

# Dell PowerMax and VMAX: Non-Disruptive and Minimally Disruptive Migration Best Practices and Operational Guide

August 2022

H17133.6

## White Paper

### Abstract

This document includes best practices and walkthrough procedures for Non-Disruptive Migration (NDM) with Dell PowerMax, VMAX All Flash, VMAX3, and VMAX family storage systems.

Dell Technologies

## Copyright

The information in this publication is provided as is. Dell Inc. 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 software described in this publication requires an applicable software license.

Copyright © 2018-2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners.

Published in the USA August 2022 H17133.6.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

# Contents

Executive summary.....	4
Migration overview.....	6
Pass-through NDM.....	19
Metro-based NDM.....	64
NDM Update: Offline migrations with minimal host downtime .....	118
Online Minimally Disruptive Migration (O-MDM).....	138
Masking enhancements (SE 9.0).....	150
Masking enhancements (SE 9.1).....	162
Appendix A: Host multipathing software notes .....	164
Appendix B: AIX, GPFS, and PowerPath with NDM .....	167
Appendix C: Consistent LUN.....	169
Appendix D: Device geometry behavior post NDM.....	170
Appendix E: VMware—VM Clone or Storage vMotion with NDM .....	171
Appendix F: NDM with ESXi virtualized cluster environment .....	172
Appendix G: NDM Update with Solaris cluster .....	173
Appendix H: Technical support and resources.....	174

## Executive summary

### Overview

This document provides best practices and walkthrough procedures for Non-Disruptive Migration (NDM) with Dell PowerMax, VMAX All Flash, VMAX3, and VMAX Family storage systems.

Prior to the introduction of NDM, migrating from a Dell VMAX storage array to a new VMAX storage array was a time-consuming and tedious process. It required detailed planning and application downtime.

Solutions Enabler Release 8.3 introduced a data migration feature to help automate the process of moving applications from a VMAX array to a VMAX3 array without application downtime in a non-disruptive fashion. NDM leverages VMAX replication technologies to move the application data to the new array, and it leverages VMAX Auto-Provisioning, in combination with host multipathing software, to manage and maintain host access to the data during the migration process.

With each release of HYPERMAX OS and PowerMaxOS along with [Solutions Enabler](#), the core functionality of NDM has been iterated in order to streamline the user experience and align the support matrix to customer needs.

This document describes both versions of NDM:

- Pass-through NDM: Source array is running 5876 code (VMAX, VMAX2)
- Metro-based NDM: Source array running 5977 or 5978 (VMAX3, VMAX AF, PowerMax)

Each method uses a similar set of commands but has significantly differing underlying architectures leveraging the source arrays abilities.

Throughout this document, the term VMAX Family is applicable to all Dell PowerMax, VMAX 250F, VMAX 450F, VMAX 850F, VMAX 950F, VMAX 100K, VMAX 200K, VMAX 400K, and VMAX arrays.

This document also describes the NDM Update migration feature, in which the source array is running any VMAX or PowerMax code level.

### Audience

This document is intended for the following audience:

- IT management and planners, storage architects, and administrators involved in evaluating, acquiring, managing, operating, or designing and implementing PowerMax, VMAX All Flash, or VMAX3 storage arrays.
- Dell Technologies field personnel and partners involved in designing the NVMe-enabled PowerMax, VMAX All Flash, or VMAX3 solutions, and involved in planning for and implementing migrations to PowerMax, VMAX All Flash, or VMAX3 arrays.



## Feature updates and code levels

The following table outlines each feature update and the code level when the feature became available. This is listed in reverse chronological order from the latest to first feature release.

**Table 1. Feature updates per code release**

Code release	Features
Q3 2020 (5978.xx.xx SE 9.2) Code name: Hickory	No new feature updates
Q3 2019 (5978.xx.xx, SE 9.1) Code name: Foxtail	<ul style="list-style-type: none"> <li>Enhanced Metro-based NDM (5977–5977, 5978–5978)</li> <li>Addition of NDM Update (orchestrated SRDF migrations)</li> <li>Ability to create DR from target while in a migrating state</li> <li>Migration support for overlapping SGs</li> <li>IBMi NDM support</li> </ul>
Q1 2018 (5978.xx.xx, SE 9.0) Code name: Elm	<ul style="list-style-type: none"> <li>Enhanced Pass-through NDM (5876–5978)</li> <li>Introduced Metro-based NDM (5977–5978)</li> <li>50 concurrent migrations</li> <li>User ability enhancements: <ul style="list-style-type: none"> <li>Auto set DRX on 5876 devices</li> <li>User selectable target PG on create</li> <li>Honor Consistent LUN setting</li> <li>Non-Disruptive SG reconfiguration</li> <li>Honor Consistent LUN addresses (<a href="#">Appendix C: Consistent LUN</a>)</li> </ul> </li> </ul>
Q2 2017 (5977.1125.1125, SE 8.4) Code name: Cypress	<ul style="list-style-type: none"> <li>Ability to have DR from target in SRDF/Synchronous Mode</li> <li>Mixed WWN and IGs in masking view</li> <li>Initiators not required to be logged in, entry in LHT sufficient</li> <li>Full Cascaded V2 IG support</li> <li>Increased support matrix including Veritas VCS and DMP</li> <li>Automatically clear device geometry on commit</li> </ul>
Q3 2016 (5977.xxx.xxx, SE 8.3) Code name: Trinity	<ul style="list-style-type: none"> <li>NDM (cutover) initial release (5876–5977): <ul style="list-style-type: none"> <li>SRDF Pass-through</li> <li>16 concurrent migrations</li> <li>Stand-alone hosts</li> <li>Clusters SCSI2 and SCSI3</li> <li>Required Initiators to be logged in</li> </ul> </li> <li>SRDF/A DR support</li> </ul>

## Revisions

Date	Description
June 2018	Initial release
February 2019	<ul style="list-style-type: none"> <li>Added document revisions section</li> <li>PowerPath 6.4 typo changed to PowerPath 6.2</li> <li>Added note on GK</li> <li>Adding warning about post Cutover Rescans to both pass through sections and Pass-through NDM overview</li> <li>Added note to Linux appendix outlining the requirement for the -a parameter and link to Red Hat KB article</li> <li>Added update to Windows 2012 appendix outline the Path Verify Enabled flag behavior</li> <li>Template and format update</li> </ul>
September 2019	<ul style="list-style-type: none"> <li>Included updates from Q2 2019 release</li> <li>Included NDM Update section</li> <li>Included Enhancements to Solutions Enabler for 9.1</li> <li>Appendixes updated for RDMS with NDM</li> <li>Template and format update</li> </ul>
September 2020	Added link to the knowledgebase article in the technical support and resources section; format updates
January 2021	Clarification added to Veritas cluster behavior with NDM
January 2022	Updated template
August 2022	Added Online Minimally Disruptive Migration (O-MDM) content

## We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

**Authors:** Richard Murphy, Mike Adams

---

**Note:** For links to other documentation for this topic, see the [PowerMax and VMAX Info Hub](#).

---

## Migration overview

### NDM overview

NDM is designed to help automate the process of migrating host applications to a PowerMax, VMAX All-Flash, or VMAX3 enterprise storage array with no downtime.

Benefits of using NDM include the following:

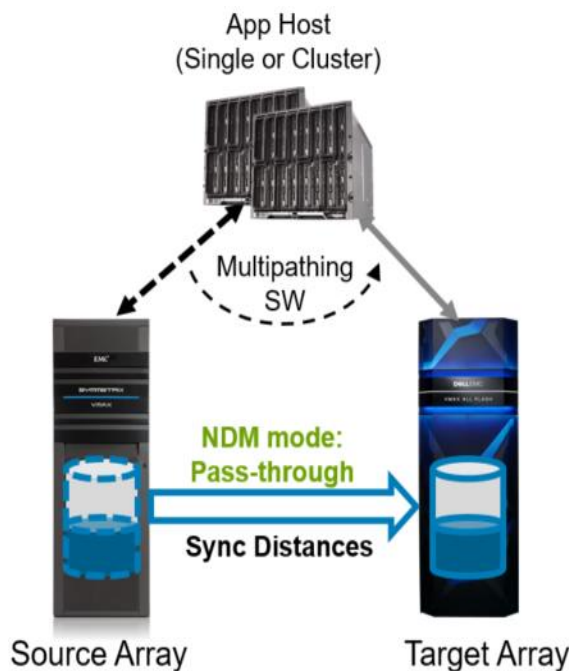
- Allows migration, with hosts and applications completely online, for the following:
  - From a VMAX array to a VMAX3, VMAX All Flash, or PowerMax array
  - From a PowerMax array to another PowerMax array
- Is designed for ease of use with control operations that automate the setup and configuration of the migration environment

- Is managed by familiar and simple user interfaces, the Solutions Enabler, and Dell Unisphere
- Allows migrations to be canceled easily and failed back to the source array prior to commit

Non-Disruptive Migration leverages VMAX SRDF replication technologies to move the application data to the new storage array. It also uses PowerMax/VMAX auto-provisioning, in combination with Dell PowerPath or a supported host multipathing solution, to manage host access to the data during the migration process.

NDM is available in the following two forms depending on the source array involved in the migration session. From a user standpoint, the process is very similar in terms of interaction, but the architecture is fundamentally different.

- **Pass-through NDM:** Source array is a VMAX or VMAX2 running 5876 code
- **Metro-based NDM:** Source array is a PowerMax, VMAX All Flash, VMAX3 running 5977 or 5978 code



**Figure 1. Pass-through NDM overview (migrating from VMAX with HYPERMAXOS 5876)**

**Note:** Reference the extensive [support matrix](#) when planning to migrate applications using NDM. In addition to the support matrix, reference the appendixes in this document for caveats on specific host-OS and multipathing combinations before attempting a migration.

### NDM operations

Since the initial release of the NDM feature, the cutover feature has been the method by which migrations from an array running 5876.xx.xx code have been undertaken. This process uses what is referred to as the three Cs: create, cutover, and commit.

The migration of an application from the source to the target array is completed using a sequence of user-initiated operations, each which is fully automated. These migrations are performed at the storage group (SG) level. The entire migration of a storage group can be accomplished with a few clicks in Unisphere or through simple and short Solutions Enabler commands.

- **Environment setup:** Setup configures the migration environment that will be required to migrate any application from the source array to the target array. It confirms that both the source and target arrays can support the NDM operations. This includes ensuring that a usable replication pathway for data migration is available between the source and target arrays. The environment setup command is run only once prior to the initial migration between two arrays. All other storage groups migrating between those arrays will use the same migration environment.
- **Create:** Solutions Enabler examines a specific application's storage on the source array and automatically provisions equivalent storage on the target array. The target devices are assigned the identity of the source devices and are configured in a pass-through mode that allows the data to be accessed from both the source and target devices.

After the create operation completes, the administrator issues a host rescan to allow the host to discover the paths to the newly created devices. Once this is complete, I/O issued by the application is directed to either the source or the target arrays through the host multipathing software. The array operating system ensures that all I/Os that are directed to the target by the host are actually serviced by the source array until the cutover.

- **Cutover:** A cutover operation moves the target devices out of pass-through mode, initiates data synchronization from the source to the target, and makes the paths to the source array inactive so that all I/Os are being serviced by the target array.
- **Commit:** After the source to target data synchronization is complete and all application data has been migrated to the target array, a commit operation is performed. During a commit operation, Solutions Enabler completes the migration by releasing temporary resources allocated to perform the migration, permanently disabling access to the source devices and assigning the target device ID to the source devices.
- **Environment remove:** This is performed after all migrations have been completed to remove the migration environment. The array-to-array connectivity configured for the data migration pathway is removed along with the rest of the infrastructure on both arrays that was configured by the array operating system to support the migrations.

Other supported NDM operations include the following:

- **Cancel:** Ends a migration that has not been committed. It removes storage provisioned for the migration on the target array, releases resources allocated by Solutions Enabler to perform the migration, and places the source devices into the state they were in before the create operation was run.
- **Cancel with revert:** This is run after a cutover operation to move the application's data access back to the source array from the target array and cancel the migration.

- **Recover:** This attempts to complete a failed migration operation. Recover is run following a failure after the cause of the failure, such as a connectivity issue, has been resolved.
- **Sync:** This controls the replication from the target side devices to the source side devices after a cutover is complete and all data has been migrated to the target side.
- **List:** This shows a list of migrations to or from a specified array, with the current status for each.

### Manipulation of device IDs and host paths

Two of the underlying processes that ensure that NDM is non-disruptive is the technologies ability to maintain device visibility at all times by spoofing and swapping devices IDs between source and target devices.

NDM can migrate data and cut over to the target array non-disruptively by both swapping device IDs between the source and target devices and manipulating the paths from the host to both arrays. The device ID contains the device's unique WWN and other information about it, such as a device identifier that the user has assigned to a device through Solutions Enabler or Unisphere. All of this information is copied to the target devices.

NDM performs the data migration and device ID swap without the host being aware. The path management changes appear as either the addition of paths or the removal of paths to the existing source device. To the host and application, there is no change in the device that it is accessing and access to the device is maintained throughout the entire migration process.

### Supported distance between source and target

NDM is supported across SRDF synchronous distances. However, because of the requirement that the host see both the source and target storage, migrations are typically performed between arrays within a data center.

### Effects on devices and existing replication sessions

Devices included in a migration session on the source array can remain in existing replication sessions throughout the migration. NDM evaluates the state of any current replication sessions before proceeding with the migration and ensures that they are in the proper state to allow the migration to succeed. By maintaining existing replication, NDM ensures there is no RPO impact during the period of the migration.

Though existing replication sessions can be left in place during the migration, replication relationships are not migrated to the target array. These replication resources need to be created on the target array, if required, at the appropriate point in the migration.

For example, SRDF replication can be configured between the target array and its remote array while in the CutoverSyncing state or after the CutoverSync state is reached. The new DR RDF pairs can then be allowed to synchronize before the commit so that DR is maintained throughout the migration. SRDF can also be set up in the CutoverNoSync state, which is reached when the sync command is used to stop replication. For local Snap/VX sessions running against the source volumes, existing sessions on the source

array can continue as normal during the migration and new sessions can also be created at the same time that the new SRDF to the DR site is configured.

Storage on the source and target arrays that is involved in the migration of an application should never be altered, and the migration resources should not be managed, outside of the NDM commands. If any changes in the migration session are detected when a migration operation is performed, the operation is blocked until the changes that were made are undone, allowing the migration operation to proceed as expected.

The following are examples of manual changes made to the NDM session that will cause the session to stop or fail:

- Storage group manipulation, or adding or removing devices
- Masking view manipulation such as changing the name, or adding or removing elements

### Configuration requirements and prerequisites

Most of the steps for configuring and unconfiguring NDM are done automatically using the environment setup and remove commands. Prior to running the setup, the following steps are required:

- SRDF ports must be configured across at least two directors.
- The source and target arrays must have SRDF directors and ports configured.
- The SRDF ports between the source and target arrays in a Fibre Channel environment must be zoned to each other.
- A Solutions Enabler or Unisphere management host that sees at least one of the arrays must be available.
- The host with the application being migrated must be zoned to the VMAX3 or VMAX All Flash array.
- For VMware, enable `consistent_lun` on the source array initiator before performing NDM operations.

---

**Note:** SRDF ports do not need to be dedicated to NDM operations. Ports involved in ongoing SRDF disaster recovery operations may be shared with NDM sessions, but analysis should be performed prior to setting up NDM to make certain there is adequate bandwidth to handle both DR and migration traffic.

---

### Restrictions

#### *NDM SRDF restrictions*

- Potential source devices cannot be R2 devices.
- Arrays must be within metro distances of each other.
- The source devices cannot be part of a concurrent RDF relationship.
- The source devices must not be enabled for RDF consistency.
- The source devices may not be part of an SRDF/Star configuration.
- The source devices may not be part of an SRDF/Metro configuration.

**Environment restrictions**

A minimum of two SRDF links (FC or GigE) are required to support an NDM session. These ports must be spread across at least two directors.

**Open replicator restrictions**

- The migration source device may not be the control device in an Open Replicator pull operation.
- Open Replicator must not be replicating data from a remote device to a control device, such as during an in-progress restore.

**TimeFinder restrictions**

- The source or target devices may not be the target of a TimeFinder copy session.
- SRDF must not be replicating data from a local replication operation, such as during an in-progress TimeFinder restore.

**RecoverPoint restrictions**

The migration source or target devices cannot be tagged for RecoverPoint use.

**Boot-from-SAN support**

Cutover NDM supports hosts that boot directly from the VMAX array. The host boot BIOS must be updated to point to the target volume so that when the host is rebooted at a later date it will find the volume containing the operating system. For details on boot drive configuration, refer to the vendor specific HBA management guide or BIOS guides.

**REST API support**

Both methods of NDM, Pass-through and Metro-based, are fully supported through the REST API.

**Recognizing the NDM SRDF group**

The RDF group created for NDM between two arrays can be identified by its label. The label follows the format of M\_XXXXYYYY. XXXX is the last four digits of the lower numbered storage array and YYYY is the last four digits of the higher numbered array. This group is used for all NDM migrations between the two arrays. This group is automatically created as part of the environment setup.

**Setting up multiple NDM environments**

Multiple-environment setup operations can be performed for a single source array, provided that a different target array is specified for each migration environment. All NDM RDF groups on a source or target array can be in use simultaneously, for concurrent migrations to or from an array.

For example, a single PowerMax, VMAX All Flash, or VMAX3 target array can have multiple NDM RDF groups, each connected to one of four different source VMAX arrays. This means that the target array can be the target of migrations from each of those four VMAX arrays in a consolidation use case.

Likewise, a single VMAX source array can have multiple NDM RDF groups, each connected to one of four different target PowerMax, VMAX All Flash, or VMAX3 arrays. This means that the VMAX array can be the source of migrations to each of those four VMAX3 or All Flash arrays.

When migrations are completed, separate environment remove operations are required for each array pair. The environment remove operation removes the NDM RDF group between the two arrays, provided that no devices on either array have an RDF mirror in the NDM RDF group.

### Masking groups and views with NDM

When NDM sessions are created, NDM configures the following items on the target array with the same names as those on the source array:

- Storage groups (SGs)
- Initiator groups (IGs)
- Port groups (PGs)
- Masking views (MVs)

Both initiator groups and port groups can exist in multiple masking views, so these groups are reused when applicable.

A host may also be attached to multiple source arrays. For example, if a storage group spans two source arrays, when the storage is migrated, the target array contains two sets of SGs, IGs, PGs, and MVs, one for each source array.

When the **first SG on the first array** is migrated to the target array, the following occurs:

- An SG is created on the target with the same name that contains the migration target devices.
- An IG is created on the target with the same name that contains the host initiators.
- A PG is created on the target based on which ports the host HBAs are logged into.

When a **second SG on the second source array** is migrated to the target array, the following rules apply:

- The SG name must be different.
- If necessary, the SG can be renamed before it is migrated.
- The IG must have the same name because an initiator can only exist in one IG.
- If the PG on the second array has the same name as the PG on the first array, the PG built by NDM during the first migration can be reused. If it has a different name, a new PG will be created with the same ports used in the PG created during the first migration.

Alternatively, you can manually create the PG on the target in advance. Then, select this as the target PG for the NDM session or create it as part of the NDM create process. This option is new for Solutions Enabler 9.1.

### Rules for masking groups and views

All migrations are performed against an SG, which is the data container that is migrated with NDM. The following rules apply:

- Only SGs contained in masking views can be migrated. If the device is mapped to a port that it is not masked to for this SG, the create operation is not permitted.



- Multiple masking views on the SG using the same IG are not allowed unless PGs on the target array already exist for each view and the ports in the PGs are selected to avoid duplicate host paths.
- If the SG is a parent, its child SGs are also migrated.
- Devices in the SG which are considered to be GKs (20 cylinders or less) are not migrated to the target array. Devices must not be masked to FCoE ports.
- Devices must not be masked to iSCSI ports.
- Devices must not be mapped to ports where the ACLX is not enabled.
- If a storage resource pool (SRP) on the target array is specified for the migration, that SRP must already exist on the target array. The names of the SGs (parent and children) that are being migrated must not exist on the target array.
- The names of the masking views that are being migrated must not exist on the target array.
- The names of the initiator groups that are being migrated may exist on the target array, the IG layout may differ on the target array to that on the source.
- The names of the port groups that are being migrated may exist on the target array, provided that the groups on the target array have the initiators logged in into at least one port in the port group.

### **NDM general considerations and limitations**

- Migrating hosts must have access and be zoned to both the source and target array
- Migrating hosts must use Fibre Channel connectivity

### **Session limits**

- 50 SGs can be migrated concurrently. Child SGs do not count towards this limit.
- Each SG can contain up to 4096 devices.

### **Hardware and software requirements**

For hardware and software requirements, refer to the [PowerMax/VMAX All Flash/VMAX3 Features Simple Support Matrix](#).

## **Metro-based NDM overview**

Metro-based NDM is built upon SRDF/Metro active/active technology with the Metro instant activate (Metro-IA) feature. For more information about SRDF/Metro, see the [SRDF/Metro Overview and Best Practices](#) document.

Benefits of using Metro-based NDM include the following:

- Allows the complete migration from a VMAX3 or VMAX All Flash arrays to VMAX All Flash or PowerMax arrays without the need for downtime
- Leverages Metro technology to reduce the required commands users need to issue by removing the need to cut over
- Is managed by familiar and simple user interfaces: Solutions Enabler and Unisphere for VMAX
- Allows migrations to be canceled easily and failed back to the source array for any reason prior to commit

- Allows users to precopy all or a large portion of data before bringing the host live to the new array. This reduces the impact to the application during the migration period.

The previous version of NDM provided with Solutions Enabler 8.3, Unisphere 8.4, and HYPERMAX OS 5977.1125 releases allowed data to be migrated from a VMAX (5876) to a VMAX3 (5977) array without application downtime.

With the release of Solutions Enabler 9.0 and HYPERMAX OS Q2 2018, the NDM feature was enhanced to help automate the process of moving applications from a VMAX3 (5977) or VMAX All Flash (5977) array to another VMAX All Flash (5978) or PowerMax (5978) array.

With the release of Solutions Enabler 9.1, full interfamily migrations are possible (5977 to 5977, and 5978 to 5978).

The source hardware is not a limiting factor. Metro-based NDM is supported from arrays running 5977 or 5978 to 5977 or 5978 code, regardless of the underlying technology.

With SRDF/Metro, the session goes active/active only after all SCSI information and application data is synchronized from R1-R2 using SRDF adaptive copy technology. The time to be fully active/active largely depends on the time it takes for the data transfer to finish.

To improve the user experience with NDM, the software is enhanced such that the SRDF/Metro session goes active/active instantly on NDM create. This only applies to underlying Metro technology for NDM and does not apply to regular SRDF/Metro for running active/active applications. This makes both sides of the SRDF/Metro active and read/write to the host within the duration of the create command.

### Metro-based NDM modes of operation


- Metro-based NDM: This is the NDM mode where the synchronization starts right after the create completes, and the create operation creates an NDM session using SRDF/Metro internally.
- Metro-based NDM with precopy: This mode of operation offers a choice to start the NDM session using adaptive copy (SRDF/ADP) which helps synchronize most of the data before moving into an active/active state. Eventually, this moves to active/active mode and synchronizes the remaining tracks. In essence, precopy allows end users to copy application data from the source array to target array while the application is still running on the source array.

### Metro-based flows

The following steps and diagram describe the **process flow** for Metro-based NDM:

1. Solutions Enabler creates an active/active SRDF group with NDM attribute.
2. Solutions Enabler activates the group.
3. SCSI information, device personality, and attributes are transferred to the target devices.
4. Masking view to the source (R2) array created.
5. Migration starts from R1 to R2 (source to target).

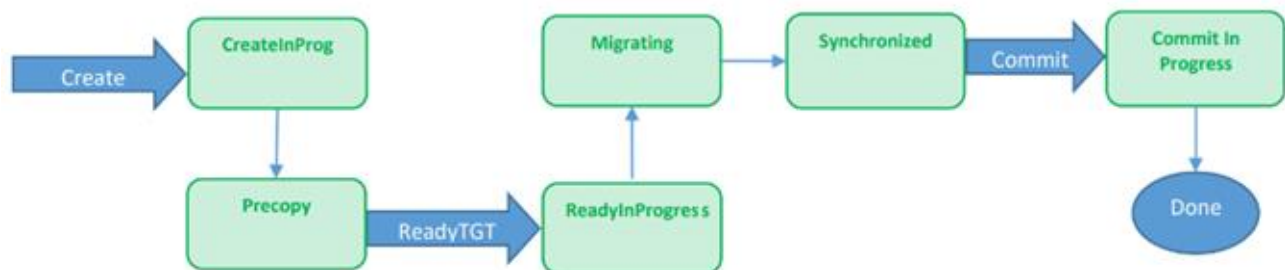
6. User rescans host for extra paths to target array.
7. Data synchronization completed.
8. Commit issued and migration completed.



**Figure 2. Process flow for Metro-based NDM**

The following steps and diagram describe the component flow for Metro-NDM with precopy:

1. Solutions Enabler creates an active/active group with NDM attribute and precopy set. The SRDF group is set to adaptive copy mode.
2. Pre-copy data sync (R1-R2)
3. Solutions Enabler sets READYTGT which will activate RDF group with Metro-IA.
4. SCSI information, device personality, and attributes are transferred to the target devices.
5. Masking view to the source array created.
6. User rescans host for extra paths to the target array.
7. Data synchronization completed.
8. Commit issued and migration completed.



**Figure 3. Component flow for Metro-based NDM with precopy**

### Metro-based NDM operations

- **Create:** This creates a Metro-based NDM session. It creates SRDF/Metro with NDM source and target attributes and puts the NDM pair in active/active mode (clear DEV-INACT on R2). This also starts data synchronization and SCSI info synchronization and moves into a migrating state. Once the data has been synchronized and the host paths to the target array have been discovered by

means of a host rescan, the NDM session reaches a synchronized state. A create results in RDF mirror invalids on R1 and local mirror invalids on R2.

- **Create with precopy:** This creates an SRDF/Metro session with NDM attributes and puts the SRDF/Metro pair into adaptive copy disk mode. It starts data synchronization from R1 to R2. Bias is enabled on the Metro-based NDM source. This creates an SRDF/Metro session with NDM attributes and puts the SRDF/Metro pair into adaptive copy disk mode. It starts syncing data from R1 to R2. Bias is on the Metro-based NDM source.
- **ReadyTGT:** This continues to synchronize the remaining invalid tracks from source target and also starts synchronization of SCSI Information. It puts the NDM pair in active/active mode (clear DEV-INACT on R2) without waiting for synchronization to finish.
- **Cancel:** This suspends the SRDF/Metro NDM session. The R2 device moves into a DEV-INACT state and all I/O is redirected to R1.
- **Sync-stop:** Once in synchronized state, the user might need to test the application performance before committing. The sync-stop command would make the NDM SRC DEV-INACT, all I/O would go through Metro-NDM target R2, and the Metro-NDM session moves into a CutoverNoSync state. This operation results in moving the bias to the Metro-NDM target.
- **Sync-start:** This moves the Metro-NDM session from CutoverNoSync to CutoverSyncing and finally to synchronized. This command should be used once the user has finished verifying the application against the NDM target after the sync-stop command and the user is ready to either commit or cancel the NDM session. The bias remains on the Metro-NDM target commit.

### Manipulation of device IDs and host paths

NDM can migrate data and cut over to the target array non-disruptively by both swapping device IDs between the source and target devices and manipulating the paths from the host to both arrays. The device ID contains the device's unique WWN and other information about it, such as a device identifier that the user has assigned to a device through Solutions Enabler or Unisphere. All of this information is copied to the target devices.

NDM performs the data migration and device ID swap without the host being aware. The path management changes appear as either the addition of paths or the removal of paths to the existing source device. To the host and application, there is no change in the device that it is accessing and access to the device is maintained throughout the entire migration process.

### NDM Update overview

Solutions Enabler 9.1 introduces the NDM Update feature to help automate the process of moving applications across arrays requiring a short application downtime.

Benefits of using NDM Update include the following:

- Allows application migration for the following:
  - From a VMAX (5876) array to a VMAX3 (5977), VMAX All Flash (5978), or PowerMax (5978) array

- From a VMAX3 (5977) array to a VMAX All Flash (5978) or PowerMax (5978) array
- Supports application (SG) level migration for open systems (FBA devices)
- Is built on existing NDM and SRDF technologies:
  - Provides the same streamlined user experience
  - Supports pre-sync data between arrays and then schedule downtime to do the cutover.
- Supports spoofing or not spoofing the LUN identity
- Requires minimal disruption of the host I/O:
  - Host application is briefly shut down when moving the host visibility to the target.
  - Application disruption time limited to its reboot time if spoofing is used.
  - Application disruption time limited to its reboot and LUN repointing time if no spoofing is used.
- Carries over existing NDM interoperability with array features to NDM Update, with all existing NDM interactions with LREP and SRDF remaining the same
- Provides the following capabilities:
  - Covers all host environments for open systems (FBA devices)
  - Addresses host environments where NDM is not supported but a high level of orchestration is preferred
  - Enables a high level of orchestration and ease of use
  - Has a similar workflow as NDM
  - Provides the same support for easy recovery of failed operations
  - Provides the same support for the easy cancellation of an ongoing migration
  - Maintains DR throughout migration
  - Places no additional hardware in the data path
  - Compresses data during migration to VMAX All Flash or PowerMax
  - Leverages SRDF/s or SRDF/Adp for data transfer
  - Is native to HYPERMAX OS with no license required (no SRDF license required)

For the latest connectivity information, see the [support matrix](#).

### NDM Update operations

Migrating an application from a source to a target array proceeds with a sequence of user-initiated operations, each fully automated. Using Solutions Enabler 9.1, the following operations are undertaken using the symdm library similar to NDM.

When initiating an NDM Update session, the **create -offline** parameter is used. Similarly, if using Unisphere for PowerMax, select the **offline** checkbox. See the walkthrough sections for examples of the following operations:

- **Migration environment setup:** This operation configures the migration environment required to migrate any application from the source array to the target array.

One or more applications can be migrated serially or in parallel from the source to the target array using the migration environment:

- **create -offline:** This operation examines a specific application's storage on the source array and automatically provisions equivalent storage on the target array. Source and target devices are configured in a mode that starts copying the data to the target devices.
- **-move\_identity:** This option assigns the target devices the identity of the source devices. If **-move\_identity** is not selected, host applications need to be directed to the new target volumes prior to rebooting the application. This is not the case if **move\_identity** is selected, much like NDM whereby the host sees what it perceives as the original volumes but with a spoofed WWN. (See the relevant NDM section for more information about WWN manipulation.)
- **-precopy:** This option modifies when the user must shut down the application to allow the source devices to be available during the time the data is copied to the target devices.
  - Without the **-precopy** option, the user must shut down the application before the **create -offline** command is given. The command will make the target devices visible to the host and the source devices **host\_inactive**.
  - With the **-precopy** option, the command will not make the target devices visible to the host and the application can continue running only on the source array while the data is copied to the target. A migration **cutover** operation is required to continue the migration. The **cutover** operation makes the target devices visible to the host and the source devices **host\_inactive**.
  - The migration **cutover** operation is only used in the **CutoverReady** state and requires the user to shut down the application before the **cutover** command is given.

After a **create -offline** or the **cutover** command completes, the administrator must perform a host rescan (or host reboot) and verify that new paths or new LUNs have been discovered prior to restarting the application. Unless the **-move\_identity** option was given, the application configuration would also need to be changed to use the new LUNs.

After the source to target data synchronization is complete and all application data has been migrated to the target array, a commit operation is performed. During a commit operation, Solutions Enabler completes the application migration by releasing resources allocated to perform the migration, which also permanently disables access to the source devices. If the **-move\_identity** option was given during the **create -offline** command, **commit** also assigns the identity the target device was created with to the source devices.

- **Environment remove:** This operation is performed after all migrations have been performed to remove any configuration created by the 'environment setup' operation.

Other supported operations provide means to the following:

- Cancel a migration, removing storage provisioned by Solutions Enabler on the target array, releasing resources allocated to perform the migration, and placing the source devices into the pre-create command state.
- Recover from a failed migration.
- List migration environments, to see an overview of the configured DM environments.
- List migrations to or from a specified array, showing the current status for each.

## Pass-through NDM


### Pass-through NDM guide plan and environment review

Pass-through NDM walkthrough guide (source running 5876 code) describes two methods for NDM:

- [Unisphere for PowerMax](#)
- [CLI \(Solutions Enabler\)](#)

The duration of this walkthrough uses the VMAX arrays in the graphic below.

Pass-through NDM migrates SGs from 000198701161 to 000197800131.

000198701161 VMAX10K   5876.309.401	5 %	<span style="color: red;">●</span> <span style="color: yellow;">●</span> <span style="color: green;">●</span>	-	MB/S	-	-	  
000197800131 VMAX250F   5978.100.100	8 %	<span style="color: red;">0</span> <span style="color: yellow;">0</span> <span style="color: green;">0</span>	-	MB/S	-	691.0:1	  

Prior to the start of a planned migration, ensure that the prerequisite checks for using NDM have been completed:

- Ensure both source and target array are RDF capable, that is RF emulation has been added to both arrays
- Ensure both arrays RDF ports are zoned to each other. There is a minimum of two connections required
- Check for the correct zoning from the target array to the application host.

Even though the guide takes you through the detailed device examination after each step and dives into the workings of each of the issued commands, the user needs to issue three commands to migrate an SG from source to target.

- **Create**, followed by a host rescan
- **Cutover**
- **Commit**, followed by a host rescan

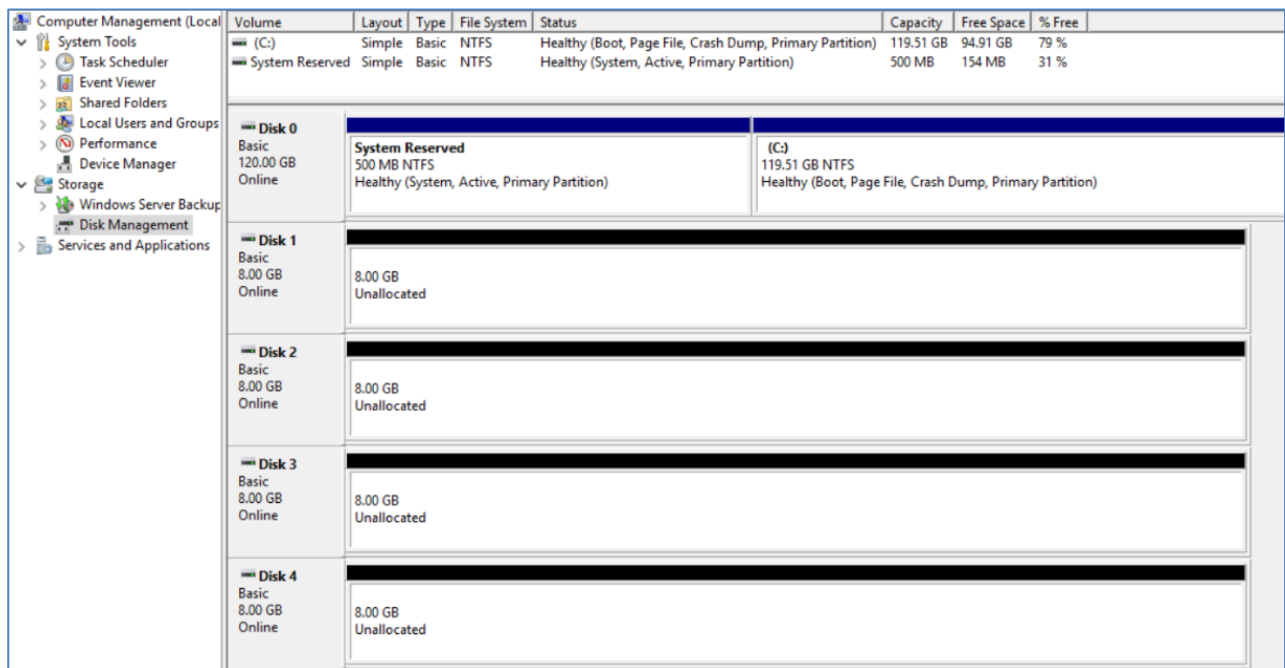
**Note:** A host rescan that will result in the permanent removal of the now inactive paths should not be undertaken post-cutover, this will limit the ability for the migration to be seamlessly canceled and normal operation reverted to the source array. In the case of multiple concurrent NDM sessions sharing the same host the same rule should apply across all sessions when issuing rescans.

**Pass-through NDM walkthrough guide (source running 5876 code)**

**Using Unisphere for PowerMax**

The following screenshots are taken from disk management on a Microsoft Windows Server 2016 host, this walkthrough intends to migrate disks 1 to 4 using NDM without any impact to the operating system or application accessing these devices. This is a virtual host with physical Raw Device Mappings.

The VMAX devices involved in this example are 1EA through 1ED added to Storage Group Uni\_Cut\_SG1, which is masked to the virtual host.





```

pplicensevmaxcse:~ # rpowermt display dev=emcpower204 host=10.
Pseudo name=emcpower204
Symmetrix ID=000198701161
Logical device ID=01EA
Device WWN=60000970000198701161533030314541
Standard UID=naa.60000970000198701161533030314541
type=Conventional; state=alive; policy=SymmOpt; queued-Ios=0
----- Host -----
### HW Path          I/O Paths    - Stor -  -- I/O Path --  -- Stats ---
                               Interf.  Mode   State   Q-Ios Errors
-----
3 vmhba6             C0:T19:L1   FA  1e:01 active  alive   0    0
3 vmhba6             C0:T8:L1    FA  2e:01 active  alive   0    0
3 vmhba6             C0:T20:L1   FA  2e:00 active  alive   0    0
3 vmhba6             C0:T17:L1   FA  1e:00 active  alive   0    0
4 vmhba2             C0:T17:L1   FA  2e:00 active  alive   0    0
4 vmhba2             C0:T13:L1   FA  1e:00 active  alive   0    0
4 vmhba2             C0:T15:L1   FA  2e:01 active  alive   0    0
4 vmhba2             C0:T2:L1    FA  1e:01 active  alive   0    0

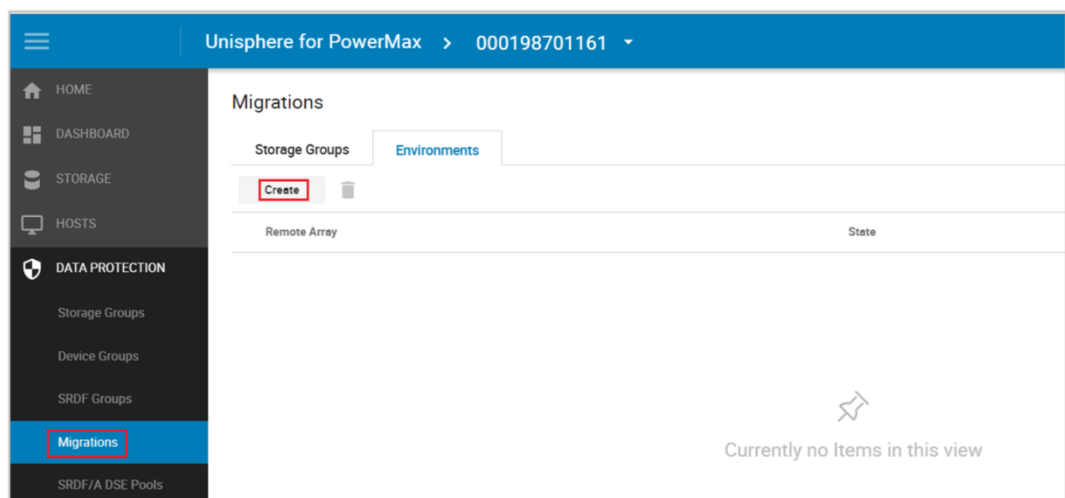
```

An example of the multipathing setup using device 1EA. It shows what the pathing looks like prior to the NDM create and the host rescan. For each of the four volumes here there are eight paths to the source array which are all alive and available for host use. At this point, there are no paths to the target array even though our zoning should be in place before the NDM create.

### ***NDM environment setup***

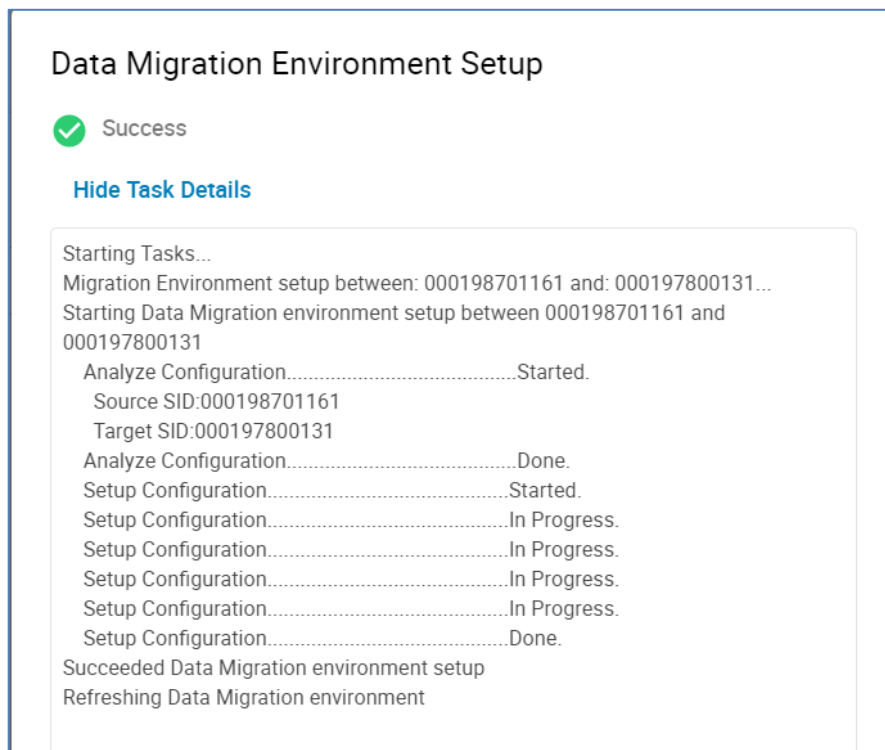
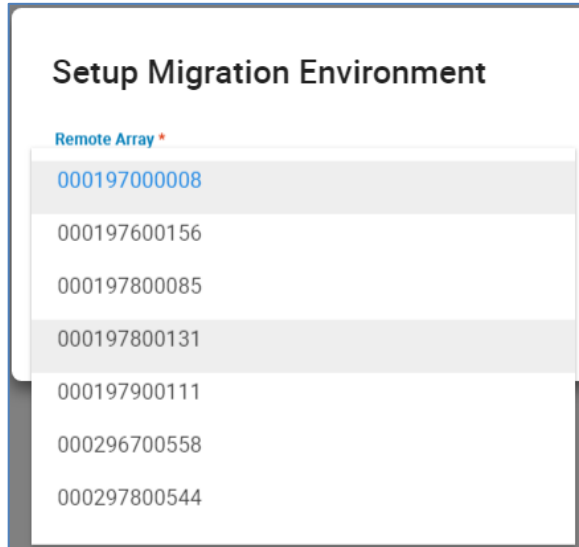
From the Unisphere Dashboard, select the Source array from the available arrays in the view. In this case, the source array is a 10K with the serial number ended in 161.

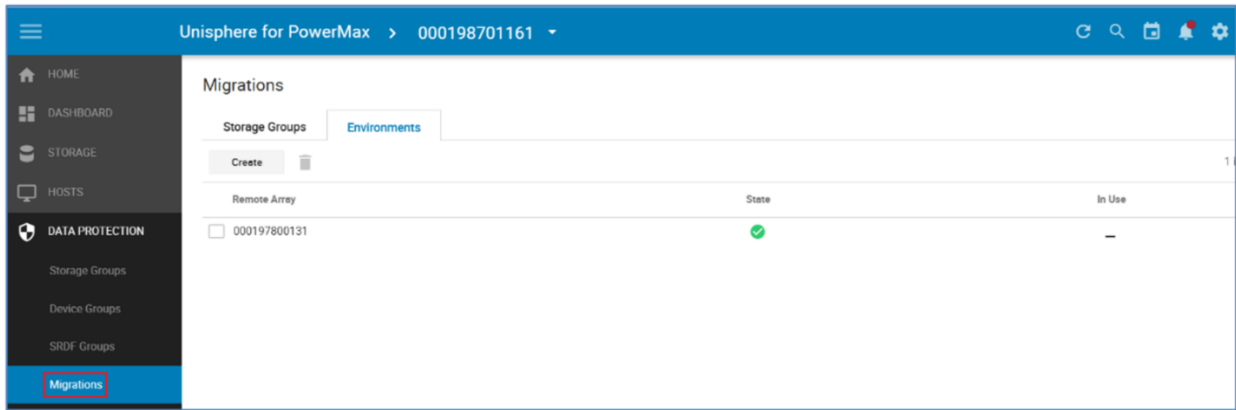
The Environment Setup configures the migration environment that will be required to migrate any application from the source array to the target array. It confirms that both the source and target arrays can support the NDM operations. This includes ensuring that a usable replication pathway for data migration is available between the source and target arrays and creating an SRDF group for the migration. The setup is run once only. When the migration pathways and SRDF group are configured, all storage groups on the source array can be migrated until all migrations from the source to the target have been completed.



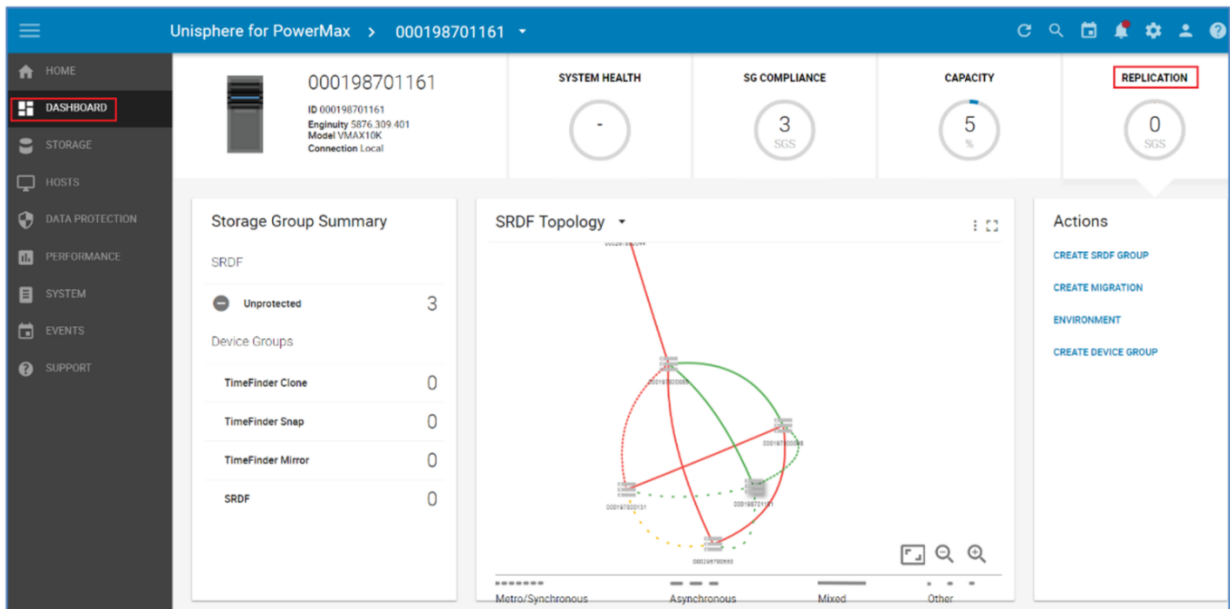
From the **Data Protection** menu select **Migrations**, Select the **Environment** tab and this will display any existing Environments already setup. The parameter **In Use** shows us if the Environment is validated and usable. The **In Use** parameter tell us if there is an active Migration using this environment.

To create an environment, select **Create**, the pop-up window below allows the choice of target array. This is populated with suitable arrays. Should the required array not be present, verify the RDF zoning and confirm the intended target array is suitable and its current code level is within the support matrix. Select the relevant array and choose **Run Now**.

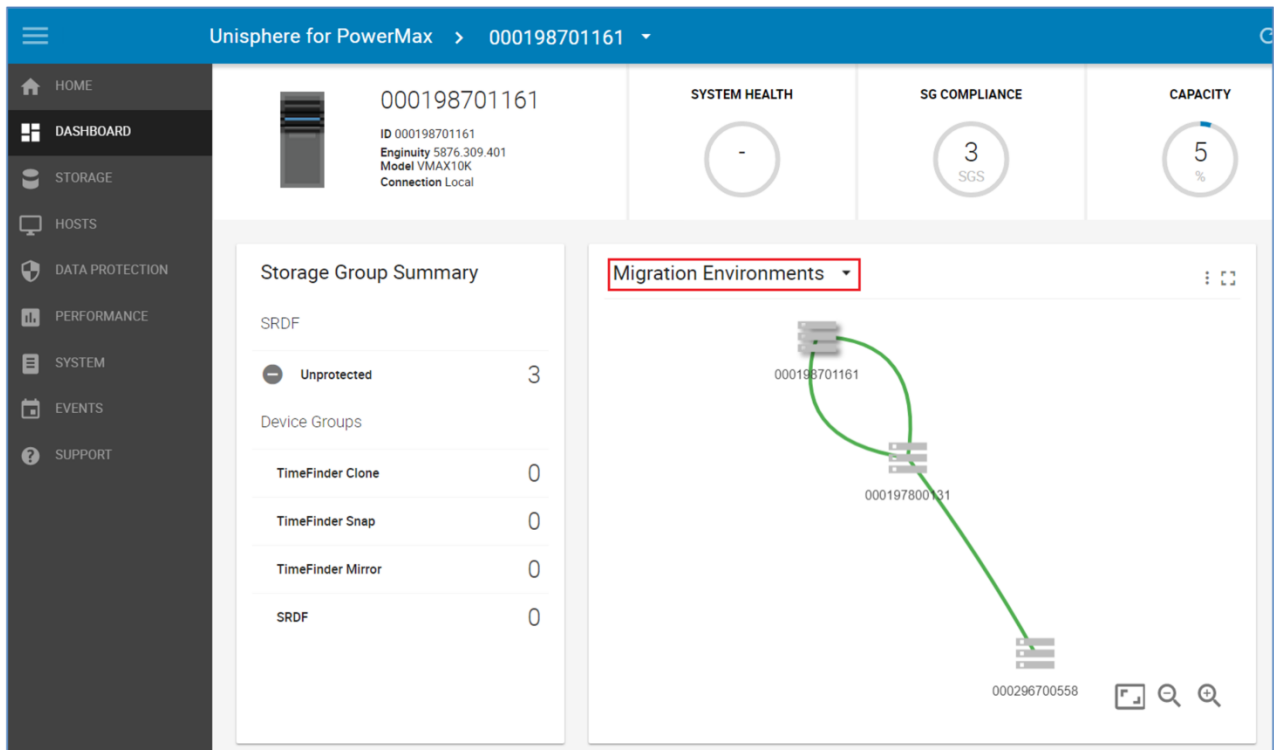




Examining the RDF environment from the new **Topology** view shows the RDF group template that has been created. Go to the **Dashboard and Replication** and hover over the line between our source and target. From here the **SRDF Groups** window appears. Select **View Groups** to display the SRDF group window highlighted below.



The dotted link between 161 and 131 suggest an “other” type of SRDF relationship. In this case, it is a migration link. The drop-down menu highlighted below allows the user to highlight just the migration relationships.

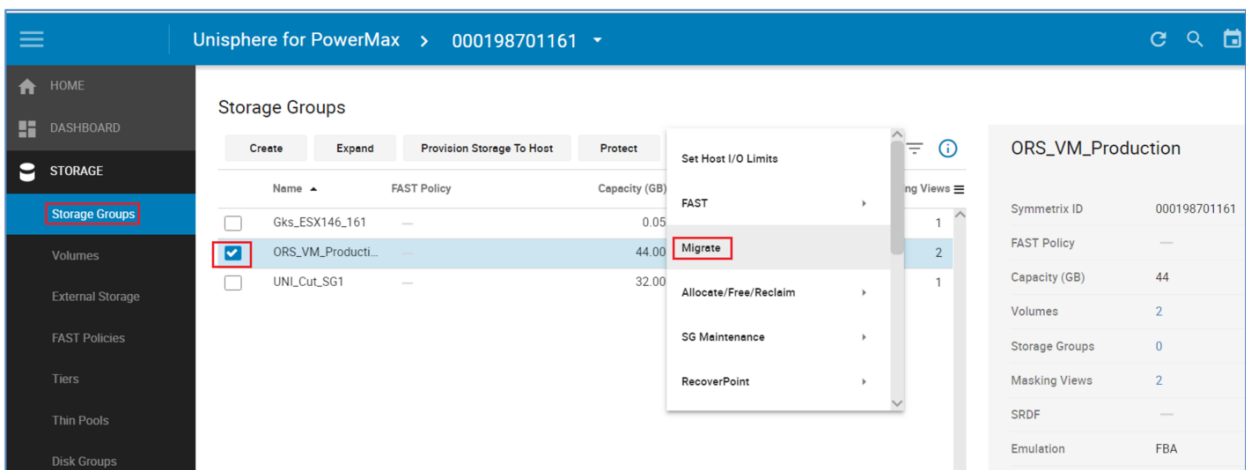


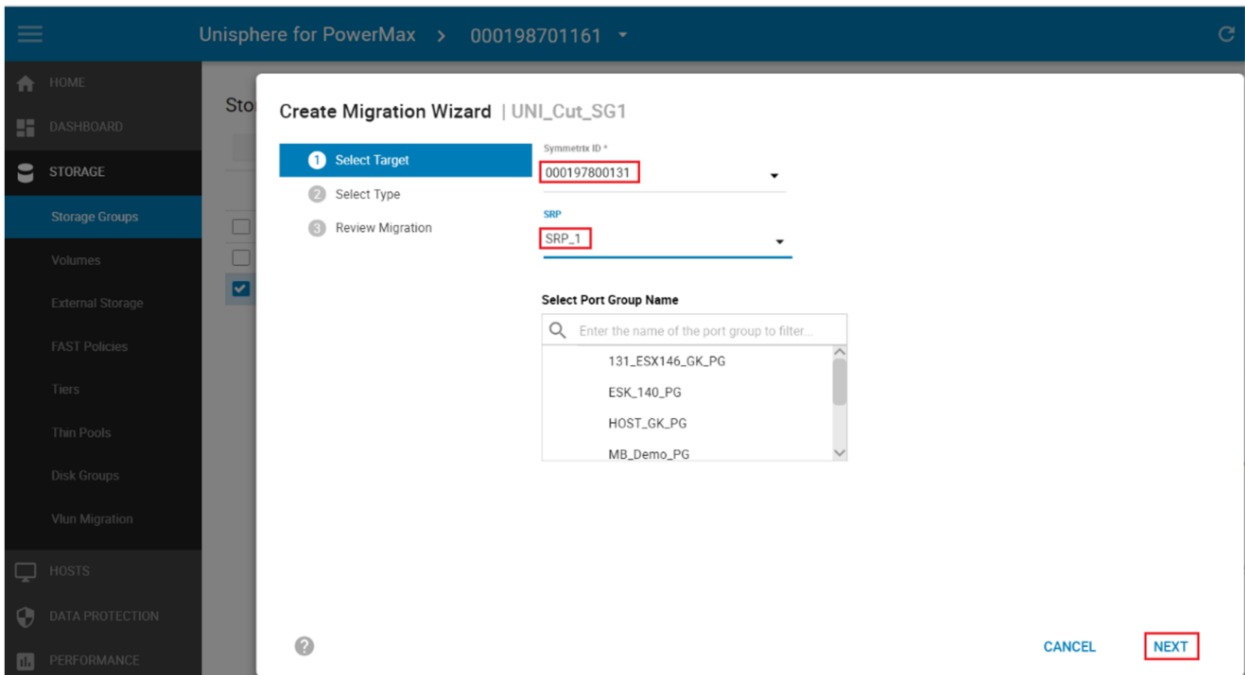
Now the Environment is in place we can continue with the *NDM Create* for the SG containing the application to be migrated.

**Note:** NDM Environment setup creates RDF links between the Source and Target using one port per Director for each zoned link. However, post-setup the user has the ability to add extra links manually using: `symrdf modifygrp -rdfg 250 -add -dir xx -remote_dir xx`

**Create migration session**


From the **Storage** tab select **Storage Groups**, from there locate the SG that is to be migrated. Set the check box and click the **More Actions** "3-Dot" icon to the right of Set Host I/O Limits. From the drop-down menu, select **Migrate**.



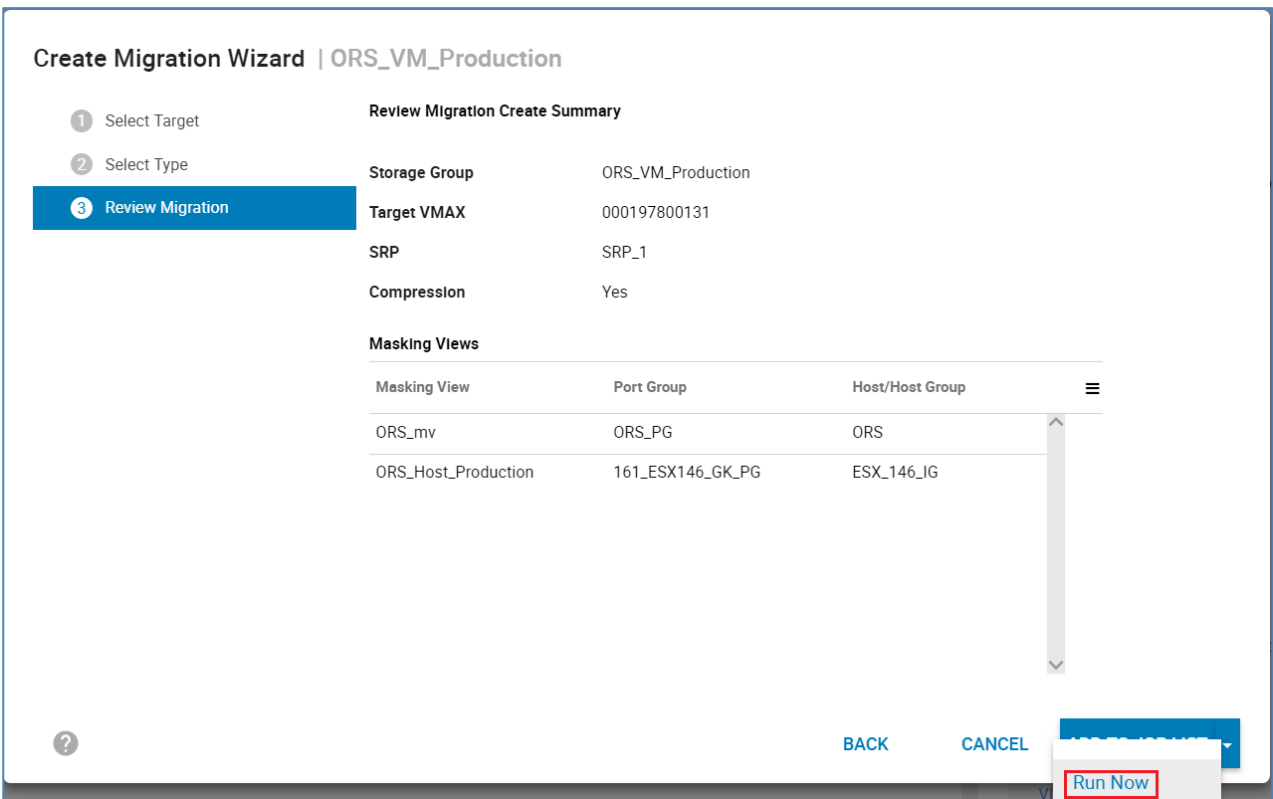


From the pop-up Create Migration Wizard, select the **Target array** (only Arrays with valid Environments setup appears on this drop-down menu). From Solutions Enabler 9.0 and later, the ability to select an existing **Port Group** on the target array is also an option. Select the target array SRP and Select **Next**.

From the next screen select Create Data Migration, from here we have the option of selecting Compression on the target SG. The Prepare Data Migration selection requires Performance data to be collected on the host. This runs a check for resources on the target array to ensure the addition of the new SG does not cause the target array to exceed any performance metrics on both FE and BE. It will also produce a spread sheet to help plan the zoning required for the host from the target array.



The final menu for the Wizard gives a final confirmation on the planned NDM session to be created. It breaks down the planned masking view elements and the NDM parameters. Select **Run Now** to continue.



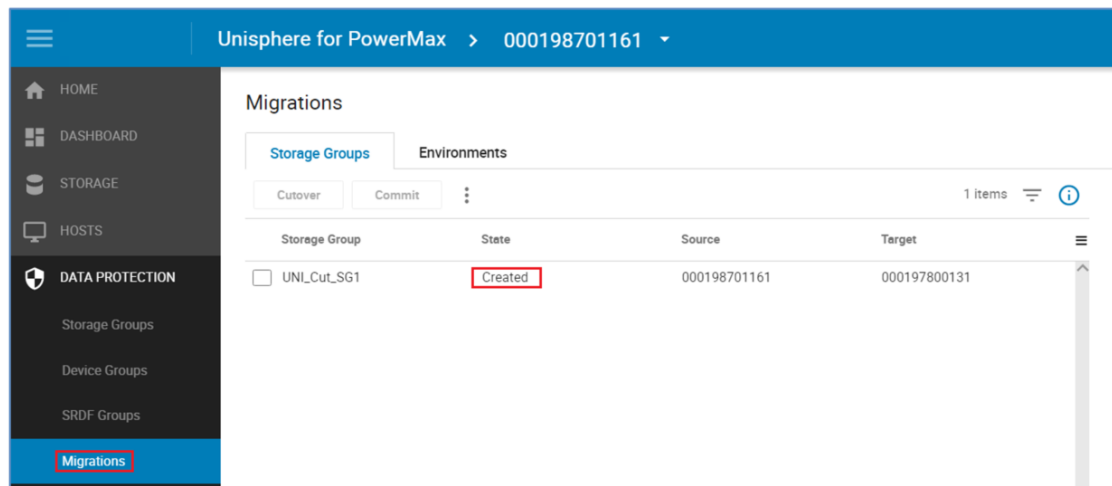
Creating the NDM session will also validate the Environment as part of the setup to ensure the migration complete successfully. The Create Command:

- Creates a Storage Group on the Target array (group name must not be in use on the target array) with the same name as the Source SG
- Creates duplicate devices on the target array to match those on the Storage Group
- Creates an initiator group using Initiators with entries in the login history table
- Creates a Port group (if one does not already exist)
- Effective (external) WWNs of the device created on the target are copied from the WWNs of the host devices
- Creates a masking view to the host from the target array

**Note:** During a Cutover NDM migration, the source of the migration is an R2 or an R21 device (if there is existing SRDF DR replication from the source device) and the target is an R1 device. This is different than basic SRDF operations and is required to allow DR protection during a migration using a cascaded SRDF configuration.

### Examine the migration session

The following screenshot shows how to examine a migration session.



### Perform a host rescan

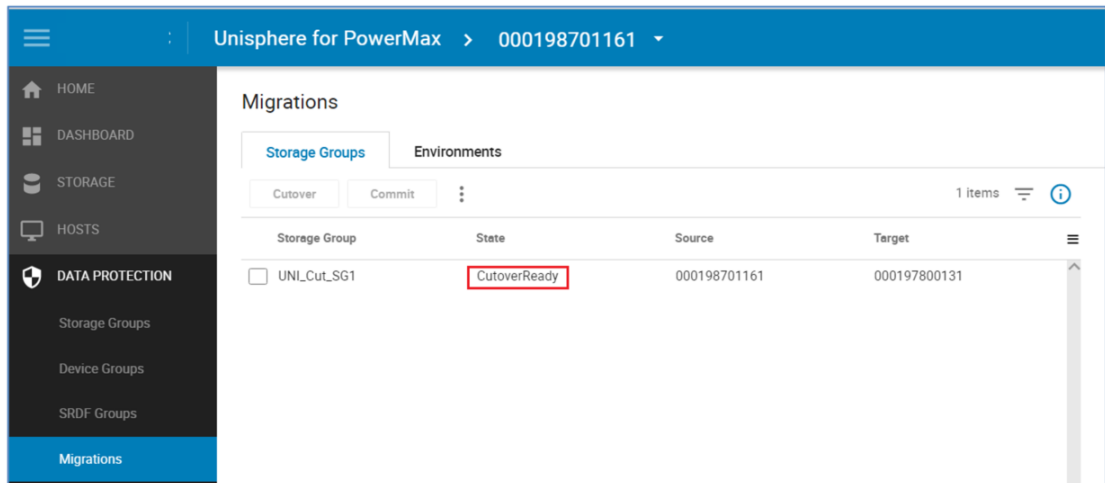
After the create operation completes the systems administrator must issue a host rescan to allow the host to discover the paths to the newly created devices. This host rescan is operating-system specific and also should include a rescan using the host multipathing software.

The NDM session goes from a Created state to a CutoverReady state after the host rescan is performed and the target devices are discoverable. After this is complete, I/O issued by the application will be directed to either the source or the target arrays through the host multipathing software. This is possible because the target devices are in pass-through mode. [Appendix A: Host multipathing software notes](#) has more details on host multipathing software settings.

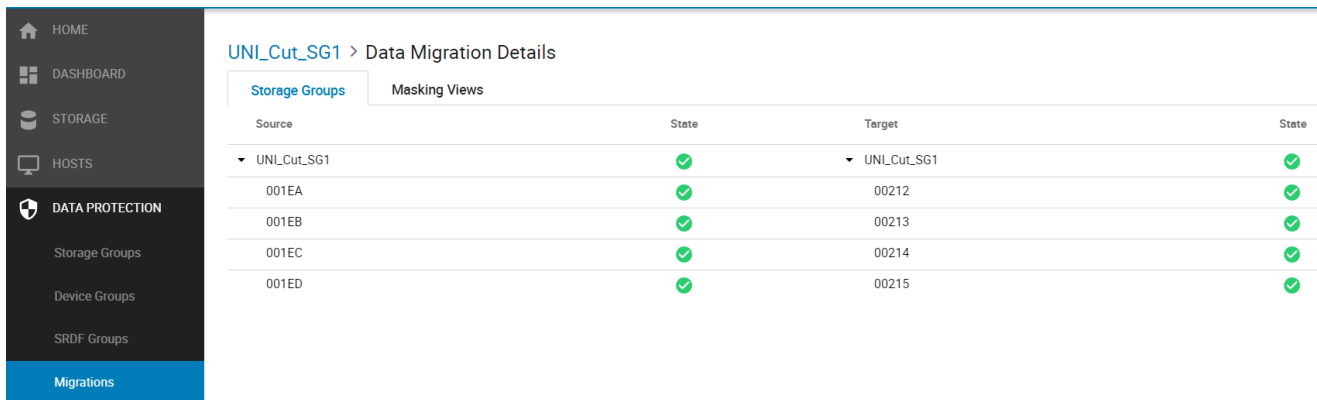
### CutoverReady and pass-through mode

Pass-through mode allows the host to write to or read from either the source or target array. Any write that is sent to the target array is sent over the SRDF link and serviced by the source array. No data is kept on the target array while in a CutoverReady state.

The CutoverReady state is a transitional state. The devices should only be in a CutoverReady state and using pass-through mode for as long as it takes to check that the Create has succeeded properly, to run the host rescan, and to run the Cutover operation.

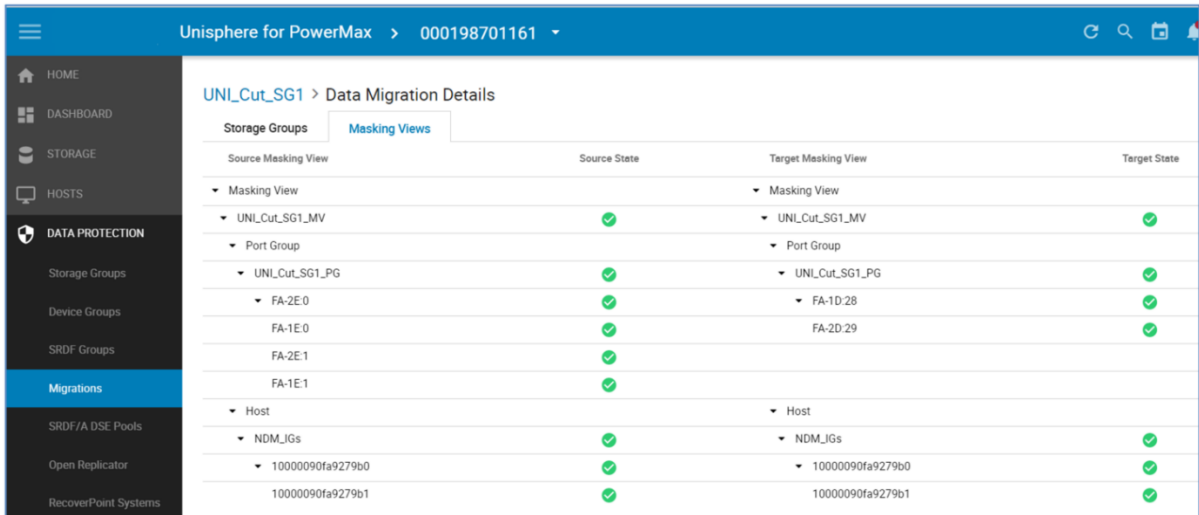


Double-clicking the NDM session displays the **Migration Details** view. This example below shows the individual devices involved in the session. Under the target tab we can see devices 200 through 203 have been created on the target side by the NDM create. The **State** also shows a live status of each of the devices involved. A device without a green tick should be investigated for potential problems before continuing.

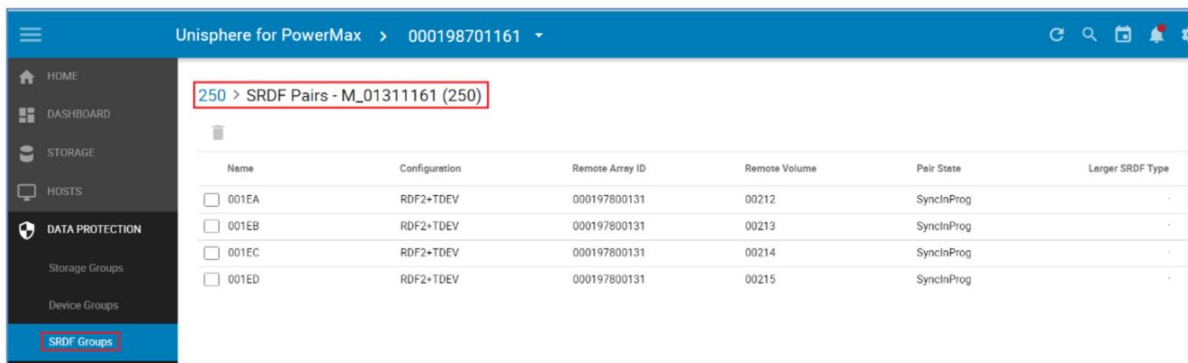


Selecting the **Masking Views** tab opens up the pane outlining the masking and masking elements involved on both sides of the NDM session. This screen is useful for troubleshooting any issues with the migration such as an unplanned or unauthorized manipulation of any of the NDM elements hindering the progress to the commit stage. This is highlighted in the **State** not showing a green tick.





The SRDF Groups tab in Data Protection shows the SRDF relationships established as part of the Create Command. The Devices 212 to 215 in the example below have been created as replicas of the source devices on the target array. The SyncInProg state can be ignored in this case as the data transfer from Source to Target has yet to start.



**View paths to new devices**

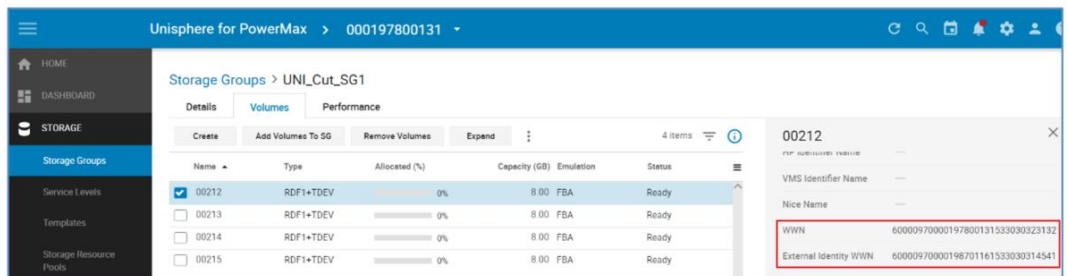
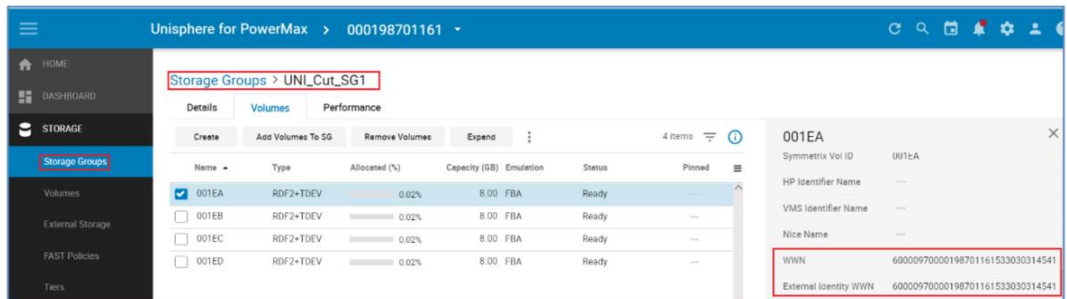
Viewing the same device (1EA) from the multipathing software **post rescan** shows the highlighted extra paths online to the target array (in this case two extra paths) this depends on the zoning setup. It also displays the source and target SIDs and the device numbers involved for these paths. This highlights the WWNs on the LUNs to appear as a single device with just extra paths. Prior to version 6.2, PowerPath was not aware of the NDM process so the dual SIDs and devices IDs were not visible.

```

pplicensevmaxcse:~ # rpowermt display dev=emcpower204 host=10.
Pseudo name=emcpower204
Symmetrix ID=000198701161, 000197800131
Logical device ID=01EA, 00212
Device WWN=60000970000198701161533030314541
Standard UID=naa.60000970000198701161533030314541
type=Conventional; state=alive; policy=SymmOpt; queued-IOs=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-IOs Errors
=====
3 vmhba6             C0:T3:L0   FA  2d:29 active  alive   0     0
4 vmhba2             C0:T12:L0  FA  1d:28 active  alive   0     0
3 vmhba6             C0:T19:L1  FA  1e:01 active  alive   0     0
3 vmhba6             C0:T8:L1   FA  2e:01 active  alive   0     0
3 vmhba6             C0:T20:L1  FA  2e:00 active  alive   0     0
3 vmhba6             C0:T17:L1  FA  1e:00 active  alive   0     0
4 vmhba2             C0:T17:L1  FA  2e:00 active  alive   0     0
4 vmhba2             C0:T13:L1  FA  1e:00 active  alive   0     0
4 vmhba2             C0:T15:L1  FA  2e:01 active  alive   0     0
4 vmhba2             C0:T2:L1   FA  1e:01 active  alive   0     0
    
```

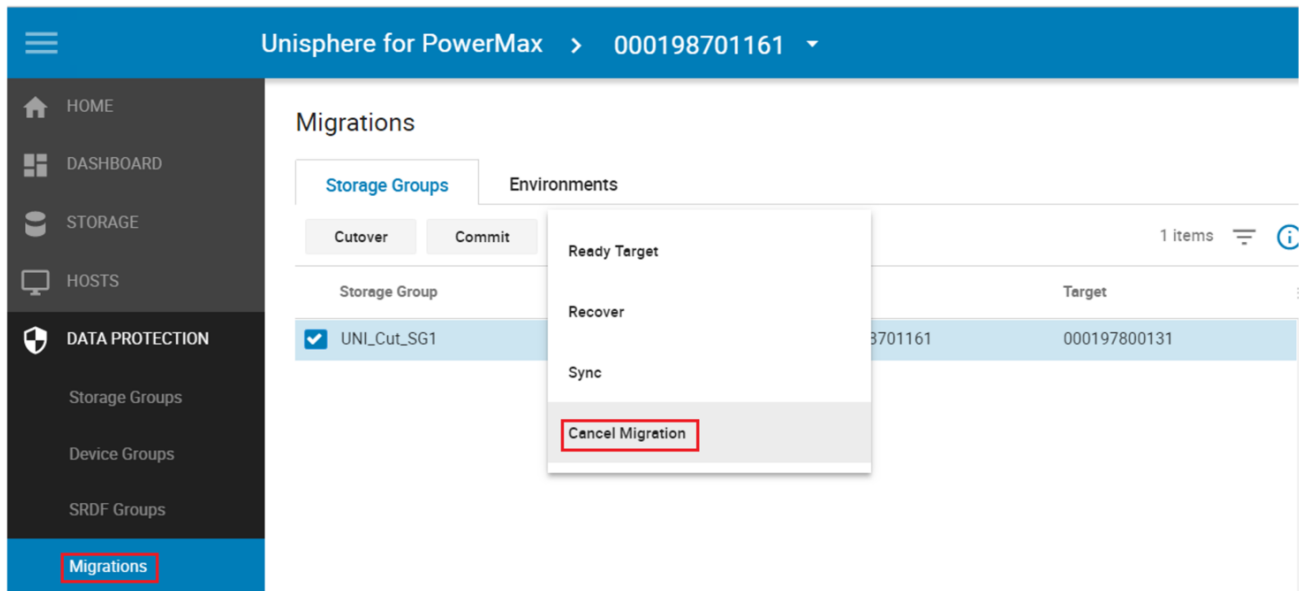
**Examine IDs of source and target**

Viewing the devices WWNs after the Create process shows that the source device has a **WWN** and **External** (host visible) WWN with the same value. However, the target's WWN and External WWNs differ. The target's External WWN has inherited the WWN of the source in order to appear logically as the same device to the host and picked up my multipathing software as an extra path to the same device.



**Cancel a migration**

At any point before a commit operation is run on a Storage Group, a migration that has not been committed can be canceled. In this example, the Cancel command is occurring before the cutover. This operation does not require the `--revert` flag because processing has not moved to the target array.



Canceling a migration removes the storage and groups provisioned for the migration on the target array, releases resources allocated by Solutions Enabler to perform the migration, and places the source devices into the state they were in before the Create operation was run. It does not affect the replication pathways put in place with the environment setup.

---

**Note:** It is best practice to run a rescan on the host after a Cancel to clear up any dead or invalid paths.

---

### *Cutover migration session*

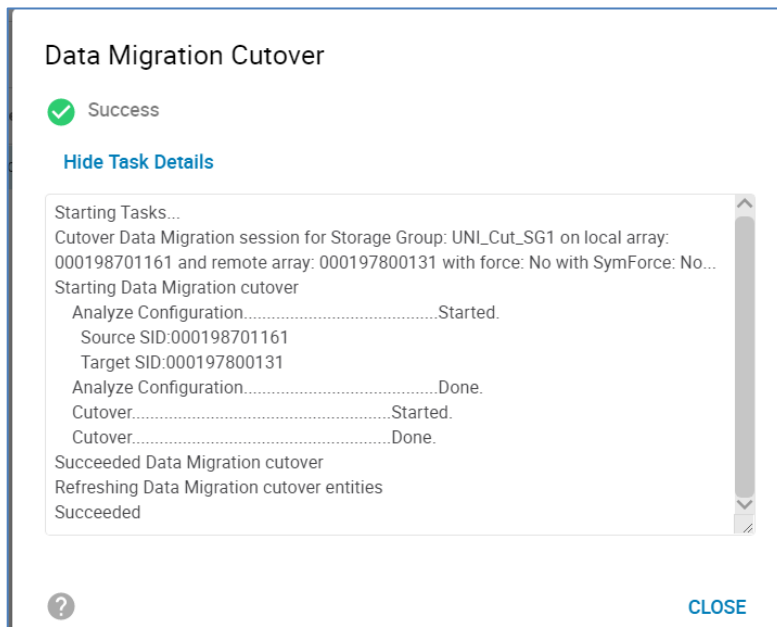
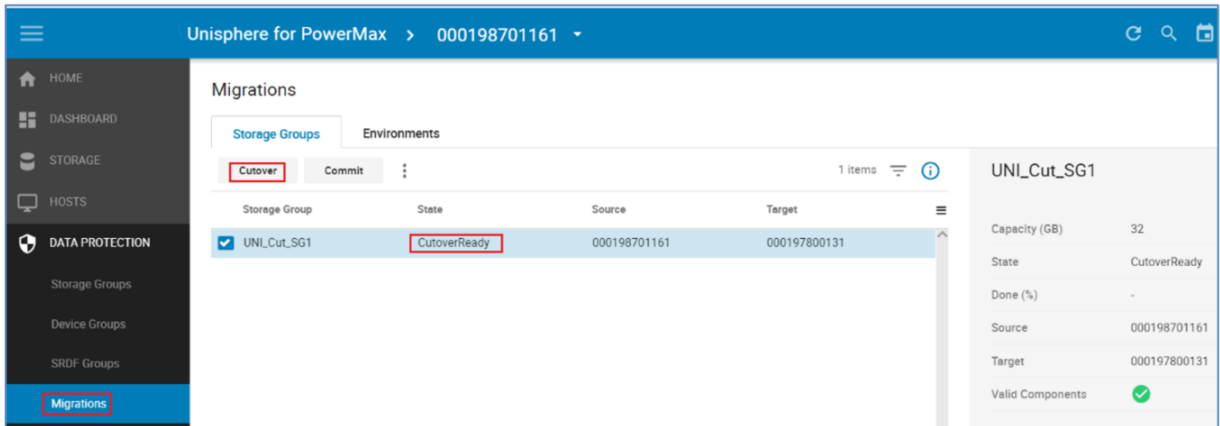
---

**Note:** A Host rescan that will result in the permanent removal of the now “inactive” paths should not be undertaken post Cutover, this will limit the ability for the migration to be seamlessly canceled and normal operation reverted to the source array. In the case of multiple concurrent NDM sessions sharing the same host the same rule should apply across all sessions when issuing rescans.

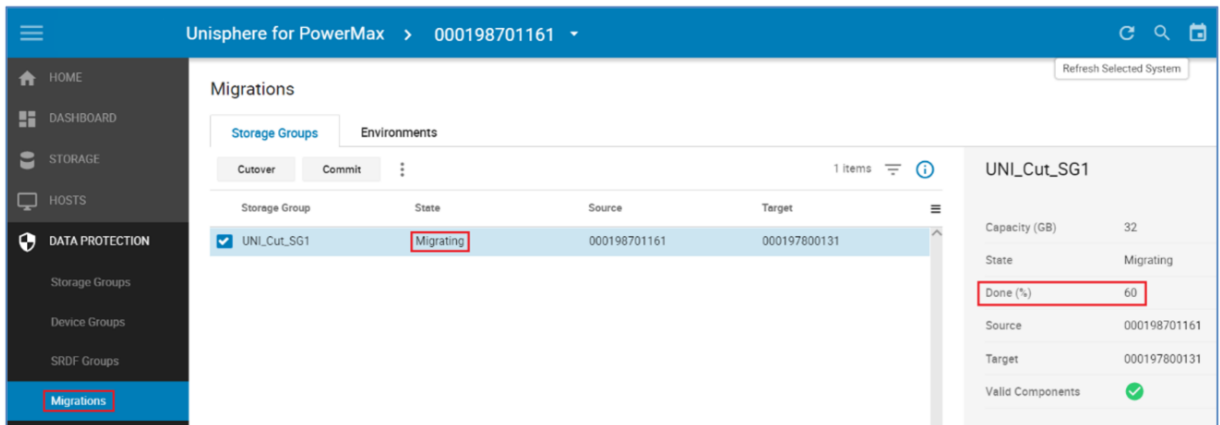
---

The normal operation following a successful Create is a Cutover.

A cutover operation moves the target devices out of pass-through mode, initiates data synchronization from the source to the target and makes the host paths to the source array inactive so that all I/Os are being serviced by the target array. From an SRDF point of view this initiates a full SRDF Restore on the devices.



When the cutover operation completes, the data copy begins. The session is in a Migrating state and will remain in that state until either the pairs are cutover to the new array or other action is taken.



In the example above, the migration session is 60% Copied. Copy time is affected by a number of factors such as:

- How busy the array is overall
- How many RDF paths are part of the NDM environment
- Whether the resources are shared between regular SRDF operations and NDM copies
- Amount of concurrent NDM session ongoing
- Amount of application I/O

The screenshot shows the 'Migrations' page in the Unisphere for PowerMax interface. The 'Storage Groups' tab is selected, and a table lists migration sessions. The session 'UNI\_Cut\_SG1' is highlighted, with its state 'CutoverSync' circled in red. The source and target storage groups are '000198701161' and '000197800131' respectively. A sidebar on the right shows details for 'UNI\_Cut\_SG1', including Capacity (32 GB), State (CutoverSync), and Done (%) (0%).

Storage Group	State	Source	Target
UNI_Cut_SG1	CutoverSync	000198701161	000197800131

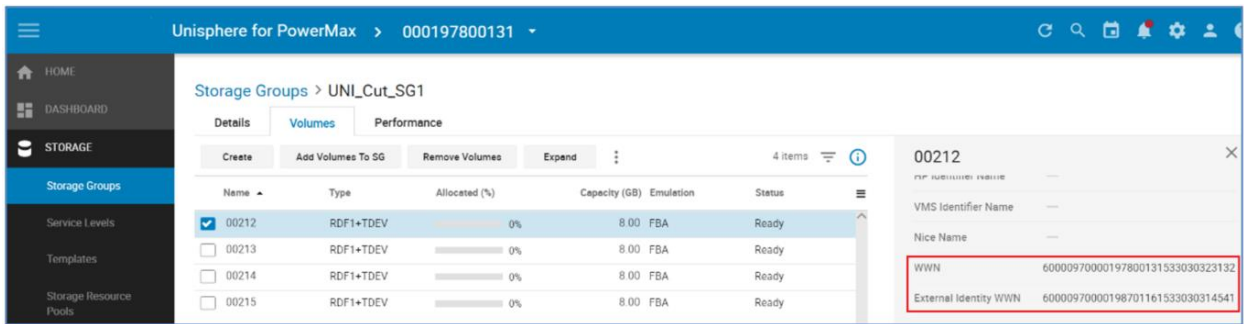
**Note:** The Done % shows 100% only for a very brief period of time. When the session transitions to a CutoverSync state it is always 100% synchronized.

### Examine devices post CutoverSync

The device IDs used on the source and target devices have not changed following the Cutover operation. The target devices are still using the effective WWN of the source devices. The source devices still have the same native and effective IDs.

The screenshot shows the 'Storage Groups > UNI\_Cut\_SG1' page in the Unisphere for PowerMax interface. The 'Volumes' tab is selected, and a table lists volumes. The volume '001EA' is highlighted, with its WWN circled in red. A sidebar on the right shows details for '001EA', including Symmetrix Vol ID (001E:A), HP Identifier Name, VMS Identifier Name, Nice Name, WWN (60000970000198701161533030314541), and External Identity WWN (60000970000198701161533030314541).

Name	Type	Allocated (%)	Capacity (GB)	Emulation	Status	Pinned
001EA	RDF2+TDEV	0.02%	8.00	FBA	Ready	--
001EB	RDF2+TDEV	0.02%	8.00	FBA	Ready	--
001EC	RDF2+TDEV	0.02%	8.00	FBA	Ready	--
001ED	RDF2+TDEV	0.02%	8.00	FBA	Ready	--



However, the host no longer has access to the source array for I/O processing. All the host I/O is being handled by the target array and is replicating by SRDF/Sync back to the source array. This means that application processing can revert non-disruptively to the source array without data loss or downtime.

Examining the multipathing following the Cutover the paths to the source array have transitioned to a Dead state. The Masking view remains to the source but the paths are in a suspended state so unavailable for host traffic.

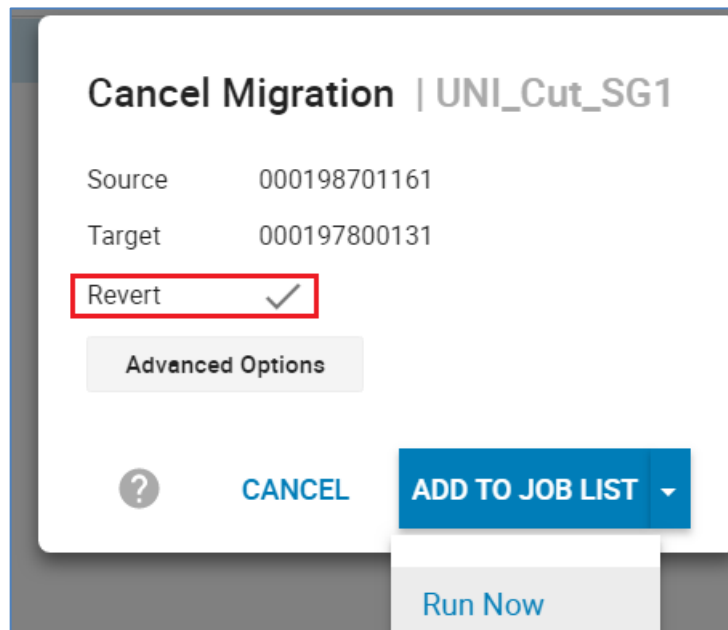
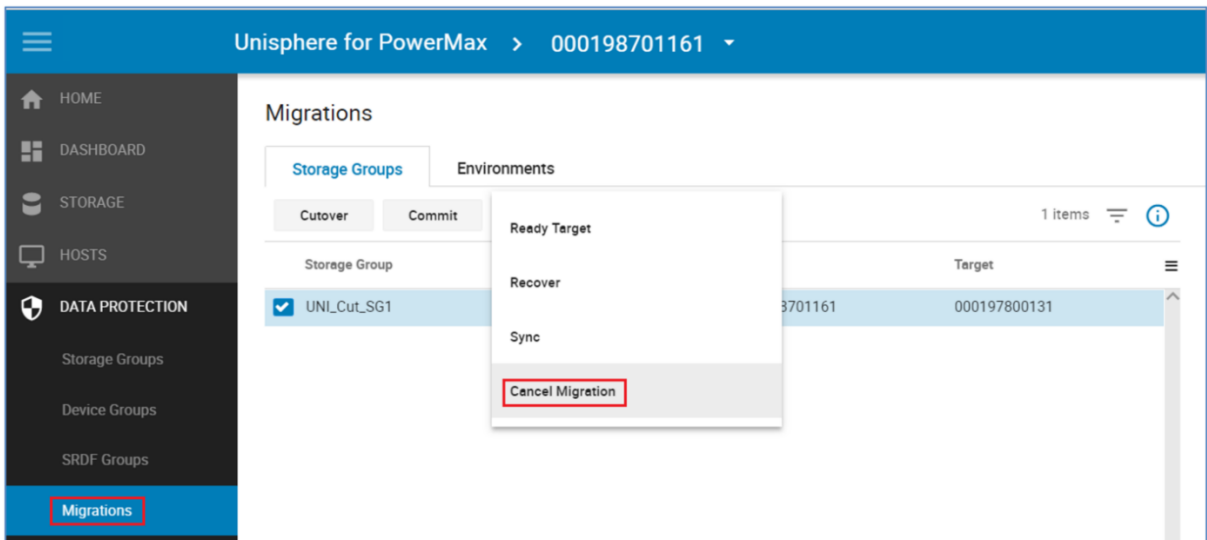
```
pplicensevmaxcse:~ # rpowermt display dev=emcpower204 host=10.60.136.146
Pseudo name=emcpower204
Symmetrix ID=000198701161, 000197800131
Logical device ID=01EA, 00212
Device WWN=60000970000198701161533030314541
Standard UID=naa.60000970000198701161533030314541
type=Conventional; state=alive; policy=SymmOpt; queued-IOs=0

----- Host -----
### HW Path          I/O Paths          - Stor -  -- I/O Path --  -- Stats ---
                               Interf.  Mode   State  Q-IOs Errors
-----
 3 vmhba6            C0:T3:L0          FA  2d:29 active  alive    0    0
 4 vmhba2            C0:T12:L0         FA  1d:28 active  alive    0    0
 3 vmhba6            C0:T19:L1         FA  1e:01 active  dead     0    3
 3 vmhba6            C0:T8:L1          FA  2e:01 active  dead     0    3
 3 vmhba6            C0:T20:L1         FA  2e:00 active  dead     0    2
 3 vmhba6            C0:T17:L1         FA  1e:00 active  dead     0    2
 4 vmhba2            C0:T17:L1         FA  2e:00 active  dead     0    2
 4 vmhba2            C0:T13:L1         FA  1e:00 active  dead     0    2
 4 vmhba2            C0:T15:L1         FA  2e:01 active  dead     0    3
 4 vmhba2            C0:T2:L1          FA  1e:01 active  dead     0    3
```

**Revert to the source array**

Because the migration is not permanent until the **commit** operation is run, after a cutover, the migration can still be canceled and reverted to the source array. To revert back to the source array following a cutover, a cancel is run with the -revert option.

The revert option moves the processing back to the source array and the cancel removes all of the target side entities created for the migration. This operation leaves the environment in the same state as it was prior to the create operation. The revert operation may take some time to run as the system waits for deallocations to complete on the target LUNs before completing. Also, as the revert is running, that the paths to the source array are active again. This is monitored by the VMAX/PowerMax, which waits for the rediscovery before proceeding.



By default, the Revert Flag is selected once the session as reached a Migrating or CutoverSync state.

#### ***Perform a host rescan***

Following the cancel revert operation, the host paths to the target array are no longer available. The host systems administrator performs a rescan to remove the dead paths to the target array.



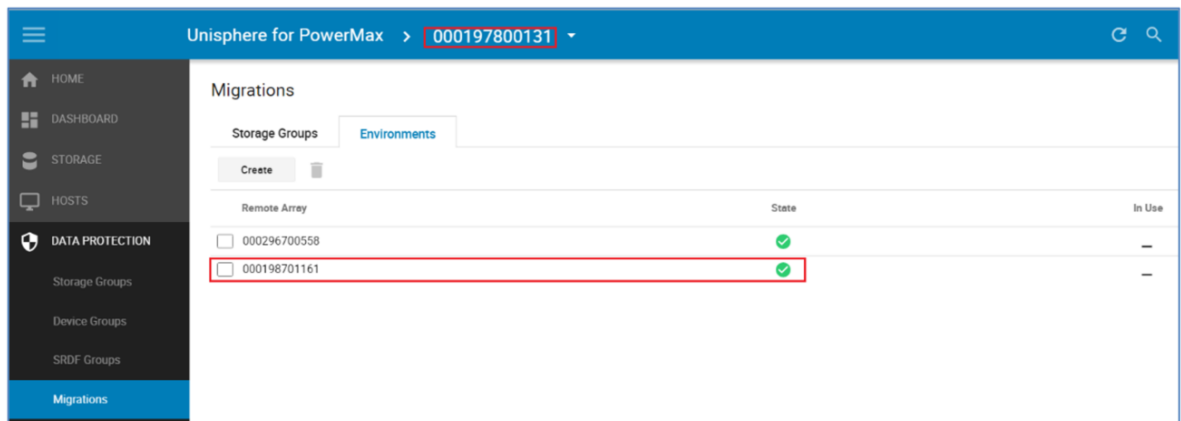
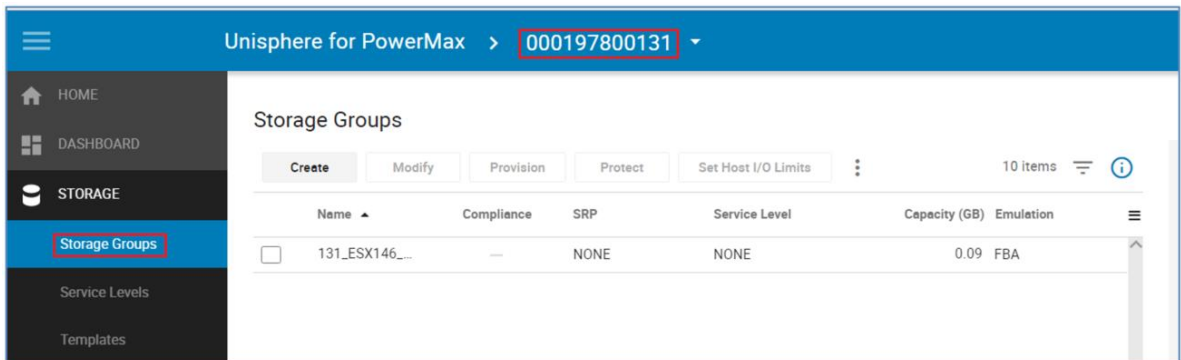
*Examine the devices post cancel with revert*

```
pplicensevmaxcse:~ # rpowermt display dev=emcpower204 host=10.60.136.146
Pseudo name=emcpower204
Symmetrix ID=000198701161
Logical device ID=01EA
Device WWN=60000970000198701161533030314541
Standard UID=naa.60000970000198701161533030314541
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
```

###	HW Path	Host	I/O Paths	Stor Interf.	I/O Mode	Path State	Stats Q-I/Os	Errors
3	vmhba6		C0:T19:L1	FA 1e:01	active	alive	0	3
3	vmhba6		C0:T8:L1	FA 2e:01	active	alive	0	3
3	vmhba6		C0:T20:L1	FA 2e:00	active	alive	0	2
3	vmhba6		C0:T17:L1	FA 1e:00	active	alive	0	2
4	vmhba2		C0:T17:L1	FA 2e:00	active	alive	0	2
4	vmhba2		C0:T13:L1	FA 1e:00	active	alive	0	2
4	vmhba2		C0:T15:L1	FA 2e:01	active	alive	0	3
4	vmhba2		C0:T2:L1	FA 1e:01	active	alive	0	3

In this example, the paths to the source array are active once again and the paths to the target array no longer exist.

The SG on the target array has also been removed but the NDM environment remains for any future NDM session between the source and target arrays.



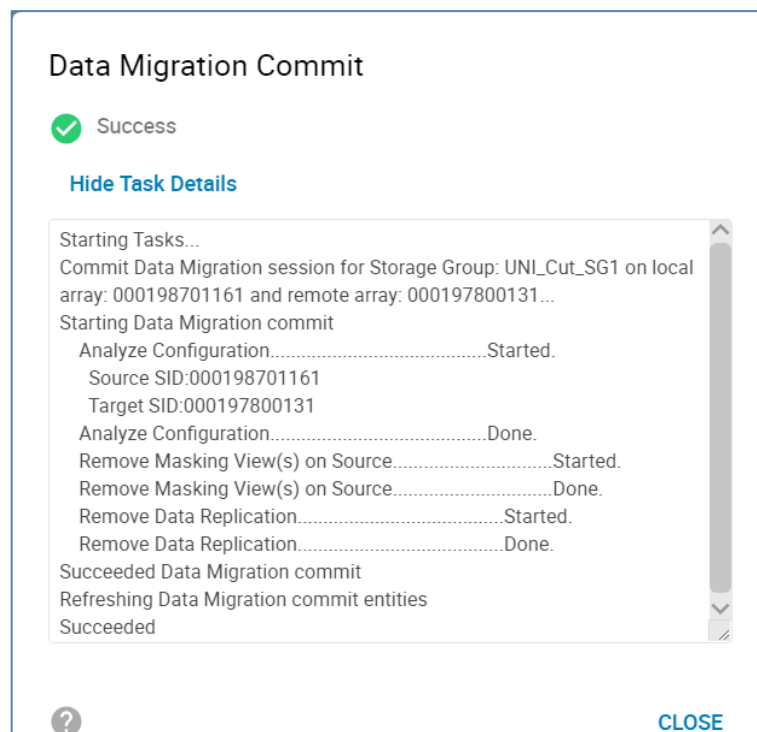
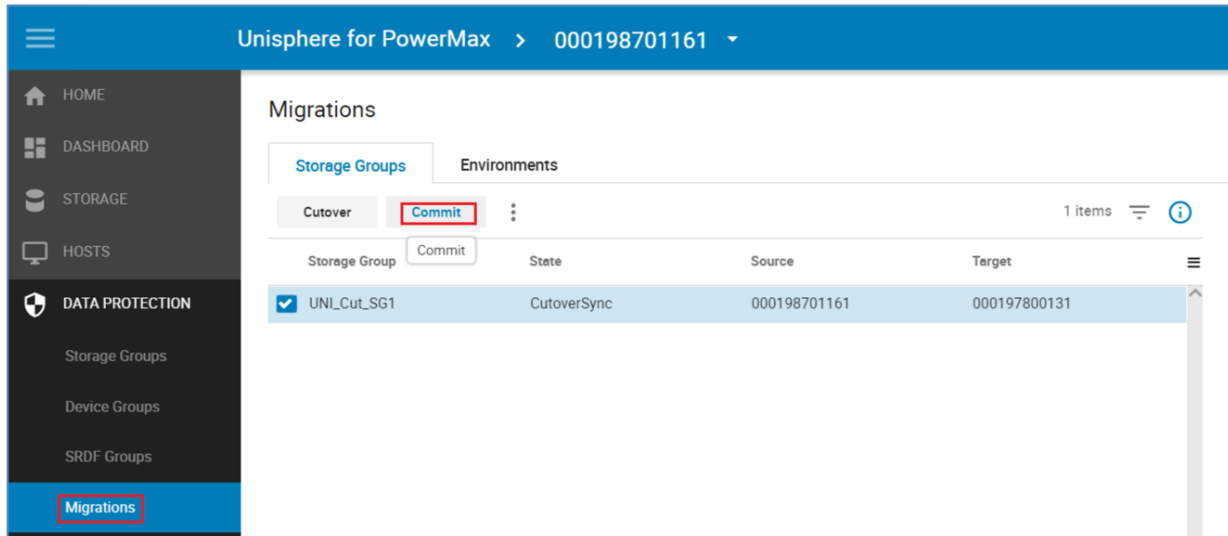


### Commit migration session

The normal operation following a successful Cutover is a Commit.

**Note:** Once the Commit has been run reverting to the Source array will not be possible non-disruptively. This is effectively the point of no return for Source array reversion.

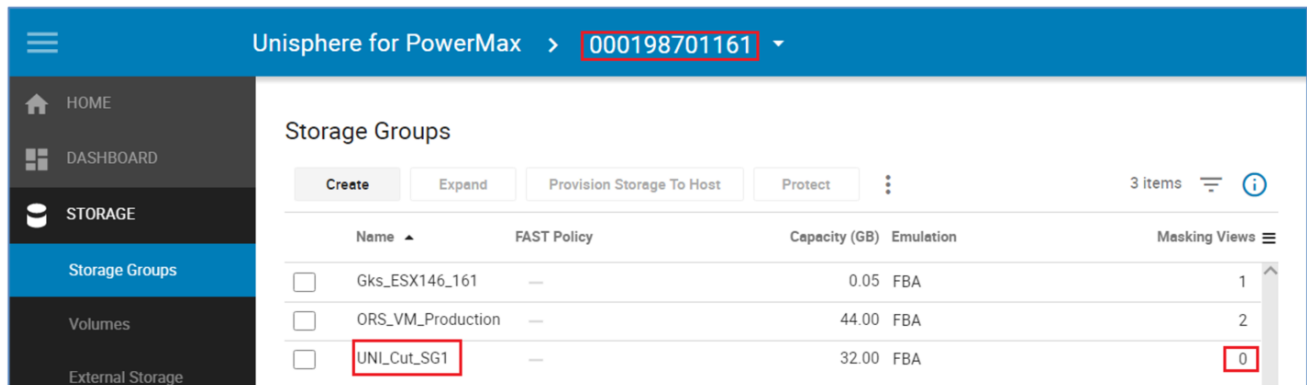
When the data copy is complete, the migration can be committed. The commit operation completes the migration by removing the migrated application resources from the source array and temporary system resources used for the migration. To commit, the state of the migration sessions must be CutoverSync or CutoverNoSync.



Once the commit is complete, replication between the source and target arrays is terminated. The source devices are no longer visible to a host because the masking has been removed. The source device IDs have also been permanently swapped with the target device IDs.

**Perform a host rescan**

After commit operation completes the systems administrator can issue a host rescan to allow the host to clean up the dead paths left by the removed paths to the source array. This host rescan is operating-system specific and also should include a rescan using the host multipathing software if it must be performed separately from the host rescan, as with PowerPath. See [Appendix A: Host multipathing software notes](#) for more details on the host multipathing software.



Because the commit completes the migration and removes all of the source side masking, there are no longer any paths seen to the source array. The logical device field only contains the target device and the Symmetrix ID contains the serial number of the target only.

```

oplicensevmaxcse:~ # rpowermt display dev=emcpower204 host=10.60.136.146
Pseudo name=emcpower204
Symmetrix ID=000197800131
Logical device ID=00212
Device WWN=60000970000198701161533030314541
Standard UID=naa.60000970000198701161533030314541
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths  Interf. Mode  State  Q-I/Os Errors
=====
  3 vmhba6            C0:T3:L0   FA  2d:29 active  alive   0    0
  4 vmhba2            C0:T12:L0 FA  1d:28 active  alive   0    0
    
```

**Compare the SGs and LUN WWNs post commit**

Following the commit operation, each device presents the opposite device ID. The source device now presents the target device ID as its external identity and the target presents the source device ID as its external identity. These changes are permanent and will persist across system power cycles and even the deletion and recreation of the devices. In other words, if device 1EA is deleted, when re-created, it will still show the identity of device 0212.

Therefore, the WWNs have effectively been reverse and the “spoofing” is permanent.

Unisphere for PowerMax > 000198701161

Storage Groups > UNI\_Cut\_SG1

Details | **Volumes** | Performance

Create | Add Volumes To SG | Remove Volumes | Expand | 4 items

Name	Type	Allocated (%)	Capacity (GB)	Emulation	Status	Pinned
<input checked="" type="checkbox"/> 001EA	TDEV	0.02%	8.00	FBA	Ready	—
<input type="checkbox"/> 001EB	TDEV	0.02%	8.00	FBA	Ready	—
<input type="checkbox"/> 001EC	TDEV	0.02%	8.00	FBA	Ready	—
<input type="checkbox"/> 001ED	TDEV	0.02%	8.00	FBA	Ready	—

001EA

VMS Identifier Name —

Nice Name —

WWN 60000970000198701161533030314541

External Identity WWN 60000970000197800131533030323132

DG Name —

Unisphere for PowerMax > 000197800131

Storage Groups > UNI\_Cut\_SG1

Details | **Volumes** | Performance

Create | Add Volumes To SG | Remove Volumes | Expand | 4 items

Name	Type	Allocated (%)	Capacity (GB)	Emulation	Status	Pinned
<input checked="" type="checkbox"/> 00212	TDEV	0.00%	8.00	FBA	Ready	—
<input type="checkbox"/> 00213	TDEV	0.00%	8.00	FBA	Ready	—
<input type="checkbox"/> 00214	TDEV	0.00%	8.00	FBA	Ready	—
<input type="checkbox"/> 00215	TDEV	0.00%	8.00	FBA	Ready	—

00212

VMS Identifier Name —

Nice Name —

WWN 60000970000197800131533030323132

External Identity WWN 60000970000198701161533030314541

DG Name —

**Using device identifiers to track migrated devices**

One example of keeping track of devices as they migrate from one array to another is to tag the LUNs. This tagging will persist throughout an NDM session.

Unisphere for PowerMax > 000198701161

Storage Groups > UNI\_Cut\_SG1

Details | **Volumes** | Performance

Create | Add Volumes To SG | Remove Volumes | Expand

Name	Type	Allocated (%)	Capacity (GB)	Emulation
<input checked="" type="checkbox"/> 001EA	TDEV	0.02%	8.00	FBA
<input type="checkbox"/> 001EB	TDEV	0.02%	8.00	FBA
<input type="checkbox"/> 001EC	TDEV	0.02%	8.00	FBA
<input type="checkbox"/> 001ED	TDEV	0.02%	8.00	FBA

001EA

Masking Info

- Emulation
- Set Volume Attributes
- Set Volume Identifier**
- Set Volume Status
- Replication QOS

### Set Volume Identifiers | 001EA

Volume Identifier Name  
Application1\_Dev1

Volume HP Identifier Name

Volume VMS Identifier Name

? CANCEL ADD TO JOB LIST ▾

Unisphere for PowerMax > 000198701161 ▾

Storage Groups > UNI\_Cut\_SG1

< Details Volumes Performance

Create Add Volumes To SG Remove Volumes Expand 4 Items ⓘ

Name	Type	Allocated (%)	Capacity (GB)	Emulation	Status	Pinned
<input checked="" type="checkbox"/> 001EA:Application1_Dev1	TDEV	0.02%	8.00	FBA	Ready	—
<input type="checkbox"/> 001EB	TDEV	0.02%	8.00	FBA	Ready	—
<input type="checkbox"/> 001EC	TDEV	0.02%	8.00	FBA	Ready	—
<input type="checkbox"/> 001ED	TDEV	0.02%	8.00	FBA	Ready	—

Unisphere for PowerMax > 000197800131 ▾

Storage Groups > UNI\_Cut\_SG1

< Details Volumes Performance

Create Add Volumes To SG Remove Volumes Expand 4 Items ⓘ

Name	Type	Allocated (%)	Capacity (GB)	Emulation	Status
<input checked="" type="checkbox"/> 00212:Application1_Dev1	TDEV	0.00%	8.00	FBA	Ready
<input type="checkbox"/> 00213	TDEV	0.00%	8.00	FBA	Ready
<input type="checkbox"/> 00214	TDEV	0.00%	8.00	FBA	Ready
<input type="checkbox"/> 00215	TDEV	0.00%	8.00	FBA	Ready

00212

Masking Info 4

Storage Groups 1

SRP 1

FBA Front End Paths 2

RDF Info 0

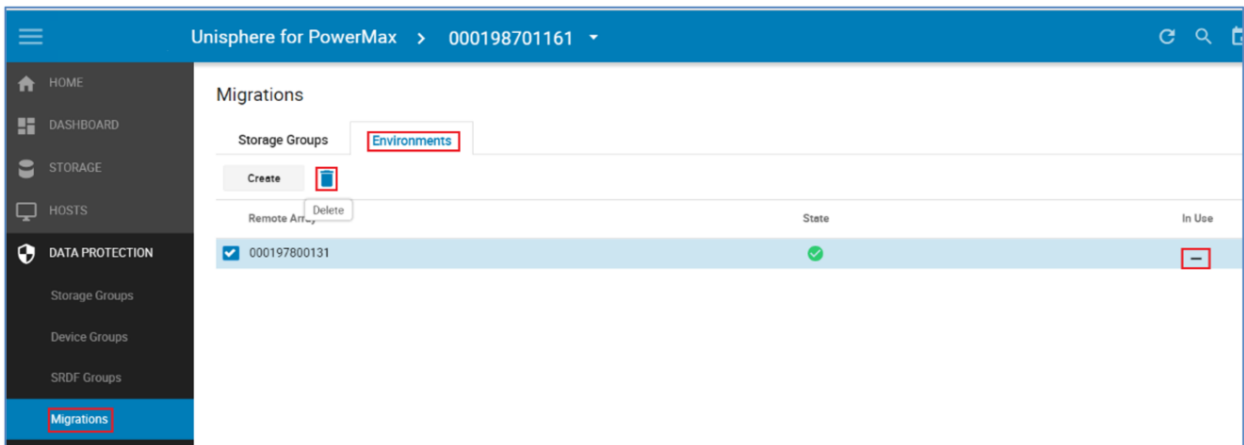
Volume Name 00212:Application1\_Dev1

The Tag or Identifier allows a quick reference of a device's origin or application.

### Remove the NDM environment

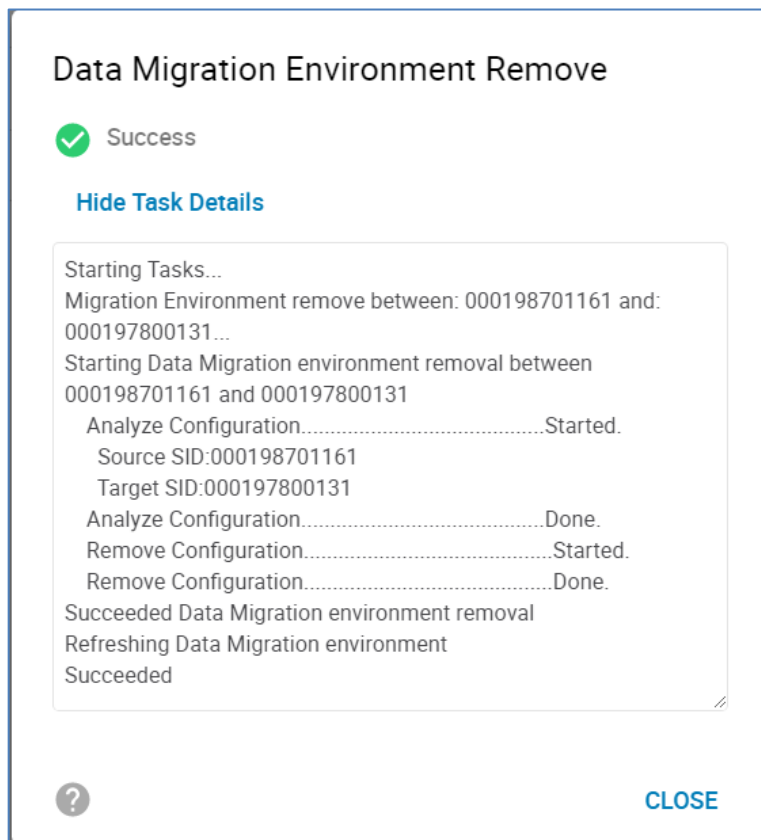
The environment remove operation removes the replication pathway configured by the environment setup operation and removes the resources that were configured to support

NDM on the source and target arrays. On successful completion of an environment remove operation, only running an environment setup operation is allowed. An environment removal operation should occur only after all migrations from the source array have been completed.



To delete the environment, select the relevant environment from the Environments tab. The In Use symbol identifies if an NDM session is active on this environment.

Once the environment remove operation is complete, the NDM process is complete.



**Using Solutions Enabler 9.x**

This example shows the migration of five devices (PHYSICALDRIVE0 to PHYSICALDRIVE4) as this syminq display shows.

These devices were previously added as RDMS to the VM with VMware vSphere.

```
C:\Program Files\EMC\SYMCLI\bin>syminq
```

Device		Product			Device	
Name	Type	Vendor	ID	Rev	Ser Num	Cap (KB)
\\.\PHYSICALDRIVE0		VMware	Virtual disk	1.0	N/A	125829120
\\.\PHYSICALDRIVE1		EMC	SYMMETRIX	5876	61001EE000	13632000
\\.\PHYSICALDRIVE2		EMC	SYMMETRIX	5876	61001EF000	13632000
\\.\PHYSICALDRIVE3		EMC	SYMMETRIX	5876	61001F1000	13632000
\\.\PHYSICALDRIVE4		EMC	SYMMETRIX	5876	61001F0000	13632000

An example of the multipathing setup using device 1EE shows what the pathing looks like prior to the NDM create and the host rescan. For each of the four volumes here there are eight paths to the source array which are all alive and available for host use. At this point, there are no paths to the target array even though our zoning should be in place before the NDM create.

```
pplicensevmaxcse:~ # rpowermt display dev=emcpower208 host=10.60.136.146
Pseudo name=emcpower208
Symmetrix ID=000198701161
Logical device ID=01EE
Device WWN=60000970000198701161533030314545
Standard UID=naa.60000970000198701161533030314545
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
```

###	HW Path	Host	I/O Paths	Stor Interf.	I/O Mode	Path State	Stats Q-I/Os	Errors
3	vmhba6		C0:T19:L1	FA 1e:01	active	alive	0	0
3	vmhba6		C0:T8:L1	FA 2e:01	active	alive	0	0
3	vmhba6		C0:T20:L1	FA 2e:00	active	alive	0	0
3	vmhba6		C0:T17:L1	FA 1e:00	active	alive	0	0
4	vmhba2		C0:T17:L1	FA 2e:00	active	alive	0	0
4	vmhba2		C0:T13:L1	FA 1e:00	active	alive	0	0
4	vmhba2		C0:T15:L1	FA 2e:01	active	alive	0	0
4	vmhba2		C0:T2:L1	FA 1e:01	active	alive	0	0

Finally, the Storage Group in which our Application due to be migrated resides consists of four LUNS of 13 GB each, Symm Devices 0x1EE through 0x1F1.

```
C:\Windows\system32>symmsg -sid 161 show SE_Cut_SG1
```

```
Devices (4):
{
-----
Sym      Device
Dev      Pdev Name      Config      Attr  Sts      Cap
(MB)
-----
001EE    N/A            TDEV        NR    13313
001EF    N/A            TDEV        NR    13313
001F0    N/A            TDEV        NR    13313
001F1    N/A            TDEV        NR    13313
}
```

### Environment setup

The Environment Setup configures the migration environment that will be required to migrate any application from the source array to the target array. It confirms that both the source and target arrays can support the NDM operations. This includes ensuring that a usable replication pathway for data migration is available between the source and target arrays and creating an SRDF group for the migration.

The setup only must be run once. When the migration pathways and SRDF group are configured, all storage groups on the source array can be migrated until all migrations from the source to the target have been completed.

```
symdm environment -src_sid <SN of Source> -tgt_sid <SN of target> -setup
```

```
C:\Windows\system32>symdm environment -src_sid 161 -tgt_sid 131 -setup
A DM 'Environment Setup' operation is in progress. Please wait...

  Analyze Configuration.....Started.
    Source SID:000198701161
    Target SID:000197800131
  Analyze Configuration.....Done.
  Setup Configuration.....Started.
  Setup Configuration.....In Progress.
  Setup Configuration.....In Progress.
  Setup Configuration.....In Progress.
  Setup Configuration.....In Progress.
  Setup Configuration.....Done.

The DM 'Environment Setup' operation successfully executed.
```

To validate the Environment is working correctly the following can be run at any point.

```
symdm environment -src_sid <SN of Source> -tgt_sid <SN of target> -validate
```

```
C:\Windows\system32>symdm environment -src_sid 161 -tgt_sid 131 -validate
A DM 'Environment Validate' operation is in progress. Please wait...

  Analyze Configuration.....Validated.

The DM 'Environment Validate' operation successfully executed.
```

To view all the environments currently configured as well as their status from all available local and remote arrays connected run the following:

```
symdm -environment list
```

```

C:\Windows\system32>symdm -environment list

Symmetrix ID: 000197600156

The migration session environment is not configured

Symmetrix ID: 000197800131

Remote SymmID  Status
-----
000198701161  OK
000296700558  OK

Symmetrix ID: 000197900111

The migration session environment is not configured

Symmetrix ID: 000198701161

Remote SymmID  Status
-----
000197800131  OK

```

Once the environment setup is complete, the migration sessions can be created.

---

**Note:** NDM Environment setup creates RDF links between the Source and Target using one port per Director for each zoned link. However, post-setup the user can add extra links manually using: `symrdf modifygrp -rdfg 250 -add dir xx -remote_dir xx`

---

### **Create migration and validate migration session**

Solutions Enabler examines a specific application's storage on the source array and automatically provisions equivalent storage on the target array. The target devices are assigned the identity of the source devices and are configured in pass-through mode which allows the data to be accessed through both the source and target devices.

Prior to running the create operation, the target array resources can be validated to ensure that the target array has the resources required to configure the migration sessions and the migration infrastructure exists on both arrays.

```

symdm create -src_sid <SN of Source> -tgt_sid <SN of target> -sg
<SG to be Migrated> -tgt_SRP <Target SRP> -validate

```

```

C:\Windows\system32>symdm create -src_sid 161 -tgt_sid 131 -sg SE_Cut_SG1 -validate

Execute 'Validate Create' operation on SG 'SE_Cut_SG1' (y/[n])? y

A DM 'Validate Create' operation is
in progress for storage group 'SE_Cut_SG1'. Please wait...

Analyze Configuration.....Validated.
Create Storage Group(s) on Target.....Validated.
Duplicate Device(s) on Target.....Validated.
Create Initiator Group(s) on Target.....Validated.
Create Port Group(s) on Target.....Validated.
Create Masking View(s) on Target.....Validated.

The DM 'Validate Create' operation successfully executed for
storage group 'SE_Cut_SG1'.

```



```
symdm create -src_sid <SN of Source> -tgt_sid <SN of target> -sg
<SG to be Migrated> -tgt_srp <target SRP> -tgt_pg <target PG>
```

```
C:\Windows\system32>symdm create -src_sid 161 -tgt_sid 131 -sg SE_Cut_SG1
Execute 'Create' operation on SG 'SE_Cut_SG1' (y/[n])? y
A DM 'Create' operation is
in progress for storage group 'SE_Cut_SG1'. Please wait...

Analyze Configuration.....Started.
  Source SID:000198701161
  Target SID:000197800131
Analyze Configuration.....Done.
Set Dynamic RDF attribute on Source Device(s).....Not Needed.
Create Storage Group(s) on Target.....Started.
Create Storage Group(s) on Target.....Done.
Duplicate Device(s) on Target.....Started.
Preparing for device create on Target.....Started.
Preparing for device create on Target.....Done.
Duplicate Device(s) on Target.....Done.
Create Initiator Group(s) on Target.....Started.
Create Initiator Group(s) on Target.....Done.
Create Port Group(s) on Target.....Started.
Create Port Group(s) on Target.....Done.
Setup Data Replication.....Started.
Setup Data Replication.....Done.
Create Masking View(s) on Target.....Started.
Create Masking View(s) on Target.....Done.

The DM 'Create' operation successfully executed for
storage group 'SE_Cut_SG1'.
```

Should the Create command fail for whatever reason (RDF link failure, Target array configuration lock), the session may partially complete some of the elements of the task. The example below shows how the RDF links had an issue just as the links were being established leaving the session in a CREATEFAILED state. The details of the create command show all the elements had been created on the target array such as Port Group, Devices, and the Storage Group.

Using the Recover Command, the NDM process tries to continue from where it left off while verifying all the completed steps are still valid. Should this fail the CANCEL command cleans up all the elements created and returns the array to the state it was before the CREATE was issued.

```
symdm -sid <SN of Source or Target> -sg <SG to be Migrated>
recover
```

```
C:\Windows\system32>symdm create -src_sid 161 -tgt_sid 131 -sg SE_Cut_SG1 -validate
Execute 'Validate Create' operation on SG 'SE_Cut_SG1' (y/[n])? y
A DM 'Validate Create' operation is
in progress for storage group 'SE_Cut_SG1'. Please wait...

The migration session is not in a valid DM state for this operation

C:\Windows\system32>symdm -sid 161 list
Symmetrix ID : 000198701161
Storage Group      Source Array      Target Array      State      Total Capacity (GB)  Done (%)
-----
SE_Cut_SG1         000198701161 000197800131 CreateFailed 52.0      N/A

C:\Windows\system32>symdm -sid 161 -sg SE_Cut_SG1 recover
Execute 'Recover' operation on SG 'SE_Cut_SG1' (y/[n])? y
A DM 'Recover' operation is
in progress for storage group 'SE_Cut_SG1'. Please wait...

Analyze Configuration.....Started.
  Source SID:000198701161
  Target SID:000197800131
Analyze Configuration.....Done.
Set Dynamic RDF attribute on Source Device(s).....Not Needed.
Create Storage Group(s) on Target.....Not Needed.
Duplicate Device(s) on Target.....Not Needed.
Create Initiator Group(s) on Target.....Not Needed.
Create Port Group(s) on Target.....Not Needed.
Setup Data Replication.....Started.
Setup Data Replication.....Done.
Recover Data Replication.....Started.
Recover Data Replication.....In Progress.
Recover Data Replication.....Done.
Create Masking View(s) on Target.....Started.
Create Masking View(s) on Target.....Done.
Update Device State.....Started.
Update Device State.....Done.

The DM 'Recover' operation successfully executed for
storage group 'SE_Cut_SG1'.
```

All active NDM sessions can be monitored using the list command, there are variations of this command that give a finer level of detail on the session such as using the `-v` parameter.

```
symdm -sid <SN of Source or Target> list
```

```
C:\Windows\system32>symdm -sid 161 list
Symmetrix ID : 000198701161
Storage Group      Source Array      Target Array      State      Total Capacity (GB)  Done (%)
-----
SE_Cut_SG1         000198701161 000197800131 Created 52.0      N/A
```

```
symdm -sid <SN of Source or Target> -sg list -v
```

```

C:\Windows\system32>symdm -sid 161 list -v

Symmetrix ID      : 000198701161
Storage Group     : SE_Cut_SG1
Source Array      : 000198701161
Target Array      : 000197800131
Migration State   : CutoverReady
Total Capacity (GB) : 52.0
Done (%)          : N/A

Source Configuration: OK
{
  Storage Groups (1) : OK
  Masking Views (1)  : OK
  Initiator Groups (1) : OK
  Port Groups (1)    : OK
}

Target Configuration: OK
{
  Storage Groups (1) : OK
  Masking Views (1)  : OK
  Initiator Groups (1) : OK
  Port Groups (1)    : OK
}

Device Pairs (4): OK

```

To explore in much finer detail the individual elements of the NDM session use the `-v` in combination with the `-detail` parameter. Cropped example:

```
symdm -sid <SN of Source or Target> -sg list -v -detail
```

```

C:\Windows\system32>symdm -sid 161 list -v -detail

Symmetrix ID      : 000198701161
Storage Group     : SE_Cut_SG1
Source Array      : 000198701161
Target Array      : 000197800131
Migration State   : CutoverReady
Total Capacity (GB) : 52.0
Done (%)          : N/A

Source Configuration: OK
{
  Storage Groups (1): OK
  {
    Name      : SE_Cut_SG1
    Status    : OK
    {
      Dev
      -----
      001EE:001F1
    }
  }
}

Masking Views (1): OK
{
  Masking View Name  Initiator Group  Port Group  Status
  -----
  SE_Cut_SG1_MV     NDM_IGs         SE_Cut_SG1_PG  OK
}

Initiator Groups (1): OK
{
  Name      : NDM_IGs
  Status    : OK
  {
    Initiator WWN
    -----
    10000090fa9279b0
    10000090fa9279b1
  }
}

Port Groups (1): OK
{
  Name      : SE_Cut_SG1_PG
  Status    : OK
  {
    Dirport Status
    -----
    02E:000 OK
  }
}

```

In summary, the Create Command undertakes the following tasks:

- Creates a Storage Group on the Target array (name must not already exist in the target array) with the same name as the Source SG.
- Creates duplicate devices on the target array to match those on the Storage Group.
- Creates an initiator group using Initiators with entries in the login history table.
- Creates a Port group. (if one does not already exist)
- Effective (external) WWNs of the device created on the target are copied from the WWNs of the host devices.
- Creates a masking view to the host from the target array.

---

**Note:** During a Cutover NDM migration, the source of the migration is an R2 or an R21 device (if there is existing SRDF DR replication from the source device) and the target is an R1 device. This is different to basic SRDF operations and is required to allow DR protection during a migration using a cascaded SRDF configuration.

---

### **Perform a host rescan**

Once the Create Command has completed the NDM session is in a Created state.

```
C:\Windows\system32>symdm -sid 161 list
Symmetrix ID : 000198701161
Storage Group      Source      Target      State      Total      Done
                   Array       Array       State      Capacity  (%)
-----
SE_Cut_SG1         000198701161 000197800131 Created      52.0     N/A
```

From this state, we cannot continue to the next step of the process without rescanning the host to pick up the new paths to the target devices.

The systems administrator must issue a host rescan to allow the host to discover the paths to the newly created devices.

Devices go into a CutoverReady state from a Created state after the host rescan is performed and the target devices are discoverable. After this is complete, I/O issued by the application will be directed to either the source or the target arrays through the host multipathing software. This is possible because the target devices are in pass-through mode. [Appendix A: Host multipathing software notes](#) has more details on host multipathing software settings.

### **CutoverReady and pass-through mode**

Pass-through mode allows the host to write to or read from either the source or target array. Any write that is sent to the target array is sent across the SRDF link and serviced by the source array. No data is kept on the target array while in a CutoverReady state.

The CutoverReady state is a transitional state. The devices should only be in a CutoverReady state and using pass-through mode for as long as it takes to check that the Create has succeeded properly, to run the host rescan, and to run the Cutover operation.

```
symdm -sid <SN of Source or Target> list
```

Storage Group	Source Array	Target Array	State	Total Capacity (GB)	Done (%)
SE_Cut_SG1	000198701161	000197800131	CutoverReady	52.0	N/A

Examining the multipathing software setup there are extra paths online to the target array (in this case two extra paths) after the rescan the number of extra paths is dependent on the zoning setup. It also displays the source and target SIDs and the device numbers involved for these paths. This highlights the fact the WWNs on the LUNs to appear are being spoofed as a single device with just extra paths. Prior to version 6.2 PowerPath was not aware of the NDM process so the dual SIDs and devices IDs were not visible.

```
pplicensevmaxcse:- # rpowermt display dev=emcpower208 host=██████████
Pseudo name=emcpower208
Symmetrix ID=000198701161, 000197800131
Logical device ID=01EE, 00216
Device WWN=60000970000198701161533030314545
Standard UID=naa.60000970000198701161533030314545
type=Conventional; state=alive; policy=SymmOpt; queued-IOs=0
```

###	HW Path	Host	I/O Paths	Stor Interf.	I/O Path Mode	State	Q-IOs	Errors
4	vmhba2		C0:T12:L4	FA 1d:28	active	alive	0	0
3	vmhba6		C0:T3:L4	FA 2d:29	active	alive	0	0
3	vmhba6		C0:T19:L1	FA 1e:01	active	alive	0	0
3	vmhba6		C0:T8:L1	FA 2e:01	active	alive	0	0
3	vmhba6		C0:T20:L1	FA 2e:00	active	alive	0	0
3	vmhba6		C0:T17:L1	FA 1e:00	active	alive	0	0
4	vmhba2		C0:T17:L1	FA 2e:00	active	alive	0	0
4	vmhba2		C0:T13:L1	FA 1e:00	active	alive	0	0
4	vmhba2		C0:T15:L1	FA 2e:01	active	alive	0	0
4	vmhba2		C0:T2:L1	FA 1e:01	active	alive	0	0

### ***Examine the device pairings and the identities following a create***

The create operation automatically configures matching volumes on the target array. These volumes are the same size and configuration, though they unlikely to have the same VMAX volume numbers. Following the create operation the four new volumes on the target array are 21A through 21D. Volume 1EE and volume 21A, for example, are paired for NDM operations.

```
symdm -sid <SN of Source or Target> -sg <SG to be Migrated> list -
v -pairs_info -detail
```

```
C:\Windows\system32>symdm -sid 161 -sg SE_Cut_SG1 list -v -pairs_info -detail
Symmetrix ID      : 000198701161
Storage Group     : SE_Cut_SG1
Source Array      : 000198701161
Target Array      : 000197800131

Migration State   : CutoverReady
Total Capacity (GB) : 52.0
Done (%)          : N/A

Device Pairs (4): OK
{
  Source          Target
  Dev    Status   Dev    Status
  -----
  001EE  OK        0021A  OK
  001EF  OK        0021B  OK
  001F0  OK        0021C  OK
  001F1  OK        0021D  OK
}
```

The RDF pairing information can also be seen in the symdev list output.

symdev -sid <SN of Source or Target> list

```
C:\Program Files\EMC\SYMCLI\bin>symdev -sid 161 list
Symmetrix ID: 000198701161
```

Device Name		Dir		Device		
Sym	Physical	SA :P	Config	Attribute	Sts	Cap (MB)
00028	Not Visible	***:***	2-Wav Mir	N/Grp'd ACLX	RW	3
001EE	\\.\PHYSICALDRIVE1	02E:000	RDF2+TDEV	N/Grp'd	RW	13313
001EF	\\.\PHYSICALDRIVE2	02E:000	RDF2+TDEV	N/Grp'd	RW	13313
001F0	\\.\PHYSICALDRIVE4	02E:000	RDF2+TDEV	N/Grp'd	RW	13313
001F1	\\.\PHYSICALDRIVE3	02E:000	RDF2+TDEV	N/Grp'd	RW	13313

```
C:\Program Files\EMC\SYMCLI\bin>symdev -sid 131 list
Symmetrix ID: 000197800131
```

Device Name		Dir		Device		
Sym	Physical	SA :P	Config	Attribute	Sts	Cap (MB)
00001	Not Visible	???:???	TDEV	N/Grp'd ACLX	RW	6
0021A	Not Visible	***:***	RDF1+TDEV	N/Grp'd	RW	13313
0021B	Not Visible	***:***	RDF1+TDEV	N/Grp'd	RW	13313
0021C	Not Visible	***:***	RDF1+TDEV	N/Grp'd	RW	13313
0021D	Not Visible	***:***	RDF1+TDEV	N/Grp'd	RW	13313

This again highlights the setup from an RDF standpoint with the target devices adopting an RDF R1 personality and the source taking on the identity of an RDF R2.

**Note:** These personalities vary based on the presence of DR from the source side to a third array in which case the R2 is an R21. The addition of DR to another array from the Target side is possible once the Cutover command has been issued. This changes the personality of the Target devices from a R1 to a R11.

Looking in more detail the effective and native (Internal and External) WWNs of the devices. This shows how “spoofing” these values allows us to manipulate the multipathing software into believing it has just had new paths to the same devices added rather than paths to a completely different array.

```
symdev -sid <SN of Source> show <Source Device>
```

```
C:\Program Files\EMC\SYMCLI\bin>symdev -sid 161 show 1EE

Device Physical Name      : \\.\PHYSICALDRIVE1
Device Symmetrix Name     : 001EE
Device Serial ID         : 61001EE000
Symmetrix ID             : 000198701161

Number of RAID Groups    : 0
Encapsulated Device      : No
Encapsulated WWN        : N/A
Encapsulated Device Flags: None

Encapsulated Array ID    : N/A
Encapsulated Device Name : N/A
Attached BCV Device      : N/A

Attached VDEV TGT Device : N/A

Vendor ID                : EMC
Product ID               : SYMMETRIX
Product Revision         : 5876
Device WWN               : 60000970000198701161533030314545
Device Emulation Type    : FBA
Device Defined Label Type: N/A
Device Defined Label     : N/A
Device Sub System Id     : 0x0001
Cache Partition Name     : DEFAULT_PARTITION
Bound Pool Name          : Sata_Pool

Device Block Size        : 512

Device Capacity
{
  Cylinders      :      14200
  Tracks         :      213000
  512-byte Blocks :    27264000
  MegaBytes     :      13313
  KiloBytes     :    13632000

  Geometry Limited : No
}

Device External Identity
{
  Device WWN      : 60000970000198701161533030314545
}
```

From the example above of a source device we see the Device WWN, the device it was born with, and the external WWN, the WWN presented to the host remains the same at this stage of the process.



```
symdev -sid <SN of Target> show <Source Device>
```

```
C:\Program Files\EMC\SYMCLI\bin>symdev -sid 131 show 21A

Device Physical Name      : Not Visible

Device Symmetrix Name    : 0021A
Device Serial ID         : N/A
Symmetrix ID             : 000197800131

Number of RAID Groups    : 0
Encapsulated Device      : No
Encapsulated WWN         : N/A
Encapsulated Device Flags: None

Encapsulated Array ID    : N/A
Encapsulated Device Name : N/A
Attached BCV Device      : N/A

Attached VDEV TGT Device : N/A

Vendor ID                 : EMC
Product ID                : SYMMETRIX
Product Revision          : 5978
Device WWN                 : 600009700001978001315330323141
Device Emulation Type     : FBA
Device Defined Label Type : N/A
Device Defined Label      : N/A
Device Sub System Id      : N/A
Cache Partition Name     : N/A
Bound Pool Name           : SRP_1

Device Block Size        : 512

Device Capacity
{
  Cylinders      :      7100
  Tracks         :     106500
  512-byte Blocks :    27264000
  MegaBytes      :      13313
  KiloBytes      :    13632000

  Geometry Limited : No
}

Device External Identity
{
  Device WWN      : 60000970000198701161533030314545
}
```

In contrast this example shows the target device. We see the WWN the device was born with and the WWN that it is presenting to the host differs. The device it is presenting to the host is inherited from the Source device thereby appearing as the same device to multipathing software.

Having both device presenting the same WWN means host I/O can use both source and target as its I/O path. However, at this point we are in Pass-Through mode so no data is stored on the target. It is merely passed through over the SRDF link to the source where the I/O is processed as usual.

---

**Note:** Due to the extra latency added to I/O that is experiencing the “double hop” if sent down a target path, it is not recommended that the migration session remains in a CutoverReady state for longer than is necessary.

---



### Cancel a migration

At any point before a Commit operation is run on a Storage Group, a migration that has not been committed can be canceled. In this example, the cancel is occurring before the cutover. This operation does not require the `-revert` flag because processing has not moved to the target array.

```
symdm -sid <SN of Source or Target> -sg <SG to be Migrated> cancel
```

```
C:\Windows\system32>symdm -sid 161 list
Symmetrix ID : 000198701161

Storage Group          Source Array      Target Array      State              Total Capacity (GB)  Done (%)
-----
SE_Cut_SG1            000198701161  000197800131  CutoverReady      52.0  N/A

C:\Windows\system32>symdm -sid 161 -sg SE_Cut_SG1 cancel
Execute 'Cancel' operation on SG 'SE_Cut_SG1' (y/[n])? y

A DM 'Cancel' operation is
in progress for storage group 'SE_Cut_SG1'. Please wait...

  Analyze Configuration.....Started.
    Source SID:000198701161
    Target SID:000197800131
  Analyze Configuration.....Done.
  Remove Masking View(s) on Target.....Started.
  Remove Masking View(s) on Target.....In Progress.
  Remove Masking View(s) on Target.....Done.
  Remove Data Replication.....Started.
  Remove Data Replication.....Done.
  Remove Port Group(s) on Target.....Started.
  Remove Port Group(s) on Target.....Done.
  Remove Initiator Group(s) on Target.....Started.
  Remove Initiator Group(s) on Target.....Done.
  Remove Duplicate Device(s) on Target.....Started.
  Wait for deallocation to complete.....Started.
  Wait for deallocation to complete.....Done.
  Remove Duplicate Device(s) on Target.....In Progress.
  Remove Duplicate Device(s) on Target.....Done.
  Remove Storage Group(s) on Target.....Started.
  Remove Storage Group(s) on Target.....Done.

The DM 'Cancel' operation successfully executed for
storage group 'SE_Cut_SG1'.
```

Canceling a migration removes the storage and groups provisioned for the migration on the target array, releases resources allocated by Solutions Enabler to perform the migration, and places the source devices into the state they were in before the Create operation was run. It does not affect the replication pathways put in place with the environment setup.

---

**Note:** It is best practice to run a rescan on the host after a Cancel to clear up any dead or invalid paths.

---

### Cutover migration session

**Note:** A Host rescan that will result in the permanent removal of the now “inactive” paths should not be undertaken post Cutover, this will limit the ability for the migration to be seamlessly canceled and normal operation reverted to the source array. In the case of multiple concurrent NDM sessions sharing the same host the same rule should apply across all sessions when issuing rescans.

Assuming the previous Cancel was not undertaken, (or having canceled the migration, a new session was created), the host was rescanned and the session reached a CutoverReady state, and the Cutover command can be issued.

A cutover operation:

- Moves the target devices out of pass-through mode.
- Initiates data synchronization from the source to the target.
- Makes the host paths to the source array inactive. The target array is now servicing all I/O requests.

```
symdm -sid <SN of Source or Target> -sg <Sg to be Migrated>
cutover
```

```
C:\Windows\system32>symdm -sid 161 list
Symmetrix ID : 000198701161

Storage Group      Source      Target      State      Total      Done
Array              Array              Capacity   (%)
-----
SE_Cut_SG1         000198701161 000197800131 CutoverReady 52.0  N/A

C:\Windows\system32>symdm -sid 161 -sg SE_Cut_SG1 cutover
Execute 'Cutover' operation on SG 'SE_Cut_SG1' (y/[n])? y
A DM 'Cutover' operation is
in progress for storage group 'SE_Cut_SG1'. Please wait...

  Analyze Configuration.....Started.
    Source SID:000198701161
    Target SID:000197800131
  Analyze Configuration.....Done.
  Cutover.....Started.
  Cutover.....Done.

The DM 'Cutover' operation successfully executed for
storage group 'SE_Cut_SG1'.
```

When the Cutover operation completes, the data copy begins. The session is in a Migrating state and remains in that state until either the pairs are cutover to the new array or other action is taken. The data movement can be monitored using symdm list command. This command has options for displaying Storage Group, Masking View, Initiator Group, Port Group, and device pairs.

```
symdm -sid <SN of source> list
```

```
C:\Windows\system32>symdm -sid 161 list
Symmetrix ID : 000198701161
Storage Group      Source Array      Target Array      State      Total Capacity (GB)  Done (%)
-----
SE_Cut_SG1         000198701161 000197800131 Migrating      52.0      18
```

The list command used in this example can be used to see the copy progress to the target array.

In the example above, the migration session is 18% Copied. Copy time is affected by a number of factors such as:

- How busy the array is overall
- How many RDF paths are part of the NDM environment
- If the resources are shared between regular SRDF operations and NDM copies
- Amount of concurrent NDM session ongoing
- Amount of application I/O

```
C:\Windows\system32>symdm -sid 161 list
Symmetrix ID : 000198701161
Storage Group      Source Array      Target Array      State      Total Capacity (GB)  Done (%)
-----
SE_Cut_SG1         000198701161 000197800131 CutoverSync 52.0      N/A
```

---

**Note:** The symdm list command shows 100% done only for a very brief period of time. When the session transitions to a CutoverSync state it is always 100% synchronized.

---

### **Examine devices at CutoverSync**

The device IDs used on the source and target devices have not changed following the Cutover operation. The target devices are still using the effective WWN of the source devices. The source devices still have the same native and effective IDs.

However, the host no longer has access to the source array for I/O processing. All the Host I/O is being handled by the target array and is replicating using SRDF/s back to the source array. This enables reversion of application processing non-disruptively to the source array without data loss or downtime.

Examining the multipathing following the Cutover, the paths to the source array have transitioned to a Dead state. The Masking view remains to the source but the paths are in a suspended state so unavailable for host traffic.

```

pplicensevmaxcse:~ # rpowermt display dev=emcpower208 host=
Pseudo name=emcpower208
Symmetrix ID=000198701161, 000197800131
Logical device ID=01EE, 00216
Device WWN=60000970000198701161533030314545
Standard UID=naa.60000970000198701161533030314545
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf. Mode   State   Q-I/Os Errors
=====
  4 vmhba2            C0:T12:L4  FA  1d:28 active  alive   0      0
  3 vmhba6            C0:T3:L4   FA  2d:29 active  alive   0      0
  3 vmhba6            C0:T19:L1  FA  1e:01 active  dead    0      1
  3 vmhba6            C0:T8:L1   FA  2e:01 active  dead    0      1
  3 vmhba6            C0:T20:L1  FA  2e:00 active  dead    0      1
  3 vmhba6            C0:T17:L1  FA  1e:00 active  dead    0      1
  4 vmhba2            C0:T17:L1  FA  2e:00 active  dead    0      1
  4 vmhba2            C0:T13:L1  FA  1e:00 active  dead    0      1
  4 vmhba2            C0:T15:L1  FA  2e:01 active  dead    0      1
  4 vmhba2            C0:T2:L1   FA  1e:01 active  dead    0      1

```

symdev -sid <Source SN> show <Device>

```

C:\Program Files\EMC\SYMCLI\bin>symdev -sid 161 show 1EE

Device Physical Name      : \\.\PHYSICALDRIVE1
Device Symmetrix Name     : 001EE
Vendor ID                 : EMC
Product ID                : SYMMETRIX
Product Revision          : 5876
Device WWN                 : 60000970000198701161533030314545
Device Emulation Type     : FBA
Device Defined Label Type : N/A
Device Defined Label      : N/A
Device Sub System Id      : 0x0001
Cache Partition Name     : DEFAULT_PARTITION
Bound Pool Name           : Sata_Pool

Device Block Size         : 512

Device Capacity
{
  Cylinders                :      14200
  Tracks                   :      213000
  512-byte Blocks          :    27264000
  MegaBytes                 :      13313
  KiloBytes                 :    13632000

  Geometry Limited         : No
}

Device External Identity
{
  Device WWN                : 60000970000198701161533030314545
}

```

symdev -sid <Target SN> show <Device>

```

C:\Program Files\EMC\SYMCLI\bin>symdev -sid 131 show 21A

Device Physical Name      : Not Visible
Device Symmetrix Name    : 0021A
Product Revision         : 5978
Device WWN                : 60000970000197800131533030323141
Device Emulation Type    : FBA
Device Defined Label Type: N/A
Device Defined Label     : N/A
Device Sub System Id     : N/A
Cache Partition Name     : N/A
Bound Pool Name          : SRP_1

Device Block Size        : 512

Device Capacity
{
  Cylinders      :      7100
  Tracks         :     106500
  512-byte Blocks :    27264000
  MegaBytes      :      13313
  KiloBytes      :    13632000

  Geometry Limited : No
}

Device External Identity
{
  Device WWN      : 60000970000198701161533030314545
}

```

The device IDs used on the source and target devices have not changed following the Cutover operation. The target devices are still using the effective WWN of the source devices. The source devices still have the same native and effective IDs.

### ***Revert to the source array***

Because the migration is not permanent until the Commit operation is run, after a Cutover, the migration can still be canceled and reverted to the source array. To revert back to the source array following a Cutover, a Cancel operation is run with the `-revert` option.

The revert option moves the processing back to the source array and the cancel removes all of the target side entities created for the migration. The operation leaves the environment in the same state as it was prior to the create operation. The revert operation may take some time to run as the system waits for deallocations to complete on the target devices before completing. Also, as the revert is running, the paths to the source array become active again. This is monitored by the source and target, which waits for the rediscovery before proceeding.

```
symdm -sid <Source or Target SN> -sg <Migration SG> cancel -revert
```

```
C:\Windows\system32>symdm -sid 161 -sg SE_Cut_SG1 cancel -revert
Execute 'Cancel Revert' operation on SG 'SE_Cut_SG1' (y/[n])? y
A DM 'Cancel Revert' operation is
in progress for storage group 'SE_Cut_SG1'. Please wait...

Analyze Configuration.....Started.
  Source SID:000198701161
  Target SID:000197800131
Analyze Configuration.....Done.
Revert Data Replication.....Started.
Revert Data Replication.....In Progress.
Wait for host path discovery on Source.....In Progress.
Wait for host path discovery on Source.....In Progress.
Wait for host path discovery on Source.....In Progress.
Wait for host path discovery on Source.....In Progress.
Revert Data Replication.....Done.
Remove Masking View(s) on Target.....Started.
Remove Masking View(s) on Target.....In Progress.
Remove Masking View(s) on Target.....Done.
Remove Data Replication.....Started.
Remove Data Replication.....Done.
Remove Port Group(s) on Target.....Started.
Remove Port Group(s) on Target.....Done.
Remove Initiator Group(s) on Target.....Started.
Remove Initiator Group(s) on Target.....Done.
Remove Duplicate Device(s) on Target.....Started.
Wait for deallocation to complete.....Started.
Wait for deallocation to complete.....Done.
Remove Duplicate Device(s) on Target.....In Progress.
Remove Duplicate Device(s) on Target.....Done.
Remove Storage Group(s) on Target.....Started.
Remove Storage Group(s) on Target.....Done.

The DM 'Cancel Revert' operation successfully executed for
storage group 'SE_Cut_SG1'.
```

**Perform a host rescan and examine the devices**

Following the Cancel operation with the Revert option, the host paths to the target array are no longer available. The host systems administrator runs a rescan to remove the dead paths to the target array.

```
apllicensevmaxcse:~ # rpowermt display dev=emcpower208 host=10.60.136.146
Pseudo name=emcpower208
Symmetrix ID=000198701161, 000197800131
Logical device ID=01EE, 00216
Device WWN=60000970000198701161533030314545
Standard UID=naa.60000970000198701161533030314545
type=Conventional; state=alive; policy=SymmOpt; queued-IOs=0
=====
--- Host ---
### HW Path I/O Paths Interf. Mode State Q-IOs Errors
=====
4 vmhba2 CO:T12:L4 FA 1d:28 active dead 0 3
3 vmhba6 CO:T3:L4 FA 2d:29 active dead 0 3
3 vmhba6 CO:T19:L1 FA 1e:01 active alive 0 1
3 vmhba6 CO:T8:L1 FA 2e:01 active alive 0 1
3 vmhba6 CO:T20:L1 FA 2e:00 active alive 0 1
3 vmhba6 CO:T17:L1 FA 1e:00 active alive 0 1
4 vmhba2 CO:T17:L1 FA 2e:00 active alive 0 1
4 vmhba2 CO:T13:L1 FA 1e:00 active alive 0 1
4 vmhba2 CO:T15:L1 FA 2e:01 active alive 0 1
4 vmhba2 CO:T2:L1 FA 1e:01 active alive 0 1
```



```

pplicensevmaxcse:~ # rpowermt display dev=emcpower208 host=10.60.136.146
Pseudo name=emcpower208
Symmetrix ID=000198701161
Logical device ID=01EE
Device WWN=60000970000198701161533030314545
Standard UID=naa.60000970000198701161533030314545
type=Conventional; state=alive; policy=SymmOpt; queued-IOs=0
=====
--- Host ---
### HW Path I/O Paths Interf. Mode State Q-IOs Errors
-----
3 vmhba6 CO:T19:L1 FA 1e:01 active alive 0 1
3 vmhba6 CO:T8:L1 FA 2e:01 active alive 0 1
3 vmhba6 CO:T20:L1 FA 2e:00 active alive 0 1
3 vmhba6 CO:T17:L1 FA 1e:00 active alive 0 1
4 vmhba2 CO:T17:L1 FA 2e:00 active alive 0 1
4 vmhba2 CO:T13:L1 FA 1e:00 active alive 0 1
4 vmhba2 CO:T15:L1 FA 2e:01 active alive 0 1
4 vmhba2 CO:T2:L1 FA 1e:01 active alive 0 1

```

The identity of the target array has been completely removed from the PowerPath device following the rescan. Before that the identity remained and the paths were showing as dead.

The SG on the target array has also been removed but the NDM environment remains for any future NDM session between the source and target arrays.

```
symmsg -sid <SN of Target> list
```

```

C:\Windows\system32>symmsg -sid 131 list

      S T O R A G E   G R O U P S

Symmetrix ID:          000197800131

Storage Group Name    Flags  Number  Number  Child
                    EFM SLC Devices GKs   SGs
-----
131_ESX146_GK_SG     F.X ...    16     16     0

```

```
symdm -sid <SN of Source or Target> list -environment
```

```

C:\Windows\system32>symdm -sid 161 list -environment

Symmetrix ID: 000198701161

Remote SymmID  Status
-----
000197800131  OK

```

### Commit migration session

When the data copy is complete, the migration can be committed. The Commit operation completes the migration by removing the migrated application resources from the source array and temporary system resources used for the migration. The Commit operation requires that the state of the migration session is CutoverSync or CutoverNoSync.

**Note:** Once the Commit has completed, reverting to the Source array will not be possible non-disruptively. `symdm -sid <SN of Source or Target> -sg <SG to be migrated> commit`

```
symdm -sid <SN of Source or Target> -sg <SG to be Migrated> commit
```

```
Symmetrix ID : 000198701161
```

Storage Group	Source Array	Target Array	State	Total Capacity (GB)	Done (%)
SE_Cut_SG1	000198701161	000197800131	CutoverSync	52.0	N/A

```
C:\Windows\system32>symdm -sid 161 -sg SE_Cut_SG1 commit
Execute 'Commit' operation on SG 'SE_Cut_SG1' (y/[n])? y
A DM 'Commit' operation is
in progress for storage group 'SE_Cut_SG1'. Please wait...

  Analyze Configuration.....Started.
    Source SID:000198701161
    Target SID:000197800131
  Analyze Configuration.....Done.
  Remove Masking View(s) on Source.....Started.
  Remove Masking View(s) on Source.....Done.
  Remove Data Replication.....Started.
  Remove Data Replication.....Done.

The DM 'Commit' operation successfully executed for
storage group 'SE_Cut_SG1'.
```

Once the Commit operation is complete, replication between the source and target array ends. The source devices are no longer be visible to a host because the masking has been removed. The source device IDs have also been permanently swapped with the target device IDs.

### **Perform a host rescan**

After the commit operation completes, the systems administrator runs a host rescan so that the host can clean up the dead paths left by the removed paths to the source array. This host rescan is OS-specific and also should include a rescan using the host multipathing software. See [Appendix A: Host multipathing software notes](#) for more details on the host multipathing software.

The commit operation completes the migration and removes all of the source side masking. Therefore, there are no longer any paths seen to the source array.

The following shows the pre-rescan status:



```
pplicensevmaxcse:~ # rpowermt display dev=emcpower208 host=10.60.136.146
Pseudo name=emcpower208
Symmetrix ID=000198701161, 000197800131
Logical device ID=01EE, 0021A
Device WWN=60000970000198701161533030314545
Standard UID=naa.60000970000198701161533030314545
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
```

###	Host	HW Path	I/O Paths	Stor Interf.	I/O Path Mode	State	Stats Q-I/Os	Errors
4	vmhba2	C0:T12:L4	FA 1d:28	active	alive	0	0	
3	vmhba6	C0:T3:L4	FA 2d:29	active	alive	0	0	
3	vmhba6	C0:T19:L1	FA 1e:01	active	dead	0	2	
3	vmhba6	C0:T8:L1	FA 2e:01	active	dead	0	2	
3	vmhba6	C0:T20:L1	FA 2e:00	active	dead	0	2	
3	vmhba6	C0:T17:L1	FA 1e:00	active	dead	0	2	
4	vmhba2	C0:T17:L1	FA 2e:00	active	dead	0	2	
4	vmhba2	C0:T13:L1	FA 1e:00	active	dead	0	2	
4	vmhba2	C0:T15:L1	FA 2e:01	active	dead	0	2	
4	vmhba2	C0:T2:L1	FA 1e:01	active	dead	0	2	

The post-rescan status is displayed as follows:

```
pplicensevmaxcse:~ # rpowermt display dev=emcpower208 host=10.60.136.146
Pseudo name=emcpower208
Symmetrix ID=000197800131
Logical device ID=0021A
Device WWN=60000970000198701161533030314545
Standard UID=naa.60000970000198701161533030314545
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
```

###	Host	HW Path	I/O Paths	Stor Interf.	I/O Path Mode	State	Stats Q-I/Os	Errors
4	vmhba2	C0:T12:L4	FA 1d:28	active	alive	0	0	
3	vmhba6	C0:T3:L4	FA 2d:29	active	alive	0	0	

### Compare the SGs and LUN WWNs post commit

Following the Commit operation, each device presents the opposite device ID. The source device now presents the target device ID as its external identity and the target presents the source device ID as its external identity. These changes are permanent and will persist across system power cycles and even the deletion and recreation of the devices. In other words, if device 1EE is deleted and then re-created, it still retains the identity of device 021A.

Therefore, the WWNs have effectively been reversed and the spoofing is permanent.

```
symdev -sid <SN of Source> show <Device to be Migrated>
```

```
C:\Windows\system32>symdev -sid 161 show 1EE

Device Physical Name      : Not Visible

Device Symmetrix Name     : 001EE
Device Serial ID         : N/A
Symmetrix ID             : 000198701161

Number of RAID Groups    : 0
Encapsulated Device      : No
Encapsulated WWN         : N/A
Encapsulated Device Flags: None

Encapsulated Array ID    : N/A
Encapsulated Device Name : N/A
Attached BCV Device      : N/A

Attached VDEV TGT Device : N/A

Vendor ID                : EMC
Product ID               : SYMMETRIX
Product Revision         : 5876
Device WWN               : 60000970000198701161533030314545

Device Capacity
{
  Cylinders      :      14200
  Tracks         :      213000
  512-byte Blocks :    27264000
  MegaBytes      :      13313
  KiloBytes      :    13632000

  Geometry Limited : No
}

Device External Identity
{
  Device WWN      : 60000970000197800131533030323141
}
```

The example shows that the Source device has inherited the WWN of the target device. This is the device WWN that it will display to a host should it be masked. This allows the customer to reuse the array and the devices previously migrated without the risk of data loss in the case where the same devices were masked into the SAN of the target devices.

```

C:\Windows\system32>symdev -sid 131 show 21A

Device Physical Name      : Not Visible

Device Symmetrix Name     : 0021A
Device Serial ID         : N/A
Symmetrix ID              : 000197800131

Number of RAID Groups    : 0
Encapsulated Device      : No
Encapsulated WWN         : N/A
Encapsulated Device Flags: None

Encapsulated Array ID    : N/A
Encapsulated Device Name : N/A
Attached BCV Device      : N/A

Attached VDEV TGT Device : N/A

Vendor ID                 : EMC
Product ID                : SYMMETRIX
Product Revision          : 5978
Device WWN                 : 60000970000197800131533030323141

Device Capacity
{
  Cylinders      :      7100
  Tracks         :     106500
  512-byte Blocks :    27264000
  MegaBytes      :      13313
  KiloBytes      :    13632000

  Geometry Limited : No
}

Device External Identity
{
  Device WWN      : 60000970000198701161533030314545
}

```

On the target device, the device retains the external WWN from the source array that it inherited during the create step. This remains the WWN after the completion of the migration.

---

**Note:** The native identities of devices can be displayed using the `-native` option on the `syminq` command.

---

### **Remove the NDM environment**

The environment remove operation removes the replication pathway configured by the environment setup operation, and it removes the resources that were configured to support NDM on the source and target arrays. On successful completion of an environment remove operation, only running an environment setup operation is allowed. An environment removal operation should occur only after all migrations from the source array have been completed.

```

symdm -src_sid <SN of Source> -tgt_sid <SN of Target> environment
-remove

```

```
C:\Windows\system32>symdm -src_sid 161 -tgt_sid 131 environment -remove
A DM 'Environment Remove' operation is in progress. Please wait...

  Analyze Configuration.....Started.
    Source SID:000198701161
    Target SID:000197800131
  Analyze Configuration.....Done.
  Remove Configuration.....Started.
  Remove Configuration.....Done.

The DM 'Environment Remove' operation successfully executed.
```

Once the environment remove operation is complete, the NDM process is complete.

## Metro-based NDM

### Introduction

This section includes a guide plan, environment overview, and walkthrough guide for Metro-based NDM.

### Metro-based NDM guide plan and environment overview

The following guide describes two methods for each version of Metro-based NDM:

- Metro-based NDM:
  - Unisphere for PowerMax: [Using Unisphere for PowerMax](#)
  - CLI (Solutions Enabler): [Using Metro NDM using Solutions Enabler 9.x](#)
- Metro-based NDM with precopy:
  - Unisphere for PowerMax: [Using Unisphere for PowerMax with precopy](#)
  - CLI (Solutions Enabler): [Using Solutions Enabler 9.x with precopy](#)

This guide uses the VMAX All Flash and VMAX3 arrays in the graphic below.

In the Metro-based NDM example SGs migrate from 000296700558 to 000197800131.

000197800131 VMAX250F   5978.37.38	9 %	<span style="color: red;">0</span> <span style="color: yellow;">0</span> <span style="color: green;">0</span>	100	MB/S	28	-			
000296700558 VMAX200K   5977.1125.1125	8 %	<span style="color: red;">0</span> <span style="color: yellow;">0</span> <span style="color: green;">0</span>	-	MB/S	-	35.3:1			

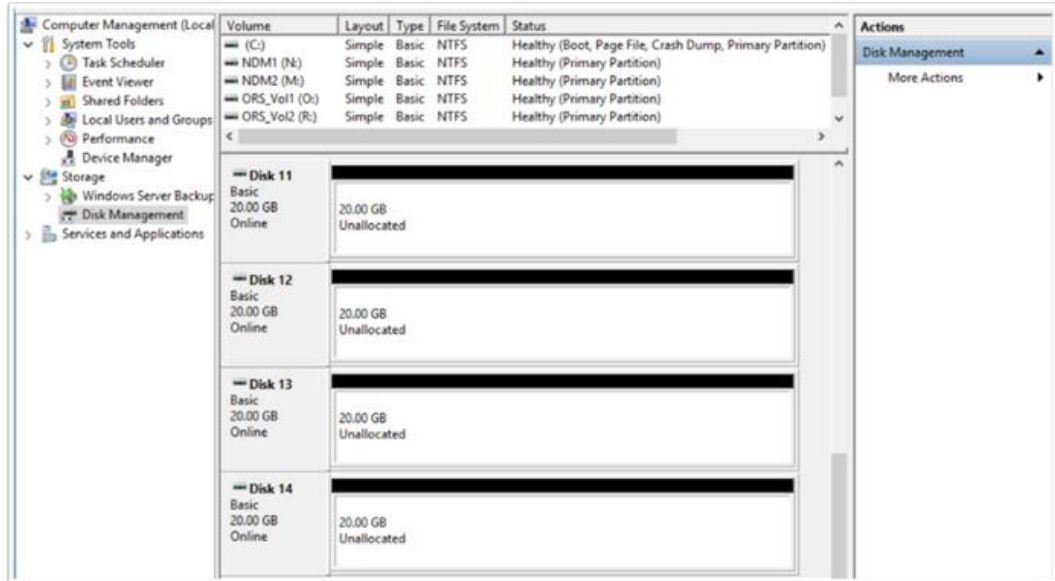
Prior to the start of a planned migration, verify the following:

- Ensure both source and target array are RDF capable (the RF emulation has been added to both arrays).
- Ensure both arrays RDF ports are zoned to each other. There is a minimum of two connections required.
- Check for the correct zoning from the target array to the application host.

**Metro-based NDM walkthrough guide (source running 5977 or 5978 code)**

**Using Unisphere for PowerMax**

Looking at the devices in this migration from the host operating system disk management (in this case, Windows Server 2016), they show as Disk 11 through Disk 14. These were previously added as RDMs to the VMs using VMware vSphere.



This is an example of the multipathing setup using device 1BA. It shows what the pathing looks like prior to the NDM Create operation and the host rescan. For each of the four volumes here there are four paths to the source array which are all alive and available for host use. At this point, there are no paths to the target array even though our zoning should be in place before the NDM create.

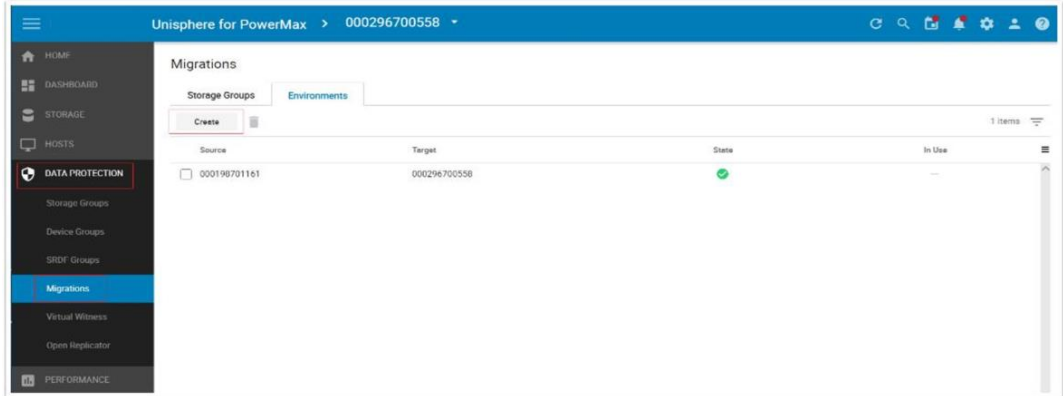
```

pplicensevmaxcse:- # ipowermt display dev=emcpower140 host=10.60.136.146
Pseudo name=emcpower140
Symmetrix ID=000296700558
Logical device ID=001BA
Device WWN=60000970000296700558533030314241
Standard UID=naa.60000970000296700558533030314241
type=Conventional; state=alive; policy=SymmOpt; queued-IOS=0
-----
### HW Path      Host ----- I/O Paths  - Stor - -- I/O Path -- -- Stats ---
                I/Path  Interf.  Mode   State  Q-I/Os  Errors
-----
3 vmhba3         C0:T2:L3  FA  1d:24 active  alive   0    0
1 vmhba4         C0:T1:L3  FA  2d:24 active  alive   0    0
1 vmhba4         C0:T7:L3  FA  1d:24 active  alive   0    0
3 vmhba3         C0:T0:L3  FA  2d:24 active  alive   0    0
    
```

**NDM environment setup**

The Environment Setup operation configures the migration environment template required to setup the Metro groups used to migrate all applications from the source array to the target array. This template is used to define the RDF groups for each migration session. Within this definition are ports used, target ports and port count. The operation also confirms that both the source and target arrays can support the NDM operations. This includes ensuring that a usable replication pathway for data migration is available between source and target arrays. Should we need a second target array from the same source, a second environment is necessary.

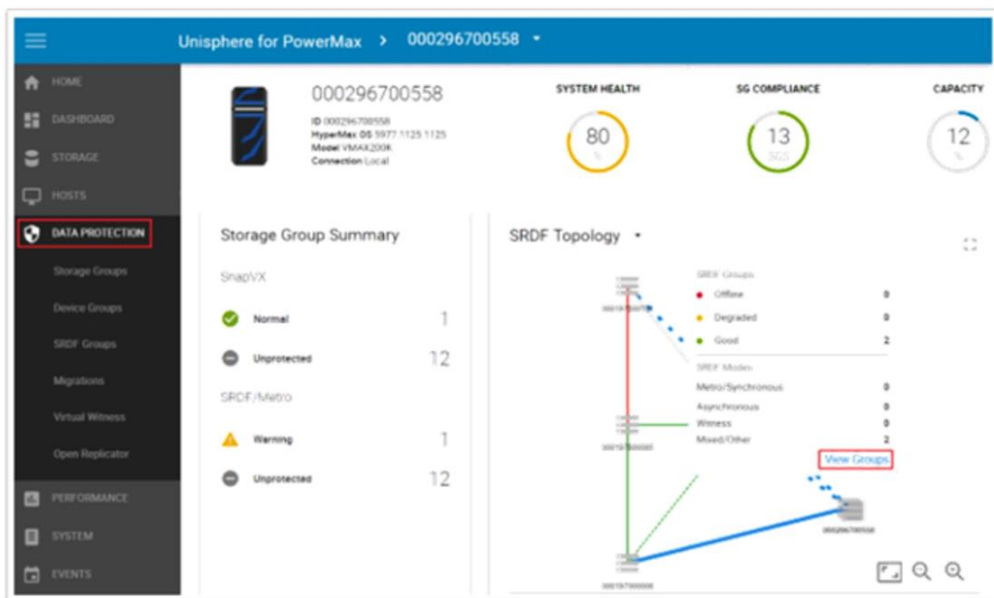
To summarize, the Environment Setup runs once only for each array relationship. This Setup operation creates a template from which all other Metro-based NDM SRDF groups are modelled. Each individual NDM session requires its own RDF/Metro group to be created unlike Pass-through NDM which used a single RDF group for all sessions between the source and target.



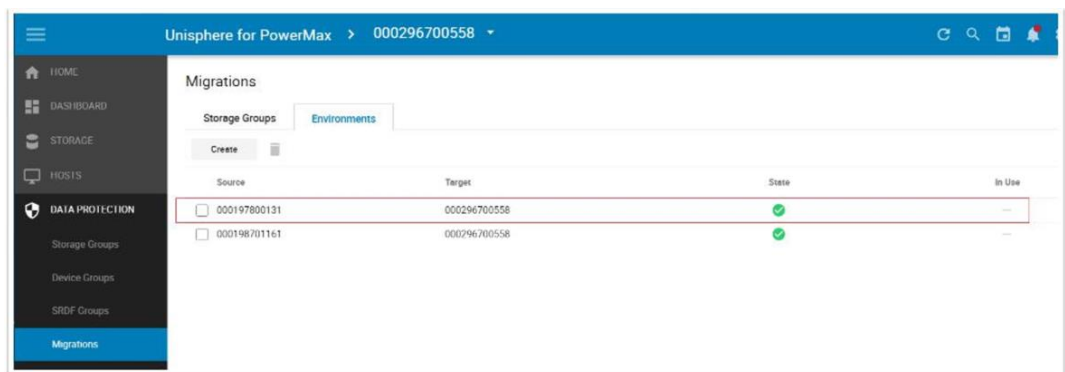
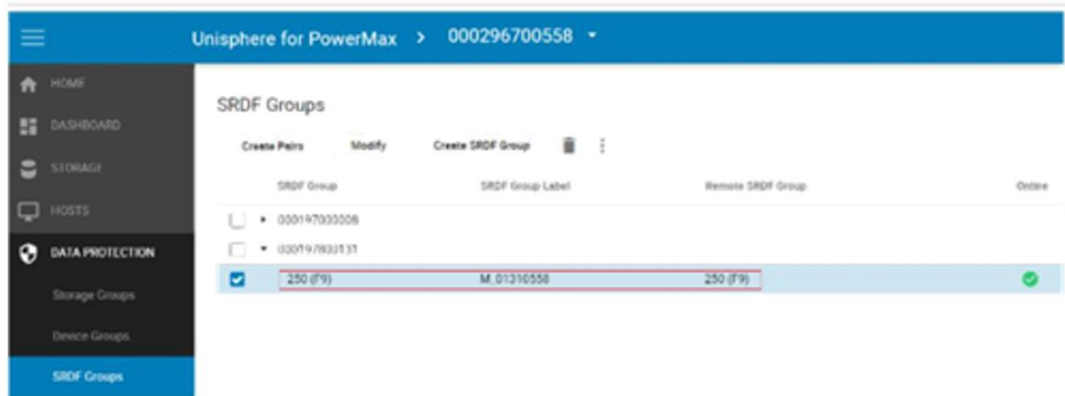
From the **Data Protection** menu select **Migrations**, Click the **Environment** tab to display any existing Environments already setup. The parameter **In Use** shows us whether the Environment is validated and usable. The **In Use** parameter also shows whether there is an active Migration using this environment.

To create an environment click **Create**, the pop-up window below appears. This contains a drop-down list of all the arrays available for migration operations. Should your array not be present, verify the RDF zoning and confirm the intended target array is suitable and its code level is within the support matrix. Select the relevant array and choose **Run Now**.

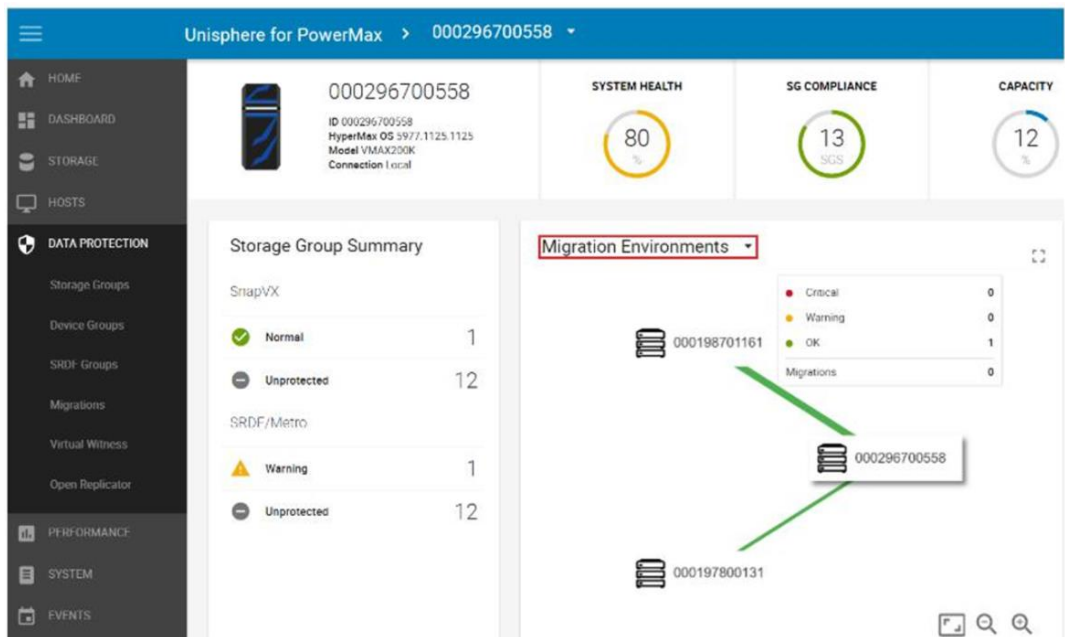
The new **Topology** view we can see the RDF group template that has been created. Go to the **DataProtection** dashboard and hover over the line between our source and target. This causes the **SRDF Groups** window to appear. Select **View Groups** to display the SRDF group window highlighted below.



From an RDF standpoint you can examine the new RDF group created to handle NDM migrations for all SG sessions between 558 and 131.



From the same Data Protection dashboard, the drop-down menu contains an option to monitor Migration Environments. A color-coded line indicates any problems. Hovering over the connection line displays the number of each connection status.

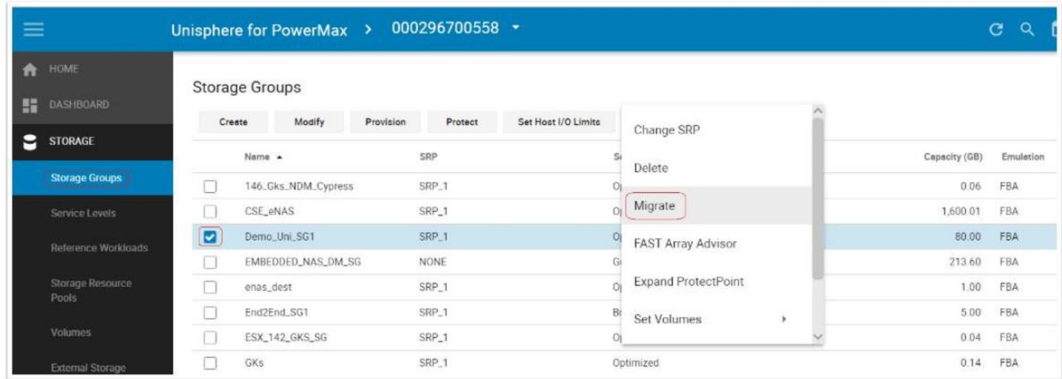




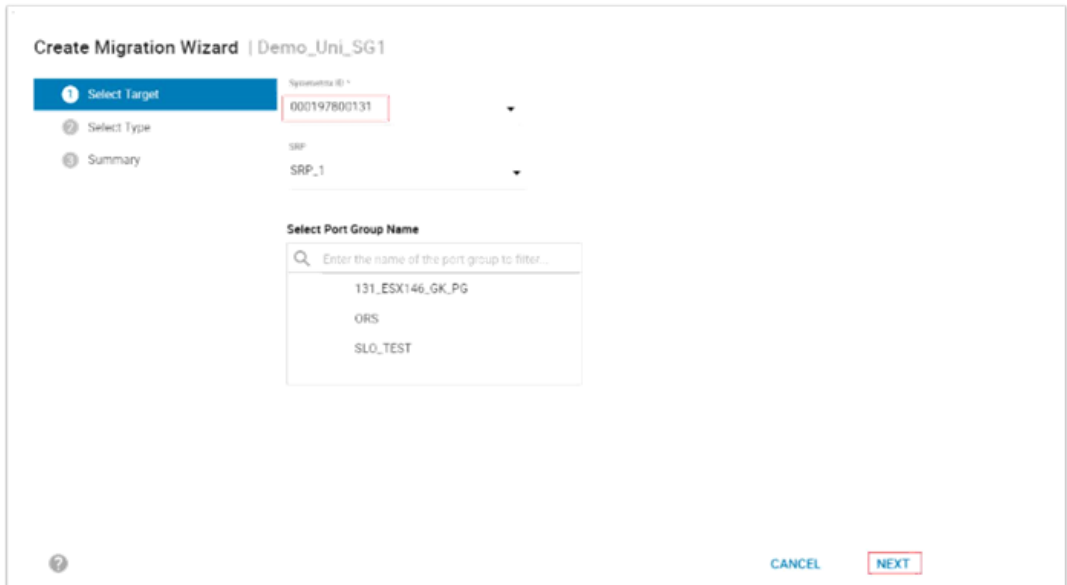
Now the Environment is in place so the NDM Create for the SG planned for migration can occur.

**Create migration session**

From the Storage tab click Storage Groups, from there locate the SG that is to be migrated. Select the check box and click the More Actions "3-Dot" icon to the right of Set Host I/O Limits. From the drop-down menu select Migrate.



In the Create Migration Wizard, select the **Target array** (only arrays with valid environments setup appear on this drop-down menu). This example does not select a Port Group. (See the relevant masking enhancements section.) Click **Next**.



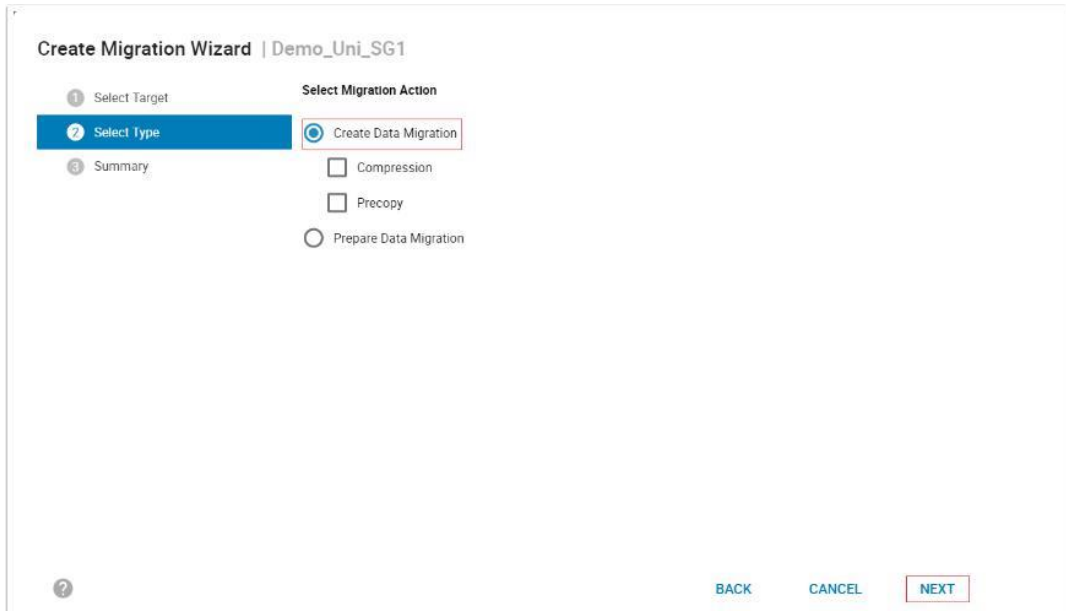
On the next screen click Create Data Migration. This provides the option of selecting Compression and precopy. The scenario that this part of the example covers does not use precopy, so that remains cleared.

(See precopy section) and select Next. The Prepare Data Migration selection requires Performance data to be collected on the host.

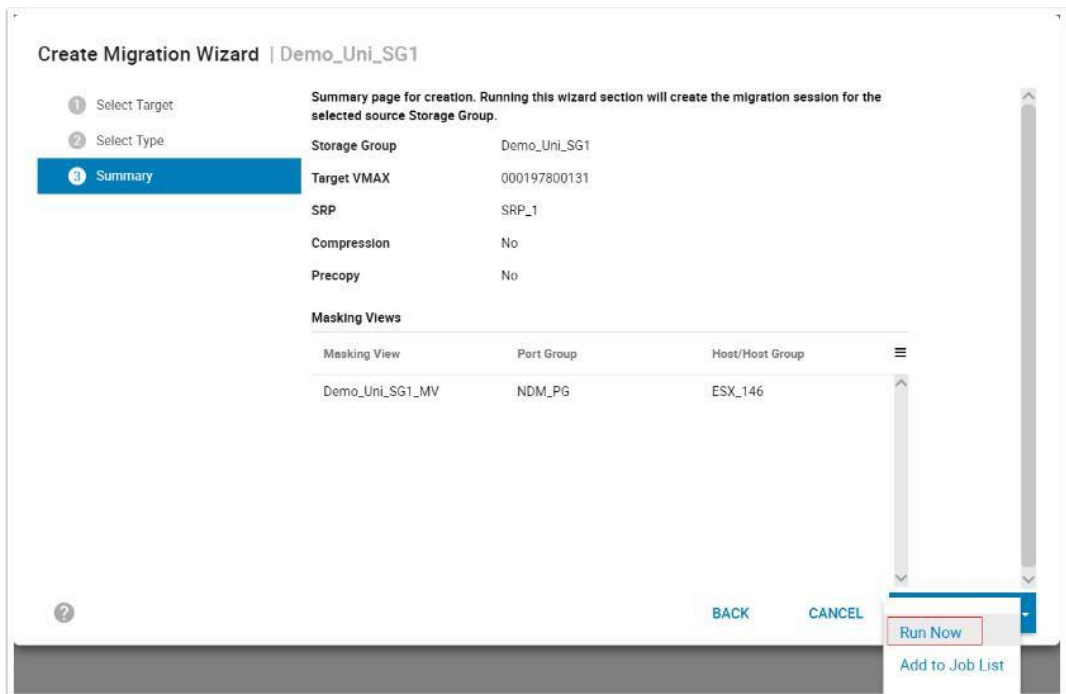
This runs a check for resources on the target array to ensure the addition of the new SG does not cause the target array to exceed any performance metrics on both FE and BE. It



also produces a spread sheet to help plan the zoning required for the host from the target array.

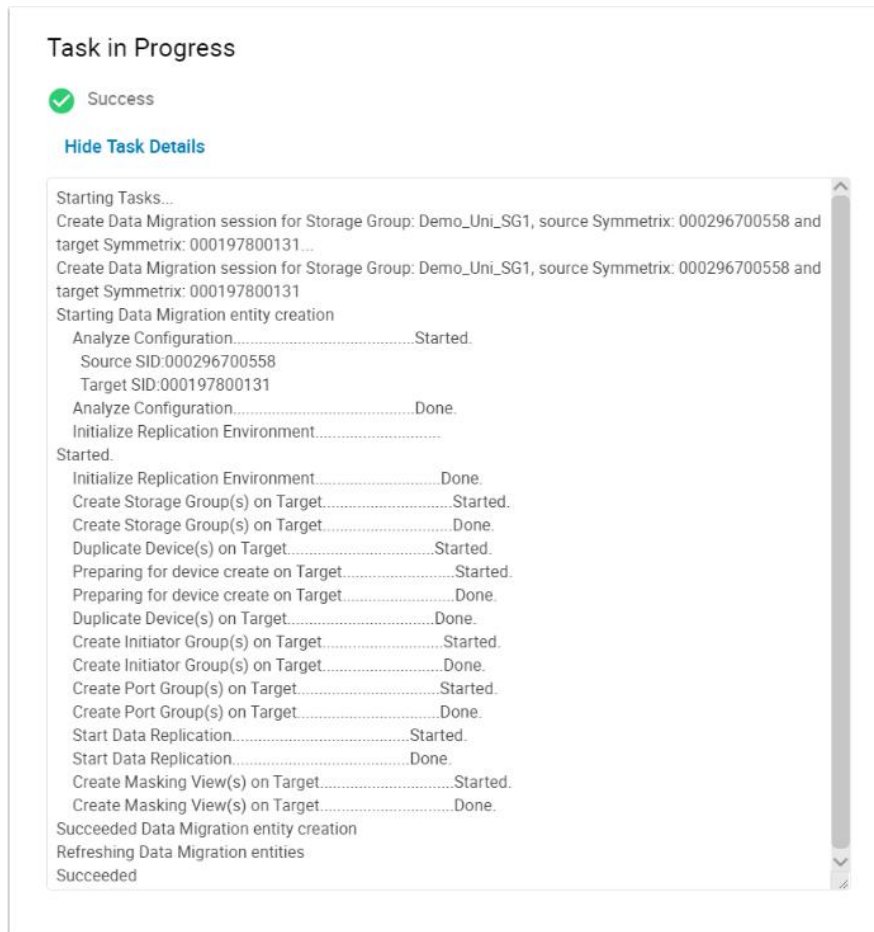


The final page of the Wizard summarizes the planned NDM session to be created. It breaks down the planned masking view elements and the NDM parameters. Select **Run Now** to continue.



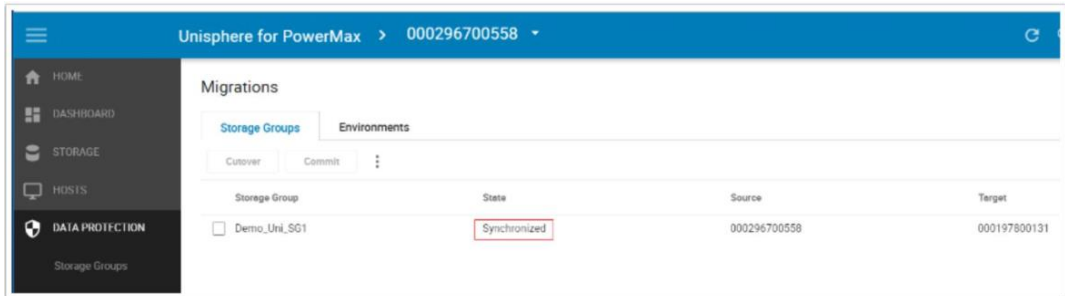
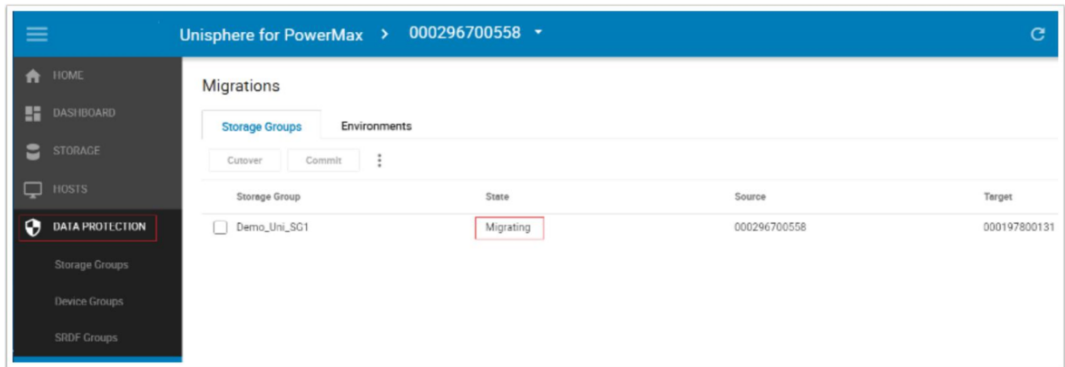
Creating the NDM session will also perform an Environment validate as part of the setup to ensure it will complete successfully. As outlined in the create command output, the Create:

- Creates a Storage Group on the Target array (name must not already exist in the target array) with the same name as the Source SG.
- Creates duplicate devices on the target array to match those on the Storage Group
- Creates an initiator group using Initiators with entries in the login history table
- Creates a Port group (if one does not already exist or has not been selected by the user, see the relevant masking enhancements section)
- Invalidates the tracks on the RDF mirror to prepare for the copy
- Starts the copy process.
- Creates a masking view to the host from the target array.

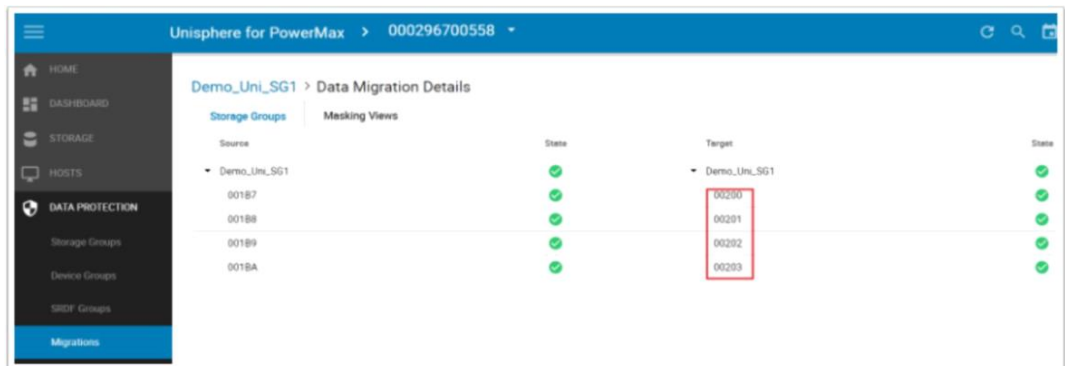


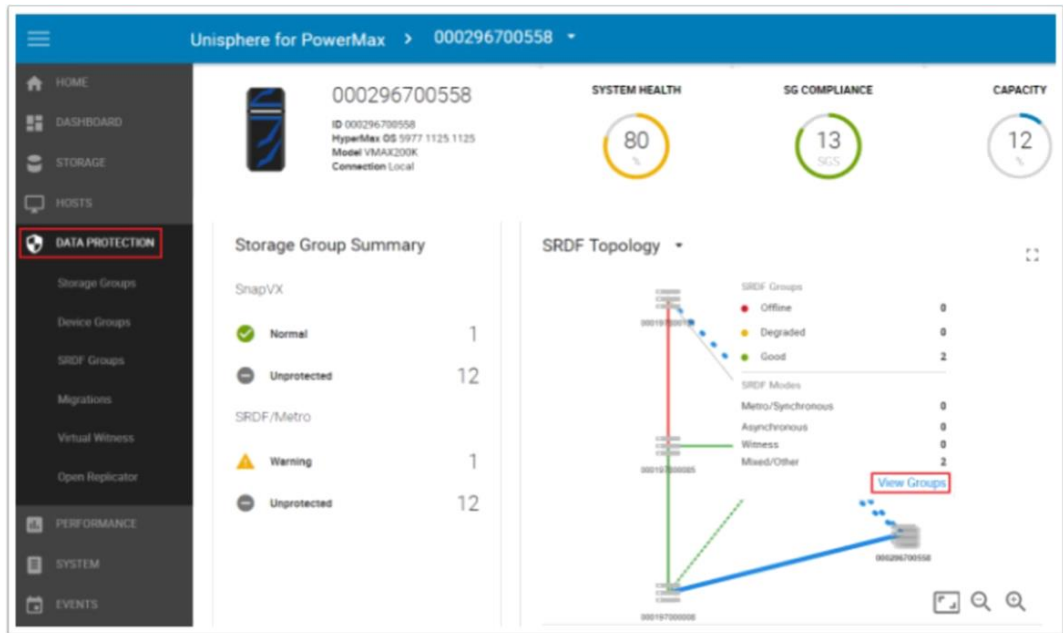
### ***Examine the created migration session***

Selecting the Data Protection tab on the left task bar and selecting Migrations in the drop-down menu, the storage groups currently involved in an NDM session are highlighted along with the current State and details on the source and target arrays. Since data transfer begins immediately following a Create operation there is no need for the Pass-Through NDM Cutover operation. Once the data is synchronizing the systems administrator runs a rescan to allow the target paths to become active to the multipathing software. At this point, both source and target arrays are involved in an active/active relationship with all I/O serviced locally.

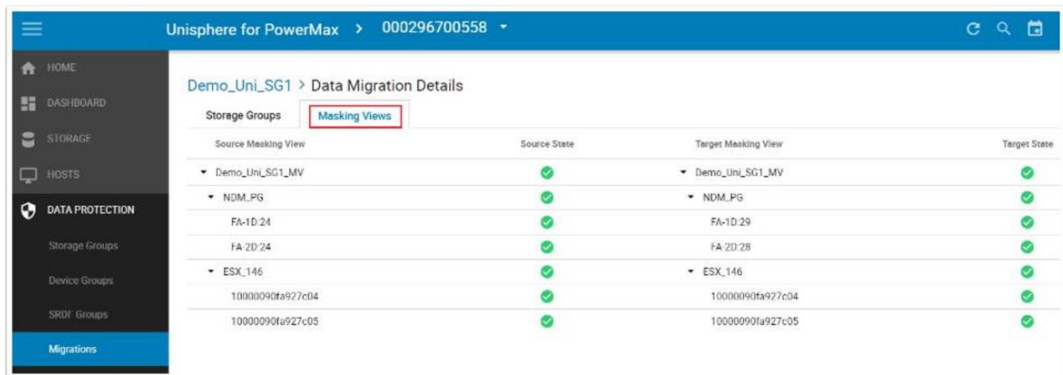


Double-clicking the NDM session will take you into the Migration Details view. In the screen above, we can see the individual devices involved in the session. Under the target tab we can see devices 200 through 203 have been created on the target side by the NDM create. The State also shows a live status of each of the devices involved. A device without a green tick should be investigated for potential problems before continuing.



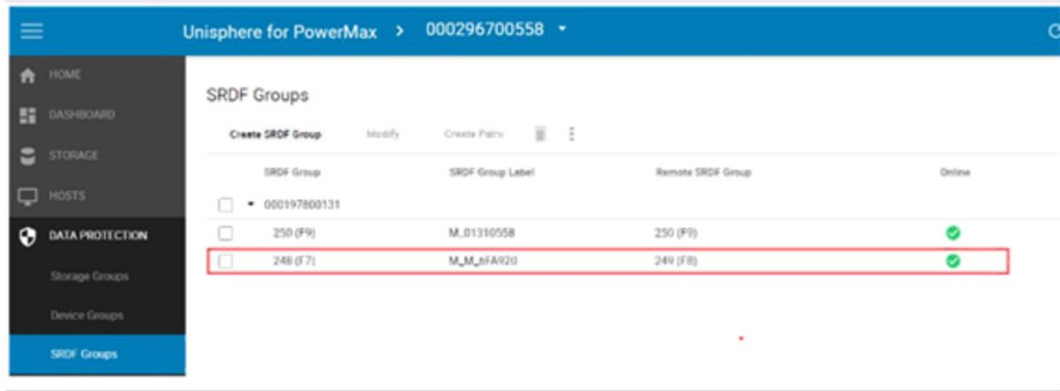


Selecting the **Masking Views** tab displays the pane outlining the masking and masking elements involved on both sides of the NDM session. This screen can be useful for troubleshooting any issues with the migration such as an unplanned or unauthorized manipulation of any of the NDM elements that hinder progress to the commit stage. In such cases the State does not contain a green tick, rather a red warning.



Examining the RDF environment from the **Topology** view shows the RDF group template that has been created. Go to the **DataProtection** dashboard and hover over the line between our source and target. This causes the **SRDF Groups** window to appear. Select **View Groups** to display the SRDF group window highlighted below.

The new SRDF group that has been created (248/249 on remote) as part of the Create command operation from the template (250). Group 248 is used for the duration of this migration.

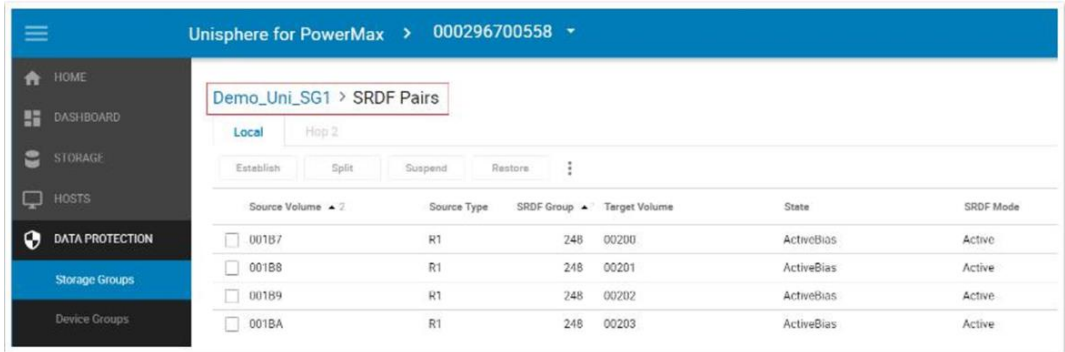


**View paths to new devices and SRDF pairs**

Viewing the multipathing software after the rescan shows the extra paths online to the target array (in this case two extra paths) The number of extra paths depends on the zoning setup. It also displays the source and target SIDs and the device numbers involved for these paths. The WWNs on the LUNs to appear as a single device with just extra paths. Prior to version 6.2 PowerPath was not aware of the NDM process so the dual SIDs and devices IDs were not visible.

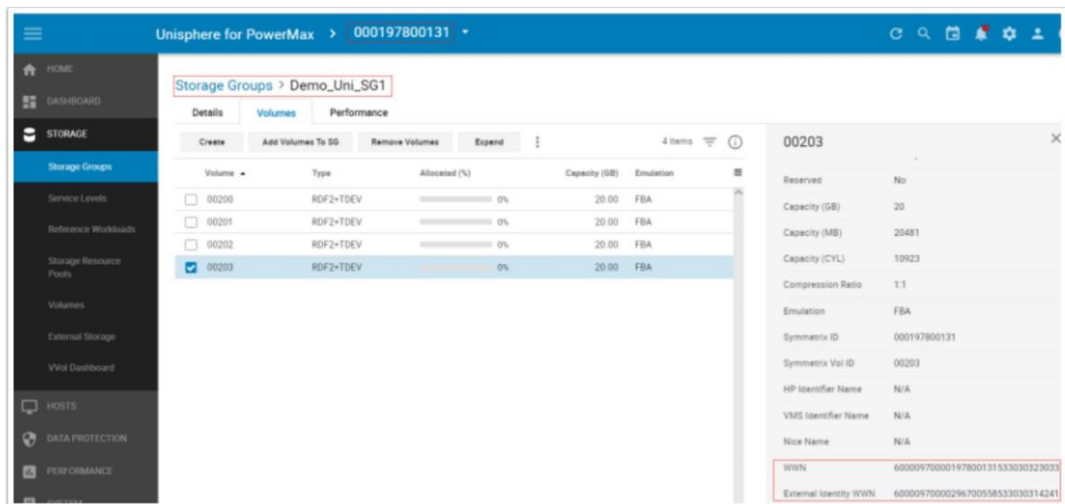
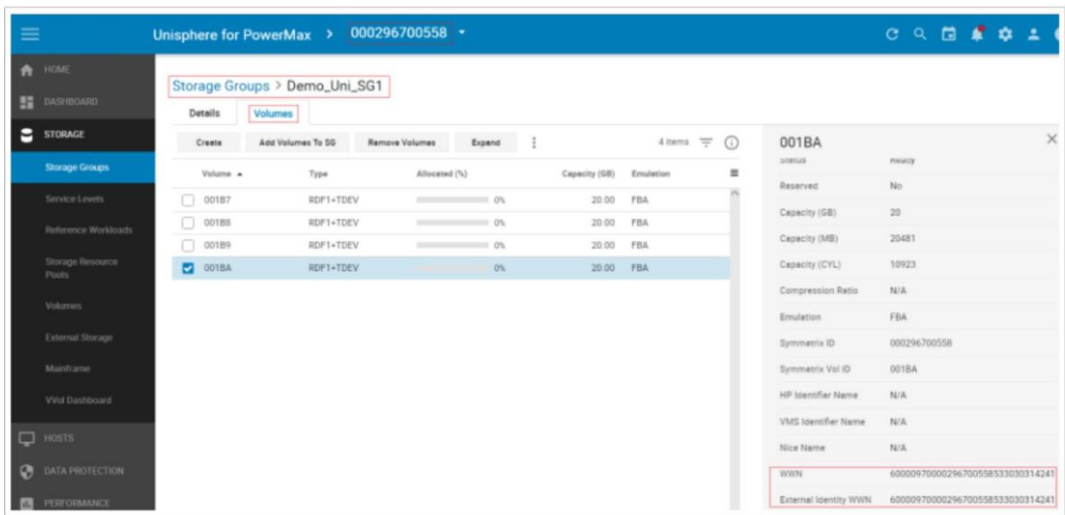
```
pplicensevmaxcse:~ # rpowermt display dev=emcpower140 host=10.60.136.146
Pseudo name=emcpower140
Symmetrix ID=000296700558, 000197800131
Logical device ID=001BA, 00203
Device WWN=60000970000296700558533030314241
Standard UID=naa.60000970000296700558533030314241
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-I/Os Errors
-----
1 vmhba4             C0:T12:L3  FA 2d:28 active  alive  0      0
3 vmhba3             C0:T7:L3   FA 1d:29 active  alive  0      0
3 vmhba3             C0:T2:L3   FA 1d:24 active  alive  0      0
1 vmhba4             C0:T11:L3  FA 2d:24 active  alive  0      0
1 vmhba4             C0:T7:L3   FA 1d:24 active  alive  0      0
3 vmhba3             C0:T0:L3   FA 2d:24 active  alive  0      0
pplicensevmaxcse:~ #
```

Once the Create operation and the scan have both completed, the migration is in an active/active functional state with I/Os being distributed to both source and target paths. Bias is set on the Source side which differs from Pass-Through NDM where the source side was the R2 in a standard R1- R2 synchronous relationship. The state in Unisphere is displayed as ActiveBias as we do not have a witness attached to this metro relationship. However, we are truly active/active from an array perspective.



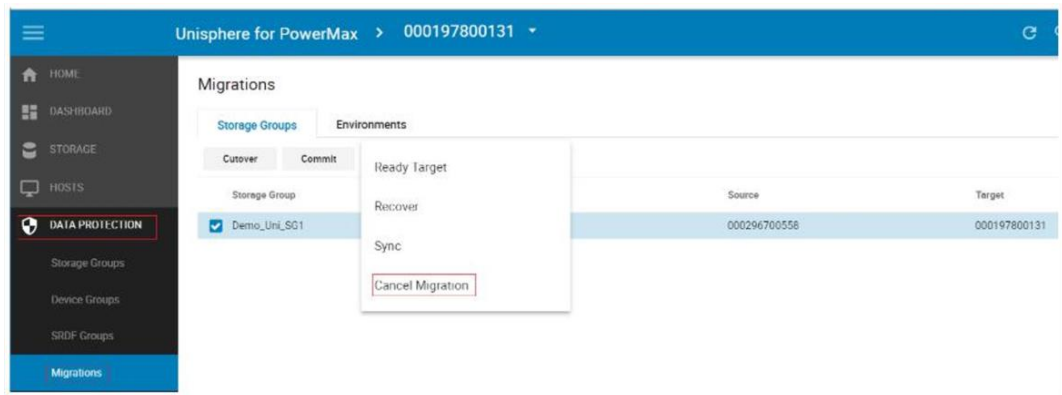
**Examine IDs of source and target**

Examining the devices WWNs post Create the source device has a WWN and External (host visible) WWN with the same value. However, the targets WWN and External WWNs differ. The target External WWN has inherited the WWN of the source in order to appear logically as the same device to the host.

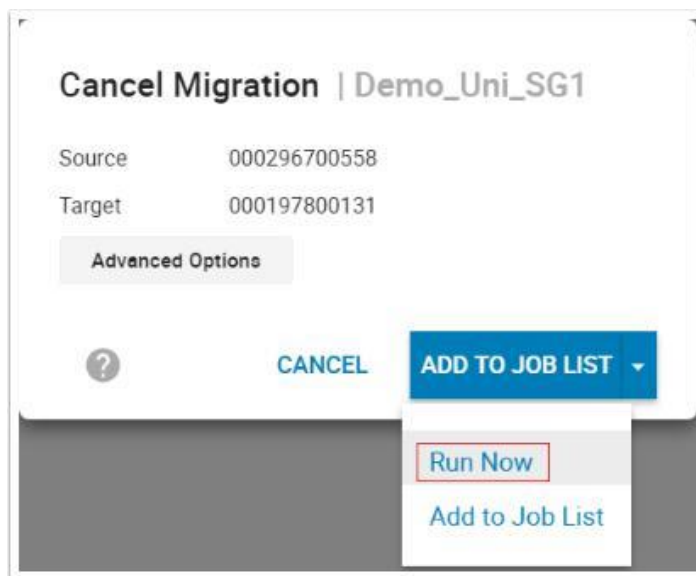


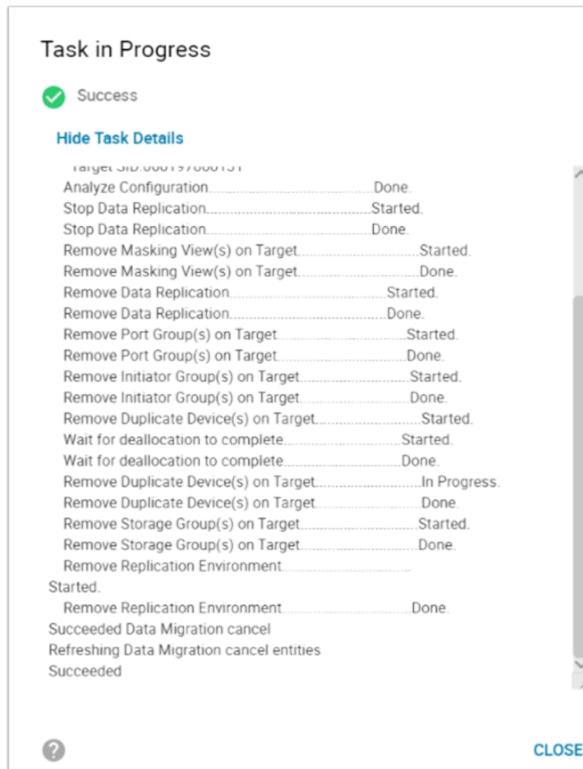
### Cancel a migration

To Cancel a migration, click the **Data Protection** tab and select **Migrations** from the drop-down menu. Select the active NDM session and click the More Actions (three Dots) menu and Cancel Migration.



Confirm the cancel operation of the SG with the correct source and target and click **Run Now**.





The cancel operation removes the storage provisioned on the target array and releases any allocations and resources allocated by the NDM create operation. It also places the source devices into the same state they were before the create operation was issued.

The cancel operation:

- Stops replication between the source and target arrays
- Removes the masking view on the target array
- Removes the RDF pairings
- Removes the port group on the target array (if not in use in another masking view)
- Removes the initiator group on the target array
- Deallocates volumes created on the target array
- Removes the devices created on the target array

---

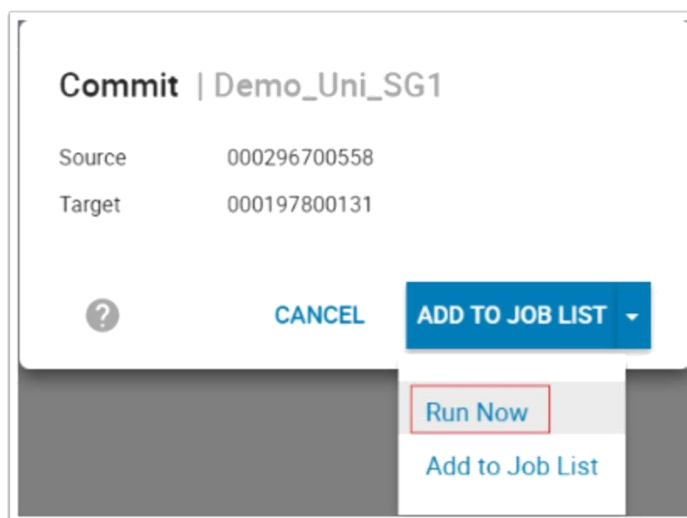
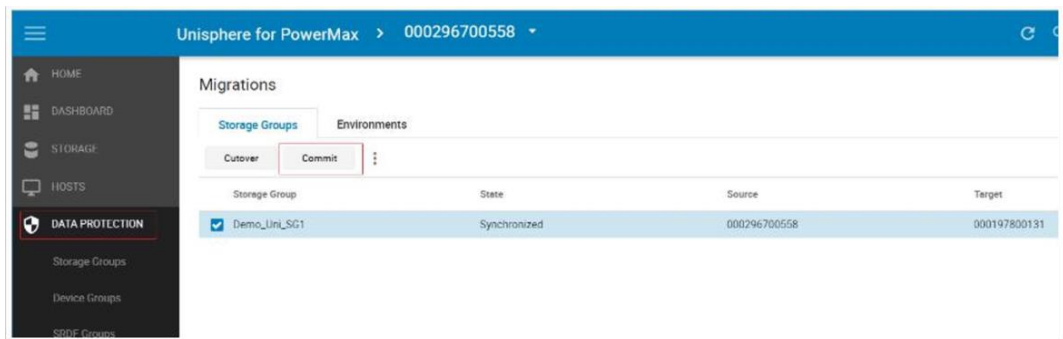
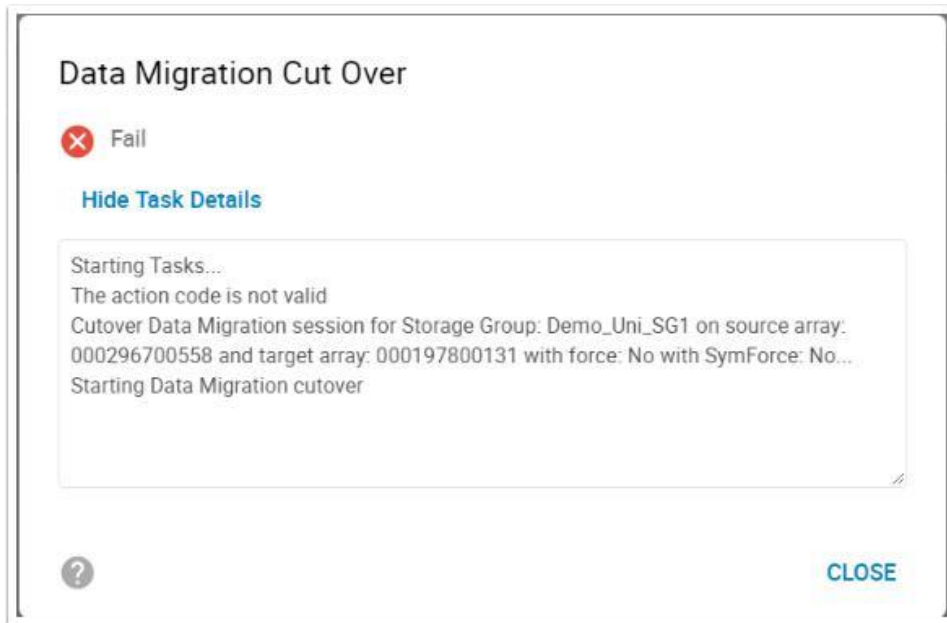
**Note:** It is best practice to run a rescan on the host should be run at this point to remove any dead or invalid paths.

---

### **Commit a migration**

If an attempt is made to issue a **cutover** on a Metro NDM session, the above error will occur. This cutover command is only for Pass-through NDM (source array running 5876 code).

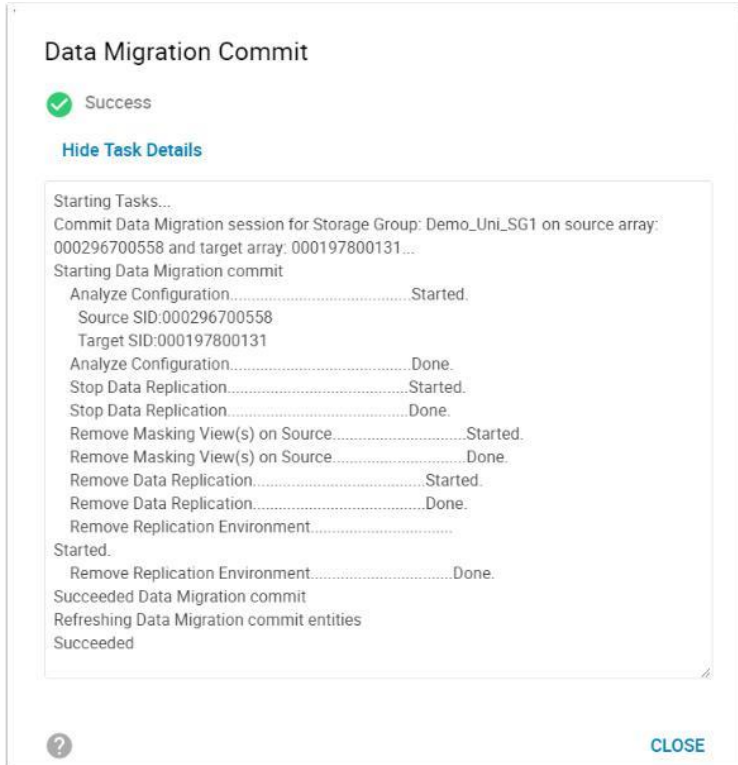




When the data commit operation is completed and the devices are synchronized, the migration can be committed. The Commit operation completes the migration by removing the migrated application resources from the source array and releases system resources used for the migration. When the commit is completed the replication relationship between the source and target devices are removed, the masking view on the source removed and

the source devices take the native (internal) WWN of the target LUN as its effective (external) WWN.

The target has an external WWN of the source and the source having an external WWN of the target. Both devices retain their native (internal) WWNs but these are not presented to the host.



**Examine paths and device post commit**

Post Commit and with the removal of the source masking view and before the host rescan the original Source paths are in a dead state. (This varies depending on your MP software) The two target paths are still active.

```

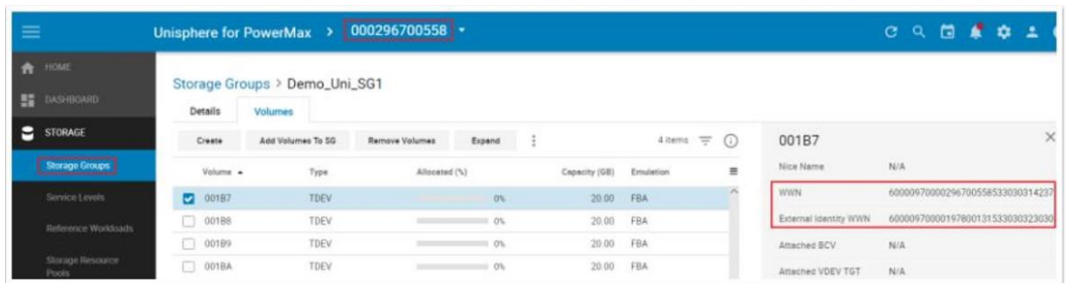
pplicensevmaxcse:~ # rpowermt display dev=emcpower140 host=10.60.136.146
Pseudo name=emcpower140
Symmetrix ID=000296700558, 000197800131
Logical device ID=001BA, 00203
Device WWN=60000970000296700558533030314241
Standard UID=naa.60000970000296700558533030314241
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-I/Os  Errors
-----
1 vmhba4             C0:T12:L3  FA  2d:28  active  alive    0    1
3 vmhba3             C0:T7:L3   FA  1d:29  active  alive    0    1
3 vmhba3             C0:T2:L3   FA  1d:24  active  dead     0    1
1 vmhba4             C0:T11:L3  FA  2d:24  active  dead     0    1
1 vmhba4             C0:T7:L3   FA  1d:24  active  dead     0    1
3 vmhba3             C0:T0:L3   FA  2d:24  active  dead     0    1
    
```

Once the host Rescan has been completed, the dead paths are now removed and the SID that of the target array. The WWN inherited from the source array is displayed as before to allow the distinction of NDM devices.

```

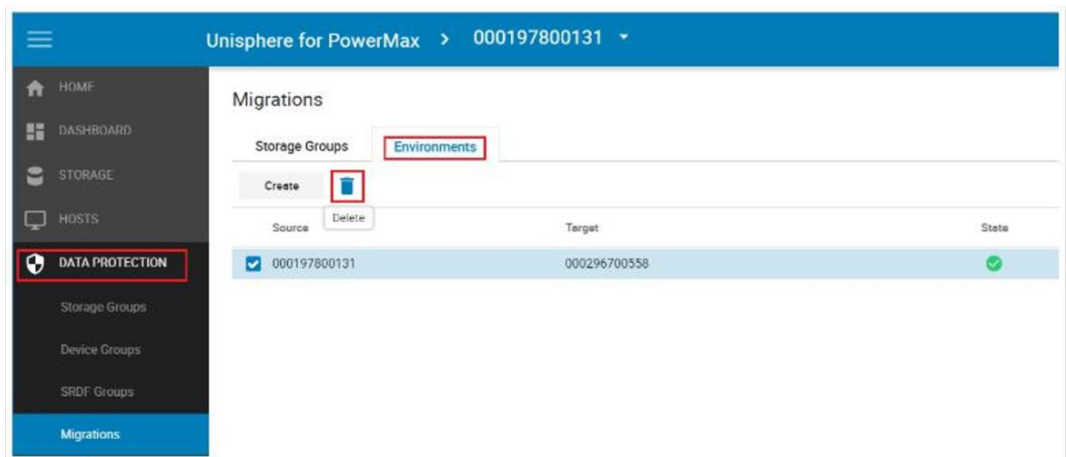
pplicensevmaxcse:- # rpowermt display dev=emcpower140 host=10.60.136.146
Pseudo name=emcpower140
Symmetrix ID=000197800131
Logical device ID=00203
Device WWN=60000970000296700558533030314241
Standard UID=naa.60000970000296700558533030314241
type=Conventional; state=alive; policy=SymmOpt; queued-IOs=0
=====
### Host ----- Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths      Interf. Mode      State      Q-IOs Errors
=====
1 vmhba4             C0:T12:L3    FA  2d:28 active   alive      0      1
3 vmhba3             C0:T7:L3     FA  1d:29 active   alive      0      1
=====
    
```

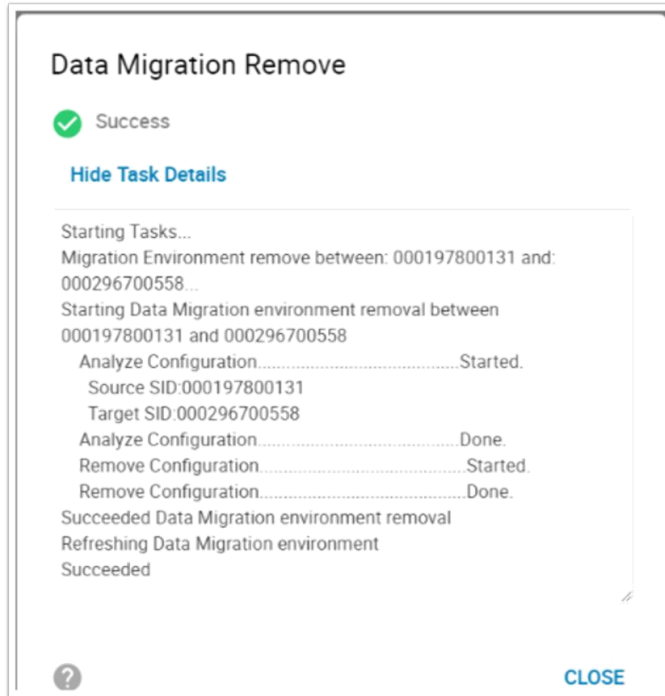
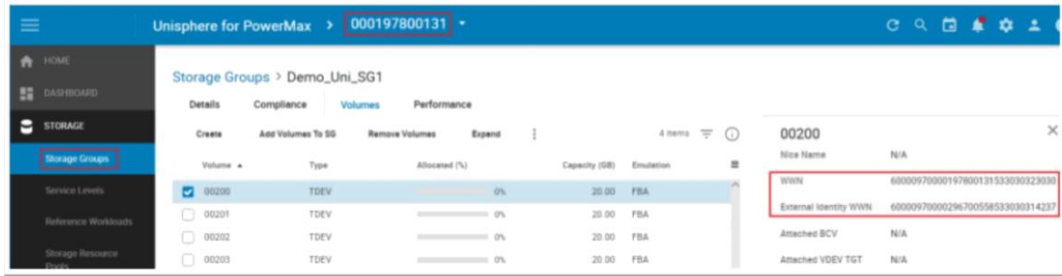
Viewing the source and target devices after the Commit operation clearly shows the WWN manipulation that has occurred. The Source device now has the target's native WWN as its effective WWN and the target device has retained the source native WWN as its effective WWN.



### Remove NDM environment

When all migrations are complete between a specific source and target, the environment can be removed. Click the Data Protection tab and click Migrations. On the Environments tab, select the Environment and click the trash icon. On the window that appears, click **Run Now** and the Environment is removed. This removes the RDF group setup and releases those resources.



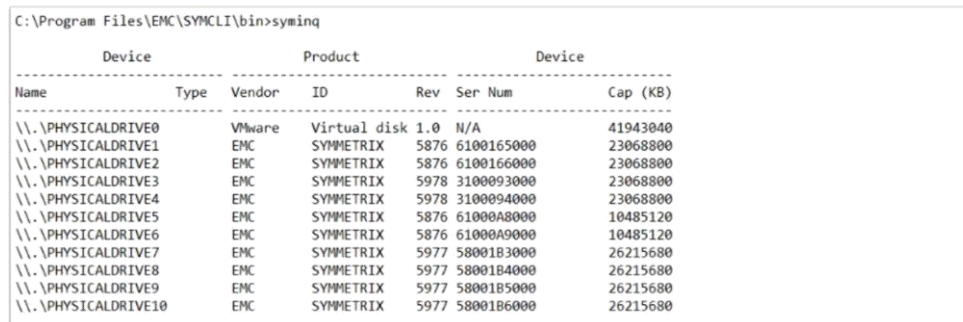


**Note:** At this point, the data migration is completed and the migration environment is removed for the specific source and target.

## Using Metro NDM using Solutions Enabler 9.x

### View devices from the host

This screenshot shows the New Devices view using the Syminq command, and the devices added to the host as PhysicalDrive7 through PhysicalDrive10.



### PowerPath view of one of the new devices

PowerPath shows what the pathing looks like prior to the NDM Create (in this case dev 1B6) and the host rescan. For each of the four volumes, there are four paths to the source array. All are alive therefore available for host use. We do not yet have any paths to the target array.

```
Pseudo name=emcpower178
Symmetrix ID=000296700558
Logical device ID=01B6
Device WWN=60000970000296700558533030314236
Standard UID=naa.60000970000296700558533030314236
type=Conventional; state=alive; policy=SymmOpt; queued-IOs=0
```

###	Host	I/O Paths	Stor Interf.	I/O Path Mode	State	Q-I/Os	Errors
3	vmhba6	C0:T11:L4	FA 1d:24	active	alive	0	0
3	vmhba6	C0:T13:L4	FA 2d:24	active	alive	0	0
4	vmhba2	C0:T6:L4	FA 2d:24	active	alive	0	0
4	vmhba2	C0:T1:L4	FA 1d:24	active	alive	0	0

### Environment setup

The Environment Setup configures the migration environment template required to create SRDF/Metro groups for the migration of any application from the source array to the target array. It confirms that both the source and target arrays can support NDM. This includes that a usable replication pathway for data migration is available between the source and target. This needs to be issued once only as the environment is to be used for all migrations between these arrays.

```
symdm -src_sid <SN of Source> -tgt_sid <SN of Target> environment
-setup
```

```
C:\Program Files\EMC\SYMCLI\bin>symdm -src_sid 558 -tgt_sid 131 environment -setup
A DM 'Environment Setup' operation is in progress. Please wait...

Analyze Configuration.....Started.
  Source SID:000296700558
  Target SID:000197800131
Analyze Configuration.....Done.
Setup Configuration.....Started.
Setup Configuration.....In Progress.
Setup Configuration.....In Progress.
Setup Configuration.....In Progress.
Setup Configuration.....In Progress.
Setup Configuration.....Done.

The DM 'Environment Setup' operation successfully executed.
```

### Validate environment

To validate the recently created environment or an existing environment to an alternative array use the - validate option to the symdm environment command.

```
symdm -src_sid <SN of Source> -tgt_sid <SN of Target> environment
-validate
```

```
C:\Program Files\EMC\SYMCLI\bin>symdm -src_sid 558 -tgt_sid 131 environment -validate
A DM 'Environment Validate' operation is in progress. Please wait...

Analyze Configuration.....Validated.

The DM 'Environment Validate' operation successfully executed.
```

### Validating and creating an NDM session

Solutions Enabler examines a specific application's storage on the source array and automatically provisions equivalent storage on the target array. The target devices are assigned the identity of the source devices. Prior to running the create operation, it is always worth running the -validate to ensure the planned migration can succeed. This allows any potential issues to be addressed leading into the migration window.

```
symdm -src_sid <SN of Source> -tgt_sid <SN of Target> -sg <SG to be Migrated> -validate
```

```
C:\Program Files\EMC\SYMCLI\bin>symdm create -src_sid 558 -tgt_sid 131 -sg NDM_Beta_SG1 -nop -validate

A DM 'Validate Create' operation is
in progress for storage group 'NDM_Beta_SG1'. Please wait...

Analyze Configuration.....Validated.
Initialize Replication Environment.....Validated.
Create Storage Group(s) on Target.....Validated.
Duplicate Device(s) on Target.....Validated.
Create Initiator Group(s) on Target.....Validated.
Create Port Group(s) on Target.....Validated.
Create Masking View(s) on Target.....Validated.

The DM 'Validate Create' operation successfully executed for
storage group 'NDM_Beta_SG1'.
```

If any of the above fails, it is worth taking a look at the SYMAPI log file, which, more often than not, points you toward an easily correctable issue within a masking view or zoning config. In the following example, one of the initiators in the source IG in an IG on the target array.

```
08/31/2017 12:41:28.688 EMC:SYMMDM validateIGEntryInMul The
initiator wwn 10000090fa927c04 is already in use in Initiator
Group 131_GKs_IG for array 000197800131
```

```
08/31/2017 12:41:28.688 Create Initiator Group(s) on
Target.....Failed.
```

Creating the NDM session also does a -validate of the environment to ensure the subsequent steps complete successfully. The Create operation:

- Creates a Storage Group on the Target array (name must not already exist in the target array) with the same name as the Source
- Creates duplicate devices on the target array to match those on the Storage Group
- Creates an initiator group using Initiators with entries in the login history table
- Creates a Port group (if one does not already exist or has not been selected by the user, see the relevant masking enhancements section)
- Invalidates the tracks on the RDF mirror to prepare for the copy
- Starts the copy process.
- Creates a masking view to the host from the target array.

```
symdm -src_sid<SN of source> -tgt_sid<SN of Source> -sg<Sg to Migrate> -tgt_SRP <SRP on Target>
```



```
C:\Program Files\EMC\SYMCLI\bin>symdm create -src_sid 558 -tgt_sid 131 -sg NDM_Beta_SG1 -nop

A DM 'Create' operation is
in progress for storage group 'NDM_Beta_SG1'. Please wait...

Analyze Configuration.....Started.
  Source SID:000296700558
  Target SID:000197800131
Analyze Configuration.....Done.
Initialize Replication Environment.....Started.
Initialize Replication Environment.....Done.
Create Storage Group(s) on Target.....Started.
Create Storage Group(s) on Target.....Done.
Duplicate Device(s) on Target.....Started.
Preparing for device create on Target.....Started.
Preparing for device create on Target.....Done.
Duplicate Device(s) on Target.....Done.
Create Initiator Group(s) on Target.....Started.
Create Initiator Group(s) on Target.....Done.
Create Port Group(s) on Target.....Started.
Create Port Group(s) on Target.....Done.
Start Data Replication.....Started.
Start Data Replication.....Done.
Create Masking View(s) on Target.....Started.
Create Masking View(s) on Target.....Done.

The DM 'Create' operation successfully executed for
storage group 'NDM_Beta_SG1'.
```

### Viewing migration sessions

As outlined in the introduction NDM from 5977 to 5978 removes the Cutover operation and need for Pass-through mode. Once the Create operation has completed, data transfer begins immediately. As this example shows, the session entered a migration state and began copying data to the target as soon as the Create operation completed.

```
C:\Program Files\EMC\SYMCLI\bin>symdm -sid 558 list

Symmetrix ID : 000296700558

Storage Group      Source      Target      State      Total
                   Array      Array      State      Capacity  Done
                   -----
NDM_Beta_SG1      000296700558 000197800131 Migrating      100.0  10
```

```
C:\Program Files\EMC\SYMCLI\bin>symdm -sid 558 list

Symmetrix ID : 000296700558

Storage Group      Source      Target      State      Total
                   Array      Array      State      Capacity  Done
                   -----
NDM_Beta_SG1      000296700558 000197800131 Synchronized  100.0  N/A
```

When the migration is complete the system administrator runs a Host Rescan. This allows the multipathing software to discover extra paths to the host. Since Metro-based NDM uses SRDF/Metro as its transmission medium I/O operations passed through the target array. All writes are saved locally and replication to the remote array is handled by SRDF/Metro.

```
Pseudo name=emcpower178
Symmetrix ID=000296700558
Logical device ID=01B6
Device WWN=60000970000296700558533030314236
Standard UID=naa.60000970000296700558533030314236
type=Conventional; state=alive; policy=SymmOpt; queued-Ios=0
=====
--- Host ---
### HW Path          I/O Paths      - Stor -  -- I/O Path --  -- Stats ---
          I/O Paths      Interf.  Mode      State      Q-Ios Errors
=====
3 vmhba6             C0:T0:L3      FA 2d:61 active  alive      0      0
4 vmhba2             C0:T7:L3      FA 1d:60 active  alive      0      0
3 vmhba6             C0:T11:L4     FA 1d:24 active  alive      0      0
3 vmhba6             C0:T13:L4     FA 2d:24 active  alive      0      0
4 vmhba2             C0:T6:L4      FA 2d:24 active  alive      0      0
4 vmhba2             C0:T1:L4      FA 1d:24 active  alive      0      0
=====
```

The example shows two additional paths to the device. These are the new paths to the target array. The number varies based on the number of paths zoned.

If target side Disaster Recovery (DR) is required, it is possible to start the process while the migration is ongoing. This can be setup using the standard symrdf commands:

```
symrdf addgrp -sid 131 -rdfg xx -dir xx -label xxxx -remote_dir xx
-remote_sid xxx - remove_rdfg xx
```

```
symrdf createpair -establish
```

***Examining the NDM session in more detail***

There are a number of ways of viewing the Session in more detail using the list -v commands. This section shows a summary of the selected session and verifies all entities are valid prior to attempting a Commit operation. Adding a -detail to this command option displays a complete breakdown of all the individual masking elements.

```
symdm -sid <SN of Source> -sg< SG to be Migrated> list -v
```



```
C:\Program Files\EMC\SYMCLI\bin>symdm -sid 558 -sg NDM_Beta_SG1 list -v

Symmetrix ID      : 000296700558

Storage Group     : NDM_Beta_SG1
Source Array      : 000296700558
Target Array      : 000197800131

Migration State   : Synchronized
Total Capacity (GB) : 100.0
Done (%)          : N/A

Source Configuration: OK
{
  Storage Groups (1) : OK
  Masking Views (1)  : OK
  Initiator Groups (1) : OK
  Port Groups (1)    : OK
}

Target Configuration: OK
{
  Storage Groups (1) : OK
  Masking Views (1)  : OK
  Initiator Groups (1) : OK
  Port Groups (1)    : OK
}

Device Pairs (4): OK
```

To view the LUN pairings, try using the `-pairs_info` parameter.

```
symdm -sid<SN of SRC or TGT> -sg <SG to be Migrated> list -v -
pairs_info -detail
```

```
C:\Program Files\EMC\SYMCLI\bin>symdm -sid 558 -sg NDM_Beta_SG1 list -v -pairs_info -detail

Symmetrix ID      : 000296700558

Storage Group     : NDM_Beta_SG1
Source Array      : 000296700558
Target Array      : 000197800131

Migration State   : Synchronized
Total Capacity (GB) : 100.0
Done (%)          : N/A

Device Pairs (4): OK
{
  Source          Target
  Dev  Status    Dev  Status
  -----
  001B3 OK        001FC OK
  001B4 OK        001FD OK
  001B5 OK        001FE OK
  001B6 OK        001FF OK
}
}
```

### **Host view of the devices**

A `symdev` from the host shows the Label of RDF1 and RDF2 against the source arrays, in contrast to pass-through NDM the source device would have been given an R2 Mirror to cater for Pass-through mode. This does not apply for Metro NDM.

```
symdev -sid <SRC or TGT SN> list
```

```
C:\Program Files\EMC\SYMCLI\bin>symdev -sid 558 list
```

```
Symmetrix ID: 000296700558
```

Device Name		Dir		Device		Cap
Sym	Physical	SA :P	Config	Attribute	Sts	(MB)
00001	Not Visible	01D:027	TDEV	N/Grp'd ACLX	RW	6
00013	Not Visible	***:***	TDEV	N/Grp'd	RW	22875
00014	Not Visible	***:***	TDEV	N/Grp'd	RW	11619
00018	Not Visible	***:***	RDF1+TDEV	N/Grp'd	RW	4155
00019	Not Visible	***:***	TDEV	N/Grp'd	RW	65543
0001A	Not Visible	***:***	TDEV	N/Grp'd	RW	2078
0001B	Not Visible	***:***	TDEV	N/Grp'd	RW	2078
...						
001B3	Not Visible	***:***	RDF1+TDEV	N/Grp'd	RW	25601
001B4	Not Visible	***:***	RDF1+TDEV	N/Grp'd	RW	25601
001B5	Not Visible	***:***	RDF1+TDEV	N/Grp'd	RW	25601
001B6	Not Visible	***:***	RDF1+TDEV	N/Grp'd	RW	25601

```
C:\Program Files\EMC\SYMCLI\bin>symdev list -sid 131
```

```
Symmetrix ID: 000197800131
```

Device Name		Dir		Device		Cap
Sym	Physical	SA :P	Config	Attribute	Sts	(MB)
00001	Not Visible	???:???	TDEV	N/Grp'd ACLX	RW	6
00002	Not Visible	???:???	TDEV	N/Grp'd	RW	10241
00003	Not Visible	???:???	TDEV	N/Grp'd	RW	10241
00004	Not Visible	???:???	TDEV	N/Grp'd	RW	6144
00005	Not Visible	???:???	TDEV	N/Grp'd	RW	5121
00006	Not Visible	???:???	TDEV	N/Grp'd	RW	8388609
...						
001FC	Not Visible	***:***	RDF2+TDEV	N/Grp'd	WD	25601
001FD	Not Visible	***:***	RDF2+TDEV	N/Grp'd	WD	25601
001FE	Not Visible	***:***	RDF2+TDEV	N/Grp'd	WD	25601
001FF	Not Visible	***:***	RDF2+TDEV	N/Grp'd	WD	25601

**Source and target post-create**

The native and effective (internal and external) devices IDs (WWNs) are the same on the source device 1B3.

```
Symdev -sid <SRC or TGT> show <SRC or TGT device>
```

```

C:\Program Files\EMC\SYMCLI\bin>symdev -sid 558 show 1b3

Device Physical Name      : Not Visible
Device Symmetrix Name     : 001B3
Device Serial ID         : N/A
Symmetrix ID             : 000296700558

Number of RAID Groups    : 0
Encapsulated Device      : No
Encapsulated WWN         : N/A
Encapsulated Device Flags: None

Encapsulated Array ID    : N/A
Encapsulated Device Name : N/A
Attached BCV Device      : N/A

Attached VDEV TGT Device : N/A

Vendor ID                : EMC
Product ID               : SYMMETRIX
Product Revision         : 5977
Device WWN                : 60000970000296700558533030314233
Device Emulation Type    : FBA
Device Defined Label Type: N/A
Device Defined Label     : N/A
Device Sub System Id     : N/A
Cache Partition Name     : DEFAULT_PARTITION
Bound Pool Name          : SRP_1

Device Block Size        : 512

Device Capacity
{
  Cylinders      :      13654
  Tracks         :      204810
  512-byte Blocks :    52431360
  MegaBytes     :      25601
  KiloBytes     :    26215680

  Geometry Limited : No
}

Device External Identity
{
  Device WWN      : 60000970000296700558533030314233
}
***

```

The native and effective (internal and external) devices IDs (WWNs) differ on the Target Array. The target device has inherited the WWN of the source and presented it as its effective or external WWN.

```

C:\Program Files\EMC\SYMCLI\bin>symdev -sid 131 show 1fc

Device Physical Name      : Not Visible

Device Symmetrix Name     : 001FC
Device Serial ID         : N/A
Symmetrix ID              : 000197800131

Number of RAID Groups    : 0
Encapsulated Device      : No
Encapsulated WWN         : N/A
Encapsulated Device Flags: None

Encapsulated Array ID    : N/A
Encapsulated Device Name : N/A
Attached BCV Device      : N/A

Attached VDEV TGT Device : N/A

Vendor ID                 : EMC
Product ID                : SYMMETRIX
Product Revision          : 5978
Device WWN                 : 60000970000197800131533030314643
Device Emulation Type     : FBA
Device Defined Label Type : N/A
Device Defined Label      : N/A
Device Sub System Id      : N/A
Cache Partition Name     : N/A
Bound Pool Name           : SRP_1

Device Block Size        : 512

Device Capacity
{
  Cylinders      :      13654
  Tracks         :      204810
  512-byte Blocks :    52431360
  MegaBytes     :       25601
  KiloBytes     :    26215680

  Geometry Limited : No
}

Device External Identity
{
  Device WWN      : 60000970000296700558533030314233
}
...

```

### Canceling a migration in progress

At any point up to the Commit, a migration can be canceled. This removes the storage provisioned on the target array and releases any allocations and resources allocated by the NDM create operation. It also returns the source devices to the state they were before the Create operation was issued.

The cancel operation carries out the following:

- Stops replication between the source and target arrays
- Removes the masking view on the target array
- Removes the RDF pairings
- Removes the port group on the target array (if not in use in another masking view)
- Removes the initiator group on the target array
- Deallocates volumes created on the target array
- Removes the devices created on the target array

Since there is no Cutover operation and therefore no pass-through state there is no need to use the - revert parameter as used in Pass-Through NDM.

```
C:\Users\Administrator>symdm -sid 558 -sg NDM_Beta_SG1 cancel -nop
A DM 'Cancel' operation is
in progress for storage group 'NDM_Beta_SG1'. Please wait...

Analyze Configuration.....Started.
  Source SID:000296700558
  Target SID:000197800131
Analyze Configuration.....Done.
Stop Data Replication.....Started.
Stop Data Replication.....Done.
Remove Masking View(s) on Target.....Started.
Remove Masking View(s) on Target.....Done.
Remove Data Replication.....Started.
Remove Data Replication.....Done.
Remove Port Group(s) on Target.....Started.
Remove Port Group(s) on Target.....Done.
Remove Initiator Group(s) on Target.....Started.
Remove Initiator Group(s) on Target.....Done.
Remove Duplicate Device(s) on Target.....Started.
Wait for deallocation to complete.....Started.
Wait for deallocation to complete.....Done.
Remove Duplicate Device(s) on Target.....In Progress.
Remove Duplicate Device(s) on Target.....Done.
Remove Storage Group(s) on Target.....Started.
Remove Storage Group(s) on Target.....Done.
Remove Replication Environment.....Started.
Remove Replication Environment.....Done.
```

### ***Committing a migration***

When the data copy is completed, and the devices are synchronized the migration can be committed. The Commit operation completes the migration by removing the migrated application resources from the source array and releases system resources used for the migration. When the Commit operation is completed, the replication relationship between the source and target devices is removed, the masking view on the source is removed, and the source devices take the native (internal) WWN of the target LUN as its effective (external) WWN.



This leaves the target having an external WWN of the source and the source having an external WWN of the target. Both devices retain their native (internal) WWNs but these are not presented to the host.

```
symdev -sid <SRC or TGT SN> -sg <SG to be Migrated> commit
```

```
C:\Users\Administrator>symdm -sid 558 -sg NDM_Beta_SG1 commit -nop

A DM 'Commit' operation is
in progress for storage group 'NDM_Beta_SG1'. Please wait...

Analyze Configuration.....Started.
  Source SID:000296700558
  Target SID:000197800131
Analyze Configuration.....Done.
Stop Data Replication.....Started.
Stop Data Replication.....Done.
Remove Masking View(s) on Source.....Started.
Remove Masking View(s) on Source.....Done.
Remove Data Replication.....Started.
Remove Data Replication.....Done.
Remove Replication Environment.....Started.
Remove Replication Environment.....Done.

The DM 'Commit' operation successfully executed for
storage group 'NDM_Beta_SG1'.
```

**Examining a device after the commit operation**

With the removal of the masking view to the source storage array the systems administrator runs a host rescan. This removes any dead paths.

```
pplicensevmaxcse:~ # rpowermt display dev=emcpower178 host=10.60.136.146
Pseudo name=emcpower178
Symmetrix ID=000296700558
Logical device ID=01B6
Device WWN=60000970000296700558533030314236
Standard UID=naa.60000970000296700558533030314236
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0

-----
--- Host ---          - Stor -  -- I/O Path --  -- Stats ---
### HW Path          I/O Paths  Interf.  Mode   State  Q-I/Os Errors
-----
  3 vmhba6           C0:T0:L3  FA  2d:61 active  alive   0    1
  4 vmhba2           C0:T7:L3  FA  1d:60 active  alive   0    1
-----
```

Viewing the source and target devices after the Commit operation shows the WWN manipulation that has occurred. The source device now has the targets native WWN as its effective WWN and the Target device has retained the source native WWN as its effective WWN. In addition, the RDF mirror has been removed from the device.

**Removing the migration environment**

Removing the environment removes the template used to create the SRDF/Metro groups for individual SG migrations. Once this template is removed another Environment Setup operation is necessary, which creates a new template, before being able to create migrations between the source and target arrays.

```

C:\Users\Administrator>symdm -sid 131 -environment list

Symmetrix ID: 000197800131

Remote SymmID  Status
-----
000296700558  OK

C:\Users\Administrator>symdm -src_sid 558 -tgt_sid 131 environment -remove

A DM 'Environment Remove' operation is in progress. Please wait...

Analyze Configuration.....Started.
Source SID:000296700558
Target SID:000197800131
Analyze Configuration.....Done.
Remove Configuration.....Started.
Remove Configuration.....Done.

The DM 'Environment Remove' operation successfully executed.

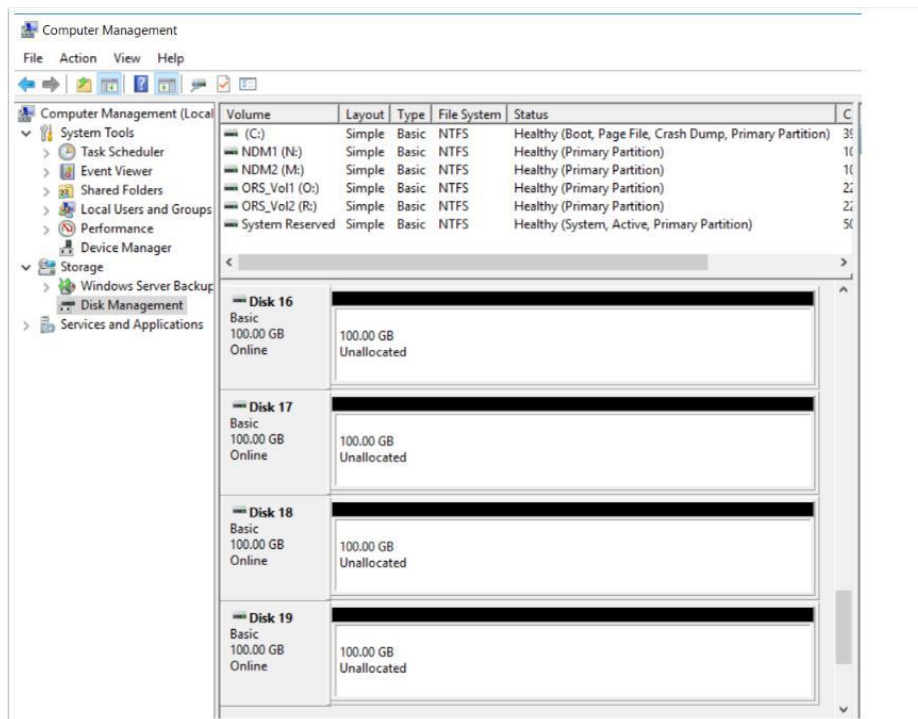
C:\Users\Administrator>symdm -sid 131 -environment list

Symmetrix ID: 000197800131

The migration session environment is not configured
    
```

**Using Unisphere for PowerMax with precopy**

Looking at the devices in this migration from the host operating system disk management (in this case, Windows Server 2016), they show as Disk 16 through Disk 19. These were previously added as RDMs to the VMs using VMware vSphere.



An example of the multipathing setup using device 137. It shows what the pathing looks like prior to the create and the host rescan. For each of the four volumes here there are four paths to the source array which are all alive and available for host use. At this point, there are no paths to the target array even though our zoning should be in place before the NDM create.

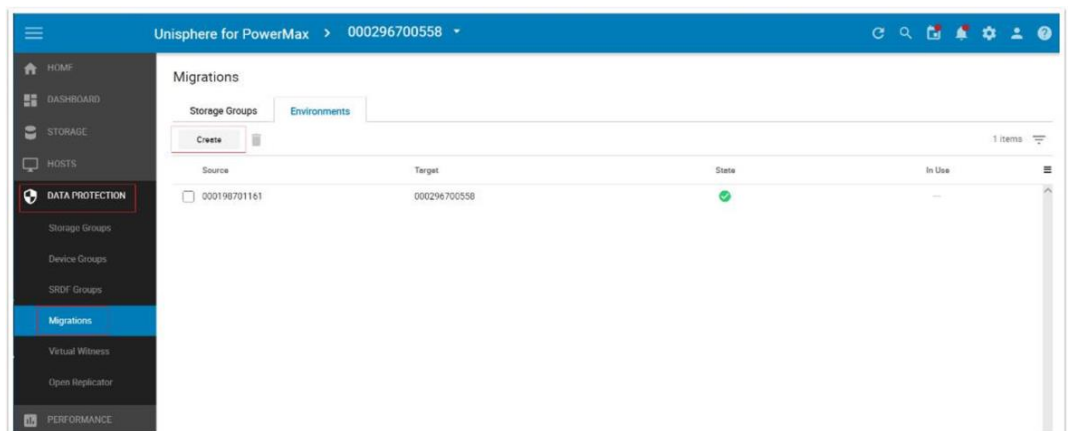
```
Pseudo name=emcpower168
Symmetrix ID=000296700558
Logical device ID=00137
Device WWN=60000970000296700558533030313337
Standard UID=naa.60000970000296700558533030313337
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
```

###	Host	I/O Paths	Stor Interf.	I/O Path Mode	State	Q-I/Os	Errors
1	vmhba4	C0:T11:L3	FA 2d:24	active	alive	0	0
1	vmhba4	C0:T7:L3	FA 1d:24	active	alive	0	0
3	vmhba3	C0:T2:L3	FA 1d:24	active	alive	0	0
3	vmhba3	C0:T0:L3	FA 2d:24	active	alive	0	0

### NDM environment setup

The Environment Setup operation configures the migration environment template required to setup the Metro groups used to migrate all applications from the source array to the target array. This template is used to define the RDF groups for each migration session. Within this definition are ports used, target ports and port count. The operation also confirms that both the source and target arrays can support the NDM operations. This includes ensuring that a usable replication pathway for data migration is available between source and target arrays. Should we need a second target array from the same source, a second environment is necessary.

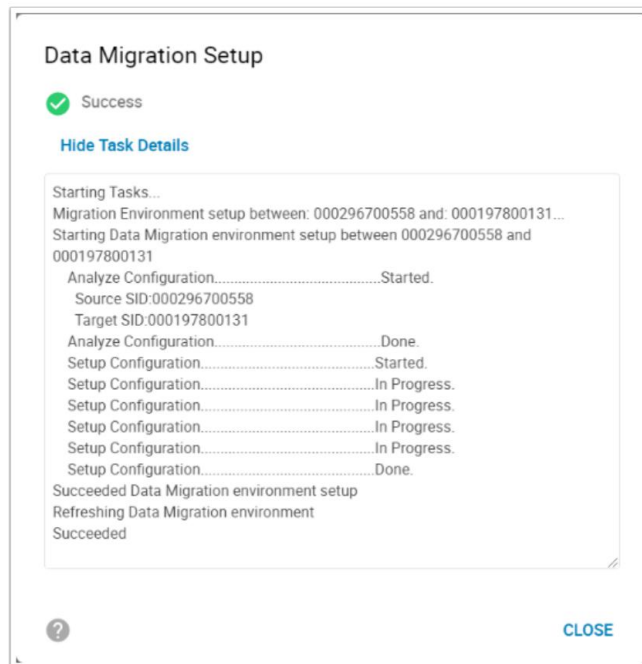
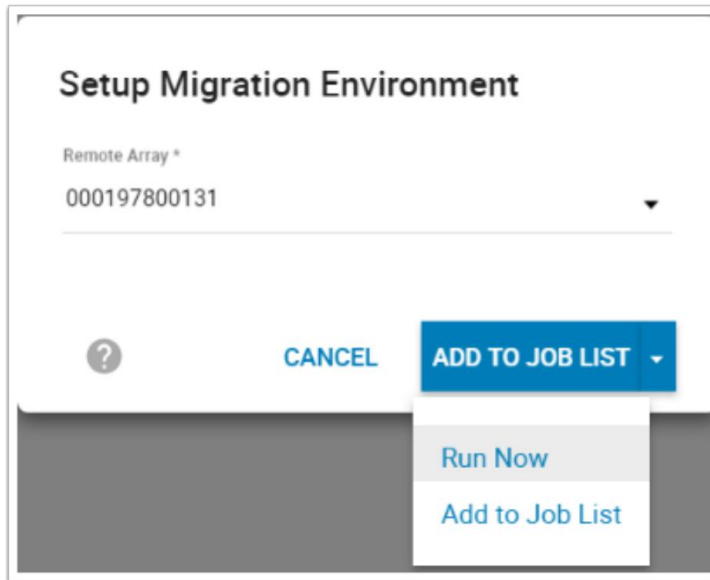
To summarize the Environment Setup runs once only for each array relationship. This Setup operation creates a template from which all other Metro-based NDM SRDF groups are modeled. Each individual NDM session requires its own RDF/Metro group to be created unlike Pass-through NDM which used a single RDF group for all sessions between the source and target.

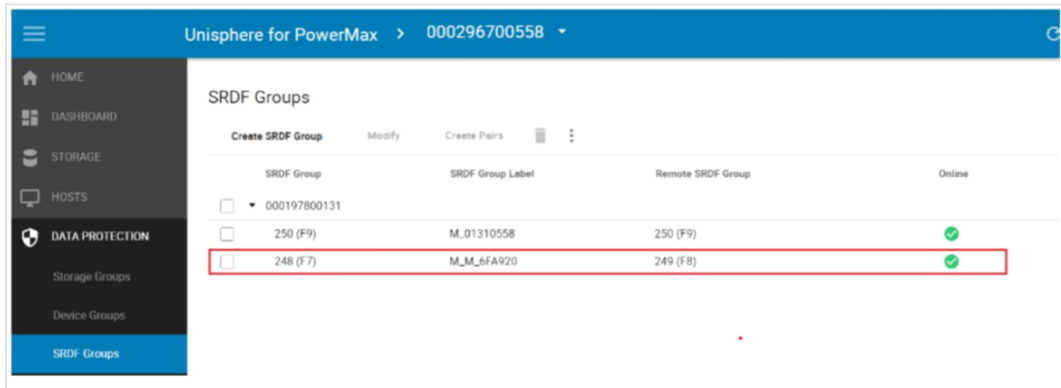
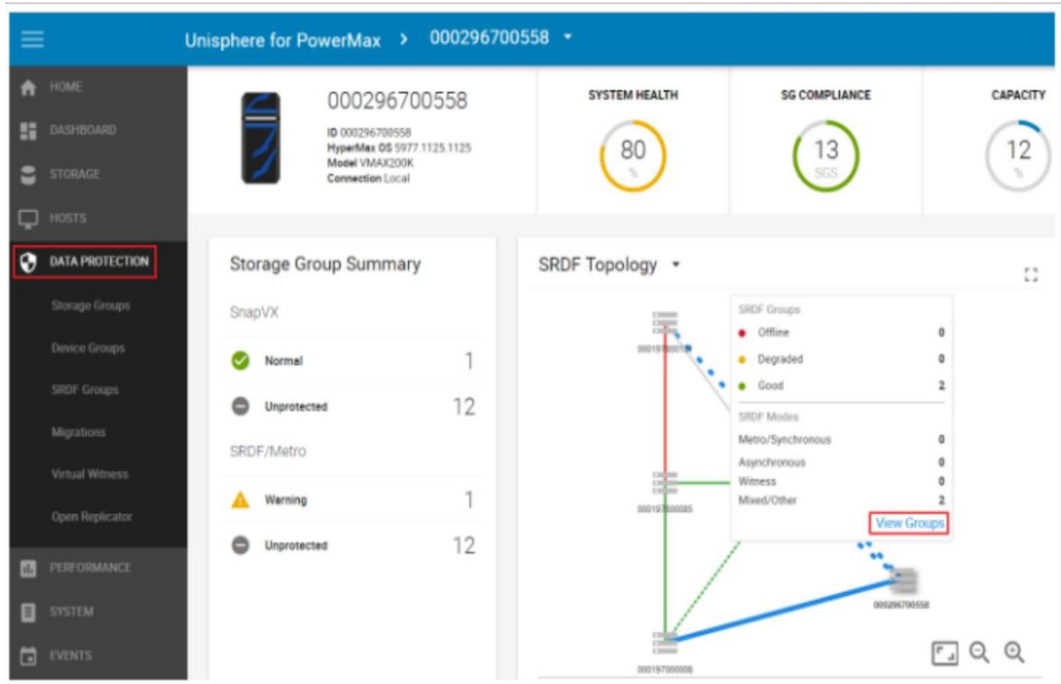


From the **Data Protection** menu select **Migrations**, Click the **Environment** tab to display any existing Environments already setup. The parameter **In Use** shows us whether the Environment is validated and usable. The **In Use** parameter also shows whether there is an active Migration using this environment.



To create an environment, click **Create**, and the window below appears. This contains a drop-down list of all the arrays available for migration operations. Should your array not be present, verify the RDF zoning and confirm the intended target array is suitable and its code level is within the support matrix. Select the relevant array and choose **Run Now**.

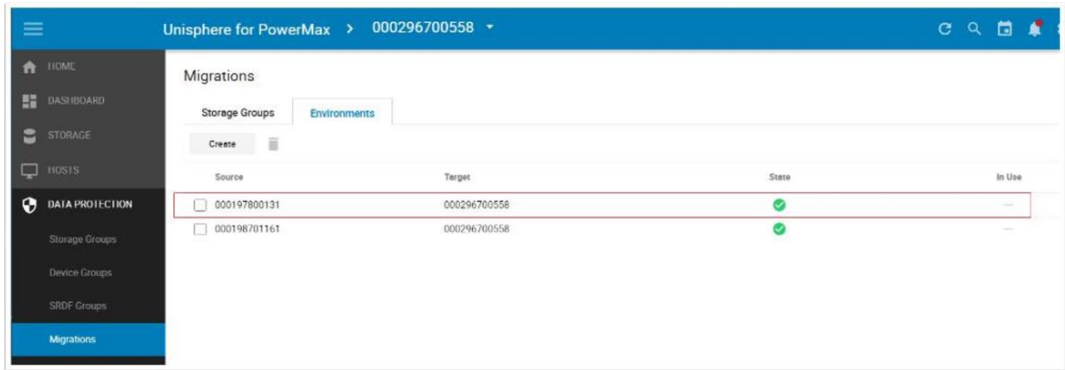




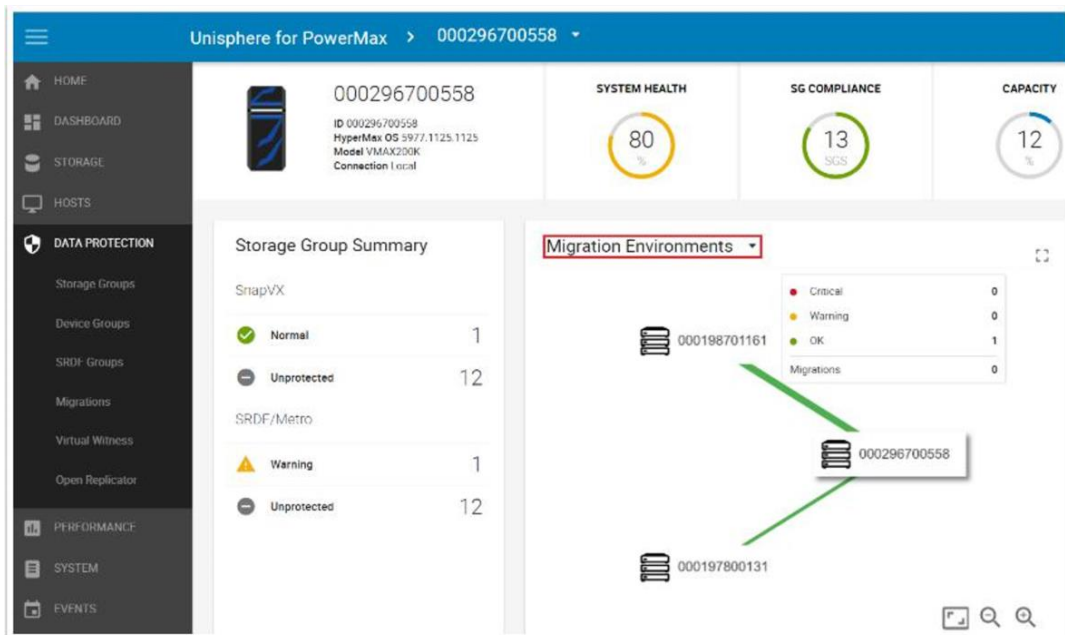
The new **Topology** view we can see the RDF group template that has been created. Go to the **DataProtection** dashboard and hover over the line between our source and target. This causes the **SRDF Groups** window to appear. Select **View Groups** to display the SRDF group window highlighted below.

The new SRDF group that has been created (248/249 on remote) as part of the Create command operation from the template (250). Group 248 is used during this migration.

From an RDF standpoint you can examine the new RDF group created to handle NDM migrations for all SG session between 558 and 131.



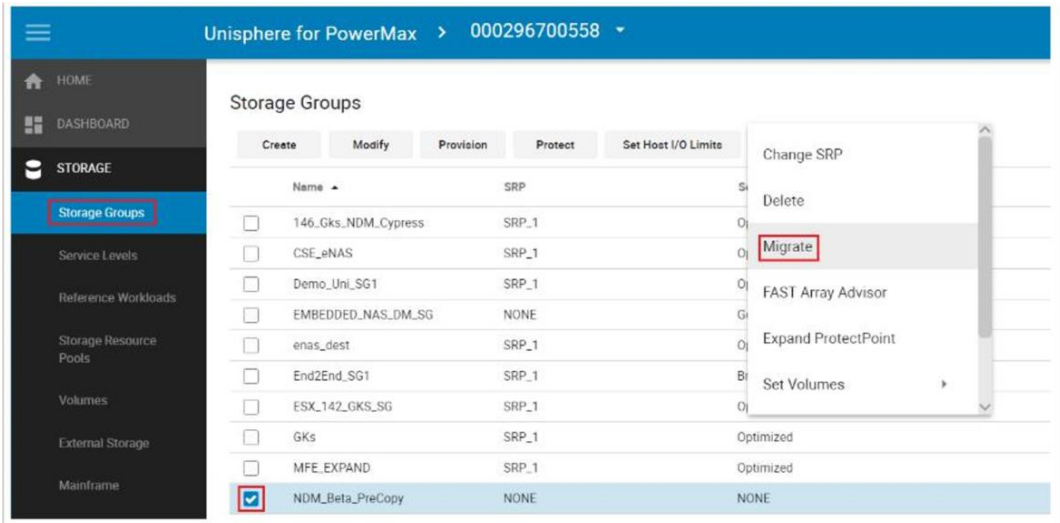
From the same Data Protection dashboard, the drop-down menu contains an option to monitor Migration Environments. A color-coded line indicates any problems. Hovering over the connection line displays the number of each connection status.



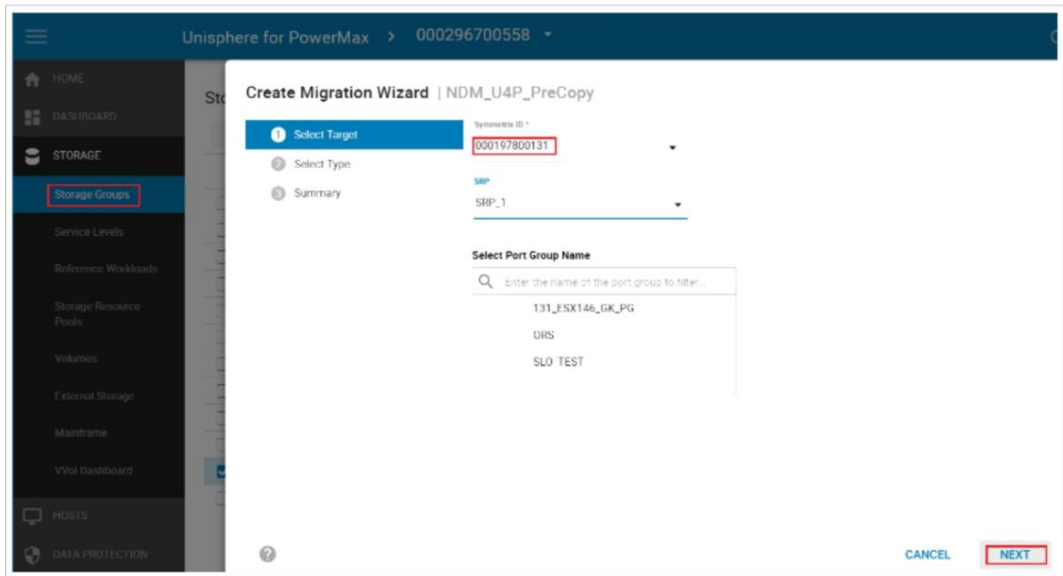
Now the environment is in place so the NDM Create for the SG planned for migration can occur.

**Create migration session with precopy**

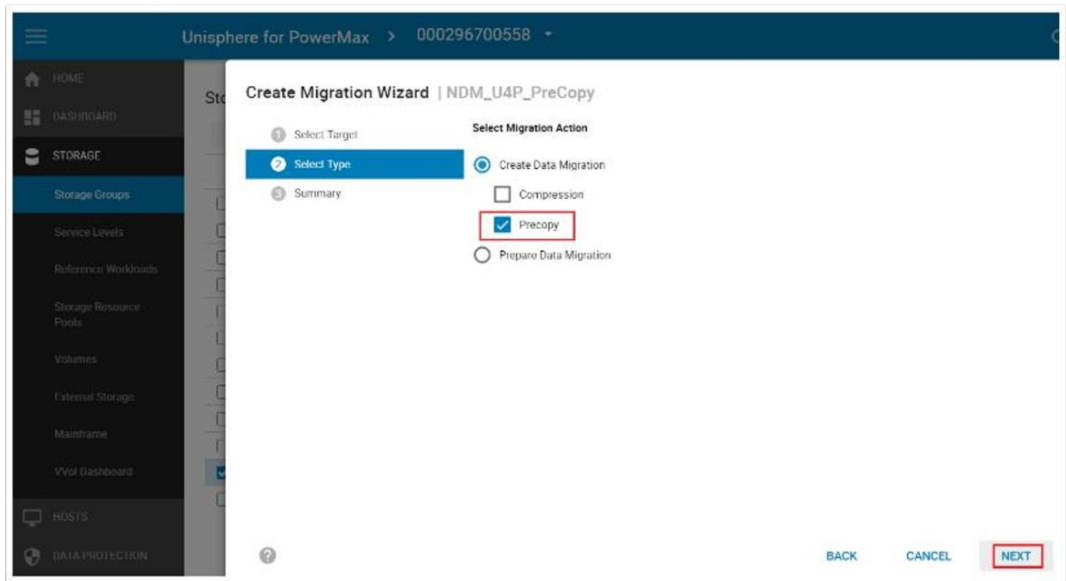
From the **Storage** tab click **Storage Groups**, from there locate the SG that is to be migrated. Select the check box and click the **More Actions** "3-Dot" icon to the right of Set Host I/O Limits. From the drop-down menu select **Migrate**.



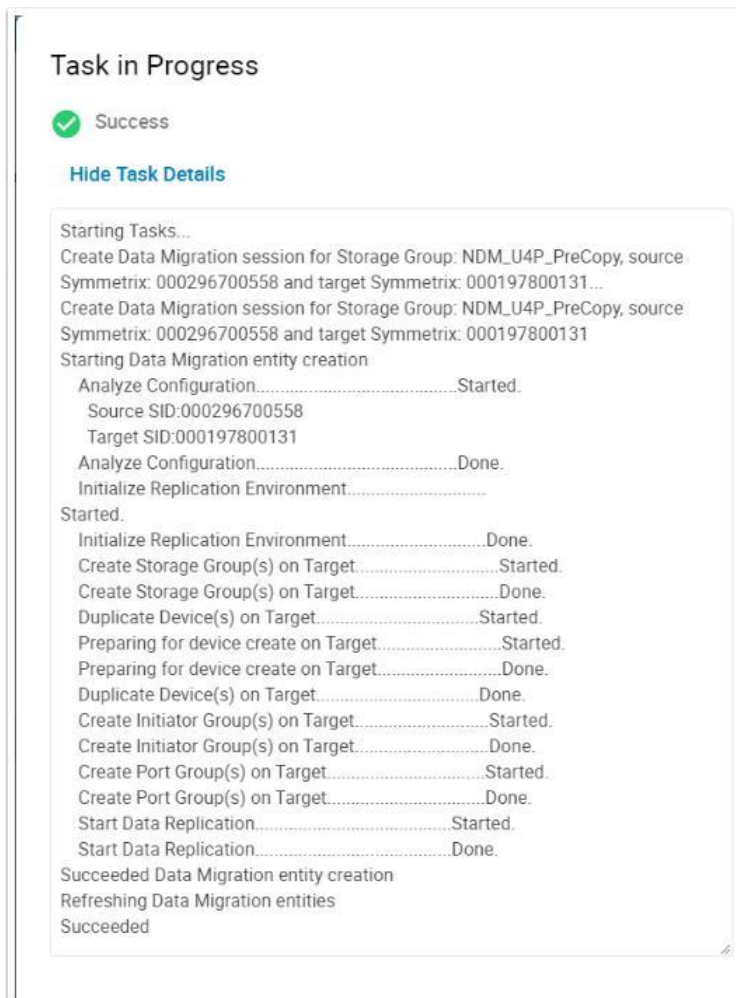
In the Create Migration Wizard, select the **Target array** (only arrays with valid environments setup appear on this drop-down menu). This example does not select a Port Group. (See the relevant masking enhancements section.) Click **Next**



On the next screen click Create Data Migration. This provides the option of selecting Compression and precopy. Check **Precopy** and select next. The Prepare Data Migration selection requires Performance data to be collected on the host. This runs a check for resources on the target array to ensure the addition of the new SG does not cause the target array to exceed any performance metrics on both FE and BE. It also produces a spread sheet to help plan the zoning required for the host from the target array.



The final page of the Wizard summarizes the planned NDM session to be created. It breaks down the planned masking view elements and the NDM parameters. Select **Run Now** to continue.



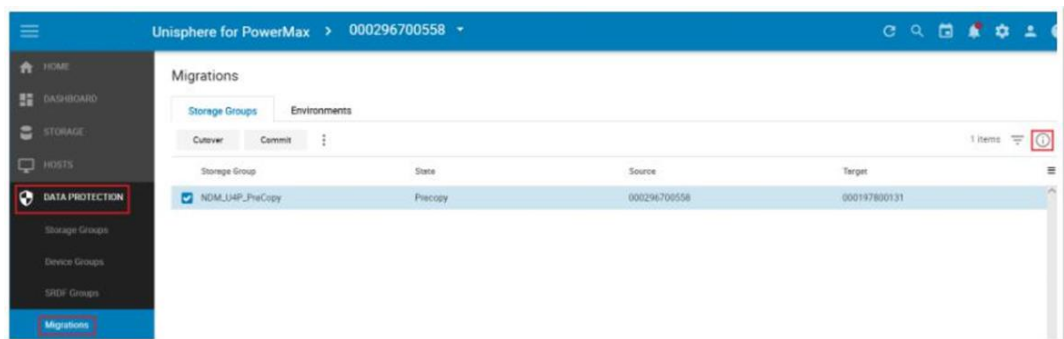
Creating the NDM session with the precopy option will also perform an environment validate as part of the setup to ensure it will complete successfully. As outlined in the create command output, the Create command completes the following:

1. Creates a storage group on the target array (name must not already exist in the target array) with the same name as the Source
2. Creates duplicate devices on the target array to match those on the storage group
3. Creates an initiator group using Initiators with entries in the login history table
4. Creates a port group (if one does not already exist or has not been selected but the user (see the relevant masking enhancements section)
5. Starts the copy process in SRDF/Adaptive Copy mode

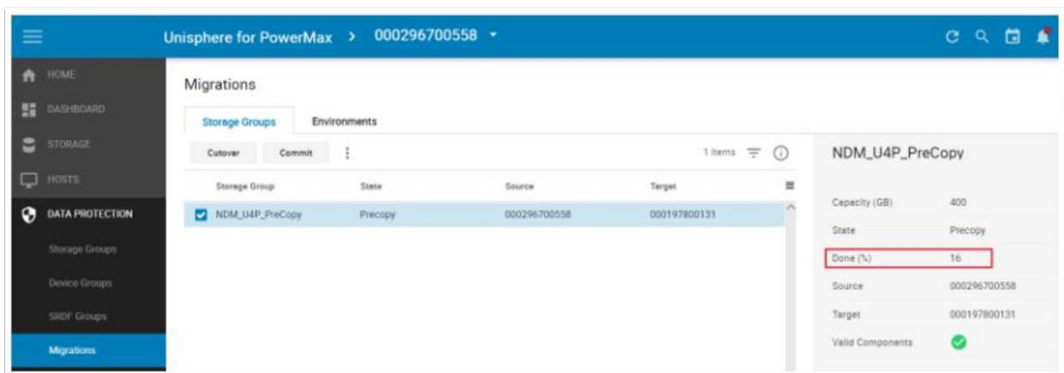
**Examine the created migration session**

Selecting the **Data Protection** tab on the left task bar and selecting **Migrations** in the drop-down menu, the storage groups currently involved in an NDM session are highlighted along with the current State and details on the source and target arrays. Since data transfer begins immediately following a Create operation, there is no need for the Pass-

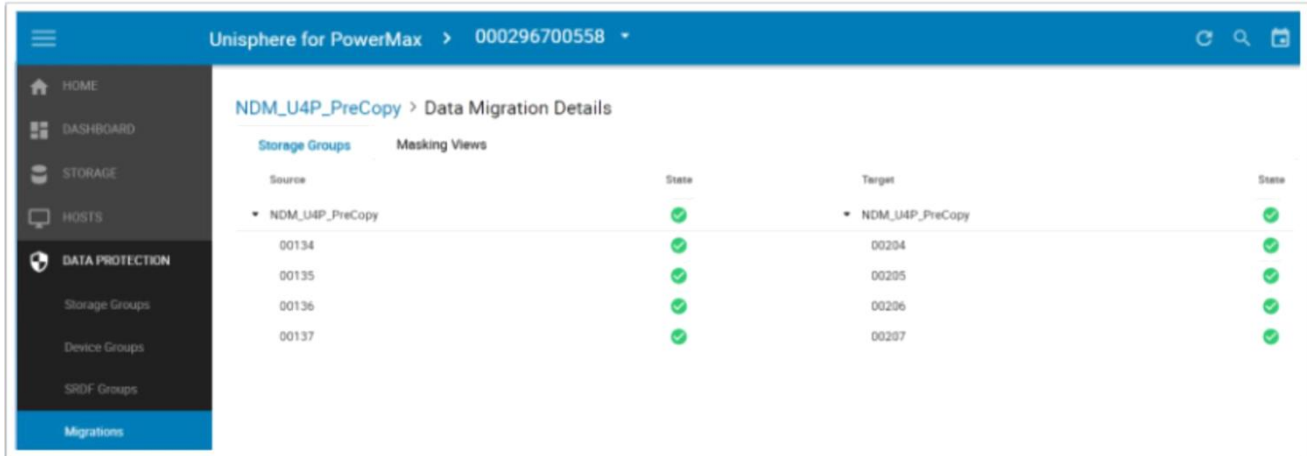
Through NDM Cutover operation. When the data is synchronizing, the systems administrator runs a rescan to allow the target paths to become active to the multipathing software. At this point, both source and target arrays are involved in an active/active relationship with all I/O serviced locally.



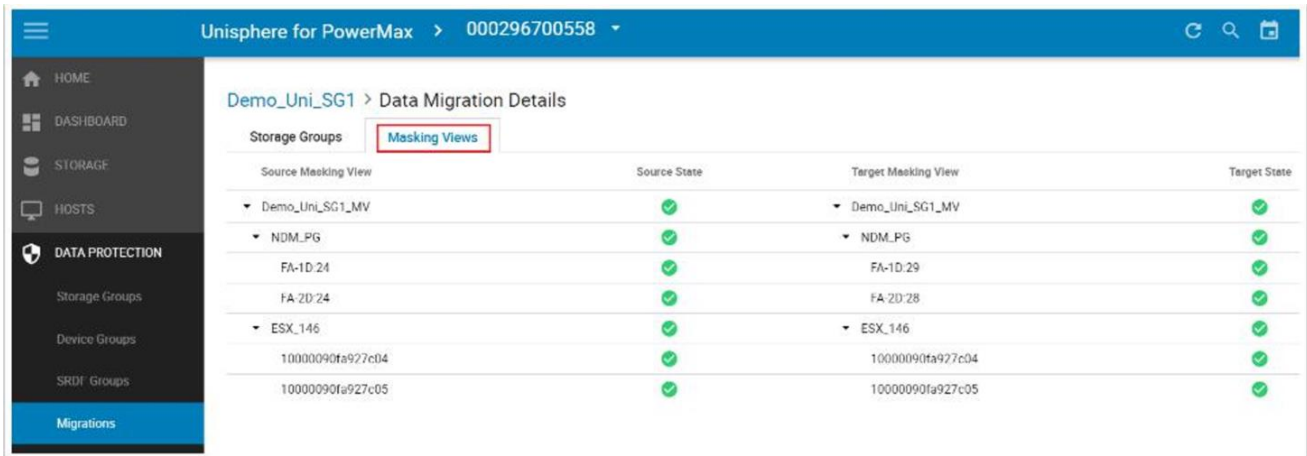
From the expanded menu in this example, the precopy is now 16% completed.



Double-clicking the NDM session will take you into the **Migration Details** view. In the screen above, we can see the individual devices involved in the session. Under the target tab, we can see devices 204 through 207 have been created on the target side by the NDM create. The **State** also shows a live status of each of the devices involved. A device without a green tick should be investigated for potential problems before continuing.

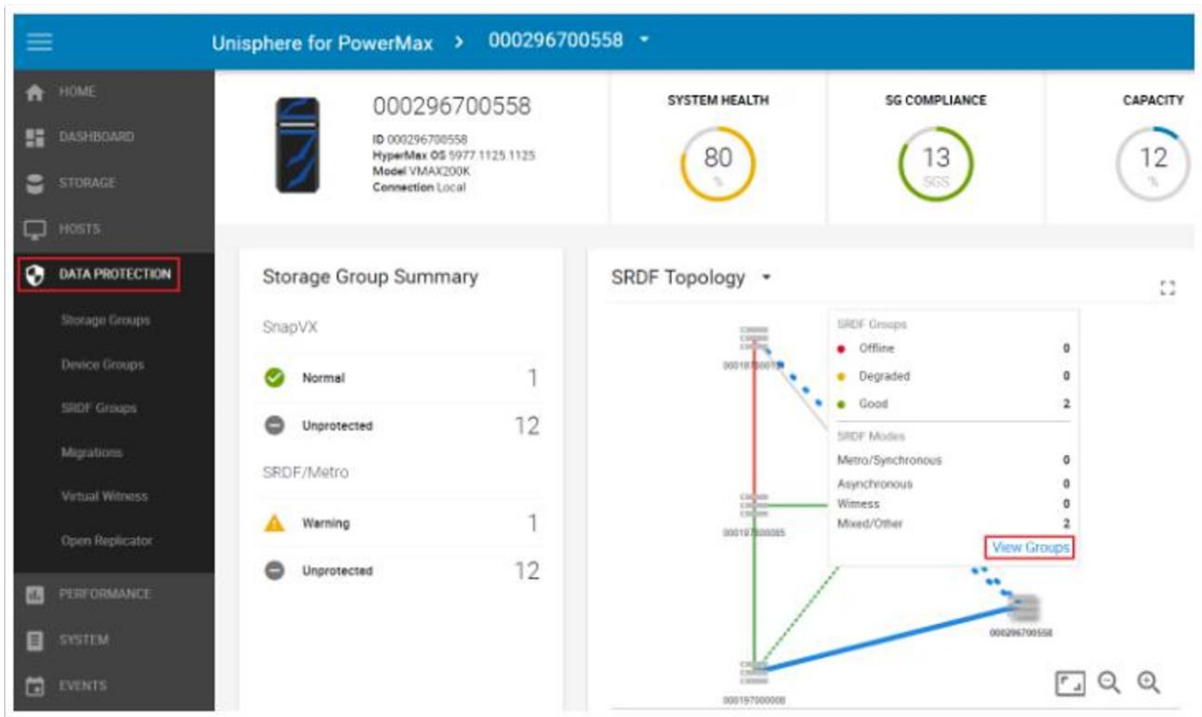


Selecting the **Masking Views** tab displays the pane outlining the masking and masking elements involved on both sides of the NDM session. This screen can be useful for troubleshooting any issues with the migration such as an unplanned or unauthorized manipulation of any of the NDM elements that hinder progress to the commit stage. In such cases the **State** does not contain a green tick, rather a red warning.



Examining the RDF environment from the Topology view shows the RDF group template that has been created. Go to the DataProtection dashboard and hover over the line between our source and target. This causes the SRDF Groups window to appear. Select View Groups to display the SRDF group window highlighted below.

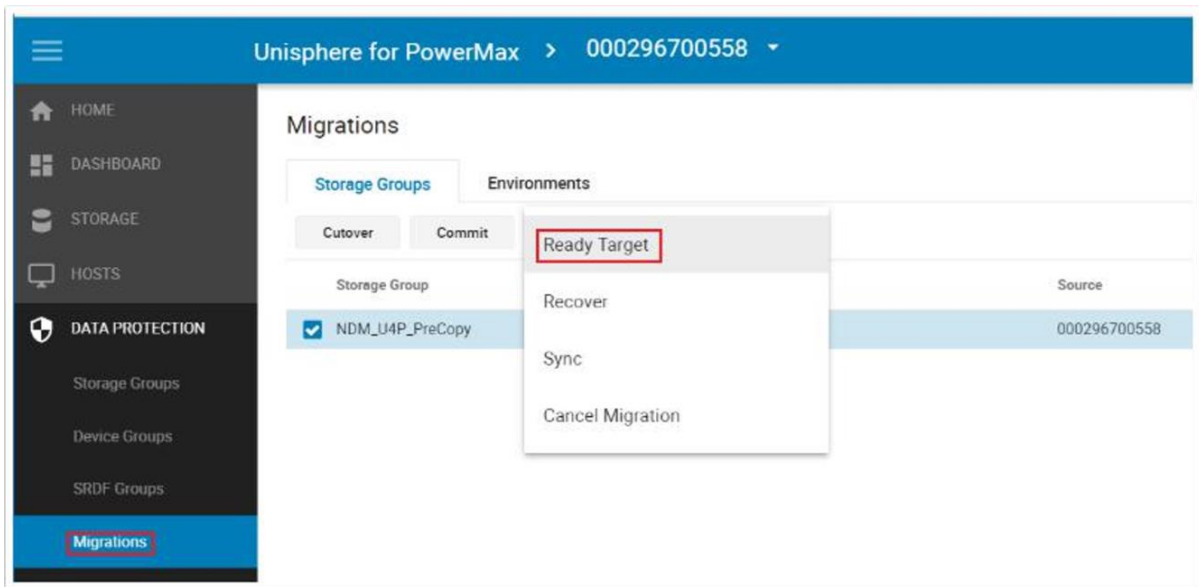




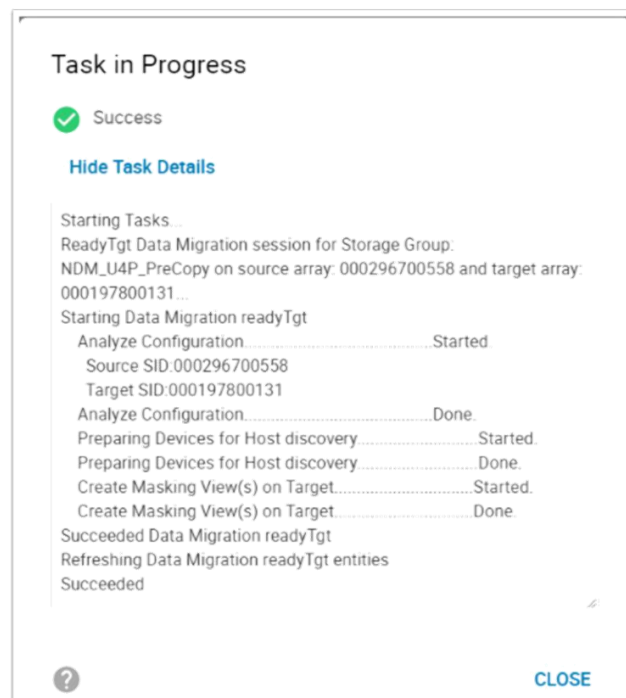
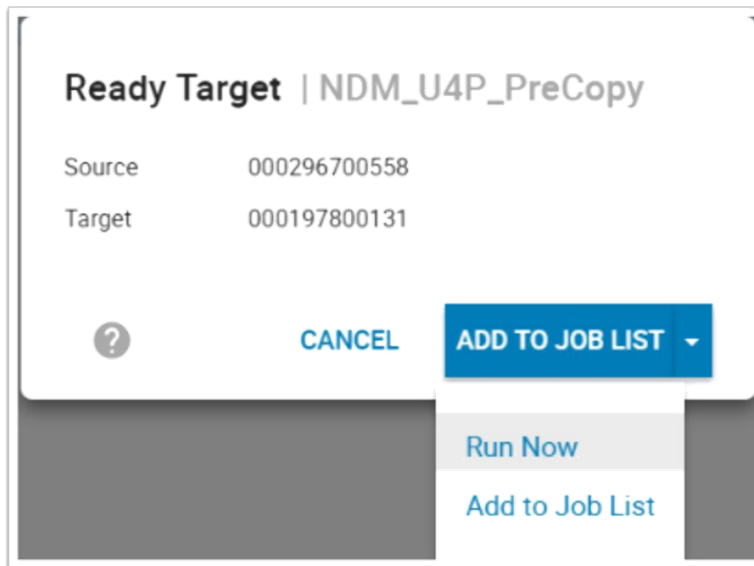
***Making the target array ready to the host***

When an adequate amount of data has been copied to the target array to negate the potential impact on the application host the target array can be made Ready to the Host.

From the active **Migrations** tab click the more actions icon (three dots) to the right of the Commit button. In the drop-down menu, select **Ready Target**. Confirm the arrays and SG are the correct combination for this migration (if multiple concurrent NDM sessions are in place) and click **Run Now**.



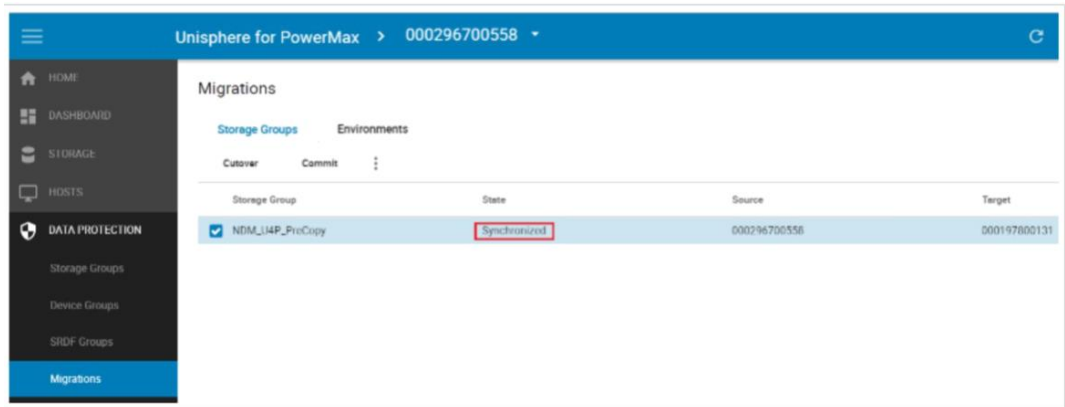




When the ready target operation begins the RDF relationship changes to Metro from adaptive copy and the masking view is created. Running a host rescan the extra paths will become available for host I/O use:

1. RDF group state moved from adaptive copy mode to Metro active
2. Target Devices are moved into a Read/Write mode
3. Masking view is created on the target array
4. SRDF state between source and target switches to active/active

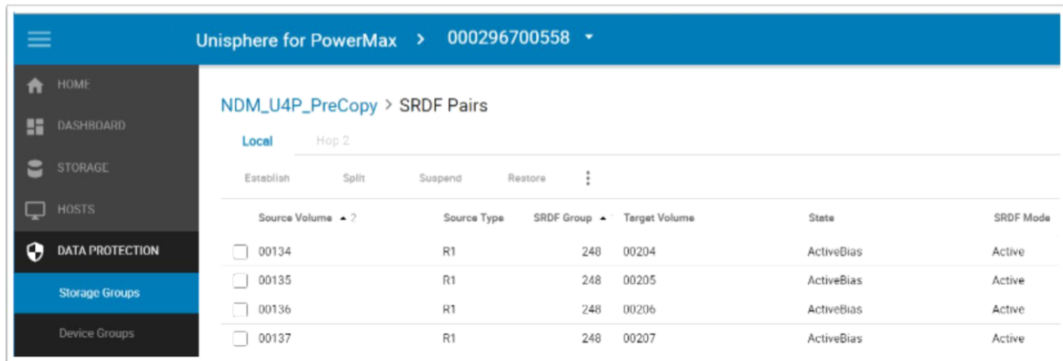
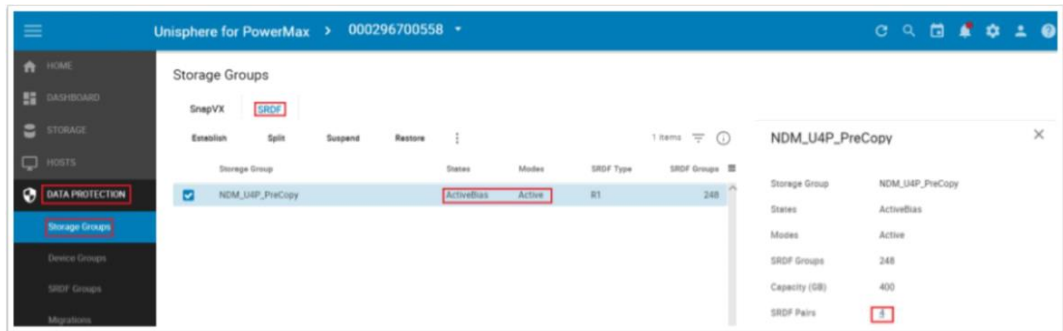
Once the ready target has been issued, we will move from a precopy state to a migrating state and eventually to a synchronized state.



At this point, the storage administrator will run a host rescan to allow the multipathing software to recognize the new paths created by issuing the Ready Target and the subsequent masking view creation on the target array. As the following figure shows, the extra paths to the target array and the dual device IDs are sharing a single effective (external) WWN. By sharing a WWN the multipathing software sees the new devices as just extra paths to the original devices.

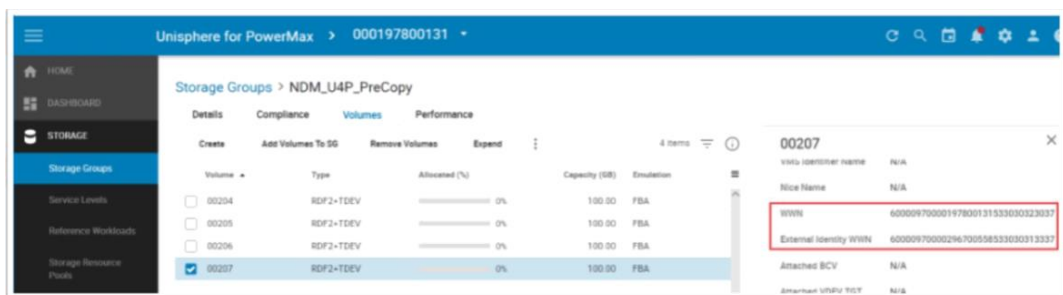
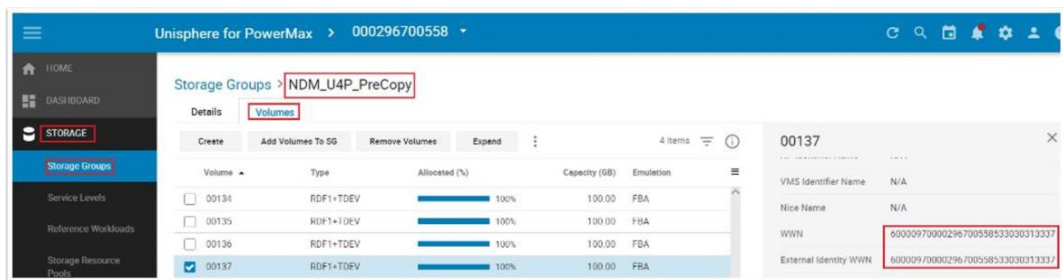
```
Pseudo name=emcpower168
Symmetrix ID=000296700558, 000197800131
Logical device ID=00137, 00207
Device WWN=60000970000296700558533030313337
Standard UID=naa.60000970000296700558533030313337
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
-----
### HW Path          Host          - Stor -  -- I/O Path --  -- Stats ---
          I/O Paths  Interf.  Mode   State   Q-I/Os  Errors
-----
1 vmhba4             C0:T12:L8  FA  2d:28 active  alive    0    0
3 vmhba3             C0:T7:L8   FA  1d:29 active  alive    0    0
1 vmhba4             C0:T11:L4  FA  2d:24 active  alive    0    0
1 vmhba4             C0:T7:L4   FA  1d:24 active  alive    0    0
3 vmhba3             C0:T2:L4   FA  1d:24 active  alive    0    0
3 vmhba3             C0:T0:L4   FA  2d:24 active  alive    0    0
```

Selecting Storage Groups from the Data Protection tab and SRDF from the window to display the SRDF relationship. In this case, we are in ActiveBias as there is no witness between the arrays. From a system standpoint, we are now processing I/Os in a Metro active/active mode with our target array being read/write to the host also.



**Examining the WWNs of source and target**

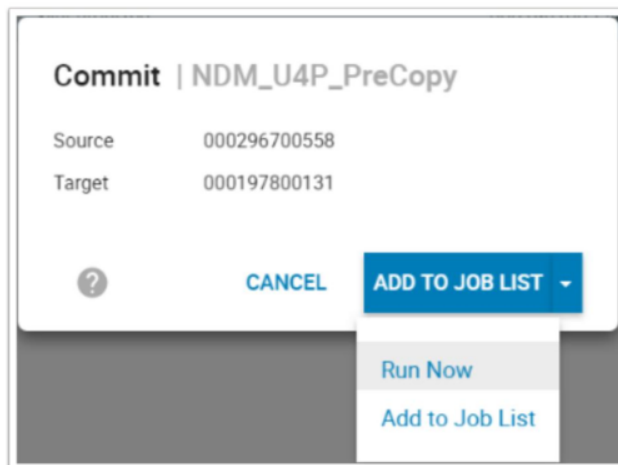
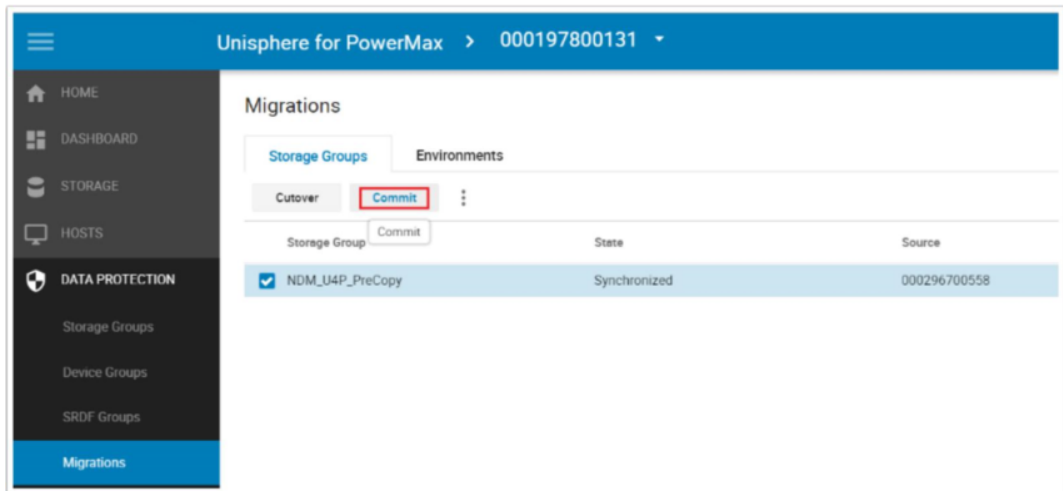
Viewing the WWN of the source device after the Create operation shows that its internal and external (that is, host visible) WWN are identical. However, in contrast, the internal WWN and external WWNs of the target device are different. Its external WWN has inherited the WWN of the source device, which means that both devices appear as a single device to the host. In effect, the multipathing software sees additional paths to the same LUN.

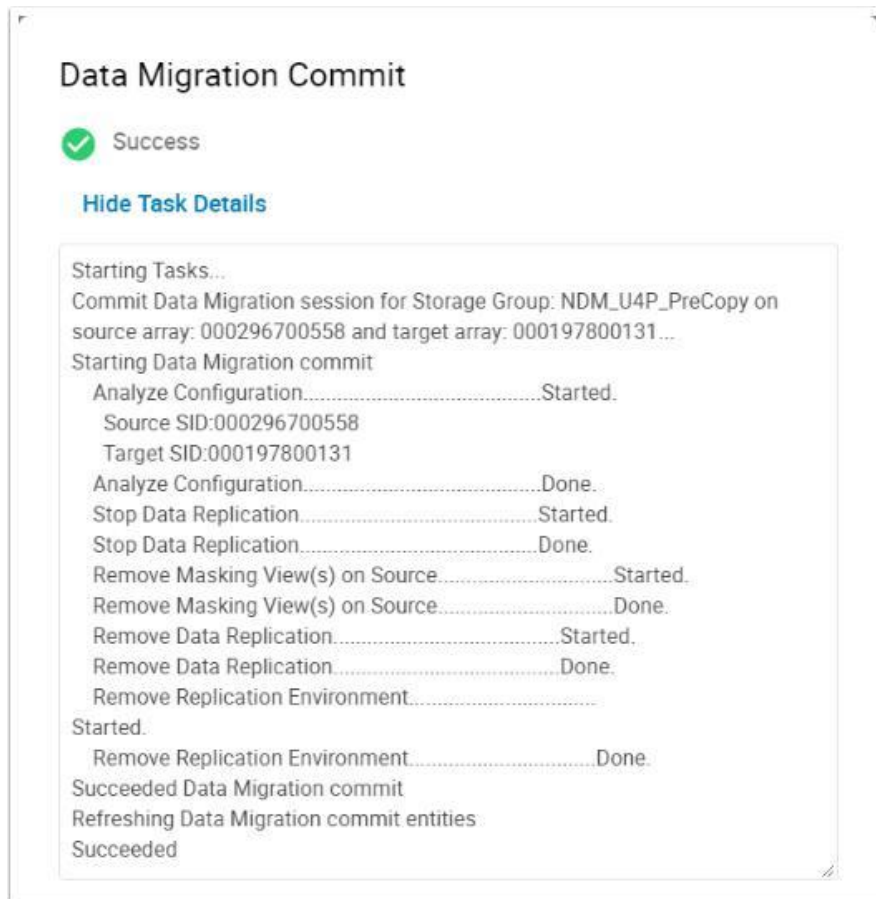


### Committing a migration

When the data copy is completed and the devices are synchronized, the migration can be committed. The Commit operation completes the migration by removing the migrated application resources from the source array and releases system resources used for the migration. When the commit is completed, the replication relationship between the source and target devices are removed, the masking view on the source is removed and the source devices take the native (internal) WWN of the target LUN as its effective (external) WWN.

The target device has the external WWN of the source and the source device has the external WWN of the target. Both devices retain their native (internal) WWNs but these are not presented to the host.





**Device paths after the commit operation**

The number of paths depends on the multipathing software in use and the zoning policy.

Carrying out a Rescan operation on the host removes the dead paths, retaining only the ones to the target devices. It also removes the SID of the original, target array, as shown in these images:

The following shows the pre-rescan status:

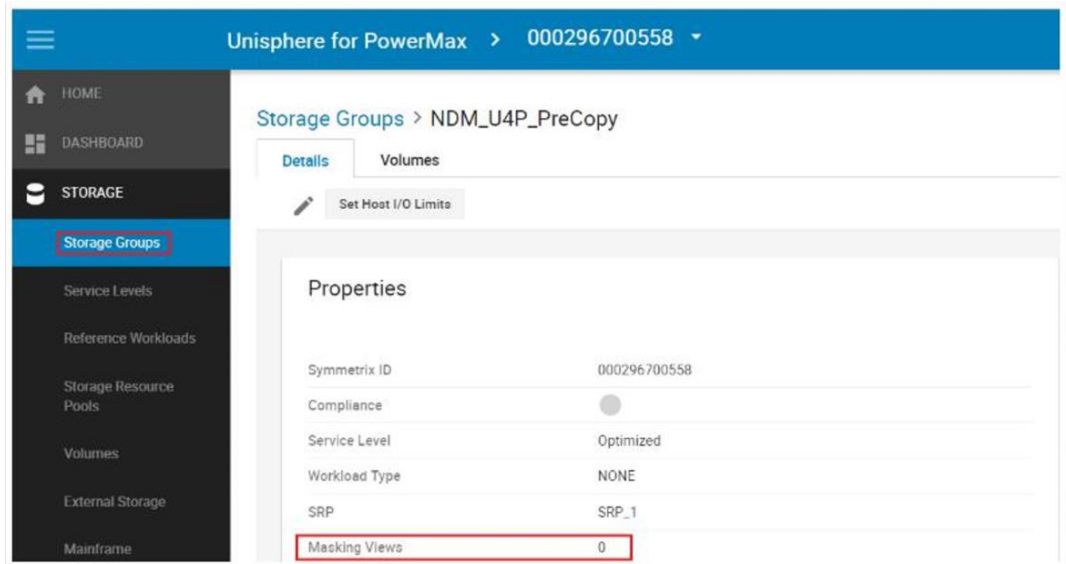
```
Pseudo name=emcpower168
Symmetrix ID=000296700558, 000197800131
Logical device ID=00137, 00207
Device WWN=60000970000296700558533030313337
Standard UID=naa.60000970000296700558533030313337
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
-----
### HW Path          Host          I/O Paths    - Stor -  -- I/O Path --  -- Stats ---
          I/O Paths  Interf.  Mode   State  Q-I/Os  Errors
-----
1 vmhba4          C0:T12:L3  FA 2d:28 active  alive    0    1
3 vmhba3          C0:T7:L3   FA 1d:29 active  alive    0    1
3 vmhba3          C0:T2:L3   FA 1d:24 active  dead     0    1
1 vmhba4          C0:T11:L3  FA 2d:24 active  dead     0    1
1 vmhba4          C0:T7:L3   FA 1d:24 active  dead     0    1
3 vmhba3          C0:T0:L3   FA 2d:24 active  dead     0    1
```

The post-rescan status is displayed as follows:

```

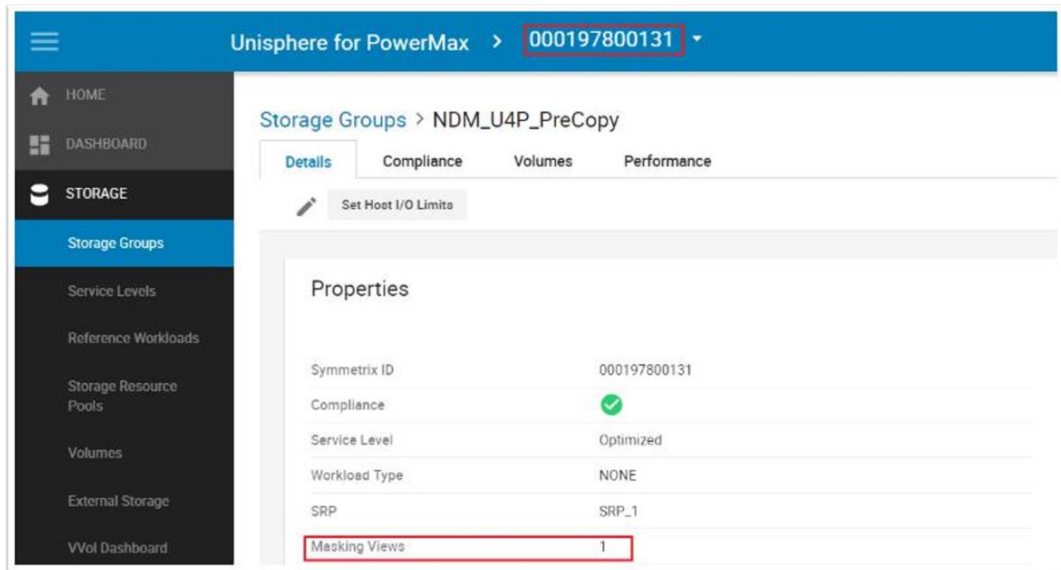
pplicensevmaxcse:- # rpowermt display dev=emcpower168 host=10.60.136.146
Pseudo name=emcpower168
Symmetrix ID=000197800131
Logical device ID=00207
Device WWN=60000970000296700558533030313337
Standard UID=naa.60000970000296700558533030313337
type=Conventional; state=alive; policy=SymmOpt; queued-IOs=0
=====
### Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf. Mode   State   Q-IOs Errors
-----
1 vmhba4             C0:T12:L8   FA 2d:28 active  alive   0      0
3 vmhba3             C0:T7:L8    FA 1d:29 active  alive   0      0
    
```

Viewing the details of the Source SG and devices demonstrate that the masking view to the host no longer exists, the RDF mirror is deleted from each of the devices and the Internal WWN from the target has been copied to the External WWN of the Source. This ensures the devices can remain on the same SAN without necessarily having to decommission the array entirely.



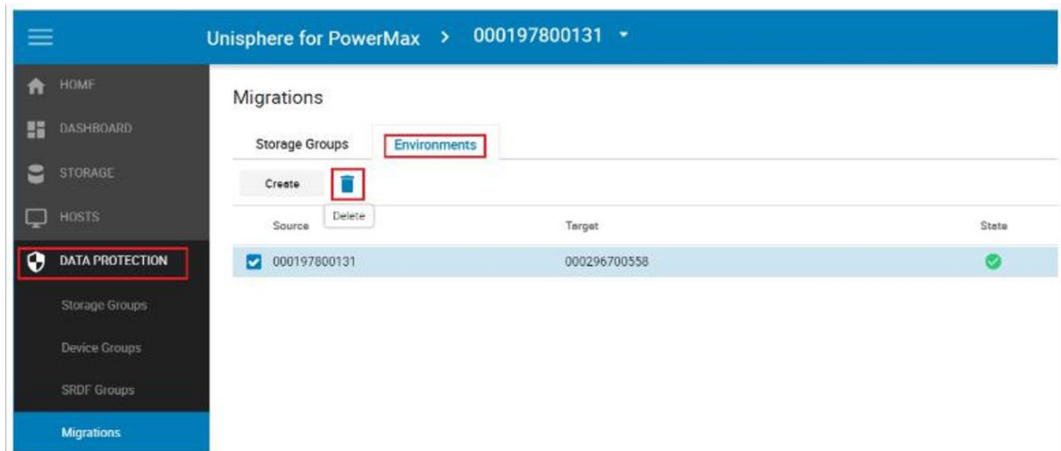
Similarly, the source device has lost its RDF mirror but retains its masking view to the host. The devices retain the Internal WWN of the Source in its External WWN identity it received at the Create stage.





### Remove NDM environment

Once all migrations are completed for a specific source and target, the migration environment can be removed. Click the **Data Protection** tab and click Migrations. On the **Environments** tab, select the environment and click on trash icon to remove. On the confirmation screen, click **Run Now**. This deletes the RDF group setup and release its resources.



### Using Solutions Enabler 9.x with precopy

The syminq command lists devices for migration. In this example, the devices are PhysicalDrive7 to PhysicalDrive10 consisting of VMAX devices 1B3 to 1B6.



```
C:\Program Files\EMC\SYMCLI\bin>syminq
```

Device		Product			Device	
Name	Type	Vendor	ID	Rev	Ser Num	Cap (KB)
\\.\PHYSICALDRIVE0		VMware	Virtual disk	1.0	N/A	41943040
\\.\PHYSICALDRIVE1		EMC	SYMMETRIX	5876	6100165000	23068800
\\.\PHYSICALDRIVE2		EMC	SYMMETRIX	5876	6100166000	23068800
\\.\PHYSICALDRIVE3		EMC	SYMMETRIX	5978	3100093000	23068800
\\.\PHYSICALDRIVE4		EMC	SYMMETRIX	5978	3100094000	23068800
\\.\PHYSICALDRIVE5		EMC	SYMMETRIX	5876	61000A8000	10485120
\\.\PHYSICALDRIVE6		EMC	SYMMETRIX	5876	61000A9000	10485120
\\.\PHYSICALDRIVE7		EMC	SYMMETRIX	5977	58001B3000	26215680
\\.\PHYSICALDRIVE8		EMC	SYMMETRIX	5977	58001B4000	26215680
\\.\PHYSICALDRIVE9		EMC	SYMMETRIX	5977	58001B5000	26215680
\\.\PHYSICALDRIVE10		EMC	SYMMETRIX	5977	58001B6000	26215680

**PowerPath view of one of the new devices**

PowerPath shows what the pathing configuration before the migration. (dev 1B6) For each of the four volumes there are four paths to the source array. All are alive and available for host use. There are no paths to the target array.

```
Pseudo name=emcpower178
Symmetrix ID=000296700558
Logical device ID=01B6
Device WWN=60000970000296700558533030314236
Standard UID=naa.60000970000296700558533030314236
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
```

Host	I/O Paths	Stor Interf.	I/O Path Mode	State	Stats Q-I/Os	Errors
3 vmhba6	C0:T11:L4	FA 1d:24	active	alive	0	0
3 vmhba6	C0:T13:L4	FA 2d:24	active	alive	0	0
4 vmhba2	C0:T6:L4	FA 2d:24	active	alive	0	0
4 vmhba2	C0:T1:L4	FA 1d:24	active	alive	0	0

**1.1.1.1 Environment setup**

Environment Setup configures the migration environment template required to create SRDF/metro groups for the migration of any application from the source array to the target array. It confirms that both the source and target arrays can support NDM. This includes that a usable replication pathway for data migration is available between the source and target. This needs to be issued once as the environment is used for all migrations between these arrays going.

```
symdm -src_sid <SRC SN> -tgt_sid <TGT SN> environment -setup
```

```
C:\Program Files\EMC\SYMCLI\bin>symdm -src_sid 558 -tgt_sid 131 environment -setup
A DM 'Environment Setup' operation is in progress. Please wait...

Analyze Configuration.....Started.
  Source SID:000296700558
  Target SID:000197800131
Analyze Configuration.....Done.
Setup Configuration.....Started.
Setup Configuration.....In Progress.
Setup Configuration.....In Progress.
Setup Configuration.....In Progress.
Setup Configuration.....In Progress.
Setup Configuration.....Done.

The DM 'Environment Setup' operation successfully executed.
```

### **Validate environment**

To validate the recently created migration environment or an existing migration environment use the `-validate` command.

```
symdm -src_sid <SRC SN> -tgt_sid <TGT SN> environment -validate
```

```
C:\Program Files\EMC\SYMCLI\bin>symdm -src_sid 558 -tgt_sid 131 environment -validate
A DM 'Environment Validate' operation is in progress. Please wait...

Analyze Configuration.....Validated.

The DM 'Environment Validate' operation successfully executed.
```

### **Validating and creating an NDM session**

Solutions Enabler examines a specific applications storage on the source array and automatically provisions equivalent storage on the target array. The target devices are assigned the identity of the source devices. Prior to running the Create command it is always worth running the `-validate` option to ensure the migration will be successful. This allows for the resolution of any issues before the migration takes place.

When issuing a Create with the `-precopy` parameter the Metro NDM session is created with the RDF relationship in SRDF/AdaptiveCopy disk mode. The data synchronization between R1 to R2 begins immediately. As with Metro-based NDM without precopy the source device is created with the R1 personality.

```
symdm -src_sid <SRC SN> -tgt_sid <TGT SN> -sg <SG to be Migrated>
- tgt_srp <SRP on TGT> -precopy -validate
```

```
symdm -src_sid <SRC SN> -tgt_sid <TGT SN> -sg <SG to be Migrated>
- tgt_srp <SRP on TGT> -precopy
```

```
C:\Program Files\EMC\SYMCLI\bin>symdm create -src_sid 558 -tgt_sid 131 -sg NDM_Beta_PreCopy -precopy -nop
A DM 'Precopy Create' operation is
in progress for storage group 'NDM_Beta_PreCopy'. Please wait...

Analyze Configuration.....Started.
  Source SID:000296700558
  Target SID:000197800131
Analyze Configuration.....Done.
Initialize Replication Environment.....Started.
Initialize Replication Environment.....Done.
Create Storage Group(s) on Target.....Started.
Create Storage Group(s) on Target.....Done.
Duplicate Device(s) on Target.....Started.
Preparing for device create on Target.....Started.
Preparing for device create on Target.....Done.
Duplicate Device(s) on Target.....Done.
Create Initiator Group(s) on Target.....Started.
Create Initiator Group(s) on Target.....Done.
Create Port Group(s) on Target.....Started.
Create Port Group(s) on Target.....Done.
Start Data Replication.....Started.
Start Data Replication.....Done.

The DM 'Precopy Create' operation successfully executed for
storage group 'NDM_Beta_PreCopy'.
```

If any stage in the validation fails, the contents of the SYMAPI log file often contain an indication of the problem. For example, a problem with an asking view or zoning conflict.

```
08/31/2017 12:41:28.688 EMC:SYMMDM validateIGEntryInMul The
initiator wwn 10000090fa927c04 is already in use in Initiator
Group 131_GKs_IG for array 000197800131
```

```
08/31/2017 12:41:28.688 Create Initiator Group(s) on
Target.....Failed.
```

Creating the NDM session with the precopy option will also perform an environment validate as part of the setup to ensure it will complete successfully. The create command performs the following:

1. Creates a storage group on the target array that has the same name as the SG on the source array (the name cannot be in use on the target array already)
2. Creates duplicate devices on the target array to match those on the storage group
3. Creates an initiator group using Initiators with entries in the login history table
4. Creates a port group (if one does not already exist or has not been selected but the user, see the relevant masking enhancements section)
5. Starts the copy process in SRDF/adaptive copy mode

**Viewing precopy status**

While the precopy is ongoing the following commands are used to monitor its progress. It should be noted since this is an R1 - R2 RDF copy all of the usual RDF query commands are valid.

```
symdm -sid <SRC or TGT SN> list
```

```
C:\Program Files\EMC\SYMCLI\bin>symdm -sid 558 list
Symmetrix ID : 000296700558
Storage Group      Source      Target      State      Total      Done
                  Array       Array       State      Capacity  (%)
                  -----
NDM_Beta_PreCopy  000296700558 000197800131 Precopy      750.0    10
```

Issuing the list command with the -v option the migration session displays a validation of the individual elements involved in the NDM session. Note the lack of masking view on the target side.

```
C:\Program Files\EMC\SYMCLI\bin>symdm -sid 558 -sg NDM_Beta_PreCopy list -v
Symmetrix ID      : 000296700558
Storage Group     : NDM_Beta_PreCopy
Source Array      : 000296700558
Target Array      : 000197800131
Migration State   : Precopy
Total Capacity (GB) : 750.0
Done (%)          : 10
Source Configuration: OK
{
  Storage Groups (1) : OK
  Masking Views (1)  : OK
  Initiator Groups (1) : OK
  Port Groups (1)   : OK
}
Target Configuration: OK
{
  Storage Groups (1) : OK
  Masking Views (0)  : N/A
  Initiator Groups (1) : OK
  Port Groups (1)   : OK
}
Device Pairs (5): OK
```

The symrdf list command shows the created pairs and the progress in terms of tracks to be copied to the R2 side. It shows the Mode (D) highlight the SRDF mode is Adaptive Copy displays the MB to track equivalent.

```
symrdf list -sid <SRC or TGT SN>
```

```
C:\Program Files\EMC\SYMCLI\bin>symrdf list -sid 558

Symmetrix ID: 000296700558

Local Device View

-----
Sym   Sym   RDF   STATUS   FLAGS   R1 Inv   R2 Inv   RDF   S T A T E S
Dev   RDev  Typ:G  SA RA LNK MTES  Tracks  Tracks  Dev  RDev Pair
-----
00018 00190 R1:13 RW RW RW S1.E      0        0 RW  WD  Synchronized
001BB 00204 R1:248 RW RW RW D1.E      0    54650 RW  WD  SyncInProg
001BC 00205 R1:248 RW RW RW D1.E      0    26587 RW  WD  SyncInProg
001BD 00206 R1:248 RW RW RW D1.E      0    53269 RW  WD  SyncInProg
001BE 00207 R1:248 RW RW RW D1.E      0    63633 RW  WD  SyncInProg
001BF 00208 R1:248 RW RW RW D1.E      0    55631 RW  WD  SyncInProg
0020A 0001C R2:13 RW WD RW S2.E      0        0 WD  RW  Synchronized

Total
Track(s)          0    253770
MB(s)             0.0  31721.3
```

The **symstat** command shows the rate at which the precopy data is copying to the target side. This can be used to estimate the time to completion for scheduling purposes. This rate will vary depending on a number of factors including RAs involved, array level of activity and distance to target. Note the RDFG will not always be 248.

```
symstat -rdfg<RDFG of Migration> -type RDF -i -sid <SRC SN>
```

```
C:\Program Files\EMC\SYMCLI\bin>symstat -rdfg 248 -type RDF -i 60 -sid 558

RDF Group Level I/O Statistics:

GRP      IO/sec      MB/sec      % Hits      IO Service Time (usec)      Q
12:01:15  READ  WRITE  READ  WRITE  RD  Min      Max      Avg      Len
12:02:15 248      0    12558      0    1594      0    295    10825    1474    3471
12:03:15 248      0    12605      0    1600      0    295    11825    1497    3964
12:04:15 248      0    12607      0    1600      0    295    11825    1468    3602
12:05:15 248      0    12604      0    1600      0    295    11825    1496    3833
12:06:15 248      0    11611      0    1474      0    295    11825    2161    3466
12:07:15 248      0    8881      0    1127      0    295    11825    2238    3458
12:08:16 248      0    8986      0    1140      0    295    11825    2056    3291
12:09:16 248      0    5823      0    739       0    295    12783    2394    0
```

During the NDM environment setup process, the first choice will be RDFG 250, descending from this number until a free group is found. In the example shown multiple NDM environments setup from array 558 so RDFG 248 was free for 558 - 131. The RDFG number does not necessarily have to be the same on both source and target.



```
C:\Program Files\EMC\SYMCLI\bin>symdm -sid 558 list
Symmetrix ID : 000296700558
```

Storage Group	Source Array	Target Array	State	Total Capacity (GB)	Done (%)
NDM_Beta_PreCopy	000296700558	000197800131	PreCopy	750.0	100

From an SRDF pair standpoint as shown in the example, we are fully synchronized to the target device.

It should be noted that it is not necessary to let the precopy to fully Synchronize before moving onto the next step and issuing the ReadyTarget command. Depending on the rate at which data is copying across and the amount of data to be copied it can be issued when the user feels comfortable I/O processing can be shared between the source and target array in an active/active configuration.

### ***Make the target array ready to the host***

Once the ReadyTarget command is issued and the systems administrator runs a rescan on the host, the migration will transition to a migrating state. If the ReadyTarget command was completed before the data has fully pre-copied the migration will enter a Migrating state until fully synchronized and then transition to a Synchronized state.

If the data has been fully pre-copied, the migration will briefly enter a migrating state to confirm data synchronization and then to a Synchronized state.

At this point, we are active/active to the host from both source and target arrays.

Issuing the ReadyTarget command performs the following:

1. Moves RDF group state from adaptive copy mode to active/active
2. Target devices are moved into a read/write mode
3. Masking view is created on the target array using the masking elements created during the create command

```
symdm -sid <SRC or TGT SN> -sg <SG to be Migrated> readytgt
```

```
C:\Program Files\EMC\SYMCLI\bin>symdm -sid 558 -sg NDM_Beta_PreCopy readytgt
Execute 'ReadyTgt' operation on SG 'NDM_Beta_PreCopy' (y/[n])? y
A DM 'ReadyTgt' operation is
in progress for storage group 'NDM_Beta_PreCopy'. Please wait...

Analyze Configuration.....Started.
  Source SID:000296700558
  Target SID:000197800131
Analyze Configuration.....Done.
Preparing Devices for Host discovery.....Started.
Preparing Devices for Host discovery.....Done.
Create Masking View(s) on Target.....Started.
Create Masking View(s) on Target.....Done.

The DM 'ReadyTgt' operation successfully executed for
storage group 'NDM_Beta_PreCopy'.
```

Following the ReadyTarget command and the host rescan the state changes to Synchronized.

```
C:\Program Files\EMC\SYMCLI\bin>symdm -sid 558 list

Symmetrix ID : 000296700558

Storage Group          Source      Target      State      Total
Array                 Array       Array       State      Capacity  Done
                        (GB)      (%)
-----
NDM_Beta_PreCopy      000296700558 000197800131 Synchronized 750.0  N/A
```

The **symrdf list** command now lists the pair state as ActiveBias signifying we are in Metro mode. And our target is read/write accessible to the host. The pairs source and target are active/active, but Solutions Enabler displays ActiveBias as there is not a witness present.

```
C:\Program Files\EMC\SYMCLI\bin>symrdf list -sid 558

Symmetrix ID: 000296700558

Local Device View
-----
Sym  Sym  RDF  STATUS  FLAGS  R1 Inv  R2 Inv  RDF  S T A T E S
Dev  RDev Typ:G  SA RA LNK MTES  Tracks  Tracks  Dev RDev Pair
-----
00018 00190 R1:13 RW RW RW S1.E  0  0 RW WD Synchronized
001BB 00204 R1:248 RW RW RW T1.E  0  0 RW RW ActiveBias
001BC 00205 R1:248 RW RW RW T1.E  0  0 RW RW ActiveBias
001BD 00206 R1:248 RW RW RW T1.E  0  0 RW RW ActiveBias
001BE 00207 R1:248 RW RW RW T1.E  0  0 RW RW ActiveBias
001BF 00208 R1:248 RW RW RW T1.E  0  0 RW RW ActiveBias
0020A 0001C R2:13 RW WD RW S2.E  0  0 WD RW Synchronized

Total
Track(s) 0 0
MB(s) 0.0 0.0
```

**Canceling a migration**

A migration can be canceled at any point up until the commit operation occurs. Cancellation removes the storage provisioned on the target array and releases any allocations and resources allocated by the NDM create –precopy operation. It also places the source deices into the state they were before the migration began.

The cancel operation performs the following:

1. Stops replication between the source and target arrays
2. Removes the masking view on the target array
3. Removes the RDF pairings
4. Removes the port group on the target array (if added as part of the create command)
5. Removes the initiator group on the target array (if not in use)



6. Deallocates volumes created on the target array
7. Removes the devices created on the target array.

Since there is no cutover step, and therefore no pass-through state, there is no need for the use of a `-revert` parameter as used in legacy NDM.

```
symdm -sid <SRC or TGT SN> -sg <SG to be Migrated> cancel
```

```
C:\Program Files\EMC\SYMCLI\bin>symdm -sid 558 -sg NDM_Beta_PreCopy cancel -nop
A DM 'Cancel' operation is
in progress for storage group 'NDM_Beta_PreCopy'. Please wait...

Analyze Configuration.....Started.
  Source SID:000296700558
  Target SID:000197800131
Analyze Configuration.....Done.
Stop Data Replication.....Started.
Stop Data Replication.....Done.
Remove Masking View(s) on Target.....Started.
Remove Masking View(s) on Target.....Done.
Remove Data Replication.....Started.
Remove Data Replication.....Done.
Remove Port Group(s) on Target.....Started.
Remove Port Group(s) on Target.....Done.
Remove Initiator Group(s) on Target.....Started.
Remove Initiator Group(s) on Target.....Done.
Remove Duplicate Device(s) on Target.....Started.
Wait for deallocation to complete.....Started.
Wait for deallocation to complete.....Done.
Remove Duplicate Device(s) on Target.....In Progress.
Remove Duplicate Device(s) on Target.....Done.
Remove Storage Group(s) on Target.....Started.
Remove Storage Group(s) on Target.....Done.
Remove Replication Environment.....Started.
Remove Replication Environment.....Done.

The DM 'Cancel' operation successfully executed for
storage group 'NDM_Beta_PreCopy'.
```

It is best practice for the storage administrator to run a rescan on the host to clear up any dead or invalid paths left over after the migration has been canceled.

### **Committing a migration**

When the data copy is completed and the devices are synchronized, the migration can be committed. The Commit operation completes the migration by removing the migrated application resources from the source array and releases system resources used for the migration. Once the commit is completed the replication relationship between the source and target devices are removed, the masking view on the source is removed and the source devices take the native (internal) WWN of the target LUN as its effective (external) WWN.

The target device has the external WWN of the source and the source device has the external WWN of the target. Both devices retain their native (internal) WWNs but these are not presented to the host.

```
symdm -sid <SRC or TGT SN> -sg <SG to be Migrated> commit
```

```
C:\Program Files\EMC\SYMCLI\bin>symdm -sid 558 -sg NDM_Beta_PreCopy commit -nop
A DM 'Commit' operation is
in progress for storage group 'NDM_Beta_PreCopy'. Please wait...

Analyze Configuration.....Started.
  Source SID:000296700558
  Target SID:000197800131
Analyze Configuration.....Done.
Stop Data Replication.....Started.
Stop Data Replication.....Done.
Remove Masking View(s) on Source.....Started.
Remove Masking View(s) on Source.....Done.
Remove Data Replication.....Started.
Remove Data Replication.....Done.
Remove Replication Environment.....Started.
Remove Replication Environment.....Done.

The DM 'Commit' operation successfully executed for
storage group 'NDM_Beta_PreCopy'.
```

**Device paths after the commit operation**

The number of paths depends on the multipathing software in use and the zoning policy.

Carrying out a Rescan operation on the host removes the dead paths, retaining only the ones to the target devices. It also removes the SID of the original, target array, as shown in these images:

```
Pseudo name=emcpower164
Symmetrix ID=000296700558, 000197800131
Logical device ID=001BF, 00208
Device WWN=60000970000296700558533030314246
Standard UID=naa.60000970000296700558533030314246
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
```

###	Host	HW Path	I/O Paths	Stor Interf.	Mode	I/O Path State	Stats Q-I/Os	Errors
1	vmhba4	C0:T12:L8	FA	2d:28	active	alive	0	1
3	vmhba3	C0:T7:L8	FA	1d:29	active	alive	0	1
1	vmhba4	C0:T11:L4	FA	2d:24	active	dead	0	1
1	vmhba4	C0:T7:L4	FA	1d:24	active	dead	0	1
3	vmhba3	C0:T2:L4	FA	1d:24	active	dead	0	1
3	vmhba3	C0:T0:L4	FA	2d:24	active	dead	0	1

Following the Commit operation but before the storage administrator runs a rescan PowerPath still shows signs of the old NDM Session in the example above. The paths are dead but still present even though the masking view has been removed from the source array during the Commit command. The Logical Device still shows the old source device number despite the RDF pairs having been broken down and the application running solely on the target array. This is why we always recommend rescanning post create, post Cancel and Commit operations.

Once the storage administrator has run a host rescan the old paths to the source array are removed as well as reference to the source array logical device and ID.

```

Pseudo name=emcpower164
Symmetrix ID=000197800131
Logical device ID=00208
Device WWN=60000970000296700558533030314246
Standard UID=naa.60000970000296700558533030314246
type=Conventional; state=alive; policy=SymmOpt; queued-I/Os=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-I/Os  Errors
-----
  1 vmhba4           C0:T12:L8  FA  2d:28 active  alive    0      1
  3 vmhba3           C0:T7:L8   FA  1d:29 active  alive    0      1

```

### Removing the migration environment

Removing the environment removes the template used to create the SRDF/Metro groups for individual SG migrations. Once this template is removed another Environment Setup operation is necessary, which creates a new template, before being able to create migrations between the source and target arrays.

```
symdm -sid <SRC or TGT SN> -environment remove
```

```

C:\Users\Administrator>symdm -sid 131 -environment list

Symmetrix ID: 000197800131

Remote SymmID  Status
-----
000296700558  OK

C:\Users\Administrator>symdm -src_sid 558 -tgt_sid 131 environment -remove

A DM 'Environment Remove' operation is in progress. Please wait...

  Analyze Configuration.....Started.
    Source SID:000296700558
    Target SID:000197800131
  Analyze Configuration.....Done.
  Remove Configuration.....Started.
  Remove Configuration.....Done.

The DM 'Environment Remove' operation successfully executed.

C:\Users\Administrator>symdm -sid 131 -environment list

Symmetrix ID: 000197800131

The migration session environment is not configured

```

## NDM Update: Offline migrations with minimal host downtime

**Introduction** This section includes a guide plan, environment overview, and walkthrough procedures for the NDM Update migration feature.

### NDM Update guide plan and environment overview

With the inclusion of NDM update for HYPERMAXOS and PowerMaxOS (released in Q3 of 2019) and Solutions Enabler 9.1, parameters available during the create step include the following:

- Offline
  - Signifies that the migration will require a short application downtime
  - Applications to be shut down prior to running (if precopy was not used)
  - Successful create requires a host rescan or reboot prior to restarting application
- Move identity
  - Propagates device identities used to access source devices to the migration target
  - If not used, the application configuration needs to be changed to reflect the new devices on the target
- Precopy
  - Application can run on source during the create while data is being migrated
  - Requires a cutover to be run once data is migrated
  - Application must be shut down prior to cutover (makes target devices visible to host and source host inactive)
  - Host rescan or reboot needed

The offline parameter must be used in order to differentiate between NDM and a migration requiring minimal downtime (NDM Update). However, the move identity (NDM Update only) and precopy parameter are optional and can be used together or as individual parameters in addition to offline. Depending on the particular environment and user needs, the offline can be used with both precopy and move identity, or with either one.

The walkthrough guide in this section describes two methods for NDM Update using the following:

- Unisphere for PowerMax, Solutions Enabler included ([Using Unisphere \(no\\_precopy, no\\_move\\_identity\)](#))
- Unisphere for PowerMax with precopy, Solutions Enabler included ([Using Unisphere \(with\\_precopy, move\\_identity\)](#))

The duration of this walkthrough uses the PowerMax arrays in the graphic below.

NDM Update migrates SGs from 000197900111 to 000197600156.

000197600156 PowerMax_8000   5978.221.221	7 %	<span style="color:red">0</span> <span style="color:orange">0</span> <span style="color:green">0</span>	-	-	-	14.4:1	
000197900111 PowerMax_2000   5978.221.221	78 %	<span style="color:red">0</span> <span style="color:orange">0</span> <span style="color:green">0</span>	-	-	-	3.5:1	

Prior to the start of a planned migration, ensure that the prerequisite checks for using NDM Update have been completed:

- Ensure both the source and target array are RDF capable (RF emulation has been added to both arrays)
- Ensure that both arrays' RDF ports are zoned to each other (minimum of two connections required)
- Check for the correct zoning from the target array to the application host.

Even though the walkthrough guide provides a detailed device examination after each step, and inspects the workings of each of the issued commands, the user needs to issue three commands to migrate an SG from source to target:

- **Create – offline** followed by a host rescan (if precopy was not used)
- **Cutover** (if precopy option used, followed by a host rescan)
- **Commit** followed by a host rescan

---

**Note:** If migrating Solaris Cluster environments using NDM Update, review [Appendix G: NDM Update with Solaris cluster](#).

---

**NDM Update walkthrough guide (source running 5977 or 5978 code)**

**Adding the migration environment: NDM or NDM Update environment setup**

Once all the pre-requisites are met, use Unisphere for PowerMax to set up the NDM environment.

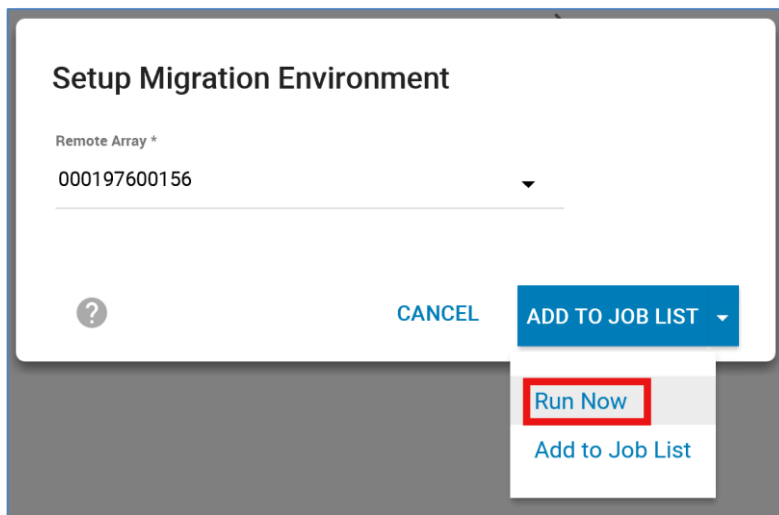
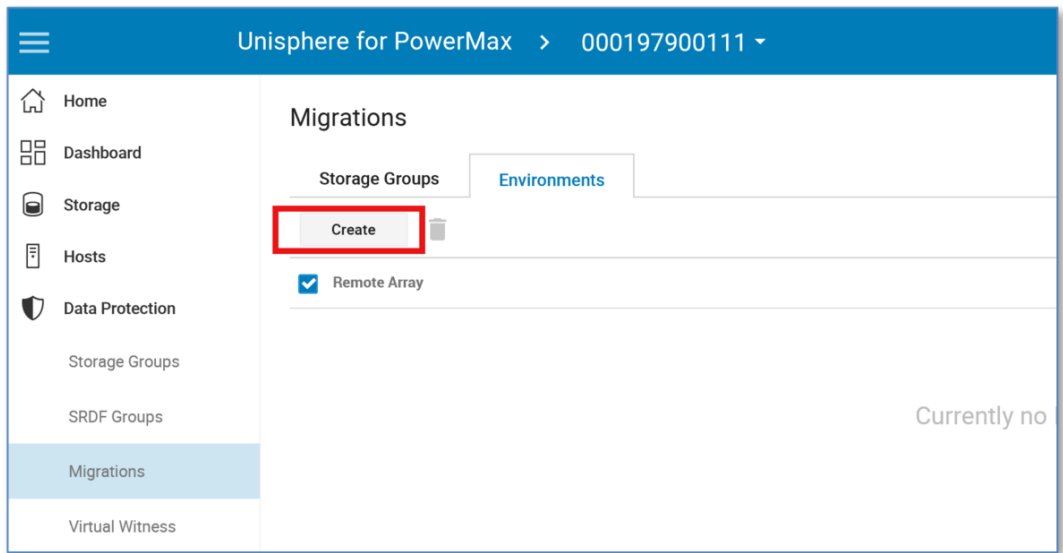
The environment setup configures the migration environment that will be required to setup the groups used to migrate all applications from the source array to the target array. It confirms that both the source and target arrays can support the NDM operations. This includes ensuring that a usable replication pathway for data migration is available between source and target arrays.

This setup only needs to be run once and will be used for all NDM migrations between the selected source and target arrays. Should we need a second target array from the same source, then a second environment will need to be configured. Should RDF group 250 not be in use, the NDM environment will set up and use this group. If it is in use, the next available RDF group descending from 250 will be used. This applies to both arrays.

From the **Data Protection** menu, select **Migrations**. Select the **Environment** tab, and this will display any existing environments already set up.

To create an environment, select **Create**, choose the target array, and select **Run Now**.






Using the CLI, run the following command:

```
symdm -src_sid 111 -tgt_sid 156 environment -setup
```

Examining further from the Replication dashboard, we can see the healthy migration environment between source array 130 and target 191.



Using the CLI, run the following commands:

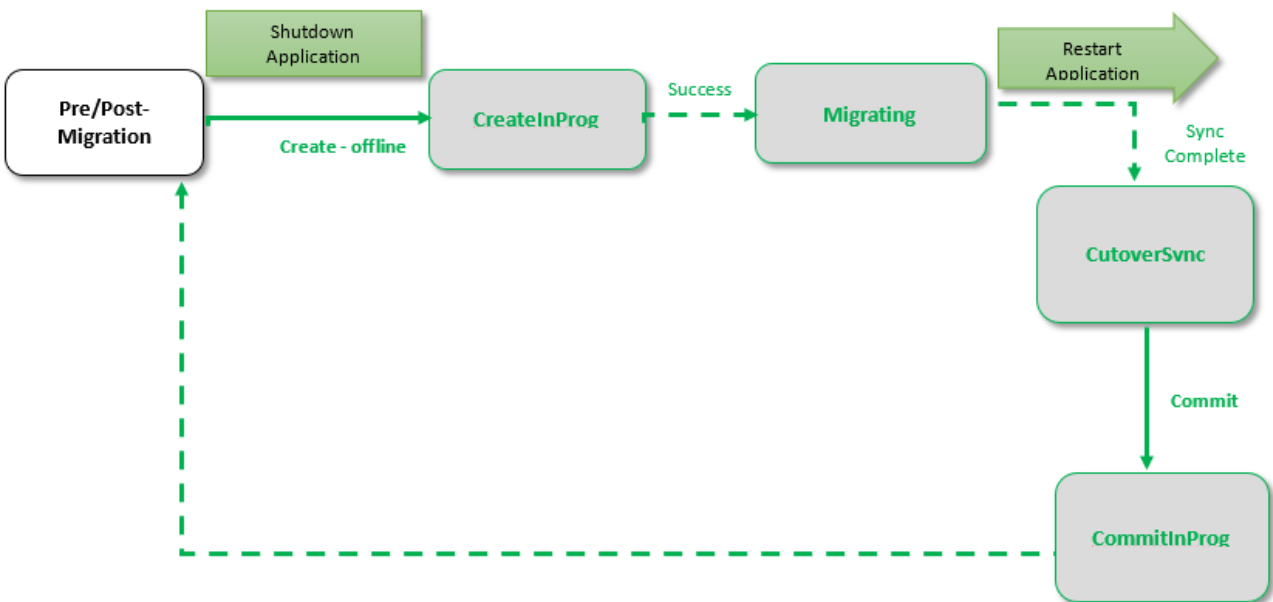
```
symdm -sid 111 -environment list
```

```
symdm -src_sid 111 -tgt_sid 156 environment -validate
```

### Using Unisphere (no\_precopy, no\_move\_identity)

This process completes in the following sequence:

1. Shut down application
2. Create offline
3. Rescan/discover paths
4. Restart application
5. Complete synchronization
6. Commit



This use case involves migrating a storage group using NDM Update. Two reasons for choosing NDM Update for migration include the following:

- The host operating system and multipathing combination is not supported (see the [support matrix](#)).
- The user's application is capable at this time of taking an outage, and the user would prefer not to spoof devices WWNs on the new array.


---

**Note:** If migrating Solaris Cluster environments using NDM Update, review [Appendix G: NDM Update with Solaris cluster](#).

---


Examining the five devices in this migration from the host operating system disk management (in this case, Windows Server 2016), they appear as Disk 2 through Disk 6. These were previously added as RDMS to the VMs using VMware vSphere.





```
C:\Users\Administrator>symlinq -wwn
Device
-----
Name          Num   Array ID      WWN
-----
\\.\PHYSICALDRIVE3  001C3  000197900111  60000970000197900111533030314333
\\.\PHYSICALDRIVE0  N/A   N/A           6000C29FC6F634EF2F656F630A7AC524
\\.\PHYSICALDRIVE2  001C2  000197900111  60000970000197900111533030314332
\\.\PHYSICALDRIVE5  001C5  000197900111  60000970000197900111533030314335
\\.\PHYSICALDRIVE1  001C1  000197900111  60000970000197900111533030314331
\\.\PHYSICALDRIVE4  001C4  000197900111  60000970000197900111533030314334
```

The volumes are contained within the storage group **Minimally\_Disrupt\_Sg**:



***NDM Update create with no\_Precopy and No Move\_Identity***

This section migrates the storage group Minimally\_Disrupt\_Sg with the offline option. The offline parameter signifies this migration will require an application outage to complete successfully.

Before migrating with the offline parameter, the host application needs to be **shut down**. Once a successful create offline is run, a host **rescan** or reboot is required.

The move identity parameter is not used for this migration, which means that the device identity is propagated from the source to the target. When the create operation is complete, the host needs to be reconfigured for the new devices.

The create command performs the following:

1. Creates a storage group on the target array (name must not already exist in the target array) with the same name as the source storage group
2. Creates duplicate devices on the target array to match those on the storage group
3. Creates an initiator group using initiators with entries in the login history table
4. Creates a port group (if one does not already exist or has not been selected by the user)

5. Invalidates the tracks on the RDF mirror to prepare for the copy
6. Starts the copy process in synchronous mode
7. Creates a masking view to the host from the target array.

At this point, you must **shut down your application**.

**Note:** There are measures in place to check for host I/O. However, this check is for a 15-second window and does **not** guarantee an application has been shut down, it is up to the user to verify manually that an application has been shut down successfully before progressing.

From the **More Actions** menu, click **Migrate**.

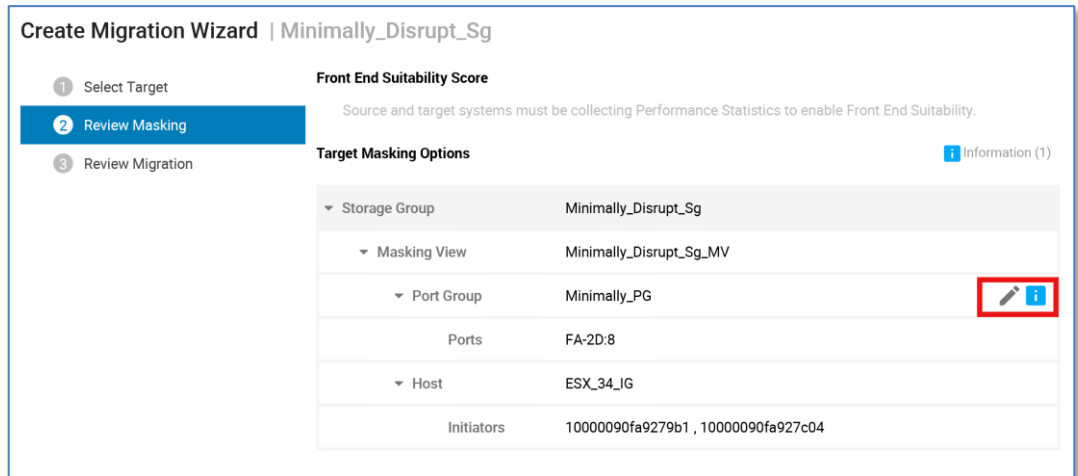
The screenshot shows the 'Storage Groups' management interface. A table lists various storage groups with columns for Name, Compliance, SRP, Service Level, and other metrics. The 'Minimally\_Disrupt\_Sg' row is selected. A context menu is open over this row, with the 'Migrate' option highlighted in red.

Name	Compliance	SRP	Service Level	Other
111_Datastore	—	SRP_1	NONE	
cse_pmax_poc_28	●	SRP_1	Diamond	
cse_pmax_poc_30	●	SRP_1	Diamond	
cse_pmax_poc_32	●	SRP_1	Diamond	
cse_pmax_poc_34	●	SRP_1	Diamond	
EMBEDDED_NAS_DM_SG	●	NONE	Diamond	
ESX34_Gks	●	SRP_1	Optimized	
Gks_SG	—	NONE	NONE	0.19 FBA 8
<input checked="" type="checkbox"/> Minimally_Disrupt_Sg	●	SRP_1	Diamond	85.00 FBA 1
robtest	●	SRP_1	Diamond	800.00 FBA 1

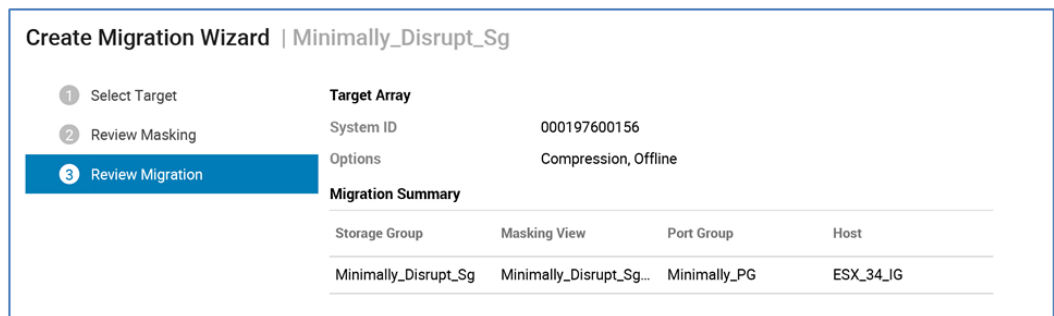
From the resulting menu, select **Offline**. This alerts Unisphere for PowerMax that NDM Update is being used to migrate this storage group. This will then highlight the need for a host reboot.

The screenshot shows the 'Create Migration Wizard' for the 'Minimally\_Disrupt\_Sg' storage group. The wizard is at the 'Select Target' step. The 'Array ID' is 000197600156 and the 'SRP' is SRP\_1. The 'Offline' checkbox is checked and highlighted with a red box. Below the checkboxes, a note states: 'Selecting offline will require a host reboot'.

From the next screen, there is the option to select a pre-existing port group from the target array, or with 9.1 and later, create a new port group as part of the create command. Click **Next**.



A breakdown of the planned migration session is displayed showing the masking elements and options selected for final approval. Click **Run Now**.




To validate and create the session using the CLI, run the following commands:

```
symdm -src_sid 111 -tgt_sid 156 -sg Minimally_Disrupt_Sg create -offline -validate
```

```
symdm -src_sid 111 -tgt_sid 156 sg Minimally_Disrupt_Sg create -offline
```

Once the create step has completed, click **OK**.

### Task in Progress

 Task in progress...

[Hide Task Details](#)

---

Starting Tasks...

Create migration for Minimally\_Disrupt\_Sg from 000197900111 to 000197600156, with compression: true, with offline selected...

Create migration for Minimally\_Disrupt\_Sg from 000197900111 to 000197600156, with compression: true, with offline selected

Starting Data Migration entity creation

Analyze Configuration.....Started.  
 Source SID:000197900111  
 Target SID:000197600156  
 Checking Source devices for IO.....Started.

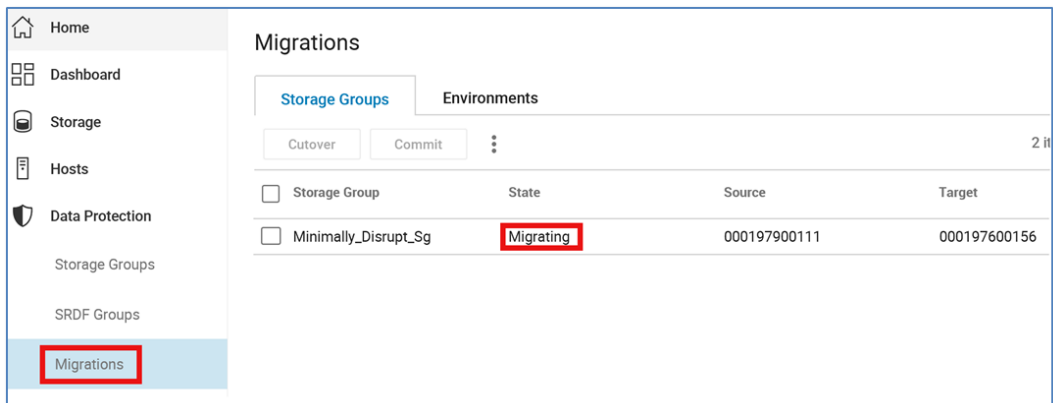
---

**Note:** Run a host rescan at this point to discover the paths to the new devices on the target array.

---

### Examine the created NDM Update session

To examine the migration session, click **Data Protection > Migrations**. The current state of the migration setting is **Migrating**.



Storage Group	State	Source	Target
<input type="checkbox"/> Minimally_Disrupt_Sg	Migrating	000197900111	000197600156

To view migration sessions using the CLI, run the following command:

```
symdm -sid 156 list
```

### Reconfigure host and commit migration

During the create, paths to the source array will become host inactive. Since move identity was not used, **the host will need to be reconfigured reflect the new devices**. In this case, the new devices were added as RDMS to the VM using vSphere, also removing the old RDMS.

From the host operating system disk management, we can see the newly added devices.

## NDM Update: Offline migrations with minimal host downtime

Volume	Layout	Type	File System	Status	Capacity	Free Space	% Free
(C:)	Simple	Basic	NTFS	Healthy (Boot, Page File, Crash Dump, Primary Partition)	119.51 GB	80.74 GB	68 %
System Reserved	Simple	Basic	NTFS	Healthy (System, Active, Primary Partition)	500 MB	156 MB	31 %
<b>Disk 1</b>							
Basic	17.00 GB Unallocated						
17.00 GB							
Online							
<b>Disk 2</b>							
Basic	17.00 GB Unallocated						
17.00 GB							
Online							
<b>Disk 3</b>							
Basic	17.00 GB Unallocated						
17.00 GB							
Online							
<b>Disk 4</b>							
Basic	17.00 GB Unallocated						
17.00 GB							
Online							
<b>Disk 5</b>							
Basic	17.00 GB Unallocated						
17.00 GB							
Online							

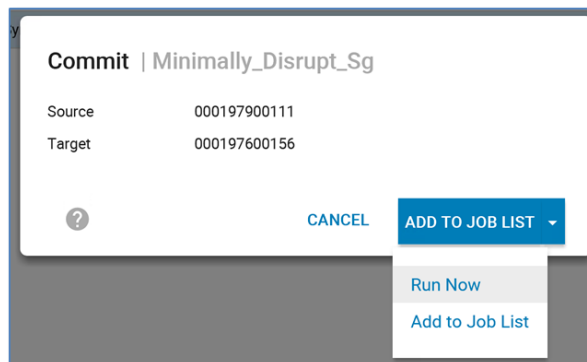
Once the devices have been reconfigured to the host and the data is migrated from the source to the target array, we can commit the migration.

Click **Data Protection > Migrations**. The migration state will show **CutoverSync**.

Since this migration did not include precopy, a cutover is not necessary.

Storage Group	State	Source	Target
<input checked="" type="checkbox"/> Minimally_Disrupt_Sg	CutoverSync	000197900111	000197600156

From here, select the **Minimally\_Disrupt\_Sg** storage group and click **Commit**. Click **Run Now**.



Upon successful completion of a commit operation, the replication between source-side and target-side devices is removed and the application migration is complete.

To commit the session using the CLI, run the following command:

```
symdm -src_sid 156 -sg Minimally_Disrupt_Sg commit
```

---

**Note:** Run a host rescan at this point to remove any potential old paths to unmasked devices.

---

### ***Canceling a migration***

The cancel operation cancels a migration that has not yet been committed. If the application was reconfigured or running I/O against the target array, the administrator must shut down the application before issuing the cancel command. Upon successful completion of a cancel operation, if the `-move_identity` option was not given as part of the create operation, the application configuration needs to be changed to use the original source LUNs.

The administrator must perform a host rescan with the option to remove dead paths to clean up the paths that were created for the target-side LUNs. The administrator must identify the source or the target array as well as the application whose migration is to be canceled.

Upon completion of a cancel operation, the application storage will be as it was prior to the start of the migration:

- Connections between source-side and target-side devices that were used to migrate data through the DM replication pathway are severed.
- Any target-side resources configured for the application that are not used by other applications on the target array are removed. This includes pre-existing IGs or PGs used by the migrated view on the target array if these are not used in other masking views or groups.

To cancel a migration using the CLI, perform the following command:

```
symdm -sid 156 -sg Minimally_Disrupt_sg cancel
```

---

**Note:** If the administrator has configured local or remote replication of the target-side devices, a cancel will be blocked. The administrator needs to remove the replication sessions before the cancel can run.

---

### ***Recovering a failed migration***

A recover operation is needed after a migration step completes with a failed state, and it is not normally required as a part of a migration. Failed states, depending on the reason for the failure, can include the following:

- CreateFailed
- CutoverFailed
- MigrateFailed
- CancelFailed
- CommitFailed

After the condition that caused a migration operation (create, cutover, commit, or cancel) to fail has been corrected, a recover operation can be invoked to continue with the migration by completing the following:

- Determining which migration operation failed
- Putting the migration session resources (connections, devices) into the appropriate state to allow the failed operation to complete
- Repeating or resuming (depending on the cause of the failure) the failed action

To recover a failed migration using the CLI (once the reason for failure has been resolved), run the following command:

```
symdm -sid 156 -sg Minimally_Disrupt_sg recover
```

---

**Note:** The administrator must identify either the source or the target array as well as the application storage from which the data migration is to be recovered.

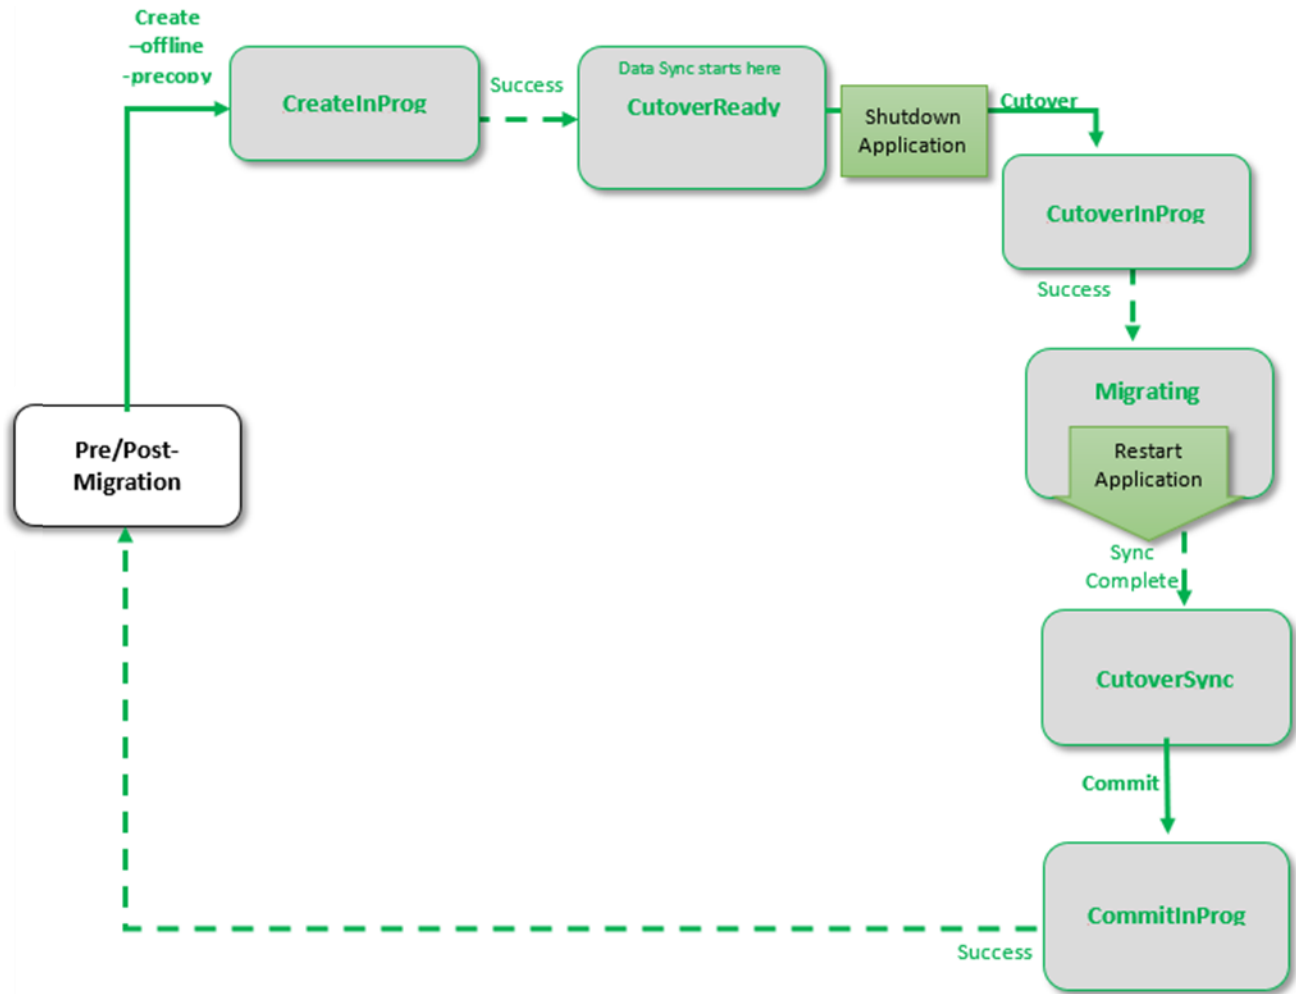
---

### Using Unisphere (with precopy, move\_identity)

This process completes in the following sequence:

1. Create offline with precopy
2. Monitor copy progress
3. Shut down application
4. Cutover
5. Rescan/discover paths
6. Restart application
7. Complete synchronization
8. Commit






This use case involves migrating a storage group using NDM Update. Two reasons for choosing NDM Update for migration include the following:

- The host operating system and multipathing combination is not supported (see the [support matrix](#)).
- The user's application is capable at this time of taking an outage and the user would prefer not to spoof device WWNs on restart

Examining the five devices in this migration from the host operating system disk management (in this case, Windows Server 2016), they show as Disk 2 through Disk 6. These were previously added as RDMs to the VMs using VMware vSphere.



The volumes are contained within a storage group called Minimally\_Disrupt\_Sg:


```
C:\Users\Administrator>syminq -wwn
```

Device		Device	
Name	Num	Array ID	WWN
\\.\PHYSICALDRIVE3	001C3	000197900111	60000970000197900111533030314333
\\.\PHYSICALDRIVE0	N/A	N/A	6000C29FC6F634EF2F656F630A7AC524
\\.\PHYSICALDRIVE2	001C2	000197900111	60000970000197900111533030314332
\\.\PHYSICALDRIVE5	001C5	000197900111	60000970000197900111533030314335
\\.\PHYSICALDRIVE1	001C1	000197900111	60000970000197900111533030314331
\\.\PHYSICALDRIVE4	001C4	000197900111	60000970000197900111533030314334

**Note:** If migrating Solaris Cluster environments using NDM Update, review [Appendix G: NDM Update with Solaris cluster](#).

***NDM Update create using precopy and choosing to move identity***

This walkthrough migrates the storage group Minimally\_Disrupt\_Sg with the offline option. The offline parameter signifies this is a migration that requires minimal disruption.



This walkthrough will also precopy the data to the target array and choose to move identity or spoof WWNs on the target.

The precopy option allows applications to continue running on the source during the create while data is being migrated. Once the data migration is complete, a cutover runs which requires the application to be shut down. The cutover also makes the target devices visible to the host and the source devices inactive.

The move identity parameter is also used, which propagates the source device identities to the target. Using the move identity does not require the host configuration to be changed to reflect new devices.

After successful completion of a create -precopy operation, the administrator can monitor the progress of the data copy and then perform a cutover operation when enough of the data copy is completed. The cutover operation switches the application to run on the target devices.

The create with precopy command completes the following:

1. Creates a storage group on the target array (name must not already exist in the target array) with the same name as the source storage group
2. Creates duplicate devices on the target array to match those on the storage group
3. Creates an initiator group using initiators with entries in the login history table
4. Creates a port group (if one does not already exist or has not been selected by the user)
5. Invalidates the tracks on the RDF mirror to prepare for the copy
6. Starts the copy process in adaptive copy mode

From the **More Actions** menu, click **Migrate**.

	Name	Compliance	SRP	Service Level		
<input type="checkbox"/>	111_Datastore	—	SRP_1	NONE		
<input type="checkbox"/>	cse_pmax_poc_28	●	SRP_1	Diamond		
<input type="checkbox"/>	cse_pmax_poc_30	●	SRP_1	Diamond		
<input type="checkbox"/>	cse_pmax_poc_32	●	SRP_1	Diamond		
<input type="checkbox"/>	cse_pmax_poc_34	●	SRP_1	Diamond		
<input type="checkbox"/>	EMBEDDED_NAS_DM_SG	●	NONE	Diamond		
<input type="checkbox"/>	ESX34_GKs	●	SRP_1	Optimized		
<input type="checkbox"/>	Gks_SG	—	NONE	NONE	0.19 FBA	8
<input checked="" type="checkbox"/>	Minimally_Disrupt_Sg	●	SRP_1	Diamond	85.00 FBA	1
<input type="checkbox"/>	robtest	●	SRP_1	Diamond	800.00 FBA	1

From the resulting menu, select **Offline**. This alerts Unisphere for PowerMax that we intend on using NDM Update to migrate this storage group. This will then highlight the need for a host reboot.

Also, for this example, choose the **Precopy** and **Move Identity** parameters.

**Create Migration Wizard | Minimally\_Disrupt\_Sg**

1 Select Target

2 Review Masking

3 Review Migration

Array ID \*  
000197600156

SRP

Compression

Precopy

Offline  Move Identity

Selecting offline will require a host reboot

On the next screen, there is the option to select a pre-existing port group from the target array, or with 9.1 and later, create a new port group as part of the create command. Click **Next**.

**Create Migration Wizard | Minimally\_Disrupt\_Sg**

1 Select Target

2 Review Masking

3 Review Migration

**Front End Suitability Score**  
Source and target systems must be collecting Performance Statistics to enable Front End Suitability.

**Target Masking Options** Information (1)

Storage Group	Minimally_Disrupt_Sg
Masking View	Minimally_Disrupt_Sg_MV
Port Group	Minimally_PG
Ports	FA-2D:8
Host	ESX_34_IG
Initiators	10000090fa9279b1 , 10000090fa927c04

A breakdown of the planned migration session is displayed showing the masking elements and options selected for final approval. Click **Run Now**.

**Create Migration Wizard | Minimally\_Disrupt\_Sg**

1 Select Target

2 Review Masking

3 Review Migration

**Target Array**

System ID 000197600156

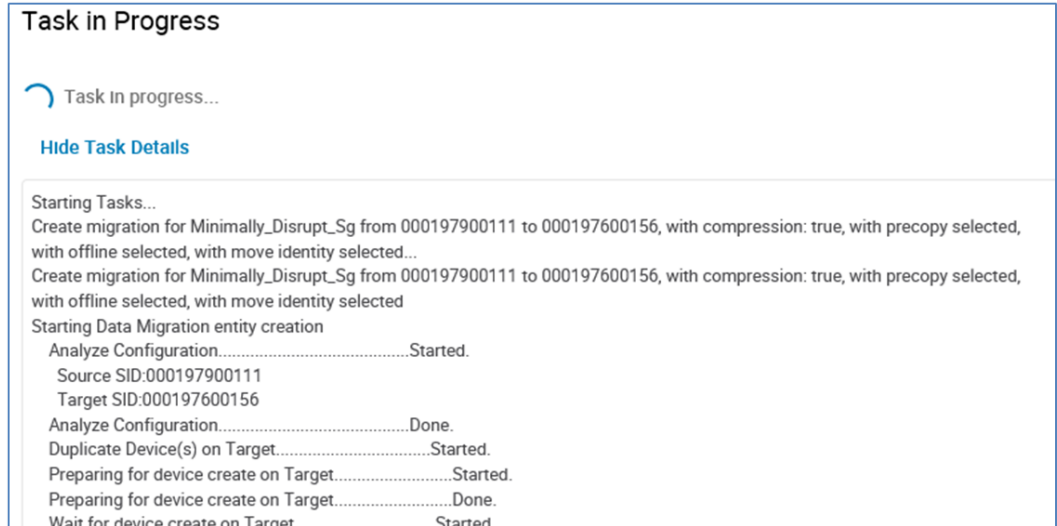
Options Compression, Precopy, Offline, Move Identity

**Migration Summary**

Storage Group	Masking View	Port Group	Host
Minimally_Disrupt_Sg	Minimally_Disrupt_Sg...	Minimally_PG	ESX_34_IG

To validate and create the session using the CLI, run the following commands:

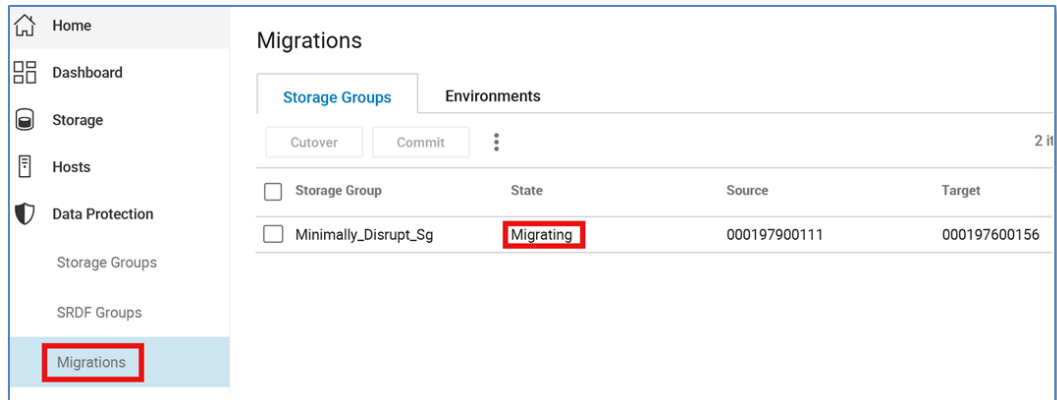
```
syndm -src_sid 111 -tgt_sid 156 -sg Minimally_Disrupt_Sg create -
offline -validate
syndm -src_sid 111 -tgt_sid 156 -sg Minimally_Disrupt_Sg create -
offline -precopy -move_identity
```



Once the create step has completed, click **OK**.

**Examine the created NDM Update session**

To examine the migration session, click **Data Protection > Migrations**. The current state of the migration setting is **Migrating**.



To view migration sessions using the CLI, run the following command:

```
syndm -sid 156 list
```

**Examine the created NDM Update session**

Refreshing the migration list view or solutions enabler command, the migration session has completed the precopy process and has now entered a CutoverReady state, meaning all data has copied and a cutting-over control to the target array is now possible.

Storage Group	State	Source	Target
<input type="checkbox"/> Minimally_Disrupt_Sg	CutoverReady	000197900111	000197600156

At this point, **shut down your application** (this example powered down the VM).

**Note:** There are checks in place to check for host I/O. However, this check is for a 15-second window and does not guarantee an application has been shut down. It is up to the user to verify manually that an application has been shut down successfully before progressing.

### Cutover the migration

This operation is only used if the -precopy option was used with the create operation.

The cutover command makes the target devices visible to the host and the source devices inactive to the host. Any updates made to data on the target array are replicated back to the source array through the RDF link.

Before issuing the command, the administrator must first shut down the application. Upon successful completion of a cutover operation, the administrator must perform a host reboot and verify that new paths or new LUNs have been discovered prior to restarting the application.

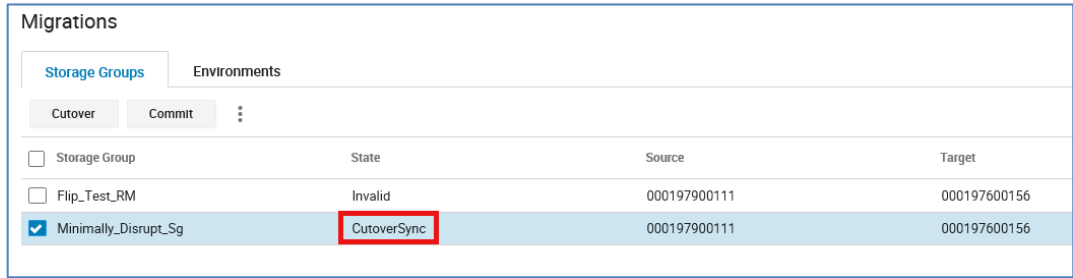
**Note:** Since the move\_identity parameter was selected for this example, a manual removal of old RDMs and the addition of new RDMs are not required. The host, upon reboot, will successfully discover the new LUNs assuming they are the original LUNs. However, before restarting an application, a visual verification of the paths should be undertaken.

Storage Group	State	Source	Target
<input type="checkbox"/> Flip_Test_RM	Invalid	000197900111	000197600156
<input checked="" type="checkbox"/> Minimally_Disrupt_Sg	CutoverReady	000197900111	000197600156

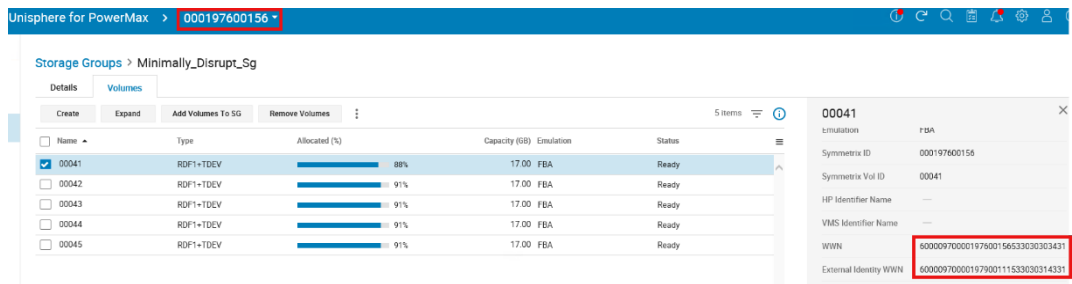
Upon a successful cutover, the migration session state will switch to CutoverSync.

To cutover the migration session using the CLI, run the following command:

```
symdm -sid 156 -sg Minimally_Disrupt_Sg cutover
```

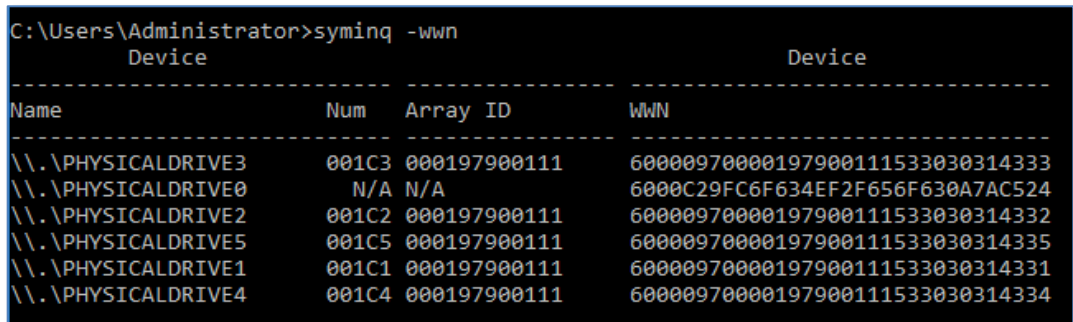


Since the move identity parameter was used, the device attributes were propagated from the source volumes to the target volumes. This can be seen below where the external identity differs from the native WWN. We have taken on the WWN of the paired source device and spoofed the target volume so it appears to the host as the same device.



At this point, **reboot the host and restart the application** (this example used a VM, requiring a rescan of the storage and restart of the VM).

Post-reboot, the device IDs remain the same to the host even though we are running on the target array 156.



### Commit the migration

Once the CutoverSync state is reached and the user confirms all data has been successfully migrated to the target array, we can complete the migration by running the commit command. This will remove the migrated application resources from the source array and releasing resources used to perform the migration.

Upon successful completion of a commit operation, the replication between source-side and target-side devices are severed. The source-side devices migrated become longer be visible to a host, and since we chose the -move\_identity option, the source-side devices



are assigned the IDs the target-side devices were created with. This ensures they are no longer used by the application that was moved to the target array during the migration.

After completing the commit operation, the application migration is complete. The application will be running only on the target array and is no longer using the devices on the source array.

<input type="checkbox"/>	Storage Group	State	Source	Target
<input checked="" type="checkbox"/>	Minimally_Disrupt_Sg	CutoverSync	000197900111	000197600156

**Data Migration Commit**

✓ Success

[Hide Task Details](#)

Starting Tasks...  
Commit migration for Minimally\_Disrupt\_Sg on array 000197900111...

Starting Data Migration commit

Analyze Configuration.....Started.  
Source SID:000197900111  
Target SID:000197600156

Analyze Configuration.....Done.

Remove Masking View(s) on Source.....Started.  
Remove Masking View(s) on Source.....Done.

Remove Data Replication.....Started.  
Remove Data Replication.....Done.

Succeeded Data Migration commit  
Refreshing Data Migration commit entities  
Succeeded

To commit the migration sessions using the CLI, run the following command:

```
syndm -sid 156 -sg Minimally_Disrupt_Sg commit
```

Once the migration session is committed, the session is removed from the migrations list view.

<input type="checkbox"/>	Storage Group	State	Source	Target
--------------------------	---------------	-------	--------	--------

### **Cancel a migration**

A cancel operation cancels a migration that has not yet been committed. If the application was reconfigured, the administrator must shut down the application before issuing the cancel command. If I/O is currently running again on the target array, the cancel operation will fail.

The administrator must perform a host rescan with the option to remove dead paths to clean up the paths that were created for the target-side LUNs. The administrator must identify the source or the target array as well as the application whose migration is to be canceled.

Upon successful completion of a cancel operation, the application storage will be as it was prior to the start of the migration:

- Connections between source-side and target-side devices that were used to migrate data through the DM replication pathway are severed.
- Any target-side resources configured for the application that are not used by other applications on the target array are removed. This includes pre-existing IGs or PGs used by the migrated view on the target array if these are not used in other masking views or groups.

### **Recovering a failed migration**

A recover operation is needed after a migration step completes with a failed state, and it is not normally required as a part of a migration. Failed states, depending on the reason for the failure can include the following:

- CreateFailed
- CutoverFailed
- MigrateFailed
- CancelFailed
- CommitFailed

After the condition that caused a migration operation (create, cutover, commit, or cancel) to fail has been corrected, a recover operation can be invoked to continue with the migration by completing the following:

- Determining which migration operation failed
- Putting the migration session resources (connections, devices) into the appropriate state to allow the failed operation to complete
- Repeating or resuming (depending on the cause of the failure) the failed action

To recover a failed migration using the CLI (once reason for failure has been resolved), run the following command:

```
symdm -sid 156 -sg Minimally_Disrupt_sg recover
```

## Online Minimally Disruptive Migration (O-MDM)

This section includes an overview, updated interfaces, and procedures for the Online Minimally Disruptive Migration (O-MDM) feature.

### Overview

O-MDM (Open Minimally Disruptive Migration) utilizes Open Replicator (ORS) as the data transfer mechanism instead of native SRDF/Metro used by NDM. ORS is a SCSI block migration feature that transfers data from an external array through a WWN designated FBA source device.

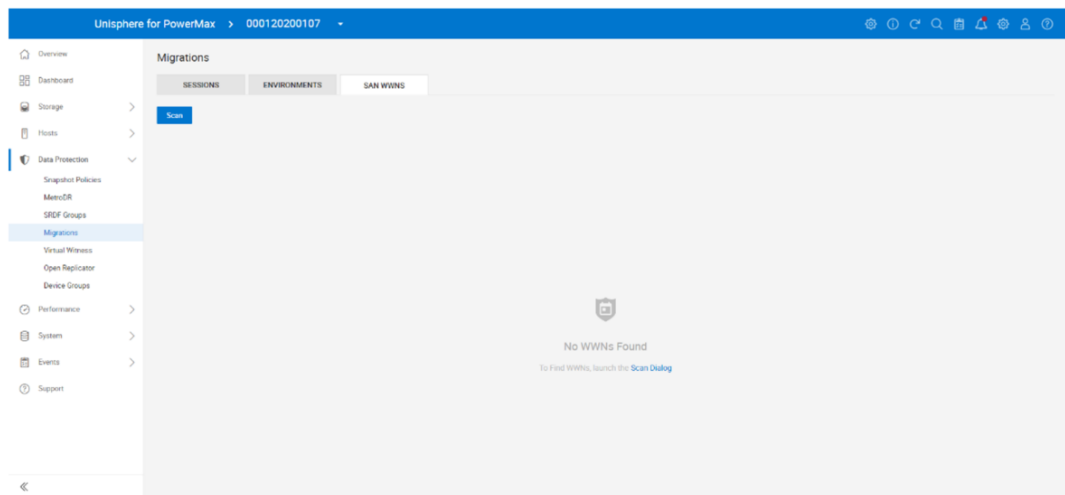
As such, it is extremely flexible in the source arrays supported:

- Currently targeting support for all arrays supported by ORS, including all shipping VMAX/PowerMax, Dell and competitive arrays.
- OS and Multi-pathing software (MPIO, PowerPath, and so on) currently supported by ORS will be supported.

ORS integration supports the familiar create/cutover/commit as a typical NDM migration. Source array devices are presented to PowerMax front end FA ports through the donor array native provisioning mechanism. The symsan SYMCLI command can be used to verify PowerMax visibility on these devices.

New symdm command syntax, `create -san -offline`, is used to initiate a new O-MDM migration session. This new command syntax requires a new option (`-file`) specifying the input of a file containing a list of donor device WWNs familiar to ORS users. Target devices, associated SGs, and front-end host mapping are then created. Target devices created on the PowerMax are rounded to the nearest cylinder size.

Similarly, Unisphere 10.0 includes a SAN WWNS tab from the Data Protection -> Migration option for initiating and managing O-MDM migrations:



Prior to the cutover command, the source application must be shut down (minimally disruptive). After a cutover command, target devices from PowerMax are made visible to the host. Applications can then be restarted now that they are sourced from the PowerMax. ORS Donor Update is then used to keep the source devices/array updated.

Commit can then be initiated to stop the migration session and associated ORS Donor Update session to source devices.

The SYMCLI `syndm` command syntax has changed, with the following new options (highlighted in bold text) to support the O-MDM functionality:

```
syndm -tgt_sid <SymmID>

[-i <Interval>] [-c <Count>] [-noprompt]
[-tgt_srp <SRPName>] [-tgt_pg <PgName>]
[-nocompression] [-validate]

create -src_sid <SymmID> -sg <SgName> [-precopy]

create -src_sid <SymmID> -sg <SgName>
-offline [-move_identity] [-precopy]

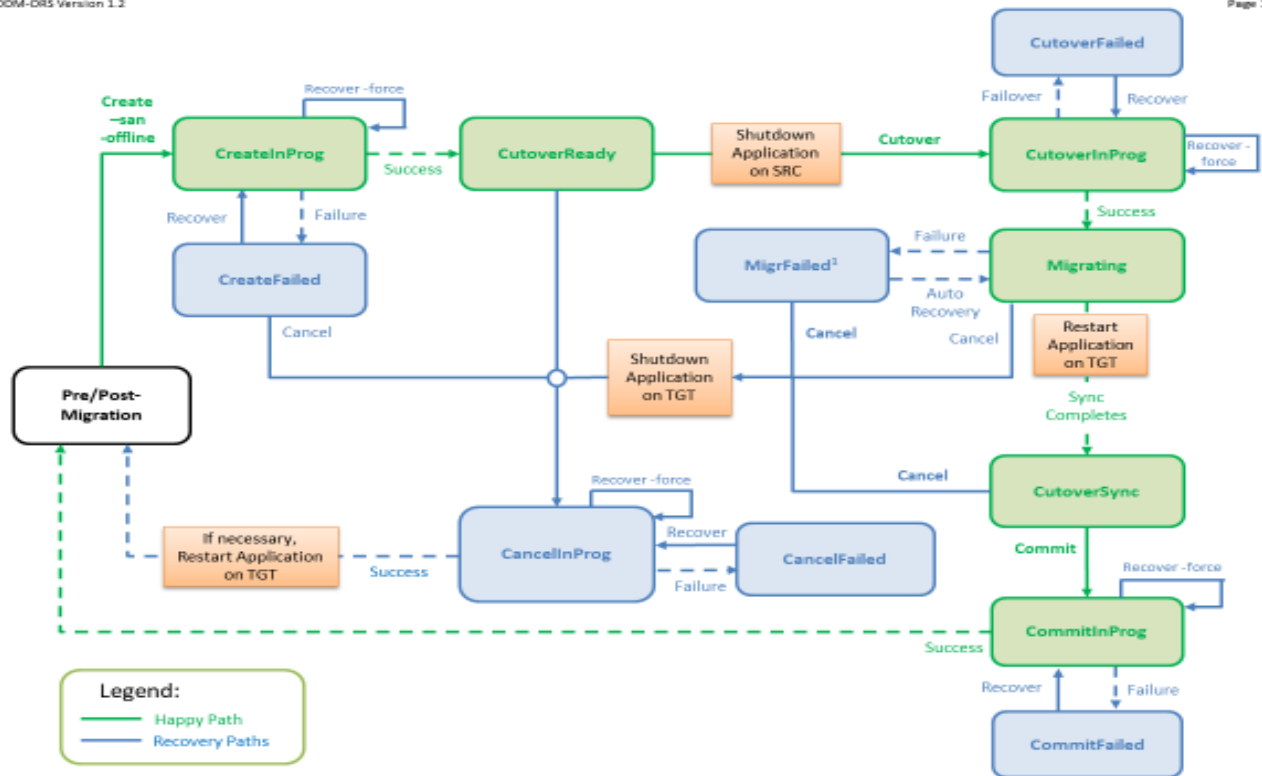
create -file <FileName> -tgt_sg <SgName>
-tgt_ig <IgName> -san -offline
```

Options:

- **-san** Specifies the migration to be an Open Data Migration (O-MDM).
- **-file** Specifies the file name that contains a list of source side device WWNs to be migrated.
- **-tgt\_sg** Specifies the target storage group name that will be created for the target devices. This name will also be used for other target side masking objects (port group, if one does not exist and masking view) created during a migration.
- **-tgt\_ig** Specifies the target side initiator group that contains host application initiators mapped to both the PowerMax and source side arrays.

## Migration flow and states

It is important to understand the migration workflows and associated states supporting the O-MDM functionality. Review the following chart and state description table with emphasis on the new create, cutover, application shutdown, cancel, and recovery states:



Note 1: A application outage will likely occur if a failure happens during the application data migration phase.

Migration state	Description
No migration	There is no migration in progress.
<b>CreateInProg</b>	The migration session is being created. The create command has not run to completion with either a success or a failed status.
CreateFailed	The create command has run to completion with a failed status.
<b>CutoverReady</b>	The create command has run to completion with a success status. Target devices are not masked to the host.
<b>CutoverInProg</b>	The cutover command is in progress and the application is being moved to the target array. The cutover command has not run to completion with either a success or a failed status.
CutoverFailed	The cutover command has run to completion with a failed status.
<b>Migrating</b>	The Cutover command succeeded. I/O's can be serviced by devices on the target array.
MigrateFailed	This state occurs when a Migrating state is interrupted, as might occur with a loss of the required DM connectivity between the source and target arrays.
<b>CutoverSync</b>	The Cutover command completed successfully. The application is running on the target array, and all data has been synchronized between the target and source arrays.
<b>CommitInProg</b>	The commit command is in progress. The commit command has not run to completion with either a success or a failed status.
CommitFailed	The commit command has run to completion with a failed status.
<b>CancelInProg</b>	The cancel command is in progress. The cancel command has not run to completion with either a success or a failed status.
CancelFailed	The cancel command has run to completion with a failed status.

### Create phase changes

This section describes the changes and new requirements in the create phase of an O-MDM migration.

The administrator must now identify the source WWNs, target array ID, and existing target array initiator group. The source WWNs are a list of hexadecimal WWNs. These WWNs can be verified using the 'symsan' CLI command.

Example contents of an 8 WWN source file: omdm\_8wwn.txt:

```
60000970000197802041533030433742
60000970000197802041533030433743
60000970000197802041533030433744
60000970000197802041533030453941
60000970000197802041533030453942
60000970000197802041533030453943
60000970000197802041533030453944
60000970000197802041533030453945
```

This list may be created and verified from the following `symsan` command as an example:

```
symsan list -sid 721 -sanluns -wnw 50000973A81FE449 -dir all -p all
```

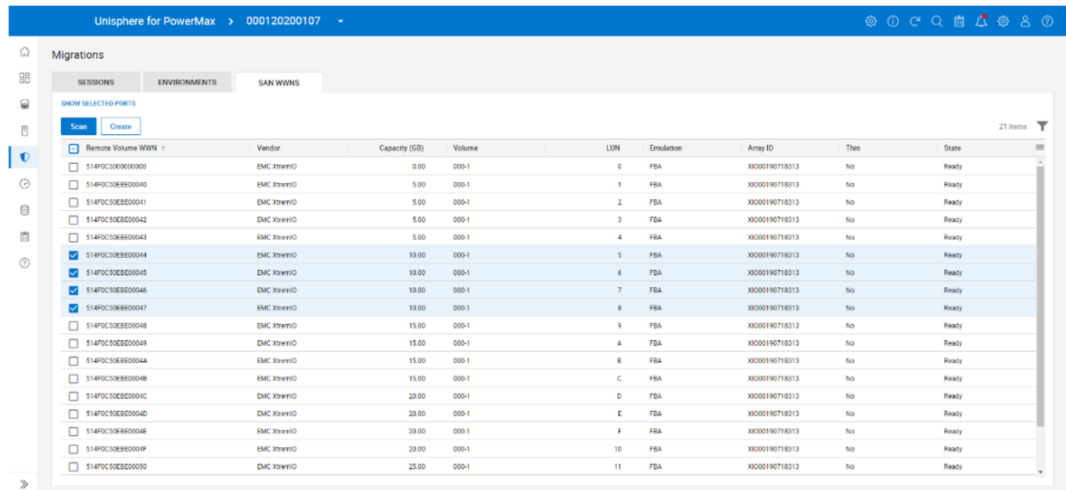
```
Symmetrix ID:      000197801721
Remote Port WWN:   50000973A81FE449
```

```

ST
A
T  Flags  Block Capacity  LUN  Dev  LUN
DIR:P  E ICR THS  Size      (MB)  Num  Num  W  WWN
-----
02D:11 RW ... FXX  512    187500    38 00C7B X 60000970000197802044533030433742
02D:11 RW ... FXX  512    187500    39 00C7C X 60000970000197802044533030433743
02D:11 RW ... FXX  512    187500    3A 00C7D X 60000970000197802044533030433744
02D:11 RW ... FXX  512    187500    3B 00E9A X 60000970000197802044533030453941
02D:11 RW ... FXX  512    187500    3C 00E9B X 60000970000197802044533030453942
02D:11 RW ... FXX  512    187500    3D 00E9C X 60000970000197802044533030453943
02D:11 RW ... FXX  512    187500    3E 00E9D X 60000970000197802044533030453944
02D:11 RW ... FXX  512    187500    3F 00E9E X 60000970000197802044533030453945

```

Similarly, the Unisphere 10.0 interface allow you to select donor devices from the WWNs that are available to the array:



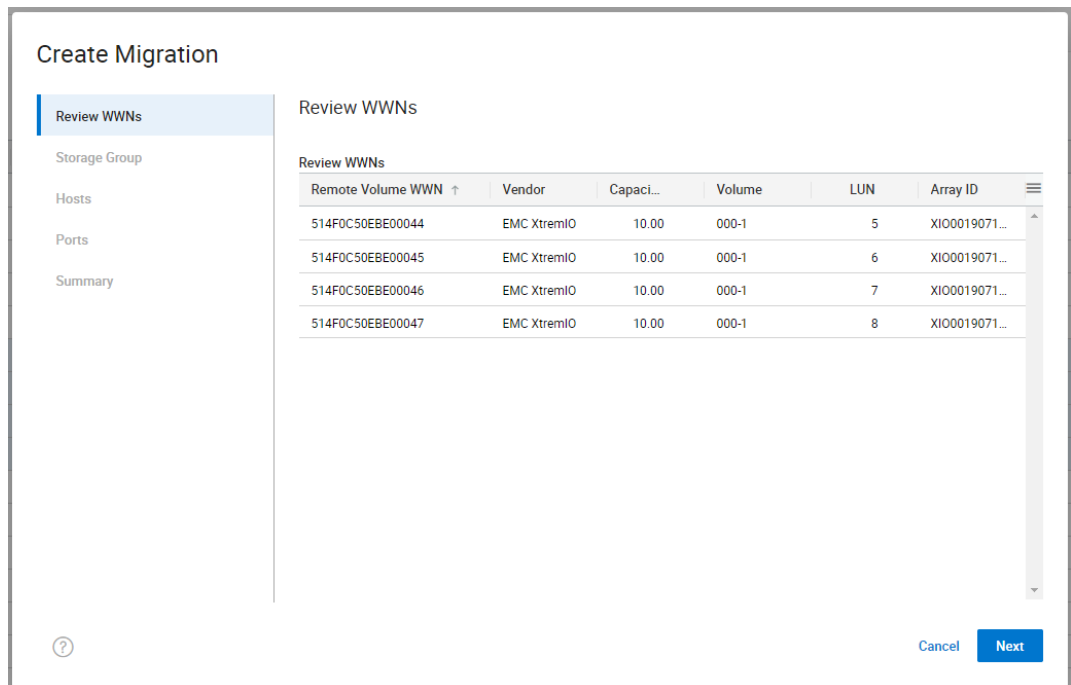
The 'create' command first verifies that the requirements and restrictions are met to ensure that the migration can proceed and verifies that the target array has sufficient resources available to duplicate the application's storage. The application's storage

configuration is then copied to the target array, creating equivalently provisioned devices. ORS activate (data copy) does not start until the following cutover action occurs.

When creating the storage configuration on the target array:

1. A device created on the target array will be rounded up to the nearest cylinder size of the capacity of the source device.
2. An SG with the name as given in the command is created and all the devices created will be added to it.
3. If an SRP is specified, that SRP will be set on the SGs created on the target array. If not specified, then no SRP will be configured on the SG created on the target array. This will result in the default SRP being used.
4. When creating a new PG on the target array, ports are selected from the Login History Table (LHT) of ACLX enabled FC ports where the host initiators were seen, ignoring the current logged in status.
5. If a PG name is specified, an existing PG on the target array will be used and the host initiators must appear in the Login History Table (LHT) of at least one port, ignoring the current logged in status.
6. The name of any port group or masking view created will be derived from the storage group name provided in the command line for the migration.

Example of a Unisphere 10.0 create dialog after the donor device WWNs have been chosen:



After the application's storage configuration has been duplicated on the target array, the create operation concludes by connecting source-side devices with the corresponding target-side devices through a DM replication pathway and starts the data migration. The migration will immediately enter the CutoverReady state. The administrator must shut down the application before a Cutover operation can be issued. A 'cancel' command can



be used to cancel the migration and undo the target-side storage provisioning done by the create operation.

If the create action fails:

- When the cause of the failure has been corrected, a recover operation can be attempted ('symdm recover') to complete creation of the migration session.
- (or)
- The migration can be canceled ('symdm cancel'), leaving the application's storage environment as it was before the start of the migration, with resources configured for the application on the target removed.

The following command creates an Open Data Migration (O-MDM) session. The following is displayed when the 'symdm create -san' command is issued:

```
symdm create -san -offline -file src_wnw_list.fil -tgt_sid 644 -tgt_sg APP1
-tgt_ig APP1_INIT_GRP
A DM 'Offline SAN Create' operation is
in progress for storage group 'APP1'. Please wait...

Analyze
Configuration.....Started.
  Source Array:APM00150519204
  Target Array:000197900644
Analyze Configuration.....Done.
Creating
Device(s).....Started.
  Creating Device(s) on Target.....Done.
  Update target SG Device(s).....Started.
  Update target SG Device(s).....Done.
  Create Storage Group(s) on Target.....Started.
  Create Storage Group(s) on Target.....Done.
  Create Port Group(s) on Target.....Started.
  Create Port Group(s) on Target.....Done.
Start Data
Replication.....Started.
  Start Data Replication.....Done.

The DM 'Offline Create' operation successfully executed for
storage group 'APP1'.
```

### Cutover and commit phase changes

This section describes the changes and new requirements in the cutover phase of an O-MDM migration.

For minimally disruptive migrations, the administrator must first shut down the application before issuing the 'cutover' command. The operation will make the target devices visible to the host. Any updates made to data on the target array will be replicated back to the source array through the DM replication pathway. Upon successful completion of a cutover operation, the administrator must perform a host rescan (or host reboot) and verify that new LUNs have been discovered prior to restarting the application. The application configuration would also need to be changed to use the new LUNs.

## Online Minimally Disruptive Migration (O-MDM)

This command is used when the migration session is in the CutoverReady state. It migrates the application's processing to run only against the target array. The following is displayed when the 'symdm cutover' command is issued:

```
symdm cutover -sid 643 -sg APP1
```

```
A DM 'Cutover' operation is  
in progress for storage group 'APP1'. Please wait...
```

```
Analyze  
Configuration.....Started.  
  Source Array:APM00150519204  
  Target Array:000197900644  
Analyze  
Configuration.....Done.  
  Preparing Devices for Host discovery.....Started.  
  
Cutover.....  
Started.  
  
Cutover.....  
Done.  
  Preparing Devices for Host discovery.....Done.  
  Create Masking View(s) on Target.....Started.  
  Create Masking View(s) on Target.....Done.
```

```
The DM 'Cutover' operation successfully executed for  
storage group 'APP1'.
```

This command completes the migration by removing application resources from the source array and releasing resources used for the migration. The following is displayed when the 'symdm commit' command is issued:

```
symdm commit -sid 643 -sg DM_APP1
```

```
Analyze Configuration.....Started.  
  Source Array:APM00150519204  
  Target Array:000197900644  
Analyze Configuration.....Done.  
Remove Data Replication.....Started.  
Remove Data Replication.....Done.
```

```
The DM 'Commit' operation successfully executed for  
storage group 'APP1'.
```

## Canceling a migration

The 'cancel' command cancels a migration that has not yet been committed. The administrator must identify the target array and the application whose migration is to be canceled. For minimally disruptive migrations, the administrator must shut down the application before issuing the cancel command. Upon successful completion of a cancel operation, the application configuration would need to be changed to use the original source LUNs before it can be restarted.

Upon successful completion of a cancel operation, the application storage will be as it was prior to the start of the migration:

- Connections between source-side and target-side devices that were used to migrate data through the DM replication pathway will be severed, and
- Any target-side resources configured for the application that are not used by other applications on the target array will be removed.
- This will include pre-existing PGs used by the migrated view on the target array, if these are not used in other masking views or groups.

The administrator must perform a host rescan with the option to remove dead paths to clean up the paths that were created for the target side LUNs.

If the administrator has configured local or remote replication of the target-side devices, then a cancel will be blocked. Administrator will need to remove the replication sessions before running the cancel operation.

The following command can be used to cancel a migration. After a successful cancel, the application will be running against its storage on the source array, and any resources configured for the application on the target array by the migration create processing will have been removed.

```
symdm cancel -sid 643 -sg APP1
  Analyze
Configuration.....Started.
  Source Array:APM00150519204
  Target Array:000197900644
  Checking Target devices for IO.....Started.
  Checking Target devices for IO.....Done.
  Analyze Configuration.....Done.
  Remove Data Replication.....Started.
  Remove Masking View(s) on Target.....Started.
  Remove Masking View(s) on Target.....Done.
  Remove Data Replication.....Done.
  Remove Port Group(s) on Target.....Started.
  Remove Port Group(s) on Target.....Done.
  Remove Storage Group(s) on Target.....Started.
  Remove Storage Group(s) on Target.....Done.
  Remove created Device(s) on Target.....Started.
Remove created Device(s) on Target.....Done.

The DM 'Cancel' operation successfully executed for storage group 'APP1'.
```

## Recovery from Failed State

A 'recover' operation is needed after a step in the migration completes with a "failed" state and is not normally required as a part of a migration.

Failed states include: CreateFailed, CutoverFailed, CancelFailed, and CommitFailed. The 'recover' command is not used for the MigrFailed state, because that state is auto recovered.

After the condition that caused a migration operation to fail has been corrected (create, cutover, commit, or cancel), a recover operation can be invoked to continue with the migration by:

- Determining which migration operation failed.
- Putting the migration session's resources (connections, devices, and so on) into the appropriate state to allow the failed operation to complete.
- Repeating or resuming (depending on the cause of the failure) the failed action.

The administrator must identify the target array and the data migration session name to be recovered.

After an error, the 'symdm recover' can be used after correcting the cause of a failed symdm action (create, commit, or cancel) to put the migration into the appropriate state, and then repeat or resume the failed action.

```
symdm recover -sid 643 -sg APP1
```

```
A DM 'Recover' operation is  
in progress for storage group 'APP1'. Please wait...
```

```
Analyze  
Configuration.....Started.  
  Source Array:APM00150519204  
  Target Array:000197900644  
Analyze Configuration.....Done.  
Creating Device(s).....Not  
Needed.  
Update target SG Device(s).....Not  
Needed.  
Create Storage Group(s) on Target.....Not Needed.  
Create Port Group(s) on Target.....Not Needed.  
Start Data  
Replication.....Started.  
Start Data  
Replication.....Done.
```

```
The DM 'Recover' operation successfully executed for  
storage group 'APP1'.
```

The following is a recovery from a failed Commit:

```
symdm recover -sid 643 -sg APP1
```

```
A DM 'Recover' operation is  
in progress for storage group 'APP1'. Please wait...
```

```
Analyze Configuration.....Started.  
  Source Array:APM00150519204  
  Target Array:000197900644  
Analyze Configuration.....Done.  
Remove Data Replication.....Started.  
Remove Data Replication.....Done.
```

The DM 'Recover' operation successfully executed for storage group 'APP1'.

The following is a recovery from a failed Cancel:

**symdm recover -sid 643 -sg APP1**

A DM 'Recover' operation is in progress for storage group 'APP1'. Please wait...

```
Analyze
Configuration.....Started.
  Source Array:APM00150519204
  Target Array:000197900644
  Checking Target devices for IO.....Not
  Needed.
  Analyze
  Configuration.....Done.
  Remove Masking View(s) on Target.....Not Needed.
  Remove Data Replication.....Started.
  Remove Data Replication.....Done.
  Remove Port Group(s) on Target.....Started.
  Remove Port Group(s) on Target.....Done.
  Remove Storage Group(s) on Target.....Started.
  Remove Storage Group(s) on Target.....Done.
  Rollback update of target SG Device(s).....Started.
  Rollback update of target SG Device(s).....Done.
  Remove created Device(s) on Target.....Started.
  Remove created Device(s) on Target.....Done.
```

The DM 'Recover' operation successfully executed for storage group 'APP1'.

**Additional query and list output changes**

This section describes the additional query and list outputs within an O-MDM migration.

The following list output displays a new flag column showing if the session is a SAN (O-MDM) migration session:

**symdm -sid 644 list**

Symmetrix ID : 000197900644

Storage Group	Flg	Source OS Array	Target Array	State	Total	
					Capacity (GB)	Done (%)
APP1	xx	APM00150519204	000197900644	CutoverSync	100.7	100
APP2	x.	000197100643	000197900644	Migrating	8.0	90

Legend:

Flags:

(O)ffline                    X = **Offline**, . = N/A  
 (S)AN                        X = **SAN**, . = N/A

## Online Minimally Disruptive Migration (O-MDM)

If the `-sg` option is given, the new SAN flag will be shown as a new field below the SG name:

```
symdm list -sid 643 -sg APP1
```

```
Symmetrix ID : 000197100643  
Storage Group: APP1  
Offline      : Yes  
SAN          : Yes
```

<b>Source</b>	Target		Total	Done
<b>Array</b>	Array	State	Capacity	(%)
-----				
<b>APM00150519204</b>	000197900644	CutoverReady	100.7	N/A

The following is an example of the verbose list output:

```
symdm list -sid 789 -sg APP3 -v -detail -pairs_info
```

```
Symmetrix ID      : 000197801789  
  
Storage Group    : APP3  
Offline         : Yes  
SAN              : Yes
```

```
Source Array     : 000197801865  
Target Array    : 000197801789
```

```
Migration State : CutoverReady  
Total Capacity (GB) : 183.1  
Done (%)          : N/A
```

```
Device Pairs (1): Failed
```

```
{  
  Source          Source      Source      Source Target      Error      Last Error Code  
  WWN            Vendor      Array      Dev   Dev   Status      Code      Timestamp  
-----  
  60000970000197900709533030384634 EMC Symmetrix 000197900709 008F4 01753 OK        NO_PATH  Fri Oct 15 10:41:05 2021  
}
```

If the Migration State is `MigrateFailed`, the Source WWN Status (Error Code) field will be one of the following different values:

- NDEF (not defined, the initial status of an O-MDM device when the session is created)
- NO\_PATH (all of the paths to the target are down)
- LUN\_NR (LUN not ready, indicates that the target/source device LUN is not ready)
- WR\_PROT (write protect set on the LUNs)

In addition, the Last Error Timestamp will note (format example, Thu Sep 19 07:03:12 2019) when the error was encountered for that specific source WWN – target device pair.

## Limitations and restrictions

The following are known limitations or restrictions regarding an O-MDM migration:

- Target array must be PowerMax 2500/8500 (V4) Hardware.
- Source or donor array must be VMAX3, VMAX All Flash, PowerMax, other Dell storage, or supported third party array.
- Data migrations are often complex operations and require careful planning and execution of predetermined procedures. Failure to identify and perform necessary steps or work within supported configurations can result in data unavailability or loss.
- Source side WWNs must be 512 byte block size
- No more than 1,024 LUNs can be migrated in a single O-MDM session
- All initiators provisioned to an application on the source array must appear in the Login History Table (LHT) of the target array provided. The name of the port group must exist on the target array and have the initiators appear on the LHT for at least one port in the port group. These ports must be ACLX enabled FC ports.
- FCoE and iSCSI ports are not supported.
- The names of the masking groups that are being created as part of the migration must not exist on the target array.
- The SRP that will be used for target-side storage, whether specified or defaulted, must have enough free capacity to support the migration.
- The target array must be capable of supporting the additional devices that the migration will create to receive the source-side data.
- The target masking configuration of the migrated application must not be added.
- Changing the masking configuration after the start of the migration can cause the O-MDM controls to be disallowed.
- The only exceptions are:
  - Changing the SLs or SRPs on the SG
  - Changing the compression attribute on SG
  - Adding Host IO limits on SG
  - Adding more ports to the PG, only FC ports (not NVMe over FC) are supported
- After the migration has started copying the data to the target array, devices can have R1 mirrors added so they can be used for remote replication. Migration controls will evaluate the states of these devices when determining if the control can proceed.
- The R1 mirrors added to the target devices can be in Asynchronous, Adaptive Copy, or Synchronous Mode but:
  - Cannot be enabled for MSC
  - Cannot be part of a Star or SQAR configuration
  - Cannot be enabled for Synchronous SRDF Consistency
  - Cannot be part of an RDF/Metro configuration



- Devices in other SRDF groups paired with target devices cannot be the target of a data synchronization operation during the lifecycle of the migration session (that is, cannot copy data to devices being migrated).
- When the migration session reaches the complete state, LREP snapshot session will be allowed to be created on O-MDM target device.
- If a device has any association with the migration session, LREP snapshot session will be blocked from being created on O-MDM target device.

## Masking enhancements (SE 9.0)

### Introduction

Solutions Enabler 9.0 introduced enhancements to manipulate masking views while maintaining host access to the array. Although the enhancements came about because of NDM, they are available for regular masking operations.

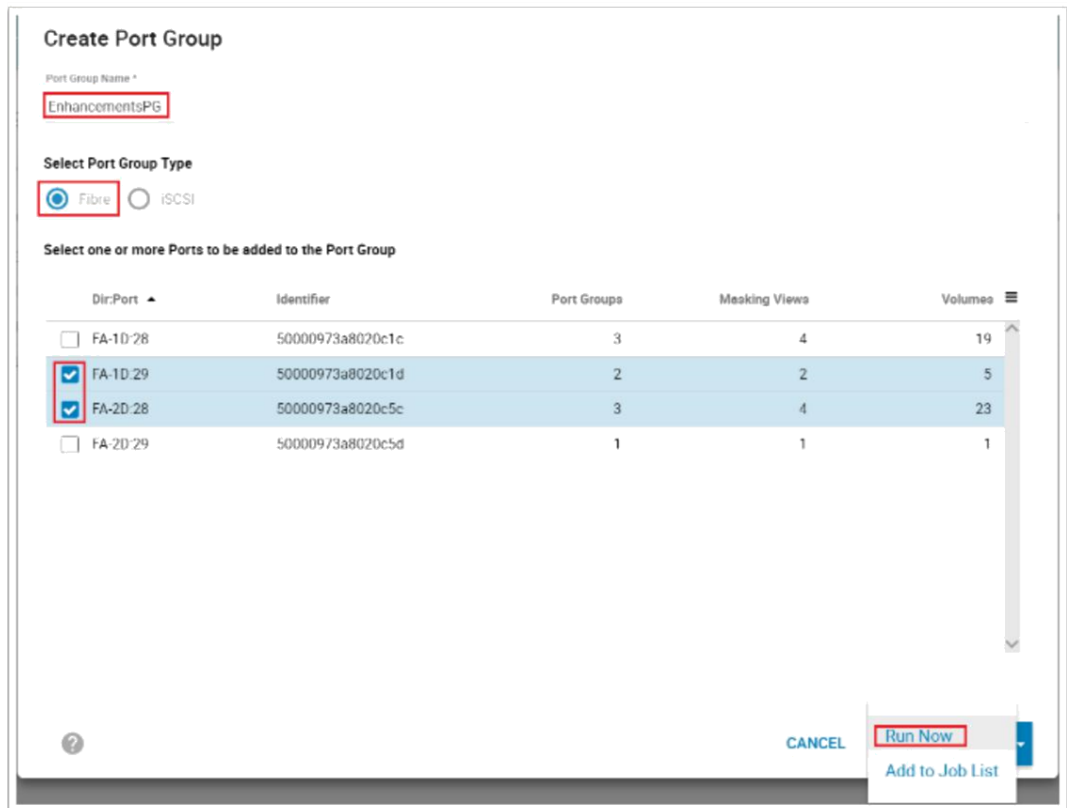
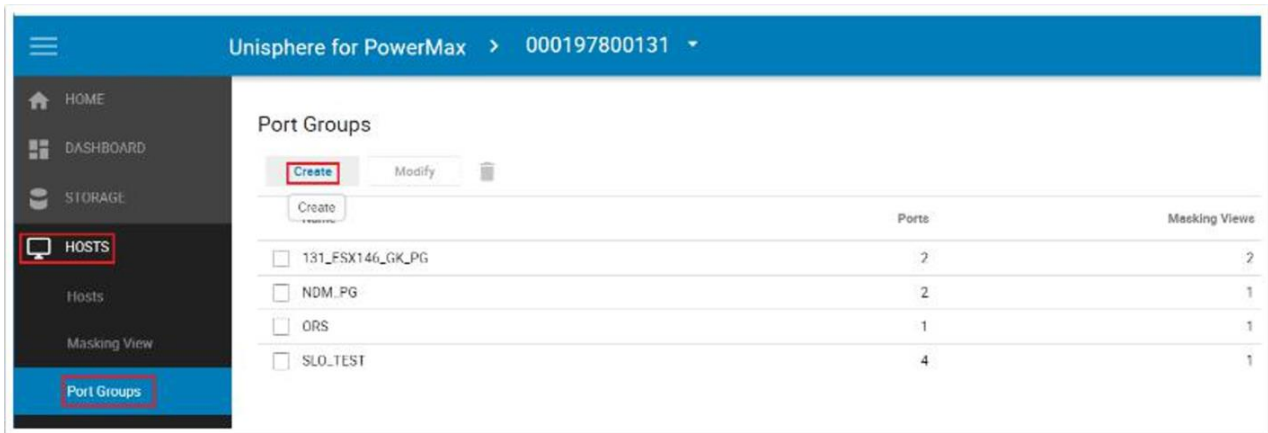
These user-experience enhancements are listed as follows:

- Added the ability to select a port group on the target array for the NDM session (see [Selecting a port group during an NDM create](#))
- Added the ability to migrate a child-only storage group (see [Migrate a child-only SG](#))
- Enabled migrating a subset of devices in a storage group (see [Migrate a subset of devices in an SG](#)).
- Allowed consolidating a single application speed across two source arrays into a single target array (see [Consolidate a single application on two arrays into a single target](#)).
- When creating an NDM session, devices are automatically set to RDF capable (see [When creating an NDM session set devices RDF capable](#)).
- Added the ability to set up DR from the target array prior to the devices synchronizing fully (see [Add DR to target SG before synchronization](#)).
- Enabled creating boot LUNs on the target array the same as the source array's LUN address (see [Create bot LUNs on the target array the same as the source array](#)).

### Selecting a port group during an NDM create

#### Creating a port group

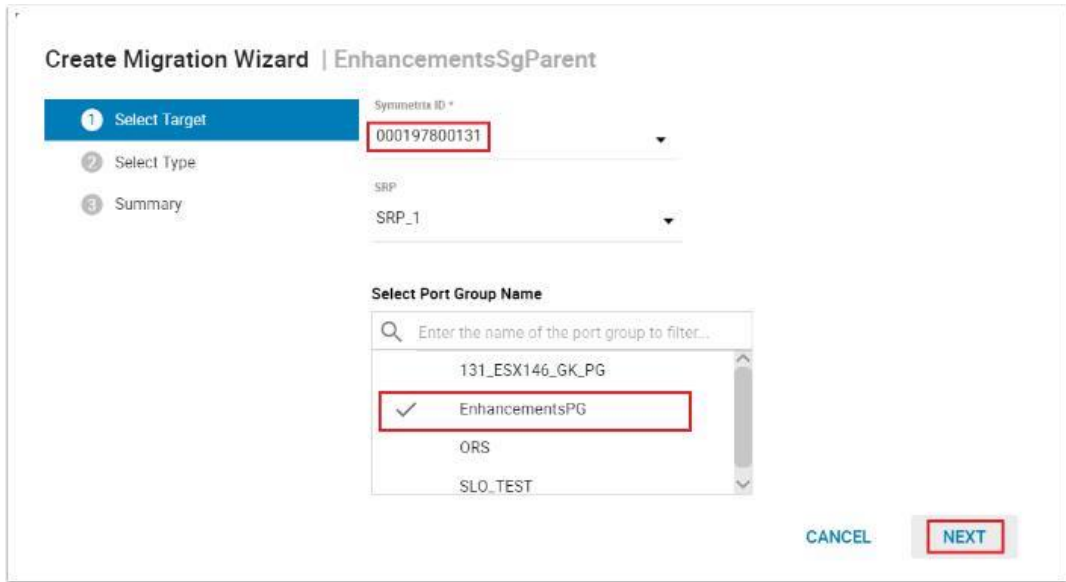
To create a port group using Unisphere, select **Hosts > Port Groups** and click **Create**. Give the group a name and select **F** or the ports for the group. Click **Run Now**.



To create a port group with Solutions Enabler, use the symaccess library, specifying the director numbers and specific ports. For example:

```
C:\Program Files\EMC\SYMCLI\bin>symaccess -sid 131 create -name EnhancementsPG -type port -dirport 1d:29,2d:28
C:\Program Files\EMC\SYMCLI\bin>
```

When using the Create Migration wizard in Unisphere, select the port group when defining the target. For example:



In Solutions Enabler, use the `-tgt_pg` option to specify the port group on the target array. For example:

```
C:\Program Files\EMC\SYMCLI\bin>symdm create -src_sid 558 -tgt_sid 131 -sg EnhancementsSgParent -tgt_pg EnhancementsPG
Execute 'Create' operation on SG 'EnhancementsSgParent' (y/[n])? y
A DM 'Create' operation is
in progress for storage group 'EnhancementsSgParent'. Please wait...

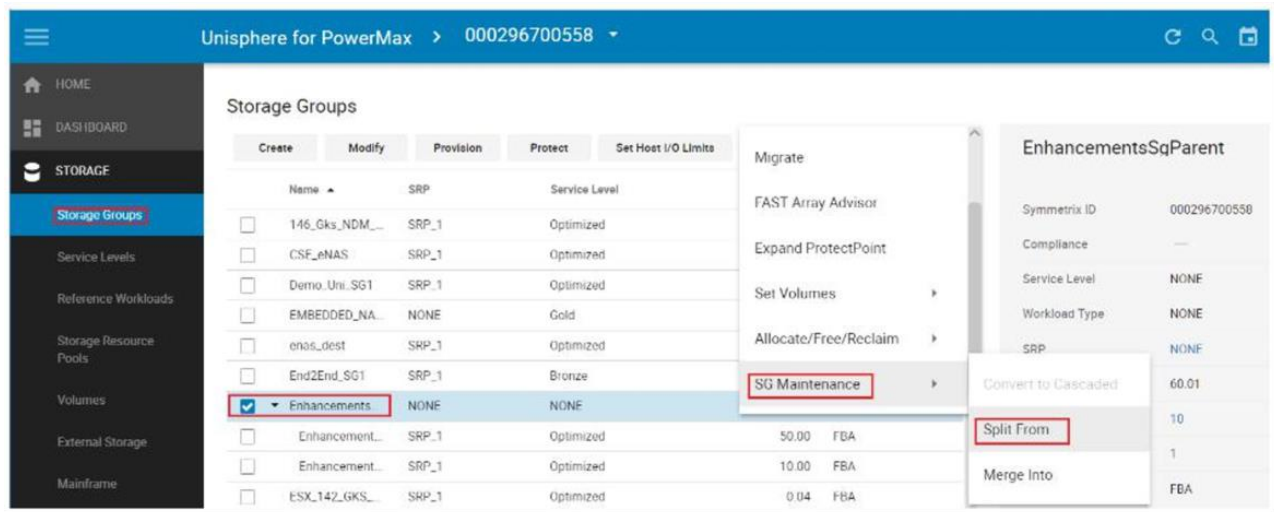
Analyze Configuration.....Started.
  Source SID:000296700558
  Target SID:000197800131
Analyze Configuration.....Done.
Initialize Replication Environment.....Started.
Initialize Replication Environment.....Done.
Create Storage Group(s) on Target.....Started.
Create Storage Group(s) on Target.....Done.
Duplicate Device(s) on Target.....Started.
Preparing for device create on Target.....Started.
Preparing for device create on Target.....Done.
Duplicate Device(s) on Target.....Done.
Create Initiator Group(s) on Target.....Started.
Create Initiator Group(s) on Target.....Done.
Create Port Group(s) on Target.....Started.
Create Port Group(s) on Target.....Done.
Start Data Replication.....Started.
Start Data Replication.....Done.
Create Masking View(s) on Target.....Started.
Create Masking View(s) on Target.....Done.

The DM 'Create' operation successfully executed for
storage group 'EnhancementsSgParent'.
```

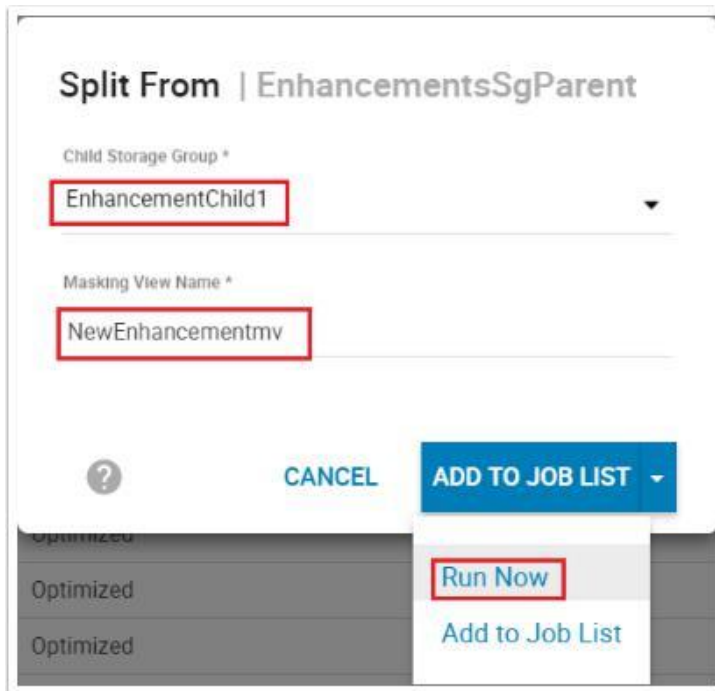
### Migrate a child-only SG

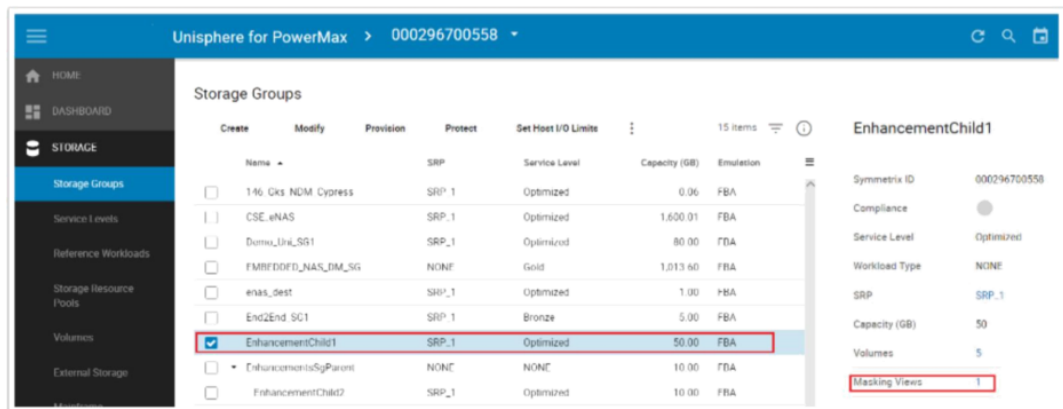
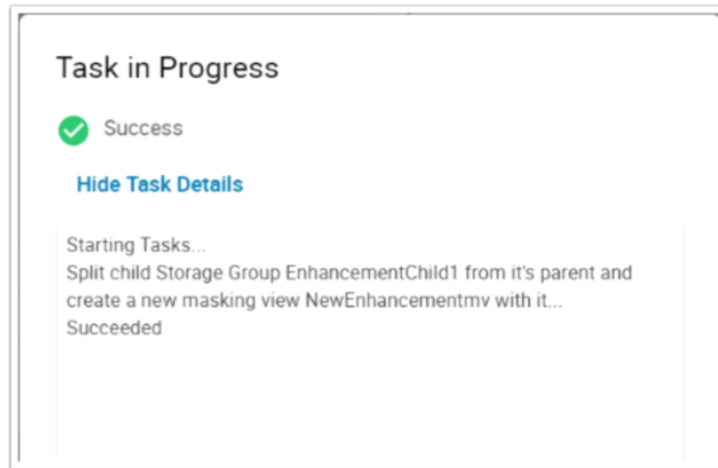
To migrate a child SG, we must first non-disruptively remove the child SG from the parent SG.

In Unisphere, select the parent SG from the **Storage Groups** dashboard. Click the **More Actions** menu and select **SG Maintenance > Split From**.



If there are multiple child SGs, select which one you would like to migrate and give a name to the masking view that will be created. Click **Run Now**.





Once the process is complete, the child storage group is promoted to tier one on the SG nesting with its own masking view to the host. For guidance on how to migrate the SG see the section on Metro-based NDM ([Using Unisphere for PowerMax](#) or [Using Metro NDM using Solutions Enabler 9.x](#)) or Metro-based NDM with precopy ([Using Unisphere for PowerMax with precopy](#) or [Using Solutions Enabler 9.x with precopy](#)).

With Solutions Enabler, use the **symsg** library:

```
symsg -sg <SgName> -sid <SymmID>
merge <SgName1>
split <SgName1> -view_name <MvName>
```

For the split operation, the -sg <SgName> option to specify the source SG, for example, EnhancementsSgParent for the split operation.

For a cascaded (parent child) source SG, the <SgName1> (EnhancementChild1 for example) option is used to specify the child SG to split from the parent. Use the view\_name option to specify the name of the new masking view to be created. The source SG must be in a single masking view.

```
C:\Program Files\EMC\SYMCLI\bin>symsg -sg EnhancementsSgParent -sid 558 split EnhancementChild1 -view_name EnhancementChild1_mv
C:\Program Files\EMC\SYMCLI\bin>symsg -sid 558 list

      S T O R A G E   G R O U P S

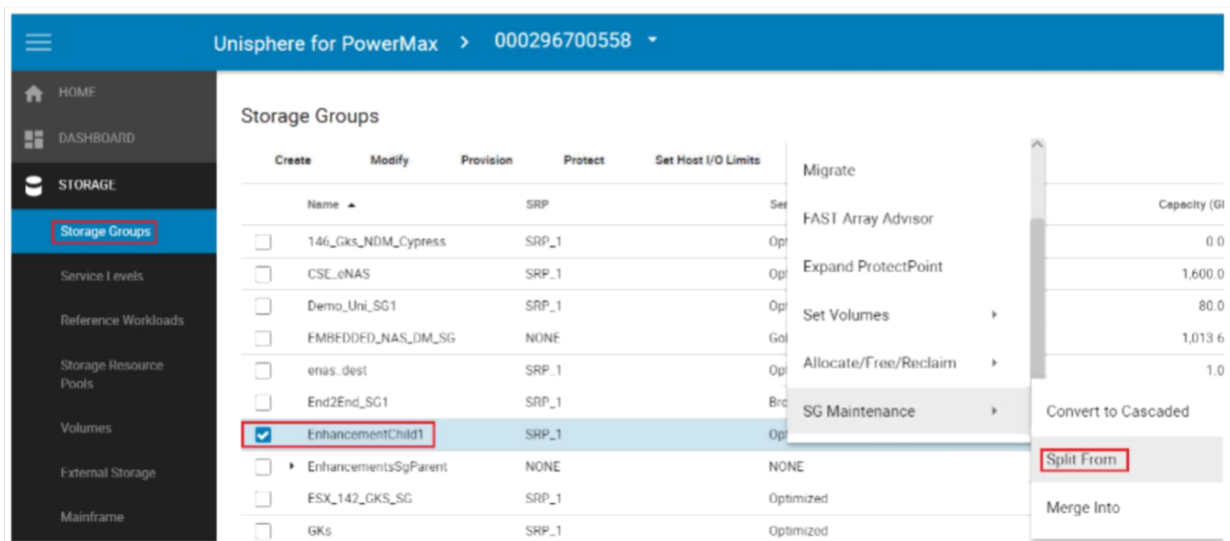
Symmetrix ID:      000296700558

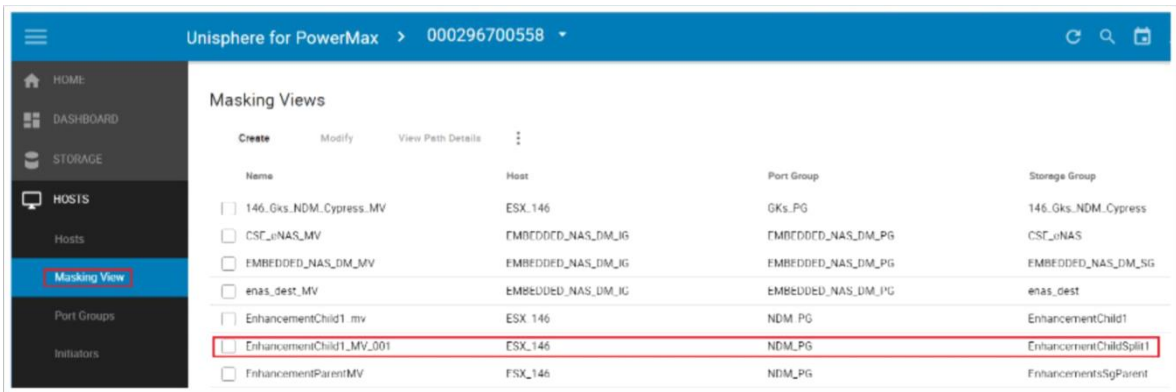
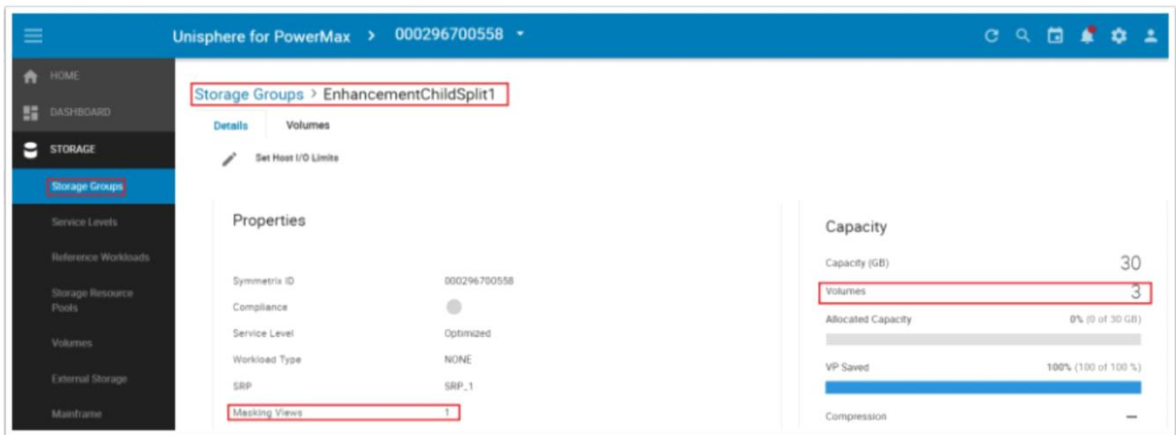
Storage Group Name      Flags      Number      Number      Child
                        EFM SLC    Devices    GKs         SGs
-----
146_Gks_NDM_Cypress     FXX ...   9           9           0
CSE_eNAS                 FXX ...   8           0           0
Demo_Uni_SG1             FX. ...   4           0           0
EMBEDDED_NAS_DM_SG      FXX ...  19           1           0
enas_dest                FXX ...   1           0           0
End2End_SG1             FX. ...   1           0           0
EnhancementChild1       FXX ...   5           0           0
EnhancementChild2       FXX C...   5           0           0
EnhancementsSgParent    F.X P...   5           0           1
ESX_142_GKS_SG          FXX ...   6           6           0
GKs                      FXX ...  24          24          0
MFE_EXPAND              9X. ...   1           0           0
NDM_Beta_PreCopy        F.. ...   5           0           0
NDM_U4P_PreCopy         FX. ...   4           0           0
OpenStack_Instances     FXX ...   8           0           0
```

**Migrate a subset of devices in an SG**

To migrate a subset of devices in an SG, divide that SG into a number of SGs. This is a non-disruptive operation.

In Unisphere, select the SG on the **Storage Group Dashboard**, click the **More Actions** menu, and select **SG Maintenance > Split From**.





The result is a stand-alone SG with its own masking view, using the same components as before the split, that is fully capable of being migrated using NDM. For guidance on how to migrate the SG, see the section on Metro-based NDM ([Using Unisphere for PowerMax](#) or [Using Metro NDM using Solutions Enabler 9.x](#)) or Metro-based NDM with precopy ([Using Unisphere for PowerMax with precopy](#) or [Using Solutions Enabler 9.x with precopy](#)).

From Solutions Enabler, use the **Symsg** library.

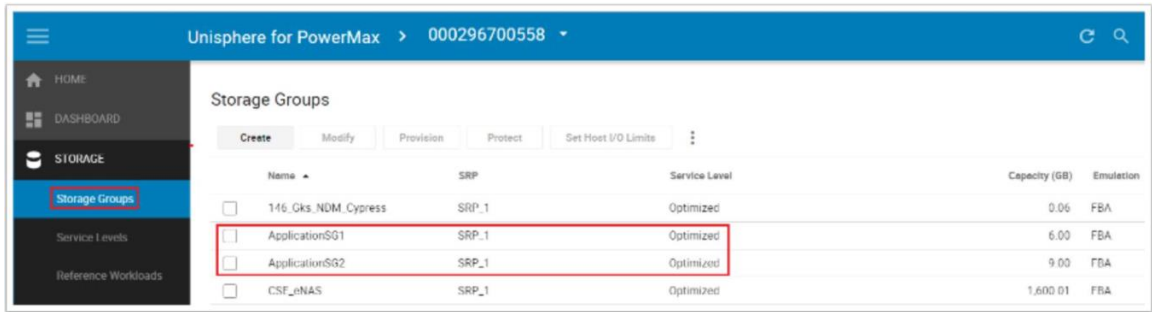
```
symsg -sg <SgName> -sid <SymmID>
merge <SgName1>
split <SgName1> -view_name <MvName> -devs
```

For the split operation, the `-sg <SgName>` option is used to specify the source SG, for example, `EnhancementsChild1` for the split operation. Specifying the name of the Storage Group create after the Split parameter (`EnhancementChild1_Split`) the devices to split with the `-devs` parameter and finally the `-view_name` option for the masking view create.

**Consolidate a single application on two arrays into a single target**

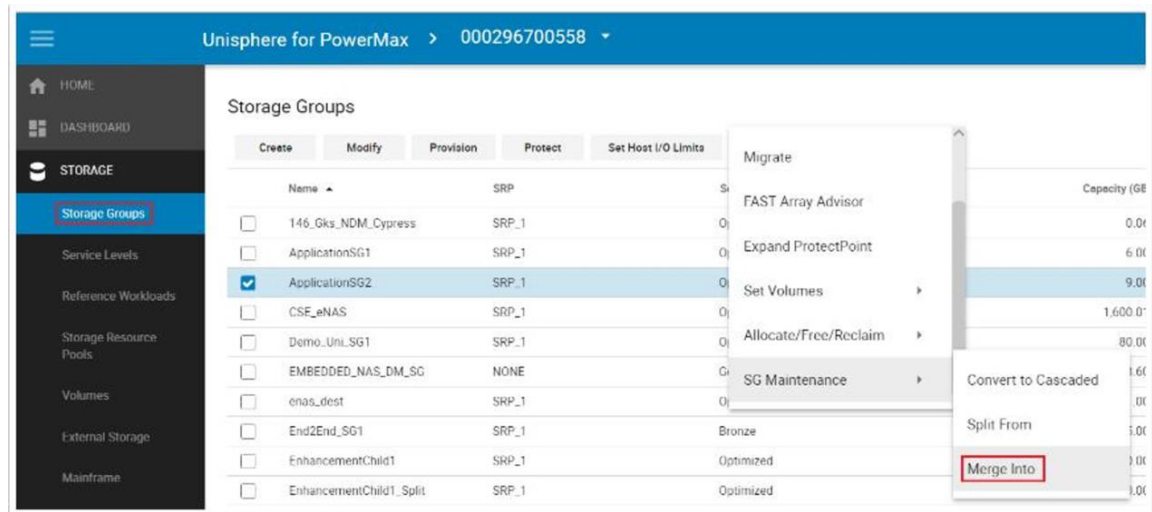
Storage arrays are more compact, have more power, and more storage as time progresses. So NDM is being used to consolidate a number of older arrays onto a single array. For various reasons, we see Storage groups spread across multiple arrays in data centers. Once the groups are moved, using NDM, to a single array they can be merged, as long as the IGs and PGs contain the same elements.





The example above shows two storage groups (ApplicationSG1 and ApplicationSG2) that have been migrated from two separate arrays. The groups use the same application host and so their IGs are identical. The target PG was selected manually by the storage administrator.

To merge the groups using Unisphere, select one of them on the Storage Groups dashboard. Click the more options (three dots) icon and select **SG Maintenance > Merge Into**.



In the pop-up box, select the storage group you want to merge the selected group into, and click **Run Now**. Only storage groups with the correct masking elements are available in this dialog.

The result is a single storage group, ApplicationSG2, that retains the original masking view.

## Masking enhancements (SE 9.0)

With Solutions Enabler, use the **symmsg** library:

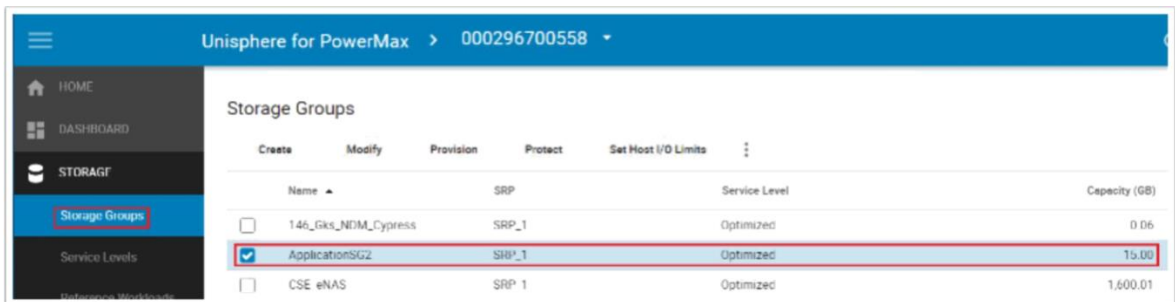
```
symmsg -sg <SgName> -sid <SymmID> merge <SgName1>  
split <SgName1> -view_name <MvName>
```

For the merge operation, the `-sg <SgName>` option is used to specify the source SG, for example, ApplicationSG1 for the merge operation. Specifying the name of the Storage Group to be merged (EnhancementChild1\_Split)

The result is a single SG called ApplicationSG1 that contains the devices from both SGs.

## When creating an NDM session set devices RDF capable

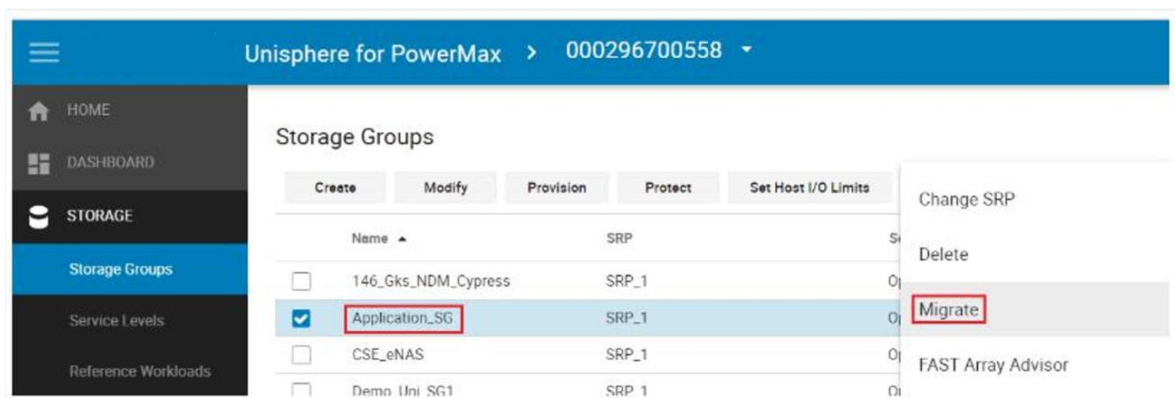
Previously it was not possible to migrate devices that did not have the Dynamic RDF compatibility. This applies to Pass-through NDM (5876 – 5977/5978) This required a configuration change and also it delayed NDM scheduling. Now, the NDM Create operation sets all source device to being Dynamic RDF capable. Therefore, there is no extra user intervention required.



## Add DR to target SG before synchronization

Initially, it was possible to set up SRDF disaster recovery (DR) on the target array using SRDF/A. A later version enabled this use of SRDF/S. In both cases, the NDM synchronization had to be complete before setting up the DR relationship.

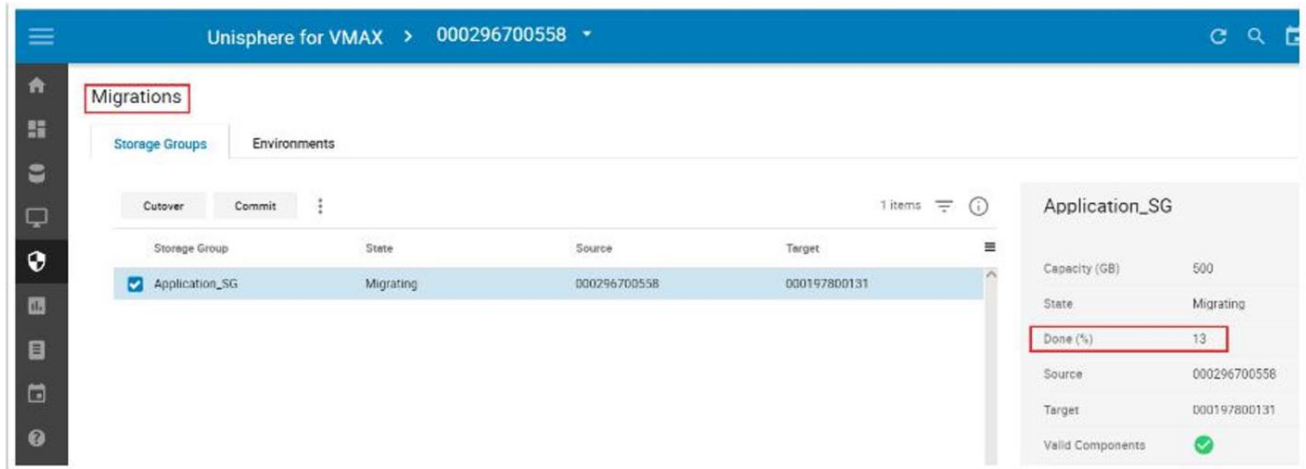
Using PowerMaxOS 5978 and Solutions Enabler 9.x, it is possible to setup SRDF/S or SRDF/A from the NDM target array to a disaster recovery site once the migration has entered a migrating state. This reduces significantly the migration process from a customer standpoint. This will also reduce the impact on response time to the host.



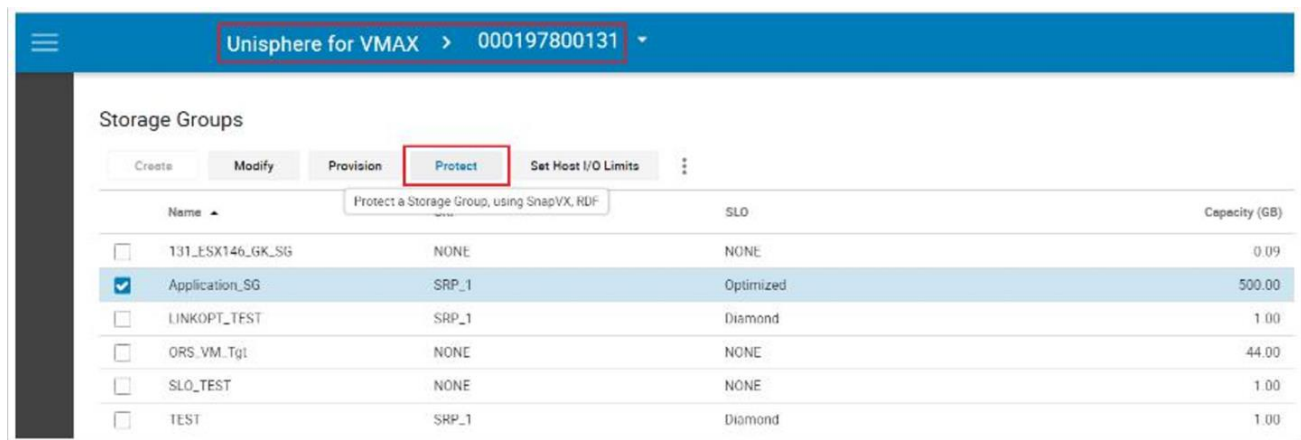
As before, select the Storage Group to be migrated and start the NDM create operation. For further guidance on how to migrate the SG, see the section on Metro-based NDM (Using Unisphere for PowerMax or Using Metro NDM using Solutions Enabler 9.x) or

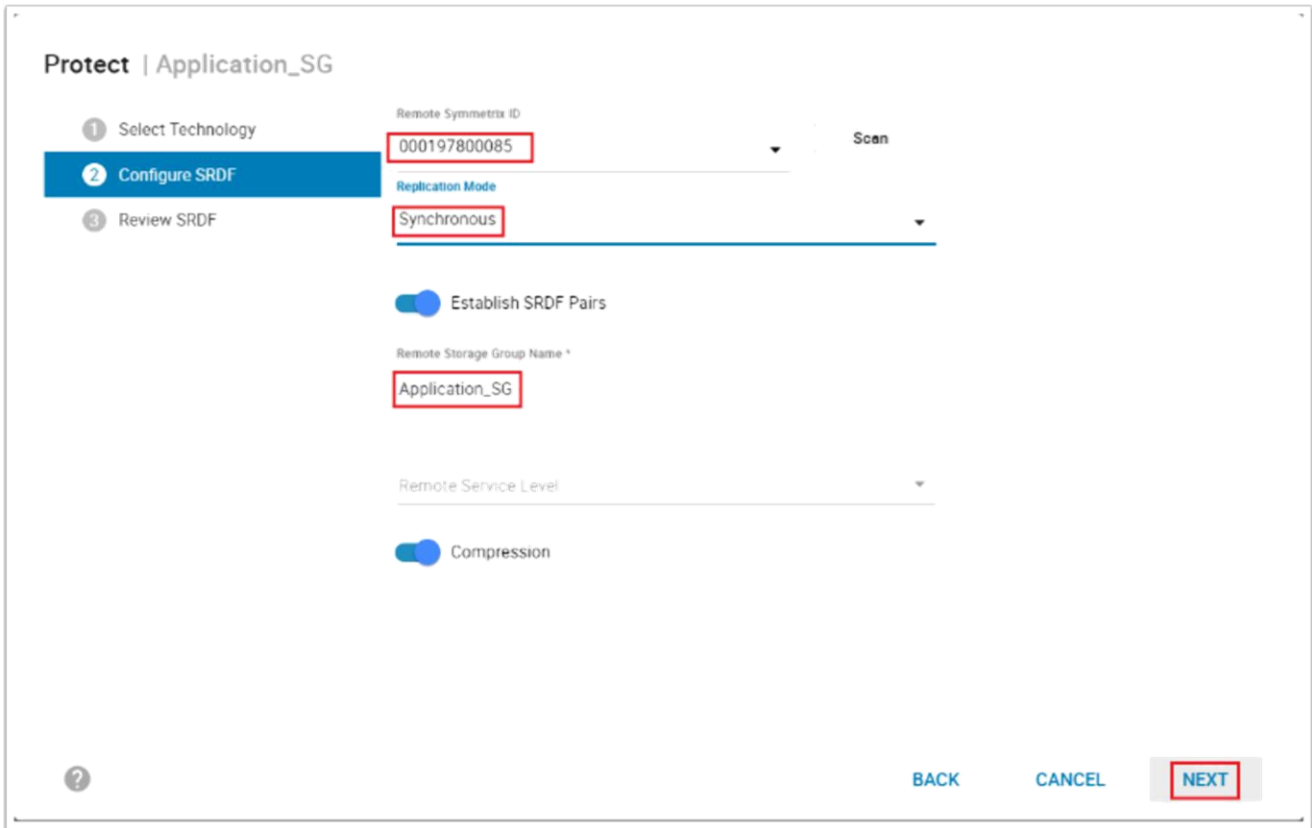
Metro-based NDM with precopy (Using Unisphere for PowerMax with precopy or Using Solutions Enabler 9.x with precopy).

In this example, the Storage Group view shows that the migration is 13% completed. In the previous version of NDM, it would have been necessary to let the migration complete to 100% before setting up DR.



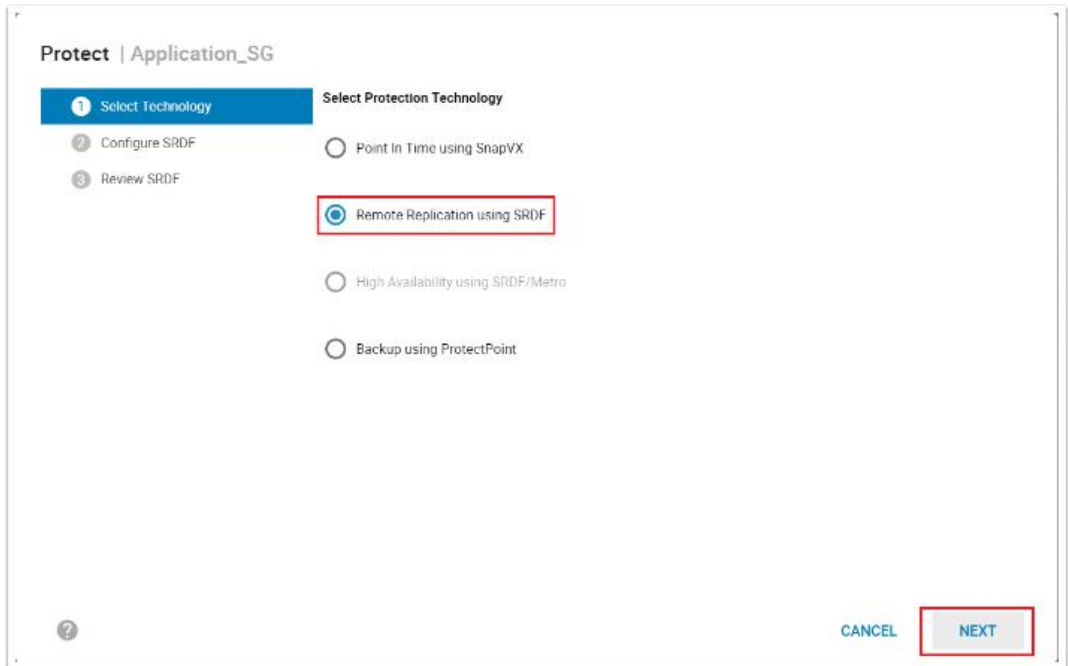
From the target array, navigate to the Storage group, highlight it, and select **Protect**.



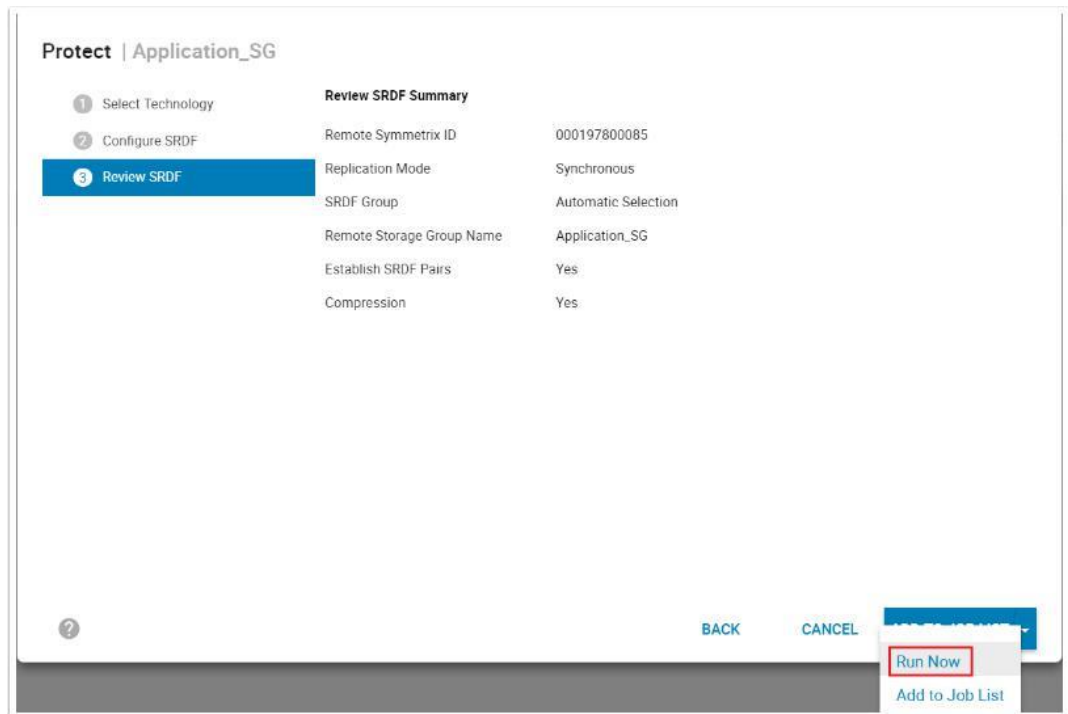


In the pop-up window, click **Remote Replication using SRDF** and click **Next**.

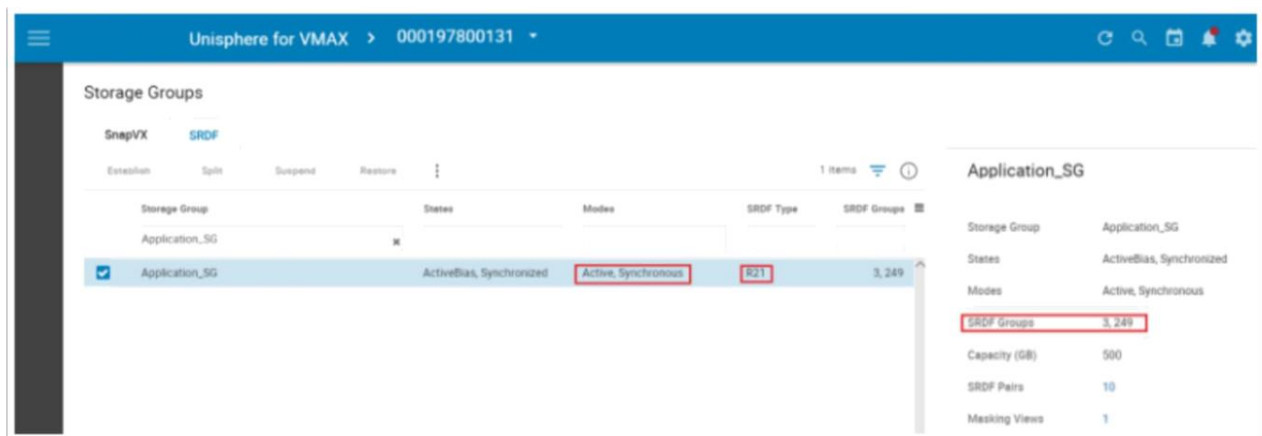
In the protection configuration window, select the target array, the SRDF mode, and the remote storage group name.



The final confirmation of the planned configuration is displayed. Review and click **Run Now**.



The result is an SG with DR in an Active/Bias and Synchronous RDF state. Essentially, this cascaded R21 with the R1 being the NDM Source and the R2 being the new DR to 085. For guidance on how to continue with the migration, see the section on Metro-based NDM ([Using Unisphere for PowerMax](#) or [Using Metro NDM using Solutions Enabler 9.x](#)) or Metro-based NDM with precopy ([Using Unisphere for PowerMax with precopy](#) or [Using Solutions Enabler 9.x with precopy](#)).



To achieve the same result with Solutions Enabler, perform the following:

1. Create an SRDF group between the NDM target and planned Disaster Recovery site:

## Masking enhancements (SE 9.1)

```
symrdf addgrp -label DrSite1 -rdfg 3 -sid 131 -dir  
1F:30,2F:30 -remote_rdfg 3 -remote_sid 085 - remote_dir  
1F:31, 2F:31
```

### 2. Create the SRDF pairings using the following:

```
symrdf -file srdf.txt -sid 131 -rdfg 3 -type r1 -establish  
createpair where the srdf.txt file contains the device  
pairings
```

In this SRDF query, the RDFG is shown as 3 which shows that it is on the second leg of the R21. The devices on the DR target are Write disabled to the host and are Synchronized. Some other information such as the 'S' under the mode (M) signifies Synchronous mode.

```
C:\Program Files\EMC\SYMCLI\bin>symrdf -sid 131 -rdfg 3 query
```

Symmetrix ID	:	000197800131	(Microcode Version: 5978)
Remote Symmetrix ID	:	000197800085	(Microcode Version: 5977)
RDF (RA) Group Number	:	3 (02)	

Source (R1) View					Target (R2) View					FLAGS
Standard	ST			LI	ST					
Logical	Sym	T R1 Inv	R2 Inv	K Sym	T R1 Inv	R2 Inv				RDF Pair
Device	Dev	E Tracks	Tracks	S Dev	E Tracks	Tracks	MCES			STATE
N/A	00006	RW	0	0 RW 000C9	WD	0	0	S..E		Synchronized
N/A	00048	RW	0	0 RW 000CE	WD	0	0	S..E		Synchronized
N/A	00058	RW	0	0 RW 00096	WD	0	0	S..E		Synchronized
N/A	00059	RW	0	0 RW 002CC	WD	0	0	S..E		Synchronized
N/A	00208	RW	0	0 RW 00095	WD	0	0	S..E		Synchronized
Total										
Track(s)		0	0			0	0			
MB(s)		0.0	0.0			0.0	0.0			

### Create bot LUNs on the target array the same as the source array

NDM now retains the LUN addresses on the target array the same as the source array when using the -consistent LUN option on the NDM Create operation. Although the LUNs address will be the same note that redirecting of the BIOS to the new boot LUN must be done by the system administrator. NDM does not have control of the individual host BIOS.

## Masking enhancements (SE 9.1)

### Migrate storage groups with shared volumes

Prior to Solutions Enabler 9.1, data migration using NDM would be blocked if devices were shared across storage groups selected for migration.

With Solutions Enabler 9.1, users can migrate storage groups that have fully or partially overlapping devices. This results in multiple SGs being migrated as part of the NDM session.

**Create port group as part of the NDM or NDM Update create command**

With the release of Solutions Enabler 9.0, the ability to target a specific port group as part of the NDM create command was introduced. This allowed users to pre-create a PG to be used for the target storage group.

Solutions Enabler 9.1 gives the user the ability to create the port group as part of the NDM or NDM Update create process. The user has the ability to name and select specific target-side ports.



## Appendix A: Host multipathing software notes

**Introduction** This section describes best practices for using multipathing software in an NDM environment. See the NDM Support Matrix for the latest operating system and multipathing software combinations.

### AIX Native Multipathing software

For Native Multipathing on AIX, best practice is to use the following settings for MPIO:

```
algorithm = round_robin    (other algorithms may also be used)
hcheck_cmd = inquiry
queue_depth = 32
reserve_policy = PR_shared*
```

\*reserve\_policy may be set to no\_reserve as long as GPFS/SCSI3 clusters are not part of the planned migration.

Check if a "PR\_key\_value" is already assigned prior to changing the "reserve\_policy" to PR\_shared.

### Linux multipathing software with LUNZ

If using address 0xf7, Using NDM with result in a failure scenario. The recommendation is to modify with other another address before attempting to use NDM.

If not greater than 0xF7 LUNs, as long as the system has a LUN with host id 0xF7, then a manual removal of the LUNZ and rescan is required:

1. Find the VMALUNZ by "ls SCSI | grep VMAXLUNZ".
2. Remove the SCSI device by "echo 1 > /sys/block/xxx/device/delete".
3. Rescan the SCSI devices by "rescan-scsi-bus.sh".
  - Some Red Hat versions require the addition of the -a parameter
  - rescan-scsi-bus.sh -a
  - See <https://access.redhat.com/solutions/1314183>

### PowerPath (version 5.7 and later)

Use the default PowerPath multipath settings.

First, run the SCSI device rescan (rescan\_scsi\_bus.sh).

To detect, rescan, and configure new paths use the powermt config command.

Use PowerPath commands/scripts to scan and remove stale paths (powermt check).

The powermt restore command can be used to detect path changes faster than PowerPath will discover them on its own.

---

**Note:** See [Appendix B: AIX, GPFS, and PowerPath with NDM](#) for important information regarding PowerPath and AIX 6.x with GPFS.

---

## PowerPath with Solaris

For Solaris with PowerPath, a cancel revert operation will end with the session in a “CancelFailed” state. The host paths to the source array running 5876 must be recovered using `cfgadm`.

For example, to recover the dead paths one by one:

```
cfgadm -c configure c5::50000973f001d109
cfgadm -c configure c5::50000973f001d105
cfgadm -c configure c4::50000973f001d109
cfgadm -c configure c4::50000973f001d105
```

To recover all paths on a controller:

```
cfgadm -c configure c4
cfgadm -c configure c5
```

Once the paths are online, the cancel revert can be resumed by performing an NDM recover operation.

## Windows Server 2012 with MPIO

Use default MPIO settings with the following parameters enabled:

- PathVerifyEnabled - Enable for optimal results with path discovery.
  - With “Path Verify Enabled” **checked**, the target V3 paths will be automatically discovered following an NDM create and a “cancel -revert”
  - If “Path Verify Enabled” is **cleared**, the target V3 paths will **not** be automatically discovered following an NDM create (host rescan discovers target V3 paths) and a “cancel -revert” may fail. Recommendation is to initiate a manual rescan during the “Wait for host path discovery on Source” step.
- PathVerificationPeriod - Set a time in seconds for automatic path detections. Dell Technologies recommends setting it to lowest allowed value between 10 and 30 seconds.

## Veritas Dynamic Multipathing

Configure the DMP tunable parameters to NDM required values. From the application host:

Check and modify the following DMP tunable parameters. If the parameter values are not set to the default values.

---

**Note:** Settings so the parameters can be returned to expected values following NDM migration.

---

DMP Tunable Parameter	Default Value	NDM Required Value
<code>dmp_path_age</code>	300	0
<code>dmp_health_time</code>	60	0
<code>dmp_restore_interval</code>	300	10

```
dmp_restore_cycles          10
10
iopolicy (per DMP node name)      MinimumQ      MinimumQ
```

Use the rescan command.

On Linux: /usr/bin/rescan\_scsi\_bus.sh followed by vxdisk scandisks to detect new paths. Use vxddmpadm to verify that the new paths are added.

[https://sort.symantec.com/public/documents/dmp/6.0/vmwareesx/productguides/html/dmp\\_admin/ch06s08.htm](https://sort.symantec.com/public/documents/dmp/6.0/vmwareesx/productguides/html/dmp_admin/ch06s08.htm)

### Veritas cluster behavior with NDM

To read the SCSI Persistent Reservation keys after the commit, you must perform a cluster failover. We recommend performing a reboot of each node in the cluster. However, this is not a mandatory action if you perform a cluster failover.

### ESXi with Native Multipathing

Use the rescan command to detect new paths, or wait for NMP to detect the paths automatically.

To reduce the delay in automatic detection, change it to 30 seconds.

To set the path polling time, login to the host and navigate to **Configuration > Advanced Settings > Disk** and update the **Disk.PathEvalTime** field.

### Solaris 10 SPARC with Solaris cluster 3\_3u2 using Pass-Through NDM, Missing Reservations

When migrating from VMAX to PowerMax, VMAX All Flash or VMAX3 using Pass-through NDM the following could potentially be observed during the Cutover command:

```
e2e-14-100242:/opt/emc/SYMCLI/bin # ./symdm cutover -sid 176 -sg
NDM176_4
```

```
Nov 29 22:10:24 soh4ser2 cl_runtime: [ID 868277 kern.warning]
WARNING: CMM: Erstwhile online quorum device /dev/did/rdisk/d61s2
(qid 2) is inaccessible now.
```

```
Nov 29 22:10:24 soh4ser2 cl_runtime: [ID 868277 kern.warning]
WARNING: CMM: Erstwhile online quorum device /dev/did/rdisk/d44s2
(qid 3) is inaccessible now.
```

This also applies when migrating from VMAX3 or VMAX All Flash to VMAX All Flash or PowerMax. Reservations

An error message such as “reservation key on the quorum device gone”

These are temporary condition with no host impact. Waiting a few minutes will result in the Reservations returning from the target array.

## Appendix B: AIX, GPFS, and PowerPath with NDM

Properly configuring multiple paths to the host with AIX 6.x, GPFS, and PowerPath requires an additional step when performing an NDM migration.

After the NDM create operation is completed, the target devices (appearing to the host as additional paths to the source devices) will be masked and available. After running `cfgmgr` to create the host native devices, the attached script, named `emc_pp_configure.sh`, must be run immediately following the completion of the `cfgmgr` command. The script will configure the new native devices on the target side into PowerPath by copying the attributes of the PowerPath `hdiskpower` pseudo devices into the new native devices and reconfigure the native devices.

---

**Note:** This script **must** be used in order to perform a migration non-disruptively in this environment. Failing to run the script following the configuration of the new native target devices can lead to data unavailability.

---

```
#!/bin/ksh
devlist=`powermt config 2>&1 | grep -p 0514-034 | grep hdiskpower | awk '{print $5}'`
for pseudo in $devlist
do
    pseudo_policy=`lsattr -El $pseudo -a reserve_policy | awk '{print $2}'`
    pseudo_prkey=`lsattr -El $pseudo -a PR_key_value | awk '{print $2}'`
    nativelist=`powermt display dev=$pseudo | grep -i hdisk| grep -v power | awk
'{print $3}'`
    echo $nativelist
    for native in $nativelist
    do
        native_policy=`lsattr -El $native -a reserve_policy | awk '{print $2}'`
        native_prkey=`lsattr -El $native -a PR_key_value | awk '{print $2}'`
        #change reserve_policy and PR_key_value of native(s), whose policy or
PR_key_value
        #vary from that of pseudo.
        if [[ $native_policy != $pseudo_policy || $native_prkey != pseudo_prkey ]]
        then
            powermt remove dev=$native
            echo changing reserve_policy,PR_key_value of $native to $pseudo_policy
and $pseudo_prkey
            chdev -l $native -a reserve_policy=$pseudo_policy -a
PR_key_value=$pseudo_prkey
            fi
        done
    done
done
powermt config
```

The script can be downloaded on the following support site: [emc\\_pp\\_configure.sh](#)

### **AIX LPM (Live Partition Mobility) with NDM**

For Solutions Enabler version 8.3, remove the passive initiator (for LPM use) from IG when they do NDM. After NDM, they need to add the passive initiator back to IG, then they can do LPM.

For Solutions Enabler version 8.4 and later, perform the LPM operation at least once. Let the passive initiator log in to the array and show up in Login History table (LHT), and after this LPM operation, perform NDM.

It is not recommended to do LPM operations during the NDM session.

## Appendix C: Consistent LUN

If the source array does not have the Consistent LUN attribute set, but the target does (if it is a pre-created IG) then there will be consistent LUN addresses on the Target array. However, there is no attempt to use the same LUN addresses as on the source even if, by chance, the LUN addresses on the source are consistent across all the paths.

- By default, VMAX and PowerMax arrays are delivered with the ACLX using LUN 0 and visible on a number of ports. Therefore, if the SHOW\_ACLX\_DEV attribute is still in effect on a port that NDM will choose and the source array is using Consistent LUN and is also using LUN 0, then we will not be able to set the same LUN addresses. This is because LUN 0 is not available on the target array for all the ports that NDM will use when it builds the Masking View.
- When the target VMAX or PowerMax array is deployed and storage is provisioned to the host and uses the LUN addresses for these new LUNs that it is also using for the 'application to be migrated' that resides currently on the source array:
  - Obviously, the existing application is using a different set of ports/paths because it is on a different array.
  - When we migrate that application, the LUN addresses we would want to use will not be available. Thus, you will get consistent, but different LUN addresses.

## Appendix D: Device geometry behavior post NDM

User Actions	Device Geometry Mode			
	User Defined		GCM	
	Pre-5978	5978	Pre-5978	5978
Not leave device geometry set post NDM Commit (when device geometry was getting set automatically during NDM)	Yes; starting with SE 8.4 (if target device size=source device size)		GCM ALWAYS set when there are odd # of cylinders on the source array	
Unset/clear geometry of devices in replication relationship*	Allowed	Allowed	Not Allowed	Not Allowed**
Unset/clear geometry of devices NOT in replication relationship*	Allowed	Allowed	Allowed	Allowed
Expansion of device (with geometry) in replication relationship*	Not Allowed	Allowed	Not Allowed	Allowed
Expansion of device (with geometry) NOT in replication relationship*	Allowed***	Allowed	Allowed***	Allowed

---

**Note:** Clearing GCM on a device without replication that is mapped is a Unisphere for PowerMax and Solutions Enabler 9.x feature only.

---



## Appendix E: VMware—VM Clone or Storage vMotion with NDM

For NDM between 5977 and 5978, Metro-based NDM a slight delay in processing VAAI instructions, namely xCopy, during the copy and sync stages. This is a known code delay will not affect normal NDM operations.

## Appendix F: NDM with ESXi virtualized cluster environment

This section describes a known procedure/timing in the VMware ESXi virtualized cluster environment:

**NDM: Create > Discover V3/PowerMax Paths > Cancel > Create**

This is the issue where we see the host/multipath assume that re-discovered V3 paths are still registered from before the cancel.

In the Microsoft virtualized cluster environment, this results in the symptom NTFS event 57 and unexpected cluster resource failover.

To Mitigate this issue Recommendation is to cleanup stale/dead host paths after a Cancel issued.

**NDM: Create > Discover V3 Paths > Create > Cutover > Cancel –revert (immediate)**

This is the issue where we introduce two NDM changes before all participating hosts have had change to recognize the first change (cutover).

We observe cases where a write can be rejected by both V2 and V3 and devices get marked PDL (permanent device loss) by VMware.

To mitigate this issue, Dell Technologies recommends that all participating hosts have realized the cutover change before running a cancel -revert.

## Appendix G: NDM Update with Solaris cluster

**Issue:** After NDM Update cutover, sc3.3 nodes on ldom panicked when checking cluster status. OS log showed cluster lost access to quorum dev which is part of the NDM Update.

**Recommendations:**

- No access to LUNs during migration, especially quorum device. Cluster related operations will trigger access to LUNs including quorum. Opening FS mount point will also trigger access to corresponding LUNs. These operations are not allowed for LUNs in migration.
- It is mandatory to always use `-move_identity` when migrating Solaris Cluster Environments with NDM Update.
- Immediately run a path rescan after the NDM Update creation command if not using the `-precopy` option.
- Paths to source will become unusable after a cancellation and require a `cfadm -c` configure to allow them to be used again.

## Appendix H: Technical support and resources

[Dell.com/support](https://www.dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage and data protection technical white papers and videos](#) provide expertise that helps to ensure customer success with Dell Technologies storage and data protection products.

### Related resources

Here is a link to the comprehensive knowledge base articles referencing some previously experienced issues while using NDM as well as guides and best practices.

<https://www.dell.com/support/kbdoc/en-us/534580>

### Operations Failure:

- Create with Validate Option `symdm create -validate`
- Create `symdm create`
- Cancel `symdm cancel`
- Cutover `symdm cutover`
- Cancel with Revert Option `symdm cancel -revert`
- Commit `symdm commit`
- Recover `symdm recover`
- Remove Environment `symdm environment -remove`

### Stuck Migration State:

- Created
- CutoverInProg (Level 50)
- Migrating
- CommitInProg
- CancelFailed
- Performance / Slow Copy Rate / Quality of Service (QoS)
- Hosts / Clusters
- PowerPath / Multipaths - Stale Group Reservations / Locks
- Device External Identity / Non-Native / Effective WWNs
- Geometry Compatibility Mode (GCM)
- SRDF Metro
- SRDF CE (Cluster Enabler)
- RecoverPoint
- Errors
- Red boxes (Level 40 + Level 50)
- General
- Guides