

# Dell EMC VPLEX: Leveraging Array Based and Native Copy Technologies

Implementation planning and best practices

## Abstract

This white paper provides best practices planning and use cases for using array-based and native replication solutions with Dell EMC™ VPLEX™ Local, VPLEX Metro along with metro node for Dell EMC PowerStore™ and Dell EMC Unity XT.

January 2021

## Revisions

Date	Description
April 2019	Revision 5
January 2021	Metro node and PowerStore volume migration considerations

## Acknowledgments

This paper was produced by the following:

Author: VPLEX CSE Team

[VPLEX\\_CSE\\_TEAM@emc.com](mailto:VPLEX_CSE_TEAM@emc.com)

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 © January 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [1/19/2021] [Best Practices] [H14293]

# Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents .....	3
Executive summary.....	4
Audience.....	4
Document organization .....	4
1 Introduction.....	6
2 Terminology.....	7
3 Array-based replication with VPLEX .....	9
3.1 Introduction.....	9
3.2 Array-based copy for VPLEX RAID-0 volumes .....	9
3.2.1 Creating VPLEX RAID-0 copies .....	11
3.2.2 Procedure: Creating VPLEX RAID-0 array based copies: .....	11
3.3 Resynchronizing (updating) local array-based copies presented through VPLEX .....	14
3.3.1 To re-synchronize (update) VPLEX virtual volumes based on array-based copies:.....	14
4 Advanced array-based copy use cases .....	17
4.1 Prerequisites.....	19
4.2 Creating VPLEX RAID-1 or distributed RAID-1 array-based copies .....	19
4.2.1 Procedure: Creating VPLEX RAID-1 or distributed RAID-1 array-based copies .....	19
4.3 Procedure: Resynchronizing (refreshing) array-based copies of RAID-1 VPLEX virtual volumes .....	21
5 Array-based restore for VPLEX RAID-0 production volumes .....	25
5.1 Procedure: Array-based restore to VPLEX RAID-0 production virtual volumes .....	25
5.2 Array-based restore of VPLEX RAID-1 production volumes.....	28
5.2.1 Prerequisites:.....	28
5.2.2 Procedure: Array-based restore to VPLEX RAID-1 production virtual volumes .....	29
6 Metro node considerations .....	34
7 PowerStore volume migration .....	35
8 Conclusion.....	36
A Technical support and resources .....	37

## Executive summary

The Dell EMC™ VPLEX™ family removes physical barriers within, across, and between data centers. VPLEX Local provides simplified management and non-disruptive data mobility across heterogeneous arrays. VPLEX Metro provides high availability and non-disruptive mobility between two data centers within synchronous distances. With unique scale-up architecture, VPLEX's advanced data caching and distributed cache coherency provides workload resiliency, automatic sharing, and balancing and failover of storage domains, and enables both local and remote data access with predictable service levels.

Preserving investments in array-based storage replication technologies like snapshots and clones or remote replication including array level replication provided by RecoverPoint is crucial in today's IT environments. This paper outlines the considerations and methodologies for array-based replication technologies within VPLEX Local and Metro environments. Procedural examples for various replication scenarios, examples of host based scripting, and VPLEX native copy capabilities are examined. The key conclusion of the paper is array-based replication technologies continue to deliver their original business value with VPLEX environments.

While the focus of this document is covering array based replication, Dell EMC does have a software product that provides management of this capability. AppSync is software that enables Integrated Copy Data Management (iCDM) with Dell EMC's primary storage systems. AppSync simplifies and automates the process of generating and consuming copies of production data by automatically discovering applications, determining the structure needed to maintain consistency, and maps them through the virtualization layers, to the underlying storage environment. AppSync then orchestrates all the activities required, from copy creation and validation, through mounting the copies at the target host and launching, or recovering, the application. Supported workflows also include refresh, expire, and restore operations through VPLEX, supported with XtremIO, Unity, Unity XT, PowerStore and VMAX AF and PowerMax back-end storage systems. Please refer to the AppSync documentation for more information.

## Audience

VPLEX is constantly evolving as a platform. The procedures and technology discussed in this white paper are only applicable to the VPLEX Local and VPLEX Metro products. Please consult with your local Dell EMC support if you are uncertain as to the applicability of these procedures to your VPLEX environment. Metro node is a new addition to the VPLEX family developed as a feature for both PowerStore and Unity XT. Due to changes in the product, the procedures and commands may not apply to metro node. Please refer to the metro node section for proper best practices that are specific to metro node with Unity XT and PowerStore or please refer to the PowerStore and Unity XT documentation.

## Document organization

This technical note is one of a set of documents that supersede the monolithic Implementation Planning and Best Practices for Dell-EMC VPLEX Technical Notes document that had previously been available. It is intended to provide more concise, per-topic information that will be easier to maintain and keep up-to-date.

All documents are available on SolVe Desktop and <https://www.dell.com/support/home/en-us>.

The following list represents the technical note best practice documents:

- Dell-EMC VPLEX Overview and General Best Practices
- Dell-EMC VPLEX SAN Connectivity

## Executive summary

- Dell-EMC VPLEX Host Multipathing
- Dell-EMC VPLEX Networking
- Dell-EMC VPLEX and RecoverPoint Interoperability
- Dell-EMC VPLEX Data Mobility and Migrations
- Dell-EMC VPLEX Upgrades

# 1 Introduction

The first section of this paper examines the technical impact of VPLEX on array-based replication productions like snapshots, clones and remote replication on any supported backend storage array. The necessary pre-copy considerations, logical checks, process changes and use case examples are provided. The examples illustrate the correct way to account for VPLEX and potential storage conditions that could impact the consistency of the copy or snapshot. Section 2 explores VPLEX's native copy capabilities for both the Unisphere for VPLEX GUI and for the VPLEX CLI. VPLEX CLI commands can be accessed via the vapi script (available from your Dell EMC account team), direct ssh login or via the VPLEX RESTful API.

---

**Note:** Using the built-in copy capability of VPLEX opens the door to creating copies of virtual volumes from heterogeneous storage arrays within and across data centers. This allows you to create and maintain the copy on a different array.

---

## 2 Terminology

Table 1 Operational definitions

Term	Definition
Storage volume	LUN or unit of storage presented by the back-end arrays.
Metadata volume	System volume that contains metadata about the devices, virtual volumes, and cluster configuration.
Extent	All or part of a storage volume.
Device	Raid protection scheme applied to an extent or group of extents.
Virtual volume	Unit of storage presented by the VPLEX front-end ports to hosts.
Front-end port	Director port connected to host initiators (acts as a target).
Back-end port	Director port connected to storage arrays (acts as an initiator).
Director	The central processing and intelligence of the VPLEX solution. There are redundant (A and B) directors in each VPLEX Engine.
Engine	Consists of two directors and is the unit of scale for the VPLEX solution.
VPLEX Cluster	A collection of VPLEX engines in one rack, using redundant, private Fibre Channel connections as the cluster interconnect.
Storage Array	An intelligent, highly available, scalable, hyper consolidated, multi-tiered block storage frame. For example, VNX, Symmetrix DMX4, VMAX, VMAX3, XtremIO, Dell EMC Unity XT, Powerstore and PowerMax

Table 2 Acronyms and abbreviations

Acronym/Abbreviation	Definition
Copy	An independent full disk copy of another device.
Source device	Standard or primary array device. Typically used to run production applications.
Target Device	A secondary array device designated to receive new data (overwritten) from the source device.
FE port	Front-end (target ports visible to hosts).
BE port	Back-end (initiator ports visible to storage arrays).



## 3 Array-based replication with VPLEX

### 3.1 Introduction

Array based copy technologies are feature rich, mature, robust, scalable, and purpose-built to provide enterprise-class replication services. Deploying VPLEX in front of storage frames with these services does not diminish their function or value. Underlying storage devices are left untouched by VPLEX and the array replication technologies can continue to provide backup, business continuity, operational recovery, or QA/Test/Dev functionality.

VPLEX enables array-based replication products like Dell EMC MirrorView, TimeFinder, SRDF, SnapView, Unity snapshot, XtremIO and PowerStore snapshots and native replication to continue to deliver full functionality and business value. Unlike other virtual storage solutions that require (or strongly recommend) slicing physical storage into small pieces and a complete reconstruction of those components into virtual storage, VPLEX has no such requirement. This is because VPLEX uses 1:1 physical to virtual storage mappings as the default methodology. When storage arrays are virtualized using VPLEX they do not lose their original replication capabilities (local or remote). Similar to traditional array based restoration or refresh, prior to restore host applications must be stopped or quiesced. Once the restore process is completed, the virtual volume can once again be accessed by the host application.

As mentioned previously, VPLEX Local and Metro use a write-through caching architecture. One of the implications of this architecture is that the VPLEX cache (not the underlying storage array) is the source for many for host read requests. Cache invalidation (for both hosts and VPLEX), therefore, becomes extremely important for array-based restoration operations because the array is modifying the underlying physical storage outside of the host and VPLEX I/O path. When this happens, the read caches maintained by the host and by VPLEX will not match the data within the array. This condition has the potential to cause data loss or corruption. For this reason, the cache invalidation must happen immediately following the array restore/refresh (prior to the resumption of host I/O). VPLEX virtual volume read cache invalidation can be accomplished in some number different ways outlined in the following sections. The version of VPLEX code on the system will dictate which methods are available. Starting with VPLEX GeoSynchrony 5.2 patch 1, an enhancement was added to facilitate direct invalidation of VPLEX cache for single virtual volumes or for entire consistency groups.

### 3.2 Array-based copy for VPLEX RAID-0 volumes

Let's start our discussion with some examples of using array-based replication (i.e. TimeFinder, SnapView, or XtremIO or Unity snapshots) to create VPLEX RAID-0 copies. Generally speaking, array-based copy technologies provide independent full disk copies (Clones or BCVs) and/or space efficient copies (Snapshots). From a VPLEX point of view, both types of copies represent normal back-end array storage volumes.

---

**Note:** Snapshots or any thinly provisioned storage volumes should be identified as thin devices during the VPLEX claiming process.

---

These copies can either be presented back through VPLEX or directly from the underlying storage array to a backup host, to a test environment, or even back to the original host. For our discussion, it is assumed the copies will be presented back through VPLEX to a host. For this basic use case example each array-based copy has a one-to-one VPLEX virtual volume configuration (VPLEX device capacity = extent capacity = array storage volume capacity) and has a single extent VPLEX RAID-0 or two extent RAID-1 device geometry. The

process for managing more complex geometry array-based copies (for example, a copy that is RAID-1 or distributed RAID-1) is covered in the advanced use case discussion at the end of this section.

Potential host, VPLEX, and storage array connectivity combinations for the basic use case are illustrated in Figure 1 and Figure 2. Figure 1 shows a single array with a single extent (RAID-0 geometry) source and single extent (RAID-0) array based copy. Both the source and the copy storage volumes pass through VPLEX and then back to hosts. The array-based copy is presented to a separate second host.




Figure 1 Array-based RAID-0 Copy of a RAID-0 Production Virtual Volume





Figure 2 Array-based RAID-0 Copy of a RAID-1 Production Virtual Volume

In Figure 2, the source volume has single extent RAID-1 geometry. Each mirror leg is a single extent that is equal in size to the underlying storage volume from each array. The difference in this topology is that there is an option to use the array-based copy services provided by array 1 and/or by array 2. The choice of which


array to copy from can be based on available capacity, array workload or, perhaps, feature licensing. The objective remains the same as with Figure 1 -- to create a copy that has single extent (RAID-0) geometry.

The following procedures address the examples shown below:

1. Creation of RAID-0 array-based copies of VPLEX RAID-0 or RAID-1 source / production virtual volumes.
2. Resynchronizing or updating RAID-0 array-based copies.



Example 1 Create / Refresh VPLEX RAID-0 Copy of RAID-0 Source



Example 2 Create / Refresh VPLEX RAID-0 Copy of RAID-1 Source

### 3.2.1 Creating VPLEX RAID-0 copies

This process augments the traditional best practices for local array-based copies using products like Dell EMC TimeFinder, Dell EMC SnapView, Unity Snapshot, and Dell EMC XtremIO snapshots and native replication. It accounts for VPLEX and for host access through VPLEX to array-based copies.

### 3.2.2 Procedure: Creating VPLEX RAID-0 array based copies:

1. For each storage array, identify source storage volumes to copy. The VPLEX storage-view will contain all VPLEX volumes that are visible to a given host or group of hosts. To determine the VPLEX to array device mapping use one of following methods:
  - Using the VPLEX CLI invoke the 'show-use-hierarchy' or 'device drill-down' command for the virtual volume(s) and underlying device(s) you wish to make a copy of. These outputs provide the virtual to physical storage volume mapping for each VPLEX volume.
  - Use the Storage Viewer feature of Virtual Storage Integrator (VSI) for VMware vCenter Server (for ESX environments) to see the virtual device to physical array device mapping.
  - Correlate the VPLEX virtual volumes to physical storage using the VPDID from the virtual volume to map to the back end storage array device. This information is available to the host via SCSI inquiry and from VPLEX via the /clusters/<cluster name>/exports/storage-view context. To get the back-end array device, you can map virtual volume VPDID back through the VPLEX and get the array LUN. From a VPLEX perspective there is a VPDID tied to the virtual volume the host sees and a different VPDID tied to the back end storage array.
2. Using the VPLEX CLI or REST API, check for valid copy conditions to ensure data consistency:

---

**Note:** Each of the following checks are typically scripted or built into code that orchestrates the overall copy process on the array.

---

- a. Confirm a VPLEX ndu is not in progress. Using the VPLEX CLI issue the 'ndu status' command and confirm that the response is 'No firmware, BIOS/POST, or SSD upgrade is in progress.'
- b. Confirm the device is healthy. Issue the 'll' command from the /clusters/<cluster name>/virtual-volumes/<virtual volume name> context for each volume(s) to be copied.
  - i. Confirm the underlying device status is not marked 'out of date' or in a 'rebuilding' state.
  - ii. Confirm Health Status is 'ok'
  - iii. Confirm Operational Status is 'ok'

---

**Note:** Consistency can't be guaranteed if Host I/O is ongoing and a R1 or DR1 is being snapped/cloned

---

- c. Confirm the source VPLEX virtual volume device geometry is not RAID-C. Device geometry can be determined by issuing ll at the /clusters/<cluster name>/devices/<device name> context.
- d. Confirm each volume is 1:1 mapped (single extent) RAID-0 or 1:1 mapped (two extent) local RAID-1. Distributed RAID-1 device legs must be a combination of RAID-0 (single extent) and/or RAID-1 (two extent) device geometries.
- e. Confirm the device is not being protected by RecoverPoint. Issue the 'll' command from the /clusters/<cluster name>/virtual-volumes/<virtual volume name> context and check 'recoverpoint-protection-at' is set to [] and 'recoverpoint-usage' is set to '-'.
- f. Confirm VPLEX volumes to be copied do not have 'remote' locality (from same VPLEX cluster). Issue the 'll' command against the /clusters/<local cluster name>/virtual volumes/<virtual volume name> context and confirm locality is 'local' or 'distributed'.
- g. Ensure virtual volumes are members of the same VPLEX consistency group and the same array-based consistency group (if available). In most cases all members of the consistency group should be copied together. Consistency group membership can be determined by issuing ll from the /clusters/<cluster name>/consistency-groups/<consistency group name> context.

---

Note: VPLEX consistency group membership should align with array based consistency group membership whenever possible. Some array replication technologies may not have a concept such as a consistency group.

---

- h. For RAID-1 or distributed RAID-1 based virtual volumes, confirm underlying storage volume status is not failed or in an error state. Issue the 'll' command from the /clusters/<cluster name>/devices context or from /distributed-storage/distributed-devices/<distributed device name>/components context.
- i. For distributed RAID-1 confirm WAN links are not down. Issue the 'cluster status' command and confirm 'wan-com' status is 'ok' or 'degraded'. If WAN links are completely down, confirm array based copy is being made at winning site.

---

**Note:** Best practice is to set the VPLEX consistency group detach rule (winning site) to match site where array based copies are made.

---

3. Follow standard array-based procedure(s) to generate desired array-based copies (i.e SnapView, TimeFinder, XtremIO User or CLI guides). The most up to date documentation for Dell EMC products is available at <https://www.dell.com/support/home/en-us>.

Once the copy is created on the storage array, it can be presented (zoned and lun masked) directly from the storage array back to a test or backup host. Most often, however, the preference is to present the array based copy back through VPLEX to a host. Since a single array contains the copy(s) / snapshot(s), the resulting virtual volume(s) will always be a local virtual volume(s). For VPLEX presented copies, perform these additional VPLEX specific steps:

4. Confirm the array-based copies are visible to VPLEX BE ports. As necessary, perform storage array to VPLEX masking/storage group modification to add the copies.
5. Perform one-to-one encapsulation through the Unisphere for VPLEX UI or VPLEX CLI:
  - a. Claim storage volumes from the storage array containing the copies
  - b. Identify the logical units from the array containing the array-based copies. One way to do this is to use the host lun number assigned by the array to each device in the array masking view (VMAX/XtremIO) or storage group (CX4/VNX).
  - c. Create VPLEX virtual volumes from the copies using the VPLEX UI or CLI:
    - i. Use the 'Create Virtual Volumes' button from the Arrays context within the UI  
or:
    - ii. Create single extent, single member RAID-0 geometry device, and virtual volume(s) with the 'storage-tool-compose' (5.3 and higher) command from the VPLEX CLI
6. Present VPLEX virtual volumes built from array-based copies to host(s)
  - a. As necessary, create VPLEX storage view(s)
  - b. Add virtual volumes built from array-based copies to storage view(s)
  - c. As necessary, perform zoning of virtual volumes to hosts following traditional zoning best practices

### 3.3 Resynchronizing (updating) local array-based copies presented through VPLEX




Figure 3 Resynchronizing a VPLEX RAID-0 copy

This process follows the traditional best practices for refreshing non-VPLEX array based copies such as those provided by products such as TimeFinder, SnapView, and XtremIO and Unity snapshots and native replication. During the resynchronization process, the storage array writes new data to the copy outside of the host and VPLEX IO path. It is very important that the host does not have the copies mounted and in-use.

---

**Note:** Failure to stop applications and to clear both the host read cache and the VPLEX read cache prior to re-synchronization may lead to data consistency issues on the copies.

---



---

**Technical Note:** These steps require using the VPLEX CLI or REST API only.

---

#### 3.3.1 To re-synchronize (update) VPLEX virtual volumes based on array-based copies:

1. Confirm the device(s) you wish to restore is/are not RecoverPoint protected. Issue the 'll' command from the /clusters/<cluster name>/virtual volumes/<virtual volume name> context. Check 'recoverpoint-protection-at' is set to [] and 'recoverpoint-usage' is set to '-'. If the virtual volume is a member of a consistency group, the consistency group context will indicate if it is RecoverPoint protected or not as well.
2. Confirm a VPLEX ndu is not in progress. Using the VPLEX CLI issue the 'ndu status' command and confirm that the response is 'No firmware, BIOS/POST, or SSD upgrade is in progress.'
3. Shut down any applications using the VPLEX virtual volumes that contain the array-based copies to be re-synchronized. Unmount the associated virtual volumes from the host. The key here is to force the host to invalidate its own read cache and to prevent any new writes during the resynchronization of the array-based copy.
4. Invalidate the VPLEX read cache. There are several options to achieve invalidation depending on your VPLEX GeoSynchrony code version:
  - a. For pre-5.2 code, remove the virtual volume(s) from all storage views. Make note of the virtual volume lun numbers within the storage view prior to removing them. You will need this information in step 6 below.

or

b. For 5.2 and higher code:

i. Use virtual-volume cache-invalidate to invalidate an individual volume:

```
VPlexcli:/> virtual-volume cache-invalidate <virtual volume>
```

Or:

Use consistency-group cache-invalidate to invalidate an entire VPLEX consistency group.

```
VPlexcli:/> consistency-group cache-invalidate <consistency group>
```

ii. Follow each command in step 1 with the command virtual-volume cache-invalidate-status to confirm the cache invalidation process has completed.

```
VPlexcli:/> virtual-volume cache-invalidate-status <virtual volume>
```

Example output for a cache invalidation job in progress:

```
cache-invalidate-status
-----
director-1-1-A status: in-progress
                result: -
                cause: -
```

Example output for a cache invalidation job completed successfully:

```
cache-invalidate-status
-----
director-1-1-A status: completed
                result: successful
                cause: -
```

---

**Technical Note 1:** For pre-5.2 code, it is recommended to wait a minimum of 30 seconds to ensure that the VPLEX read cache has been invalidated for each virtual volume. This can be done concurrently with Step 3.

---



---

**Technical Note 2:** The virtual-volume cache-invalidate commands operate on a single virtual-volume at a time. This is the case even when the consistency group command is used. The consistency-group cache-invalidate command will fail if any member virtual volume doesn't invalidate. Invalidation will fail if host I/O is still in progress to the virtual volumes. The invalidation process can take a non-trivial amount of time (usually not more than a few seconds), so the use of the virtual-volume cache-invalidate-status command is recommended to confirm completion of invalidation tasks. With GeoSynchrony 5.4 SP1+ code the cache invalidate command runs in the foreground and does not require a follow up status check. Once the command completes the invalidation is complete.

---



---

**Technical Note 3:** Cache-invalidate command must not be executed on the RecoverPoint enabled virtual-volumes. This means copying should be done using RecoverPoint or using array-based copies, but not both for a given virtual volume. The VPLEX Clusters should not be undergoing a NDU while this command is being executed.

---

5. Identify the source storage volumes within the array you wish to resynchronize. Follow your normal array resynchronization procedure(s) to refresh the desired array-based copies.
6. Confirm the IO Status of storage volumes based on array-based copies is “alive” by doing a long listing against the storage-volumes context for your cluster.

For example:

```
Vplexcli:/> ll /clusters/cluster-1/storage-elements/storage-volumes/
/clusters/cluster-1/storage-elements/storage-volumes:
Name                               VPD83 ID                               Capacity  Use      Vendor  IO Status  Type
-----
Symm1554Tdev_061D                  VPD83T3:60000970000192601554533030363144 100G      used     EMC      alive     normal
Symm1554Tdev_061E                  VPD83T3:60000970000192601554533030363145 100G      claimed  EMC      alive     normal
Symm1554Tdev_061F                  VPD83T3:60000970000192601554533030363146 100G      claimed  EMC      alive     normal
Symm1554Tdev_0620                  VPD83T3:60000970000192601554533030363230 100G      claimed  EMC      alive     normal
Symm1554Tdev_0621                  VPD83T3:60000970000192601554533030363231 100G      claimed  EMC      alive     normal
Symm1554Tdev_0622                  VPD83T3:60000970000192601554533030363232 100G      used     EMC      alive     normal
Symm1554Tdev_0623                  VPD83T3:60000970000192601554533030363233 100G      used     EMC      alive     normal
Symm1554Tdev_0624                  VPD83T3:60000970000192601554533030363234 100G      used     EMC      alive     normal
Symm1554Tdev_0625                  VPD83T3:60000970000192601554533030363235 100G      used     EMC      alive     normal
Symm1554Tdev_0626                  VPD83T3:60000970000192601554533030363236 100G      claimed  EMC      alive     normal
```

7. Confirm VPLEX back-end paths are healthy by issuing the “connectivity validate-be” command from the VPLEX CLI. Ensure that there are no errors or connectivity issues to the back-end storage devices. Resolve any error conditions with the back-end storage before proceeding.

Example output showing desired back-end status:

```
Vplexcli:/clusters/cluster-1> connectivity validate-be
Summary
Cluster cluster-2
  This cluster has 0 storage-volumes which are dead or unreachable
  This cluster has 0 storage-volumes which do not have dual paths
  This cluster has 0 storage-volumes which are not visible from all directors
Cluster cluster-1
  This cluster has 0 storage-volumes which are dead or unreachable
  This cluster has 0 storage-volumes which do not have dual paths
  This cluster has 0 storage-volumes which are not visible from all directors
```

8. For pre-5.2 code, restore access to virtual volumes based on array-based copies for host(s). If you removed the virtual volume from a storage view, add the virtual volume back, specifying the original LUN number (noted in step 2) using the VPLEX CLI:

```
Vplexcli:/> /clusters/<cluster name>/exports/storage-views> addvirtualvolume
-v storage_view_name/ -o (lun#, virtual_volume_name) -f
```

9. As necessary, rescan devices and restore paths (for example, powermt restore) on hosts.
10. As necessary, re-mount VPLEX virtual volumes on each host.
11. Restart applications.

---


**Technical Note:** VPLEX does not alter operating system limitations relative to array-based copies. When presenting an array-based copy back to the same host, you will need to confirm both the host operating system and the logical volume manager support such an operation. If the host OS does not support seeing the same device signature or logical volume manager identifier, VPLEX will not change this situation.

---



## 4 Advanced array-based copy use cases

When an array-based copy (e.g. a Snapshot or a full disk copy) is a component of a VPLEX RAID-1 virtual volume, it is important to properly account for both mirror legs to avoid potential data inconsistency. VPLEX is providing the mirroring, but during the copy and resynchronization process the array is writing to one or more of the mirror legs outside of the VPLEX IO path. The key here is that VPLEX is not able to mirror writes unless they pass through its IO stack. For this reason, an additional step must be added to the array-based copy process to temporarily detach the mirror leg that will not be updated by the array. Once the array-based copy process completes, the final step would be to back the leg that was detached in order to trigger a VPLEX RAID-1 rebuild (as shown in Example 3 below). This step must be done during both the creation and the resynchronization (update) processes.



Example 3 Create VPLEX RAID-1 Copy or Resynchronize VPLEX RAID-1 Copy

This section assumes VPLEX is providing RAID-1 device geometry for the array-based copy. The VPLEX copy virtual volume consists of one mirror leg that contains an array-based copy and a second mirror leg that may (Figure 5) or may not (Figure 4) contain array based copy.

---

**Note:** Array based copies on two different arrays are feasible, but consistency across disparate arrays cannot be assumed due to differences in things such as copy timing and copy performance. That said, some array management software interfaces do have a concept of consistency across arrays. See individual array administrative guides for further details.

---

This configuration applies to both array-based copies within a local RAID-1. For some use cases, it may be desirable to create copies of each mirror leg to protect against the loss of an array or loss of connectivity to an array. In addition, the second copy can be located within the same array or a separate second array. Figure 4 illustrates a local VPLEX device protected using an array-based copy with RAID-1 geometry. device and to array-based (for the non-copy leg) and to ensure each virtual volume's copies that make up distributed (VPLEX Metro) RAID-1 devices. The steps that follow are critical to ensure proper mirror synchronization read cache is properly updated.




Figure 1 RAID-1 array-based copies of VPLEX volumes

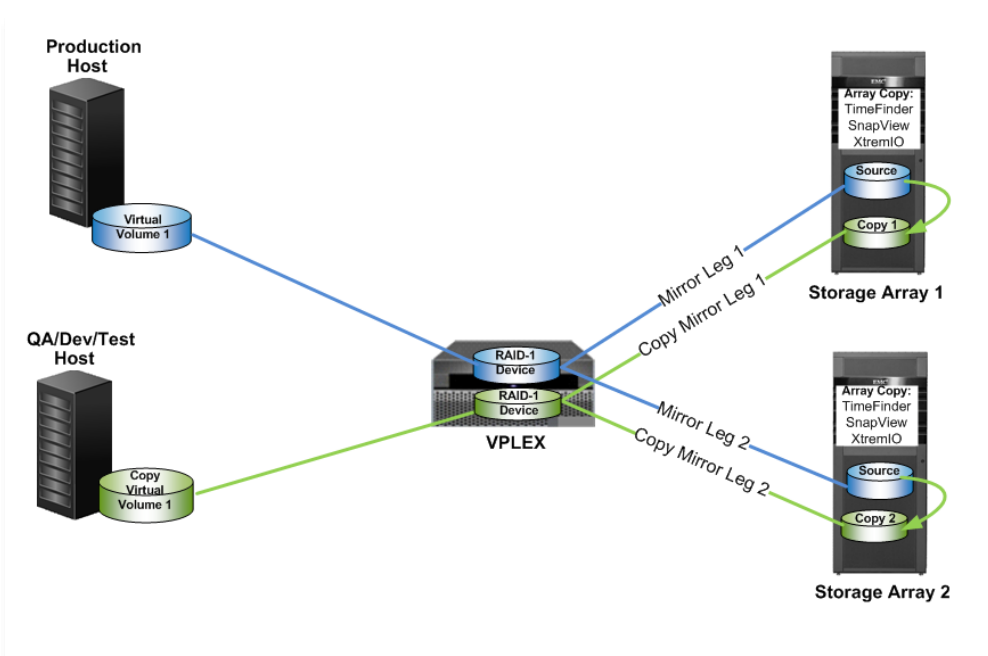


Figure 2 RAID-1 source and array-based Copies of VPLEX volumes

## 4.1 Prerequisites

This section assumes you are using or planning to use distributed or local RAID-1 VPLEX virtual volumes built with at least one mirror leg that is an array-based copy. In addition, the VPLEX virtual volumes must possess both of the following attributes:

- Be comprised of devices that have a 1:1 storage volume pass-through configuration to VPLEX (device capacity = extent capacity = storage volume capacity).
- Have a single device with single-extent RAID-1 (two single-extent devices being mirrored) geometry.

## 4.2 Creating VPLEX RAID-1 or distributed RAID-1 array-based copies

This process augments the traditional best practices for non-VPLEX array-based copies using TimeFinder and SnapView. It accounts for VPLEX and for host access through VPLEX to array-based copies. During the copy process, the storage array will be writing new data to the array-based copy outside of the VPLEX and the host IO path.

### 4.2.1 Procedure: Creating VPLEX RAID-1 or distributed RAID-1 array-based copies

1. For each storage array, identify source storage volumes to copy. The VPLEX storage-view will contain all VPLEX volumes that are visible to a given host or group of hosts. To determine the VPLEX to array device mapping use one of following methods:
  - Using the VPLEX CLI invoke the 'show-use-hierarchy' or 'device drill-down' command for the virtual volume(s) and underlying device(s) you wish to make a copy of. These outputs provide the virtual to physical volume mapping for each VPLEX volume.
  - Use the Storage Viewer feature of Virtual Storage Integrator (VSI) for VMware vCenter Server (for ESX environments) to see the virtual device to physical array device mapping.
  - Correlate the VPLEX virtual volumes to physical storage using the VPDID from the virtual volume to map to the back end storage array device. This information is available to the host via SCSI inquiry and from VPLEX via the /clusters/<cluster name>/exports/storage-view context. To get the back-end array device, you can map virtual volume VPDID back through the VPLEX and get the array LUN. From a VPLEX perspective there is a VPDID tied to the virtual volume the host sees and a different VPDID tied to the back end storage array.
2. Using the VPLEX CLI or REST API, check for valid copy conditions to ensure data consistency:

---

**Note:** Each of the following checks are typically scripted or built into code that orchestrates the overall copy process on the array.

---

- a. Confirm a VPLEX ndu is not in progress. Using the VPLEX CLI issue the 'ndu status' command and confirm that the response is 'No firmware, BIOS/POST, or SSD upgrade is in progress.'
  - b. Confirm the device is healthy. Issue the 'll' command from the /clusters/<cluster name>/virtual-volumes/<virtual volume name> context for each volume(s) to be copied.
    - i. Confirm the underlying device status is not marked 'out of date' or in a 'rebuilding' state.
    - ii. Confirm Health Status is 'ok'
-

- iii. Confirm Operational Status is 'ok'
- c. Confirm the underlying device geometry is not RAID-c. Device geometry can be determined by issuing `ll` at the `/clusters/<cluster name>/devices/<device name>` context.
- d. Confirm each volume is 1:1 mapped (single extent) RAID-0 or 1:1 mapped (two extent) local RAID-1. Distributed RAID-1 device legs must be a combination of RAID-0 (single extent) and/or RAID-1 (two extent) device geometries.
- e. Confirm the device is not being protected by RecoverPoint. Issue the `ll` command from the `/clusters/<cluster name>/virtual-volumes/<virtual volume name>` context and check `'recoverpoint-protection-at'` is set to `[]` and `'recoverpoint-usage'` is set to `'-'`.
- f. Confirm VPLEX volumes to be copied have the same locality (from same VPLEX cluster).
- g. Ensure virtual volumes are members of the same VPLEX consistency group and the same array based consistency group (if available). In most cases all members of the consistency group should be copied together. Consistency group membership can be determined by issuing `ll` from the `/clusters/<cluster name>/consistency-groups/<consistency group name>` context. The array CLI or API should be used to determine the array's consistency group membership.

---

**Note:** VPLEX consistency group membership should align with array based consistency group membership whenever possible.

---

- h. For RAID-1 or distributed RAID-1 based virtual volumes, confirm underlying storage volume status is not failed or in an error state. Issue the `ll` command from the `/clusters/<cluster name>/devices` context or from `/distributed-storage/distributed-devices/<distributed device name>/components` context.
- i. For distributed RAID-1 confirm WAN links are not down. Issue the `'cluster status'` command and confirm `'wan-com'` status is `'ok'` or `'degraded'`. If WAN links are completely down, confirm array based copy is being made at winning site.

---

**Note:** Best practice is to set the consistency group detach rule (winning site) to match site where array based copy is made.

---

- 3. Follow standard array-based procedure(s) to generate desired array-based copies. For example, use the TimeFinder, SnapView, or XtremIO Snapshots cli documentation. The most up to date versions of these documents are available at <https://www.dell.com/support/home/en-us>.

#### Perform VPLEX specific steps:

- 4. Confirm the array-based copy (Clone/Snapshot) devices are visible to VPLEX. As necessary, perform storage array to VPLEX LUN masking and storage array to VPLEX SAN zoning.
- 5. Perform one-to-one encapsulation through the Unisphere for VPLEX UI or VPLEX CLI:
  - a. Claim storage volumes from the storage array containing the copies
  - b. Identify the logical units from the array containing the array-based copies. One way to do this is to use the host lun number assigned by the array to each device in the array masking view (Symmetrix/VMAX) or storage group (CX4/VNX).
  - c. Create VPLEX virtual volumes from the copies using the VPLEX UI or CLI:
    - i. Use the 'Create Virtual Volumes' button from the Arrays context within the UI  
Or:
    - ii. Create single extent, single member RAID-0 geometry device, and virtual volume(s) from the VPLEX CLI

- d. Add local or remote mirrors to virtual volumes created in the previous step using the Unisphere for VPLEX UI or the VPLEX CLI. See <https://www.dell.com/support/home/en-us> for the latest VPLEX CLI and Users guides for further details.
6. Present VPLEX virtual volumes built from array-based copies to host(s)
    - a. As necessary, create VPLEX storage view(s)
    - b. Add virtual volumes built from array-based copies to storage view(s)
    - c. As necessary, perform zoning of virtual volumes to hosts following traditional zoning best practices

### 4.3 Procedure: Resynchronizing (refreshing) array-based copies of RAID-1 VPLEX virtual volumes

1. Confirm the device(s) you wish to resynchronize is/are not RecoverPoint protected. Issue the 'll' command from the /clusters/<cluster name>/virtual volumes/<virtual volume name> context. Check 'recoverpoint-protection-at' is set to [] and 'recoverpoint-usage' is set to '-'. If the virtual volume is a member of a consistency group, the consistency group context will indicate if it is RecoverPoint protected or not as well.
2. Confirm a VPLEX ndu is not in progress. Using the VPLEX CLI issue the 'ndu status' command and confirm that the response is 'No firmware, BIOS/POST, or SSD upgrade is in progress.'
3. Shut down any applications accessing the VPLEX virtual volumes that contain the array-based copies to be resynchronized. This must be done at both clusters if the virtual volumes are being accessed across multiple sites. If necessary, unmount the associated virtual volumes from the host(s). The key here is to force the hosts to invalidate their read cache and to prevent any new writes during the resynchronization of the array-based copy.
4. Invalidate the VPLEX read cache. There are two options to achieve invalidation depending on your VPLEX GeoSynchrony code version:
  - a. For pre-5.2 code, remove the virtual volume(s) from all storage views.
    - i. If the virtual volume is built from a local RAID-1 device and/or is a member of a single storage view, using the VPLEX CLI run:

```
Vplexcli:/> /clusters/<cluster name>/exports/storage-views>
removevirtualvolume -v storage_view_name -o virtual_volume_name -f
```

- ii. If the virtual volume is built from a distributed device and is a member of storage views in both clusters, using the VPLEX CLI run:

```
Vplexcli:/> /clusters/<local cluster name>/exports/storage-views>
removevirtualvolume -v storage_view_name -o distributed_device_name_vol -f
```

```
Vplexcli:/> /clusters/<remote cluster name>/exports/storage-views>
removevirtualvolume -v storage_view_name -o distributed_device_name_vol -f
```

Make note of the virtual volume lun numbers within the VPLEX storage view prior to removing them. You will need this information in step 7 below.

Or:

b. For 5.2 and higher code,

i. Use virtual-volume cache-invalidate to invalidate an individual volume:

```
VPlexcli:/> virtual-volume cache-invalidate <virtual volume>
```

Or:

Use consistency-group cache-invalidate to invalidate an entire VPLEX consistency group.

```
VPlexcli:/> consistency-group cache-invalidate <consistency group>
```

ii. Follow each command in step 1 with the command virtual-volume cache-invalidate-status to confirm the cache invalidation process has completed.

```
VPlexcli:/> virtual-volume cache-invalidate-status <virtual volume>
```

Example output for a cache invalidation job in progress:

```
cache-invalidate-status
-----
director-1-1-A    status: in-progress
result: -
cause: -
```

Example output for a cache invalidation job completed successfully:

```
cache-invalidate-status
-----
director-1-1-A    status: completed
result: successful
cause: -
```

---

**Technical Note 1:** For pre-5.2 code, it is recommended to wait a minimum of 30 seconds to ensure that the VPLEX read cache has been invalidated for each virtual volume. This can be done concurrently with Step 3.

---

---

**Technical Note 2:** The virtual-volume cache-invalidate commands operate on a single virtual-volume at a time. This is the case even when the consistency group command is used. The consistency-group invalidate-cache command will fail if any member virtual volume hasn't been invalidated. Invalidation will fail if host I/O is still in progress to the virtual volumes. The invalidation process can take a non-trivial amount of time, so the use of the virtual-volume cache-invalidate-status command is recommended to confirm completion of

---

invalidation tasks. With GeoSynchrony 5.4 SP1+ code the cache invalidate command runs in the foreground and does not require a follow up status check. Once the command completes the invalidation is complete.

**Technical Note 3:** Cache-invalidate command must not be executed on the Recover Point enabled virtual-volumes. This means using either RecoverPoint or array-based copies, but not both for any given virtual volume. The VPLEX Clusters should not be undergoing a NDU while this command is being executed.

5. Detach the VPLEX device mirror leg that will not be updated during the array-based replication or resynchronization processes:

```
device detach-mirror -m <device_mirror_to_detach> -d
<distributed_device_name> -i -f
```

6. Perform resynchronization (update) of array-based copies using standard TimeFinder, SRDF, SnapView, MirrorView, or XtremIO administrative procedures. See <https://www.dell.com/support/home/en-us> for the latest copies of Dell EMC product documentation.
7. Confirm the IO Status of storage volumes based on array-based copies is “alive” by doing a long listing against the storage-volumes context for your cluster.

For example:

```
VPlexcli:/> ll /clusters/cluster-1/storage-elements/storage-volumes/
/clusters/cluster-1/storage-elements/storage-volumes:
Name                               VPD83 ID                               Capacity Use      Vendor  IO Status  Type
-----
Symm1554Tdev_061D                  VPD83T3:60000970000192.601554533030363144 100G    used     EMC     alive     normal
Symm1554Tdev_061E                  VPD83T3:60000970000192.601554533030363145 100G    claimed  EMC     alive     normal
Symm1554Tdev_061F                  VPD83T3:60000970000192.601554533030363146 100G    claimed  EMC     alive     normal
Symm1554Tdev_0620                  VPD83T3:60000970000192.601554533030363230 100G    claimed  EMC     alive     normal
Symm1554Tdev_0621                  VPD83T3:60000970000192.601554533030363231 100G    claimed  EMC     alive     normal
Symm1554Tdev_0622                  VPD83T3:60000970000192.601554533030363232 100G    used     EMC     alive     normal
Symm1554Tdev_0623                  VPD83T3:60000970000192.601554533030363233 100G    used     EMC     alive     normal
Symm1554Tdev_0624                  VPD83T3:60000970000192.601554533030363234 100G    used     EMC     alive     normal
Symm1554Tdev_0625                  VPD83T3:60000970000192.601554533030363235 100G    used     EMC     alive     normal
Symm1554Tdev_0626                  VPD83T3:60000970000192.601554533030363236 100G    claimed  EMC     alive     normal
```

In addition, confirm VPLEX back-end paths are healthy by issuing the “connectivity validate-be” command from the VPLEX CLI. Ensure that there are no errors or connectivity issues to the back-end storage devices. Resolve any error conditions with the back-end storage before proceeding.

Example output showing desired back-end status:

```
VPlexcli:/clusters/cluster-1> connectivity validate-be
Summary
Cluster cluster-2
  This cluster has 0 storage-volumes which are dead or unreachable
  This cluster has 0 storage-volumes which do not have dual paths
  This cluster has 0 storage-volumes which are not visible from all directors
Cluster cluster-1
  This cluster has 0 storage-volumes which are dead or unreachable
  This cluster has 0 storage-volumes which do not have dual paths
  This cluster has 0 storage-volumes which are not visible from all directors
```

8. Reattach the second mirror leg:

RAID-1:

```
device attach-mirror -m <2nd mirror leg to attach> -d /clusters/<local cluster
name>/devices/<existing RAID-1 device>
```

**Distributed RAID-1:**

```
device attach-mirror -m <2nd mirror leg to attach> -d /clusters/<cluster
name>/devices/<existing distributed RAID-1 device>
```

---

**Technical Note:** The device you are attaching is the non-copy mirror leg. It will be overwritten with the data from the copy. Depending on the amount of data being resynchronized this process can take quite some time. This should be factored into the RTO consideration.

---

9. For pre-5.2 GeoSynchrony code, restore host access to VPLEX volume(s).


If the virtual volume is built from a local RAID 1 device:

```
/clusters/<local cluster name>/exports/storage-views> addvirtualvolume -v
storage_view_name/ -o (lun#,device_Symm0191_065_1_vol/) -f
```

If the virtual volume is built from a distributed RAID 1 device:

```
/clusters/<remote cluster name>/exports/storage-views> addvirtualvolume -v
storage_view_name/ -o (lun#, distributed_device_name_vol)-f
```

```
/clusters/<local cluster name>/exports/storage-views> addvirtualvolume -v
storage_view_name/ -o (lun#, distributed_device_name_vol)-f
```

The lun# is the previously recorded value from step 2 for each virtual volume.

---

**Technical Note 1:** For pre-5.2 GeoSynchrony code, Dell EMC recommends waiting at least 30 seconds after removing access from a storage view to restore access. Waiting ensures that the VPLEX cache has been cleared for the volumes. The array-based resynchronization will likely will take 30 seconds, but if you are scripting, be sure to add a pause prior to performing this step.

---



---

**Technical Note 2:** Some hosts and applications are sensitive to LUN numbering changes. Use the information you recorded in step 2 to ensure the same LUN numbering for the host(s) when the virtual volume access is restored.

---



---

**Technical Note 3:** It is not necessary to perform full mirror synchronization prior to restoring access to virtual volumes. VPLEX will synchronize the second mirror leg in the background while using the first mirror leg as necessary to service reads to any unsynchronized blocks.

---

10. Rescan devices and restore paths (powermt restore) on hosts.
11. Mount devices (if mounts are used).
12. Restart applications.



## 5 Array-based restore for VPLEX RAID-0 production volumes

The array-based restore or recovery process with VPLEX is similar to the previous array-based resynchronization use cases. The primary difference is that the array is now writing data from the copy/snapshot back to the production or source volume. In this section we review the array-based restore process. It is assumed the array-based copy / snapshot is located on the array containing the source device (restore target). For the basic local restore use case, each source (restore target) device must have a one-to-one storage volume pass-through configuration (device capacity = extent capacity = storage volume capacity) through VPLEX and a RAID-0 (single extent only) device geometry. This is the most basic restore use case for restore and recovery. Example 4 illustrates the case where data is being written within a storage array from a copy (Snapshot or Full disk copy) or backup media to a storage volume (standard device) that is used by VPLEX.




Figure 4 Example 4 Restore From Array-based Copy to VPLEX RAID-0 Production Volume

### 5.1 Procedure: Array-based restore to VPLEX RAID-0 production virtual volumes

1. Confirm the production storage volume(s) you wish to restore is/are not RecoverPoint protected. Issue the 'll' command from the /clusters/<cluster name>/virtual volumes/<virtual volume name> context. Check 'recoverpoint-protection-at' is set to [] and 'recoverpoint-usage' is set to '-'. If the virtual volume is a member of a consistency group, the consistency group context will indicate if it is RecoverPoint protected or not as well.
2. Confirm a VPLEX ndu is not in progress. Using the VPLEX CLI issue the 'ndu status' command and confirm that the response is 'No firmware, BIOS/POST, or SSD upgrade is in progress.'
3. Confirm consistency group membership and dependent virtual volume relationships. In other words, confirm that you are restoring all of the volumes that the application requires to run and not a subset of all volumes.
4. Shut down any host applications using the source VPLEX volume(s) that will be restored. As necessary, unmount the associated virtual volumes on the host. The objectives here are to prevent host access and to clear the host's read cache.
5. Invalidate VPLEX read cache on the source virtual volume(s). There are several options to achieve VPLEX read cache invalidation depending on your VPLEX GeoSynchrony code version:
  - a. For pre-5.2 code, remove the source virtual volume(s) from all storage views. Make note of the virtual volume lun numbers within the storage view prior to removing them. You will need this information in step 5 below.

Or:

  - b. For 5.2 and higher code:

- i. Use virtual-volume cache-invalidate to invalidate the individual source volume(s):

```
Vplexcli:/> virtual-volume cache-invalidate <virtual volume>
```

Or:

Use consistency-group cache-invalidate to invalidate an entire VPLEX consistency group of source volumes.

```
Vplexcli:/> consistency-group cache-invalidate <consistency group>
```

- ii. Follow each command in step 1 with the command virtual-volume cache-invalidate-status to confirm the cache invalidation process has completed.

```
Vplexcli:/> virtual-volume cache-invalidate-status <virtual volume>
```

Example output for a cache invalidation job in progress:

```
cache-invalidate-status
-----
director-1-1-A    status: in-progress
result: -
                 cause: -
```

Example output for a cache invalidation job completed successfully:

```
cache-invalidate-status
-----
director-1-1-A    status: completed
result: successful
                 cause: -
```

---

**Technical Note 1:** For pre-5.2 code, it is recommended to wait a minimum of 30 seconds to ensure that the VPLEX read cache has been invalidated for each virtual volume. This can be done concurrently with Step 3.

---

---

**Technical Note 2:** The virtual-volume cache-invalidate commands operate on a single virtual-volume at a time. This is the case even when the consistency group command is used. The consistency-group cache-invalidate command will fail if any member virtual volume doesn't invalidate. Invalidation will fail if host I/O is still in progress to the virtual volumes. The invalidation process can take a non-trivial amount of time, so the use of the virtual-volume cache-invalidate-status command is recommended to confirm completion of invalidation tasks. With GeoSynchrony 5.4 SP1+ code the cache invalidate command runs in the foreground and does not require a follow up status check. Once the command completes the invalidation is complete.

---

---

**Technical Note 3:** Cache-invalidate command must not be executed on the RecoverPoint enabled virtual volumes. This means using either RecoverPoint or array-based copies, but not both, for any given virtual volume. The VPLEX Clusters should not be undergoing a NDU while this command is being executed.

---

6. Identify the copy (BCV/Clone/Snapshot) to source device pairings within the array(s). For Dell EMC products, follow the TimeFinder, SnapView, SRDF, MirrorView, XtremIO and Unity Snapshot restore procedure(s) to restore data to the desired source devices. See <https://www.dell.com/support/home/en-us> for the latest Dell EMC product documentation for TimeFinder, SnapView, SRDF, MirrorView, or XtremIO and Unity Snapshots.
7. Confirm the IO Status of the source storage volumes within VPLEX is “alive” by doing a long listing against the storage volumes context for your cluster.

For example:

```

Vplexcli:/> ll /clusters/cluster-1/storage-elements/storage-volumes/
/clusters/cluster-1/storage-elements/storage-volumes:
Name                               VPD83 ID                               Capacity Use      Vendor  IO Status  Type
-----
Symm1554Tdev_061D                   VPD83T3:60000970000192601554533030363144 100G    used     EMC     alive     normal
Symm1554Tdev_061E                   VPD83T3:60000970000192601554533030363145 100G    claimed  EMC     alive     normal
Symm1554Tdev_061F                   VPD83T3:60000970000192601554533030363146 100G    claimed  EMC     alive     normal
Symm1554Tdev_0620                   VPD83T3:60000970000192601554533030363230 100G    claimed  EMC     alive     normal
Symm1554Tdev_0621                   VPD83T3:60000970000192601554533030363231 100G    claimed  EMC     alive     normal
Symm1554Tdev_0622                   VPD83T3:60000970000192601554533030363232 100G    used     EMC     alive     normal
Symm1554Tdev_0623                   VPD83T3:60000970000192601554533030363233 100G    used     EMC     alive     normal
Symm1554Tdev_0624                   VPD83T3:60000970000192601554533030363234 100G    used     EMC     alive     normal
Symm1554Tdev_0625                   VPD83T3:60000970000192601554533030363235 100G    used     EMC     alive     normal
Symm1554Tdev_0626                   VPD83T3:60000970000192601554533030363236 100G    claimed  EMC     alive     normal

```

In addition, confirm VPLEX back-end paths are healthy by issuing the “connectivity validate-be” command from the VPLEX CLI. Ensure that there are no errors or connectivity issues to the back-end storage devices. Resolve any error conditions with the back-end storage before proceeding.

Example output showing desired back-end status:

```

Vplexcli:/clusters/cluster-1> connectivity validate-be
Summary
Cluster cluster-2
  This cluster has 0 storage-volumes which are dead or unreachable
  This cluster has 0 storage-volumes which do not have dual paths
  This cluster has 0 storage-volumes which are not visible from all directors
Cluster cluster-1
  This cluster has 0 storage-volumes which are dead or unreachable
  This cluster has 0 storage-volumes which do not have dual paths
  This cluster has 0 storage-volumes which are not visible from all directors

```

8. For Pre 5.2 code, restore access to virtual volume(s) based on source devices for host(s): Add the virtual volume back to the view, specifying the original LUN number (noted in step 2) using VPLEX CLI:

```

Vplexcli:/> /clusters/<cluster name>/exports/storage-views> addvirtualvolume
-v storage_view_name/ -o (lun#, virtual_volume_name) -f

```

9. As necessary, rescan devices and restore paths (for example, powermt restore) on hosts.

10. As necessary, mount devices.
11. Restart applications.

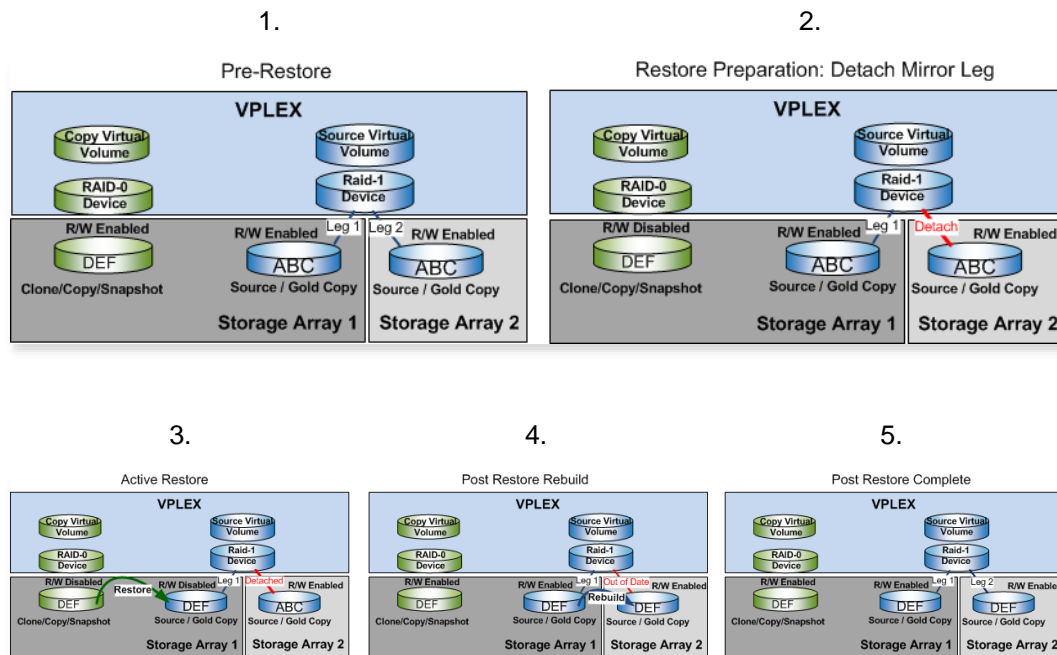
---

**Note:** Even though applications are able to run during resync, the standard RTO is not delivered until both mirror legs are completely synchronize.

---

## 5.2 Array-based restore of VPLEX RAID-1 production volumes

When VPLEX virtual volumes have RAID-1 geometry, the restore process must take into account the second (non-restored) mirror leg. This applies for both local RAID-1 and distributed (VPLEX Metro) RAID-1 VPLEX devices. The typical array-based source device restore procedure will only restore one of the two mirror legs of a VPLEX RAID 1 device. In order to synchronize the second VPLEX device, users need to add a step to resynchronize the 2nd mirror leg. These steps are critical to ensure proper synchronization of the second VPLEX RAID-1 mirror leg (the one that is not being restored by the array) and to ensure each virtual volume's read cache is properly updated. Example 5 (below) illustrates the 5-step workflow when an array-based copy is used to restore to a RAID-1 or distributed RAID-1 source volume.



Example 5 Array-based Restore of VPLEX RAID-1 Production Volumes

---

**Technical Note:** This same set of steps can be applied to remote array-based copy products like SRDF or MirrorView. For example, an SRDF R2 or MirrorView Secondary Image is essentially identical in function to local array-based copy. The remote copy, in this case, can be used to do a restore to a production (R1/Primary Image) volume.

---

### 5.2.1 Prerequisites:

This section assumes users have existing distributed or local RAID-1 VPLEX virtual volumes built from the array source devices being restored. In addition, the VPLEX virtual volumes must possess both of the following attributes:

- The volumes must comprise of devices that have a one-to-one storage volume pass-through VPLEX configuration (device capacity = extent capacity = storage volume capacity).
- The volumes must have a single-extent RAID-1 (two single extents being mirrored) geometry. Distributed, local, and distributed on top of local RAID-1 all meet this pre-requisite.

## 5.2.2 Procedure: Array-based restore to VPLEX RAID-1 production virtual volumes

1. Using the VPLEX CLI or REST API, check for valid restore conditions to ensure data consistency:

---

**Note:** Each of the following checks are typically scripted or built into code that orchestrates the overall restore process on the array.

---

- a. Confirm a VPLEX ndu is not in progress. Using the VPLEX CLI issue the 'ndu status' command and confirm that the response is 'No firmware, BIOS/POST, or SSD upgrade is in progress.'
- b. Confirm the restore target device is healthy. Issue the 'll' command from the /clusters/<cluster name>/virtual-volumes/<virtual volume name> context for each volume(s) to be copied.
  - i. Confirm the underlying device status is not marked 'out of date' or in a 'rebuilding' state.
  - ii. Confirm Health Status is 'ok'
  - iii. Confirm Operational Status is 'ok'
- c. Confirm the underlying restore target device geometry is not RAID-c. Device geometry can be determined by issuing ll at the /clusters/<cluster name>/devices/<device name> context.
- d. Confirm each volume is 1:1 mapped (single extent) RAID-0 or 1:1 mapped (two extent) local RAID-1. Distributed RAID-1 device legs must be a combination of RAID-0 (single extent) and/or RAID-1 (two extent) device geometries.
- e. Confirm the restore target device is not being protected by RecoverPoint. Issue the 'll' command from the /clusters/<cluster name>/virtual-volumes/<virtual volume name> context and check 'recoverpoint-protection-at' is set to [] and 'recoverpoint-usage' is set to '-'.
- f. Confirm VPLEX volumes to be restore have the same locality (from same VPLEX cluster) and are in the same array.
- g. Ensure virtual volumes are members of the same VPLEX consistency group and, when possible, part of a similar array-based consistency group construct within the array. In most cases all members of the consistency should be copied together. Consistency group membership can be determined by issuing ll from the /clusters/<cluster name>/consistency-groups/<consistency group name> context. Array based consistency group membership can be determined using the appropriate array cli and/or api.
- h. For RAID-1 or distributed RAID-1 based virtual volumes, confirm underlying storage volume status is not failed or in an error state. Issue the 'll' command from the /clusters/<cluster name>/devices context or from /distributed-storage/distributed-devices/<distributed device name>/components context.
- i. For distributed RAID-1 confirm WAN links are not down. Issue the 'cluster status' command and confirm 'wan-com' status is 'ok' or 'degraded'. If WAN links are completely down, confirm array based restore is being done at winning site.

---

**Note:** Best practice is to set the consistency group detach rule (winning site) to match site where array based restore is being done.

---

2. Shut down any host applications (both local and remote with DR1) using the source VPLEX volume(s) that will be restored. Unmount the associated virtual volumes on the host. The objectives here are to prevent host access and to clear the host's read cache.
3. Invalidate VPLEX read cache on the source virtual volume(s). There are several options to achieve VPLEX read cache invalidation depending on your VPLEX GeoSynchrony code version:
  - a. For pre-5.2 code, remove the source virtual volume(s) from all storage views. Make note of the virtual volume lun numbers within the storage view prior to removing them. You will need this information in step 7 below.

Or:

- b. For 5.2 and higher code,
  - i. Use virtual-volume cache-invalidate to invalidate the individual source volume(s):

```
Vplexcli:/> virtual-volume cache-invalidate <virtual volume>
```

Or:

Use consistency-group cache-invalidate to invalidate an entire VPLEX consistency group of source volumes.

```
Vplexcli:/> consistency-group cache-invalidate <consistency group>
```

- ii. Follow each command in step 1 with the command virtual-volume cache-invalidate-status to confirm the cache invalidation process has completed.

```
Vplexcli:/> virtual-volume cache-invalidate-status <virtual volume>
```

Example output for a cache invalidation job in progress:

```
cache-invalidate-status
-----
director-1-1-A    status: in-progress
result: -
cause: -
```

Example output for a cache invalidation job completed successfully:

```
cache-invalidate-status
-----
director-1-1-A    status: completed
result: successful
cause: -
```

---

**Technical Note 1:** For pre-5.2 code, it is recommended to wait a minimum of 30 seconds to ensure that the VPLEX read cache has been invalidated for each virtual volume. This can be done concurrently with Step 3.

---

**Technical Note 2:** The virtual-volume cache-invalidate commands operate on a single virtual-volume at a time. This is the case even when the consistency group command is used. The consistency-group invalidate-cache command will fail if any member virtual volume doesn't invalidate. Invalidation will fail if host I/O is still in progress to the virtual volumes. The invalidation process can take a non-trivial amount of time, so the use of the virtual-volume cache-invalidate-status command is recommended to confirm completion of invalidation tasks. With GeoSynchrony 5.4 SP1+ code the cache invalidate command runs in the foreground and does not require a follow up status check. Once the command completes the invalidation is complete.

**Technical Note 3:** Cache-invalidate command must not be executed on the Recover Point enabled virtual-volumes. This means using either RecoverPoint or array-based copies, but not both, for a given virtual volume. The VPLEX Clusters should not be undergoing a NDU while this command is being executed.

4. Detach the VPLEX device RAID-1 or Distributed RAID-1 mirror leg that will not be restored during the array-based restore processes. If the virtual volume is a member of a consistency group, in some cases the virtual volume may no longer have storage at one site which may cause the detach command to fail. In this case the virtual volume will need to be removed from the consistency group \*before\* it the mirror leg is detached. Use the detach-mirror command to detach the mirror leg(s):

```
device detach-mirror -m <device_mirror_to_detach> -d
<distributed_device_name> -i -f
```

**Note:** Depending on the raid geometry for each leg of the distributed device, it may be necessary to detach both the local mirror leg and the remote mirror leg. This is because only 1 storage volume is being restored and there are up to 3 additional mirrored copies maintained by VPLEX (1 local and 1 or 2 remote). For example, if the VPLEX distributed device mirror leg being restored is, itself, a RAID-1 device then both the non-restored local leg and the remote leg must be detached.

5. Identify the copy (BCV/Clone/Snapshot) to source device pairings within the array(s). Follow the TimeFinder, SnapView, SRDF, MirrorView, XtremIO Snapshot restore procedure(s) to restore data to the desired source devices. See <https://www.dell.com/support/home/en-us> for the latest Dell EMC product documentation for TimeFinder, SnapView, SRDF, MirrorView, or XtremIO for more details.
6. Confirm the IO Status of storage volumes based on array-based clones is "alive" by doing a long listing against the storage volumes context for your cluster.

For example:

```
VPlexcli1: /clusters/cluster-1/storage-elements/storage-volumes/
/custers/cluster-1/storage-elements/storage-volumes:
Name                               VPD83 ID                               Capacity Use      Vendor  IO Status  Type
-----
Symm1554Tdev_061D                  VPD83T3: 60000970000192.6015545330303 63144 100% used    EMC      alive     normal
Symm1554Tdev_061E                  VPD83T3: 60000970000192.6015545330303 63145 100% claimed EMC      alive     normal
Symm1554Tdev_061F                  VPD83T3: 60000970000192.6015545330303 63146 100% claimed EMC      alive     normal
Symm1554Tdev_0620                  VPD83T3: 60000970000192.6015545330303 63230 100% claimed EMC      alive     normal
Symm1554Tdev_0621                  VPD83T3: 60000970000192.6015545330303 63231 100% claimed EMC      alive     normal
Symm1554Tdev_0622                  VPD83T3: 60000970000192.6015545330303 63232 100% used    EMC      alive     normal
Symm1554Tdev_0623                  VPD83T3: 60000970000192.6015545330303 63233 100% used    EMC      alive     normal
Symm1554Tdev_0624                  VPD83T3: 60000970000192.6015545330303 63234 100% used    EMC      alive     normal
Symm1554Tdev_0625                  VPD83T3: 60000970000192.6015545330303 63235 100% used    EMC      alive     normal
Symm1554Tdev_0626                  VPD83T3: 60000970000192.6015545330303 63236 100% claimed EMC      alive     normal
```

In addition, confirm VPLEX back-end paths are healthy by issuing the "connectivity validate-be" command from the VPLEX CLI. Ensure that there are no errors or connectivity issues to the back-end storage devices. Resolve any error conditions with the back-end storage before proceeding.

Example output showing desired back-end status:

```

Vplexcli:/clusters/cluster-1> connectivity validate-be
Summary
Cluster cluster-2
  This cluster has 0 storage-volumes which are dead or unreachable
  This cluster has 0 storage-volumes which do not have dual paths
  This cluster has 0 storage-volumes which are not visible from all directors
Cluster cluster-1
  This cluster has 0 storage-volumes which are dead or unreachable
  This cluster has 0 storage-volumes which do not have dual paths
  This cluster has 0 storage-volumes which are not visible from all directors

```

#### 7. Reattach the second mirror leg:

##### RAID-1:

```
device attach-mirror -m <2nd mirror leg to attach> -d /clusters/<local cluster name>/devices/<existing RAID-1 device>
```

##### Distributed RAID-1:

```
device attach-mirror -m <2nd mirror leg to attach> -d /clusters/<cluster name>/devices/<existing distributed RAID-1 device>
```

---

**Note:** The device you are attaching in this step will be overwritten with the data from the newly restored source device.

---

#### 8. For Pre 5.2 code, restore host access to the VPLEX volume(s).

If the virtual volume is built from a local RAID 1 device:

```
/clusters/<local cluster name>/exports/storage-views> addvirtualvolume -v storage_view_name/ -o (lun#,device_Symm0191_065_1_vol/) -f
```

If the virtual volume is built from a distributed RAID 1 device:

```
/clusters/<remote cluster name>/exports/storage-views> addvirtualvolume -v storage_view_name/ -o (lun#, distributed_device_name_vol)-f
```

```
/clusters/<local cluster name>/exports/storage-views> addvirtualvolume -v storage_view_name/ -o (lun#, distributed_device_name_vol)-f
```

The `lun#` is the previously recorded value from step 2 for each virtual volume.

---

**Technical Note:** Dell EMC recommends waiting at least 30 seconds after removing access from a storage view to restore access. This is done to ensure that the VPLEX cache has been cleared for the volumes. The array-based restore will likely will take 30 seconds, but if you are scripting be sure to add a pause.

---



---

**Technical Note:** Some hosts and applications are sensitive to LUN numbering changes. Use the information you recorded in step 3 to ensure the same LUN numbering when you restore the virtual volume access.

---

## 6 Metro node considerations

Metro node is an add-on feature for both PowerStore and Unity XT. Please reference the PowerStore or Unity XT documentation for the general best practices for metro node.

---

Note: Metro node as mentioned here is built on the R640 platform and the code starts at 7.0. This release is sold specifically as a greenfield deployment and the feature is purchased along with the purchase of a PowerStore or Unity XT array. If VPLEX VS6 is added to an existing environment, it may also be referred to as metro node but it is based on GeoSynchrony 6.x code and this section does not apply.

---

The metro node solution is based on GeoSynchrony 7.0 code but has been modified for use on the R640 hardware platform. Starting with 7.0, the internal caching has been removed so the solution is strictly read/write through only. Additionally, 7.0 has simplified device creation allowing only 1:1 Virtual Volume to Storage Volume mapping.

The combination of these two feature changes align to more easily support native array replication without having to take additional steps to invalidate cache. The process now is the same as direct host attach to the storage array. While the caching has been removed, commands such as cache invalidate will still return a command complete successful even though it didn't have to do anything.

RAID-1 devices for both Local and Distributed must still have the remote mirror leg(s) removed when the source mirror leg is being used as the target of a restore operation. The mirror leg(s) is/are then reattached after the restore has completed. This step is necessary in order to fully resynchronize the data with the mirror leg(s) after the restore.

RecoverPoint splitter is no longer integrated with the metro node code but RecoverPoint is supported at the back-end array level under metro node. Restoring a RecoverPoint protected device has the same considerations as array level replication services mentioned here.

## 7 PowerStore volume migration

The PowerStore hardware platform is created with a focus on performance, reliability, and modular flexibility. Designed from the ground up to leverage next-generation innovations such as NVMe and Intel® Optane™ Storage Class Memory (SCM) to deliver higher IOPS at lower latencies for real-world workloads. This offers all the performance and expansion headroom that is needed to ensure long-term value through multiple solution life cycles. PowerStore helps simplify and consolidate diverse infrastructures with high-performance multi-protocol network I/O options and single-architecture support for block, file, and VMware® vSphere® Virtual Volumes™ (vVols), transforming operations for both traditional and modern workloads.



Figure 5 PowerStore Appliance

PowerStore volume migration has the same considerations whether it is behind a VPLEX or a metro node. Please understand that metro node is implied when VPLEX is mentioned within this section.

PowerStore “T” model has the unique capability of clustering appliances together to help balance the workload across the appliances. This is done by the software automatically selecting which appliance is most optimal to provision a volume from. This will help spread the I/O out evenly. You do have the option of manually selecting which appliance to provision it on. Also, PowerStore can move volumes between appliances. A target appliance is selected to migrate the selected volume to but the target volume is inaccessible until the data transfer has completed. This presents a problem to VPLEX as the volume may retain its original identity but the paths accessing it have changed. Anything that changes outside of VPLEX control is not immediately recognized by VPLEX therefore you are in a potential DU situation, at least temporarily.

The storage behind a PowerStore Appliance is specific to that appliance only. It is not accessible from another appliance until migrated. VPLEX will not pick up the migrated volume automatically therefore you must perform an array rediscover after the migration. This approach would require a host outage as the device would be inaccessible during this phase.

---

Note: A better approach would be to provision a new device from the target appliance and use the VPLEX Data Mobility feature to move the data to that new volume. This approach would be completely non-disruptive. You would then follow the teardown procedure to remove the source device on the original appliance to free up the capacity.

---

## 8 Conclusion

Array-based replication continues to provide business value when Dell EMC VPLEX or metro node is added to your existing storage infrastructure. While modifications to existing procedures may be necessary, the changes are relatively simple to implement. In addition to using array-based copy technologies, VPLEX and metro node customers can also use native mirroring capabilities within VPLEX or metro node to create full disk copies of existing virtual volumes. These copies can be done within and across heterogeneous storage products in a single datacenter or across datacenters. This is a truly differentiating capability that individual frame based copy technologies cannot deliver. VPLEX and metro node UI, CLI, and RESTful API as well as Ansible provide a variety of methods to continue to deliver high business value for array based copy technologies.

## A 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 documents and videos](#) provide expertise to ensure customer success with Dell EMC storage and data protection products.