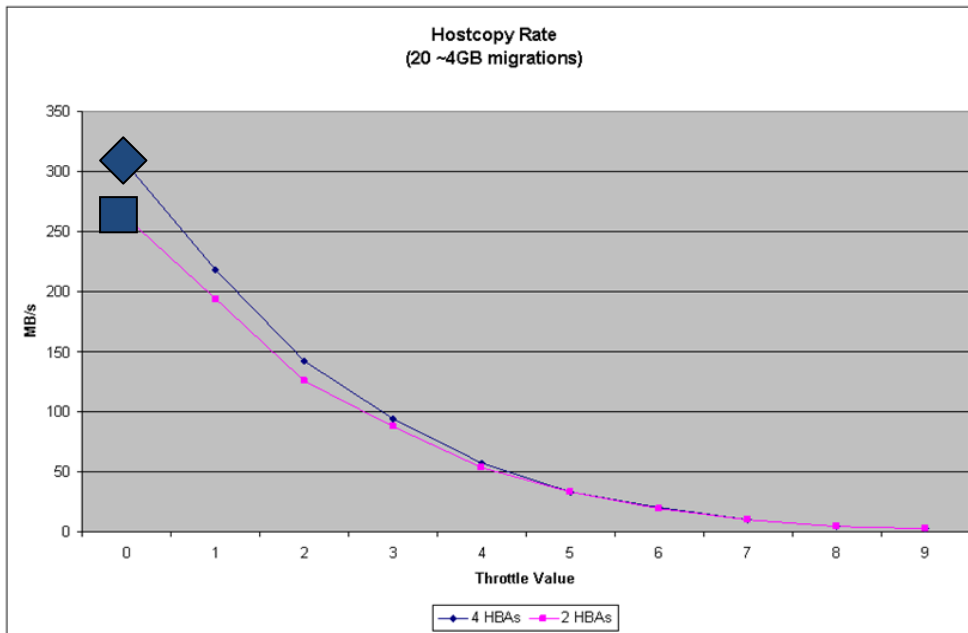


Figure 17. Twenty Host Copy throughput tests using each throttle value

Figure 17 shows the decreasing rate of bulk data copying as higher throttle values are used. As the throttle value increased, the bulk data copy rate decreased. See Figure 8 on page 26 for throttle values. For the two HBA and four HBA test sets, the copy rate converged as the throttle value increased.

In this example, the throttle is unchanged for each individual test. However, an administrator can choose to increase or decrease the copy rate by modifying the throttle value during the Syncing state of the migration.

---

Note: The figure represents the speed at which data is migrated. The actual I/O rate of the migration is double because migrating a block of data requires both a read (from the source device) and a write (to the target device) operation. The chart focuses on bulk data copying. These tests did not include application write I/O. Host Copy throughput may be reduced if I/O-intensive applications are running on the host during a migration. The amount of reduction will vary among environments.

---

## Installing PowerPath Migration Enabler

PowerPath Migration Enabler binaries are components of the PowerPath installation. On UNIX operating systems, a full installation of PowerPath Multipathing and advanced features occurs on new installs and upgrades.

On Windows operating systems, the user has three options in the InstallShield Wizard: Typical, Custom, and Complete. Consult the *EMC PowerPath and PowerPath/VE for Microsoft Windows Installation and Administration Guide* for more information.

Since **Typical Install** is the default on Windows 2003 and Windows 2008 (including Hyper-V), it is possible that PowerPath Migration Enabler binaries are not present. To determine if PowerPath Migration is installed, at a command prompt, type **powermig** and press **<Enter>**.

If PowerPath Migration Enabler commands are not displayed, then it is not installed and the output appears as follows:

```
C:\>powermig
```

```
'powermig' is not recognized as an internal or external command, operable program or batch file.
```

The InstallShield Wizard must be run again. The administrator should choose **Modify**.

After selecting **Next**, the administrator can now choose the components to be installed. The components are PowerPath Migration Enabler, PowerPath Encryption with RSA, and Array Support. After clicking **Migration Enabler**, select either option in Figure 18. Click **Next**. Both options install the same set of binaries.

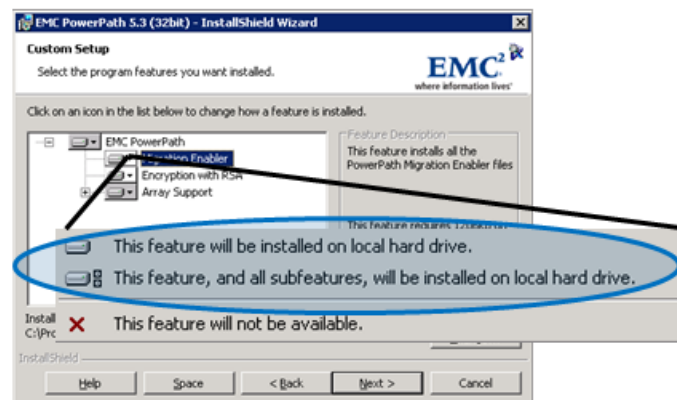


Figure 18. Select Migration Enabler files

---

Note: After modifying a PowerPath installation in Windows, the host must be rebooted for the changes to take effect. No modification is required with UNIX-based operating systems because the full PowerPath package is installed by default.

---

## Licensing PowerPath Migration Enabler

PowerPath Migration Enabler is available at no charge for new and existing customers with PowerPath Multipathing licenses. As long as the customers have a current software maintenance agreement with EMC, they are entitled to receive a PowerPath Migration Enabler license key. The request can be for any of the supported technologies: Host Copy, Open Replicator, TimeFinder/Clone, or Encapsulation. PowerPath Migration Enabler (for EMC Open Replicator for Symmetrix), EMC TimeFinder/Clone (for Symmetrix), or PowerPath Migration Enabler (for EMC Invista®) require separate licenses (or hardware in some cases).

## Obtaining a license

To obtain a free PowerPath Migration Enabler license, an EMC account representative will confirm that the customer has a valid maintenance contract. After that, the EMC representative will email the EMC License Keys group to obtain the appropriate license.

It is not necessary to have a separate license for every server. A single license for the customer for each technology (if more than one technology is needed) is sufficient to support a customer's migration needs. Depending on the requirements of the migration, the customer can use multiple PowerPath Migration Enabler licenses on the same host at the same time.

EMC personnel can contact [licensekeys@emc.com](mailto:licensekeys@emc.com) to obtain licenses.

## Licensing migration of encrypted devices

PowerPath Encryption with RSA leverages PowerPath Migration Enabler Host Copy to migrate encrypted (cipher-text) devices. The PowerPath Encryption with RSA license supports data migration to or from encrypted devices without the need to install the PowerPath Migration Enabler Host Copy license.

If the administrator plans to migrate only plain-text devices (no encrypted devices as the source or the target), then the PowerPath Migration Enabler Host Copy license is a requirement.

## Applying a license

During the installation of PowerPath, the administrator can add one or more license keys before activating PowerPath. However, it is not necessary to add the license keys at installation. The new license can be added at any time without disrupting the applications or operating system. Consult the appropriate PowerPath Installation and Administration Guide for more information on applying licenses.

## Restrictions

PowerPath Migration Enabler does *not* support:

- Migrating paging or swapping devices
- Migrating boot devices
- Uninstalling or upgrading PowerPath while a migration is in progress

The target logical unit in a migration:

- Must be the same size as or larger than the source logical unit. When Host Copy is used to migrate data from a plain-text source to an encrypted target, the target must be 65 KB larger than the source
- Cannot be under the control of a volume manager
- Should not have I/O directed to it

## Conclusion

EMC PowerPath Migration Enabler is a free migration tool that enables non-disruptive or minimally disruptive data migration between storage systems or between logical units within a single storage system. Migration Enabler resides on the host and allows applications to have continued data access throughout the migration process. When pseudo device names are used, the migration can be non-disruptive. When native device names are used, a migration is minimally disruptive because applications must be reconfigured to use the new target device name for the device that contains the migrated data.

Host Copy is one of four supported technology types. It is the most flexible technology because it supports migrating to virtually provisioned devices, migrating to encrypted devices, and migrating EMC and third-party arrays.

## References

For more information on PowerPath Migration Enabler, refer to the following material:

- *EMC PowerPath Migration Enabler User Guide*
- *EMC PowerPath Migration Enabler Release Notes*
- *EMC PowerPath Migration Enabler Support Matrix*