

EMC RECOVERPOINT ON-DEMAND OPERATIONAL RECOVERY WITH EMC VPLEX

Abstract

This white paper discusses using EMC® RecoverPoint continuous data protection (CDP), remote replication (CRR), and Concurrent Local and Remote (CLR) data protection with EMC VPLEX™ Local and EMC VPLEX Metro. RecoverPoint adds local and remote point-in-time recovery capability to VPLEX Local and Metro solutions.

February 2012

Copyright © 2012 EMC Corporation. All Rights Reserved.

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

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

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

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

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

Part Number h8146

Table of Contents

Executive summary	4
Document scope and limitations	4
Introduction	5
Audience.....	5
VPLEX terminology.....	5
EMC VPLEX Technology	6
EMC VPLEX architecture.....	7
EMC VPLEX Family Overview	8
VPLEX architecture	8
VPLEX VS2 hardware platform	8
EMC RecoverPoint technology	9
RecoverPoint architecture.....	10
RecoverPoint appliance.....	10
Solution Topologies for EMC VPLEX and RecoverPoint Array-Based Splitters	13
Single-Site VPLEX and RecoverPoint CDP Topology	13
Two-site VPLEX and RecoverPoint CDP Topology	14
Replicating VPLEX virtual volumes with RecoverPoint CDP	14
Restoring VPLEX virtual volumes with RecoverPoint CDP	16
Replicating VPLEX virtual volumes with RecoverPoint CRR	19
Restoring VPLEX virtual volumes with RecoverPoint CRR	21
Replicating and Restoring VPLEX virtual volumes with RecoverPoint CLR	24
VPLEX data mobility considerations with RecoverPoint	25
Conclusion	26
References	26

Executive summary

The 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 and Geo provide data access and mobility between two VPLEX clusters within synchronous and asynchronous distances respectively. With a unique scale-up architecture, VPLEX's advanced data caching and distributed cache coherency provide workload resiliency, automatic sharing, and balancing and failover of storage domains, and enable both local and remote data access with predictable service levels.

EMC® RecoverPoint is an enterprise-scale solution designed to protect application data on heterogeneous SAN-attached servers and storage arrays. RecoverPoint runs on a dedicated appliance and combines industry-leading continuous data protection technology with a bandwidth-efficient, no-data-loss replication technology, allowing it to protect data both locally and remotely.

Innovative data change journaling and application integration capabilities enable customers to address their pressing business, operations, and regulatory data protection concerns. Customers implementing RecoverPoint will see dramatic improvements in application protection and recovery times as compared to traditional host and array snapshots or disk-to-tape backup products.

This white paper is designed to give technology decision-makers a deeper understanding of VPLEX and RecoverPoint design, features, and functionality. This paper highlights the key technical considerations for implementing RecoverPoint with VPLEX technology to achieve local and/or remote point in time operational recovery. The integration of the two products results in a solution that builds on the strengths of RecoverPoint and VPLEX to provide the highest levels of protection and capability for today's internal and external enterprise storage environments.

Document scope and limitations

This document applies to EMC VPLEX Local and Metro with RecoverPoint CDP, CLR, and CRR running with supported array-based splitters. The details provided in this white paper are based on the following configurations:

- RecoverPoint 3.4.1 and higher
- CLARiiON® CX4, VNX (FLARE® OS 04.29.000.5.012) or higher
- VMAXe (Enginuity 5875 and higher)
- CX4, VNX, or VMAXe array based splitter
- VPLEX Local and Metro (VPLEX Geo is not supported)

The use cases and procedures described are applicable for new and existing VPLEX customers planning to implement RecoverPoint or existing RP customers planning to implement VPLEX.

Please consult with your local EMC Support representative if you are uncertain as to the applicability of these procedures to your VPLEX environment.

Introduction

Today's businesses are faced with an ever-increasing amount of data that threatens to undermine their existing storage management solutions. Data protection is no longer the simple copying of yesterday's changed files to tape. Critical data changes occur throughout the day, and to protect this data customers are frequently turning to new technology such as continuous remote replication. This white paper reviews the technical impact of VPLEX on RecoverPoint remote replication practices.

The primary use cases for this type solution are:

- VPLEX Local customers seeking point-in-time protection with RecoverPoint CDP and/or remote point-in-time protection with CRR or both with CLR
- VPLEX Metro customers seeking point-in-time protection with RecoverPoint CDP and/or third site replication with point-in-time protection using RecoverPoint CRR or both with CLR
- RecoverPoint customers introducing VPLEX into their environment

This paper reviews the technical considerations and the necessary procedural adjustments when using VPLEX with RecoverPoint. The key takeaway for each of these use cases is RecoverPoint and VPLEX can be combined into a solution that delivers on the strengths of both products.

This white paper is divided up into the following sections:

- [EMC VPLEX technology](#)
- [EMC RecoverPoint technology](#)
- [Solution topologies for VPLEX and RecoverPoint](#)
- [Replicating VPLEX virtual volumes with RecoverPoint](#)
- [Restoring VPLEX virtual volumes with RecoverPoint](#)
- [VPLEX Data Mobility considerations with RecoverPoint](#)

Audience

This white paper is intended for technology architects, storage administrators, and EMC professional services partners who are responsible for architecting, creating, managing, and using IT environments that utilize EMC VPLEX technologies. The white paper assumes that the reader is familiar with EMC VPLEX and RecoverPoint replication technologies.

VPLEX 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	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
VPLEX Metro	A cooperative set of two VPLEX clusters, each serving their own storage domain

Table 2. Acronyms and abbreviations

Acronym/Abbreviation	Definition
Clone	An independent RecoverPoint full disk copy of another device
Source device	Standard or primary array device. Typically used to run production applications
FE port	Front-end (target ports visible to hosts)
BE port	Back-end (initiator ports visible to storage arrays)

EMC VPLEX Technology

VPLEX encapsulates traditional physical storage array devices and applies three layers of logical abstraction to them. The logical relationships of each layer are shown in Figure 1.

Extents are the mechanism VPLEX uses to divide storage volumes. Extents may be all or part of the underlying storage volume. EMC VPLEX aggregates extents and applies RAID protection in the device layer. Devices are constructed using one or more extents and can be combined into more complex RAID schemes and device structures as

desired. At the top layer of the VPLEX storage structures are virtual volumes. Virtual volumes are created from devices and inherit the size of the underlying device. Virtual volumes are the elements VPLEX exposes to hosts using its FE ports. Access to virtual volumes is controlled using storage views. Storage views are comparable to Auto-provisioning Groups on EMC Symmetrix® or to storage groups on EMC CLARiiON®. They act as logical containers determining host initiator access to VPLEX FE ports and virtual volumes.




Figure 1. EMC VPLEX logical storage structures

EMC VPLEX architecture

EMC VPLEX represents the next-generation architecture for data mobility and information access. The new architecture is based on EMC's more than 20 years of expertise in designing, implementing, and perfecting enterprise-class intelligent cache and distributed data protection solutions.

As shown in Figure 2, VPLEX is a solution for federating both EMC and non-EMC storage. VPLEX resides between servers and heterogeneous storage assets and introduces a new architecture with these unique characteristics:

- Scale-up clustering hardware, which lets customers start small and grow big with predictable service levels
- Advanced data caching, which utilizes large-scale SDRAM cache to improve performance and reduce I/O latency and array contention
- Distributed cache coherence for automatic sharing, balancing, and failover of I/O across the cluster
- A consistent view of one or more LUNs across VPLEX clusters separated either by a few feet within a data center or across synchronous distances, enabling new models of high availability and workload relocation




Figure 2. Capability of an EMC VPLEX system to federate heterogeneous storage

EMC VPLEX Family Overview

The EMC VPLEX family today consists of:

- **VPLEX Local** for managing data mobility and access within the data center using a single VPLEX cluster.
- **VPLEX Metro** for mobility and access across two locations separated by inter-site RTT of up to 5 ms. VPLEX Metro uses two VPLEX clusters and includes the unique capability where a remote VPLEX Metro cluster can present LUNs without the need for physical storage for those LUNs at the remote cluster. It also supports synchronous distributed volumes that mirror data between the two clusters using write-through caching.
- **VPLEX Geo**, which also uses two VPLEX clusters, for access between two sites over extended asynchronous distances with RTT latencies up to 50 ms. VPLEX Geo distributed volumes support AccessAnywhere distributed mirroring using write-back caching. **Note: VPLEX Geo is not currently supported with RecoverPoint.**

VPLEX architecture

VPLEX VS2 hardware platform

A VPLEX VS2 system with GeoSynchrony 5.1 is composed of one or two VPLEX clusters: one cluster for VPLEX Local systems and two clusters for VPLEX Metro and VPLEX Geo systems. These clusters provide the VPLEX AccessAnywhere capabilities.

Each VPLEX cluster consists of:

- A VPLEX Management Console
- One, two, or four engines
- One standby power supply for each engine

In configurations with more than one engine, the cluster also contains:

- A pair of Fibre Channel switches

- An uninterruptible power supply for each Fibre Channel switch

VPLEX Metro and VPLEX Geo systems optionally include a VPLEX Witness. The VPLEX Witness is implemented as a virtual machine and is deployed in a separate fault domain from two VPLEX clusters. The VPLEX Witness is used to improve application availability in the presence of site failures and inter-cluster communication loss.

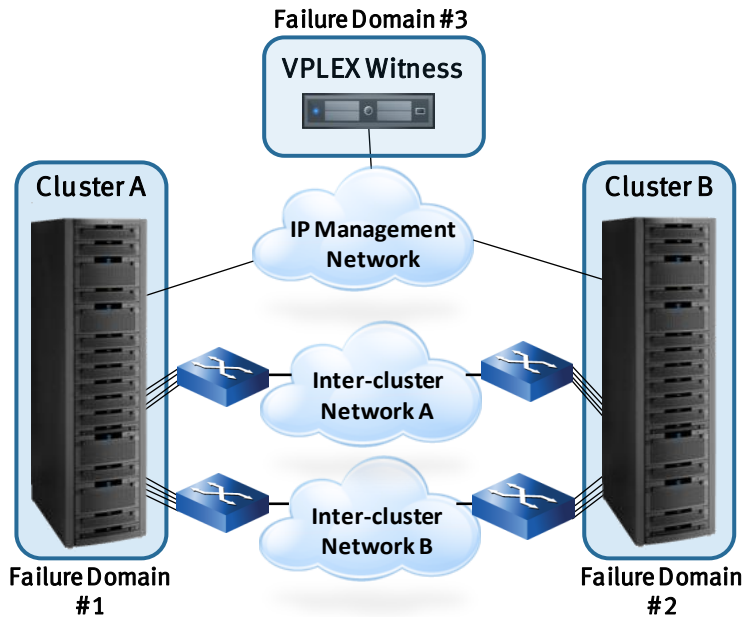


Figure 3. A VPLEX system with the VPLEX Witness

EMC RecoverPoint technology

RecoverPoint CRR provides remote replication between pairs of storage volumes residing in two different sites. For asynchronous remote replication multiple writes are collected at the local site, compressed, and sent across periodically to the remote site where they are uncompressed, written to a journal, and then distributed to the target volumes. For synchronous remote replication every write is collected and written to a remote journal and then distributed to the target volumes. The journal data provides the ability to roll back the target volume to any point in time.

RecoverPoint is designed to replicate changes at a block level on one or more SAN volumes (*source volumes*) residing in one or more storage arrays. It allows the replicated targets (*target volumes*) to reside in one or more heterogeneous storage arrays.

RecoverPoint maintains transactional consistent journals for each application (*consistency group*) defined within a RecoverPoint system. The journal allows convenient rollback to any point in time, enabling instantaneous recovery for application environments.

RecoverPoint architecture

The specific components of EMC RecoverPoint are shown in

Figure . This configuration shows the source, local copy, and remote copy volumes residing across multiple storage arrays in two different sites. There are three types of replication shown: local replication using CDP, remote replication using CRR, and both local and remote replication using the combination of CDP and CRR called concurrent local and remote. Details on the components are described in the next section.

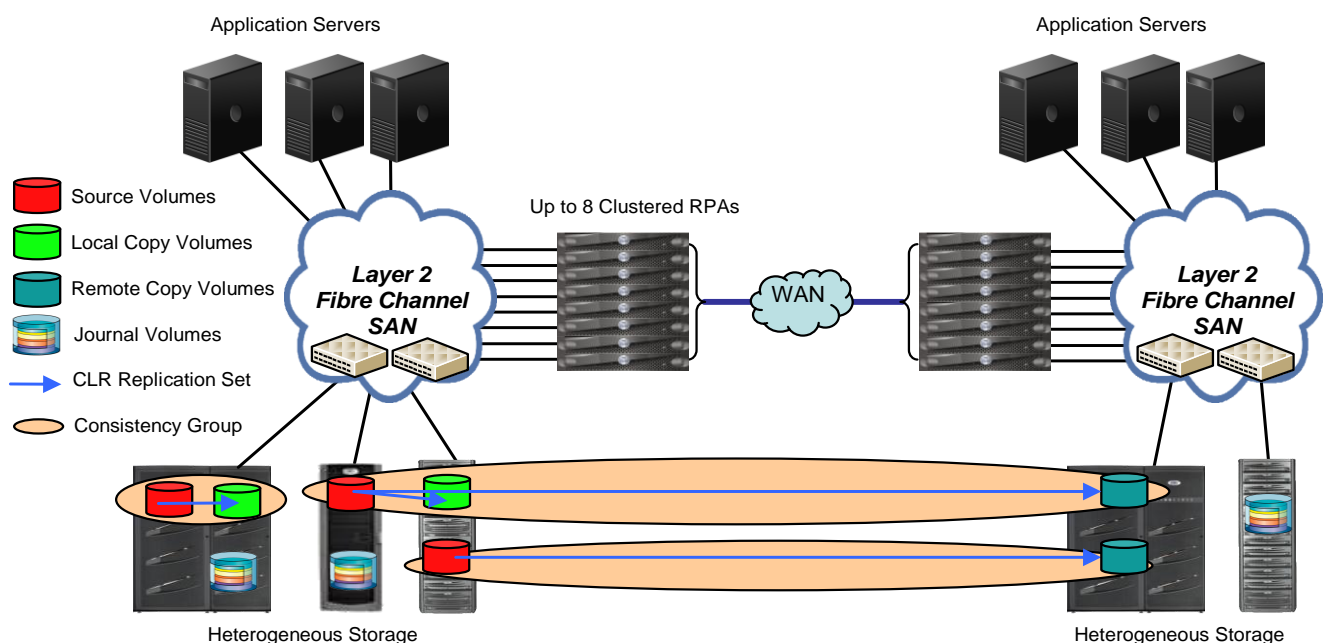


Figure 4. EMC RecoverPoint architecture

RecoverPoint appliance

The RecoverPoint appliance (RPA) runs the RecoverPoint software on top of a custom 64-bit Linux kernel inside a secure environment built from an industry-standard server platform. An RPA manages all aspects of data protection for a storage group, including capturing changes, maintaining the images in the journal volumes, and performing image recovery. Moreover, one appliance can manage multiple storage groups, each with differing policies.

There are at least two active RPAs per site that constitute a RecoverPoint cluster. Physically, the RecoverPoint cluster is located in the same facility as the host and storage subsystems. All RPAs in a cluster have identical functionality. In normal operation, all RPAs are active all of the time. Consequently, if one of the RPAs in a cluster goes down, EMC RecoverPoint supports immediate *switchover* of the functions of that appliance to one or more of the remaining RPAs.

Write Splitters

RecoverPoint monitors writes using technology called write splitters, which ensure that a copy of all host writes to a protected volume are tracked and sent to the local RecoverPoint appliance. RecoverPoint supports 3 types of write splitters - the array write splitter, the intelligent fabric-based write splitter, and the host-based write splitter. RecoverPoint/CL supports all 3 types of write splitters; RecoverPoint/EX supports the array-based write splitter for Symmetrix VMAXe, VNX series, and CLARiiON CX3 and CX4 series arrays; and RecoverPoint/SE supports the array-based write splitter for the VNX series and CLARiiON CX3 and CX4 series arrays, and the Windows host-based write splitter. A RecoverPoint configuration requires at least one type of splitter at each site. Consult the RecoverPoint release notes and product documentation for the most up to date support statements on other combinations of splitters.

Array write splitter

RecoverPoint supports an array-based write splitter for the Symmetrix VMAXe, as well as array-based write splitters that run inside the storage processors on EMC VNX series and CLARiiON CX3 and CX4 arrays. In this case, the *splitter* function is carried out by the storage processor; a KDriver is not installed on the host. The array-based write splitter is supported with Symmetrix VMAXe and VNX series arrays and with CLARiiON CX3 and CX4 series arrays running FLARE® 03.26, 04.28, 04.29, and 04.30. The array-based write splitter requires the installation of the no-charge RecoverPoint enabler. Unlike the host and fabric splitters, the array-based write splitter supports

LUNs up to 32 TB in size; the aforementioned splitters are limited to LUNs up to 2 TB in size.

Array-based write splitters enable RecoverPoint to support hosts such as AIX, HP-UX, Linux, OpenVMS, Solaris, VMware, and Windows. Additionally, multiple clusters can share the array-based write splitter, enabling up to four RecoverPoint clusters to use a single Symmetrix VMAXe, VNX series, CLARiiON CX3 series or CX4 series array. A LUN cannot span clusters, which means that more than one cluster cannot use the same LUN.

Array-based splitting technology is used for this white paper. This array-based splitter carries out write splitting inside each array. RecoverPoint supports two other methods of write splitting that will not be covered in this white paper, host-based splitter, sometimes referred to as a kdriver and fabric splitters.

Repository volume

The repository volume is a SAN-attached volume provisioned only to RecoverPoint. The repository volume is used to maintain the configuration and communication between RPAs in a cluster. Similar to a cluster server's quorum volume, the repository volume contains the status of the overall RecoverPoint system and acts as a resource arbitrator during RPA failover and recovery operations. There is no user-accessible information stored on the repository volume.

Journal volume

Journal volumes hold data waiting to be distributed to target replica volumes and also retain copies of the data previously distributed to the target volumes to facilitate operational recovery to any point in time that is retained in the journal history. Each consistency group has its own journal volumes, which allows for differing retention periods across consistency groups. Each consistency group has two or three journal volumes, one assigned to the local copy volumes (if present), one assigned to the remote copy volumes (if present), and one assigned to the production or source volumes. Journal volumes are used for the source and both copies in order to support production failover from the current active source volume to either the local or remote copy volume. A given copy journal can consist of one or more storage devices. RecoverPoint will stripe the data across the number of devices provisioned for a given journal.

Storage efficiency is realized in the history journal by retaining only the changes between journal entries when snapshot consolidation is enabled. Additionally, the journal volume can also be compressed, resulting in even more storage savings. These options can be turned on or off on a per-consistency group basis. Source and target images on the Replica volumes are always consistent upon completing the distribution of each write change.

Consistency groups

Figure shows the architecture of RecoverPoint consistency groups. Volumes are grouped into consistency groups, and replicated to target volumes in the local and/or

remote site. Writes to source volumes will be replicated in the same order on the local and/or remote volumes, which ensures that transactional consistency is maintained.

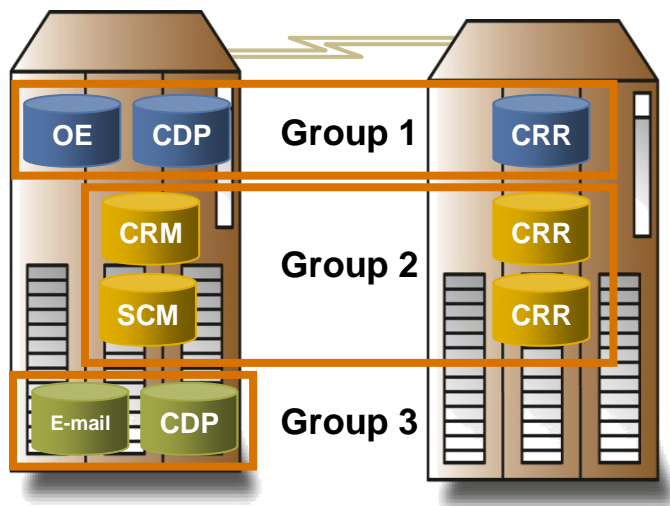


Figure 5. RecoverPoint consistency groups

If one of the source volumes in a consistency group goes offline, such as what may occur in a *rolling* disaster where a partial loss affects only a few of the source volumes, then the replication for all volumes in the consistency group will be paused until the source volume comes back online. In this way, consistency groups prevent dependent writes from getting out of sync, thus ensuring the integrity and consistency of the data at the remote site.

Solution Topologies for EMC VPLEX and RecoverPoint Array Splitters

Deployment topologies for VPLEX with RecoverPoint will vary based on solution requirements, data center design, and geographic distance between sites. In general, the deployment topologies can be broken into three main types of configurations consisting of one, two, or three sites (for CLR and CRR). There are many different possible combinations of arrays, sites, and replications solutions with VPLEX and RecoverPoint. All possible combinations will not be illustrated in this document, but they are still supported provided they meet the guidelines set forth by the VPLEX and RecoverPoint support matrices.

Single-Site VPLEX and RecoverPoint CDP Topology

Figure shows an example of a single-site topology containing two arrays with VPLEX and RecoverPoint. A single-site configuration provides continuous data protection for one or more arrays within a site. Replicas (copies of production devices) can be made within the same array or on a separate array.

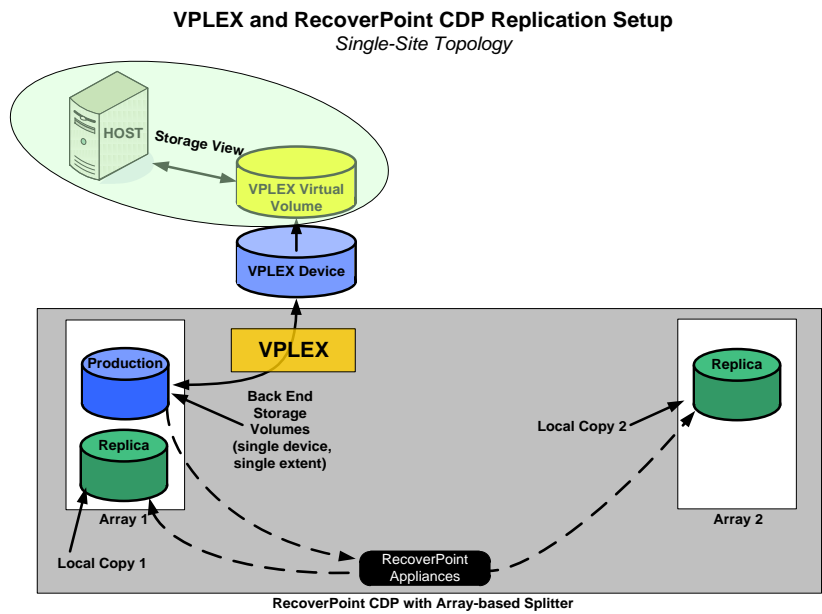


Figure 6. VPLEX and RecoverPoint CDP in a Single-Site Topology

Two-site VPLEX and RecoverPoint CDP Topology

Figure 7 (below) shows a two-site topology with four arrays. Each of the arrays provides a mirror leg for a VPLEX Distributed RAID 1 (DR1) device. This DR1 device consists mirror legs built from a single extent created from a single storage volume (a 1:1 encapsulated volume) provided by one array at each location. Replicas can be created within the same array or on other arrays at each site. The maximum distance between Site A and Site B will vary based on the based on the recovery objectives, inter-site bandwidth, latency, and other limitations outlined in the *EMC Support Matrix* for RecoverPoint 3.4.1

VPLEX Metro + RecoverPoint CDP with Array-based Splitter 2-Site Topology

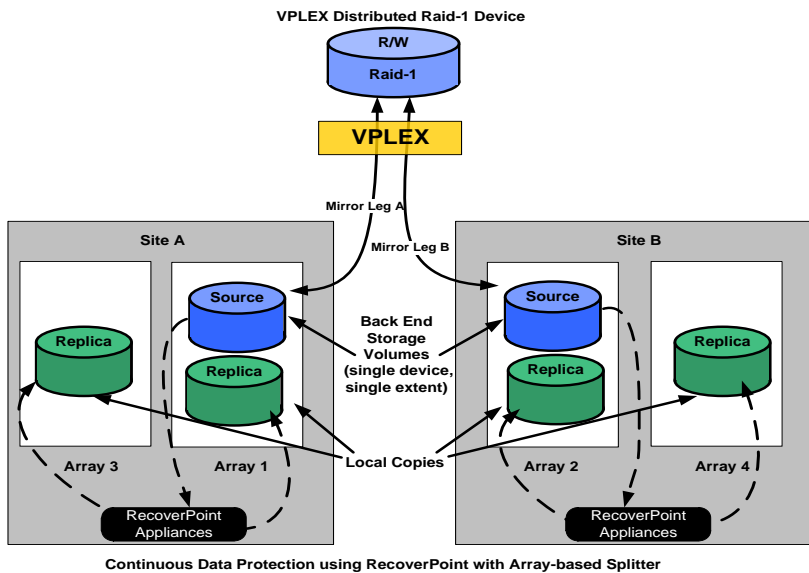


Figure 7. VPLEX and RecoverPoint CDP in a two-site topology

Replicating VPLEX virtual volumes with RecoverPoint CDP

EMC VPLEX Local and Metro's write-through caching model ensures host write order and consistency to back-end storage arrays and, when using array-based splitting, to RecoverPoint appliances. By mapping storage volumes provided by the source array in their entirety (1:1 mapping) through the VPLEX virtualization stack EMC RecoverPoint can be leveraged to provide remote replicas of production VPLEX virtual volumes. From a RecoverPoint perspective the virtual volumes are simply source volumes and the one-to-one encapsulation ensures that the host view of the SCSI LUN it is writing to matches what is provided by the array. This configuration method ensures the underlying storage devices are left unchanged by VPLEX, enabling RecoverPoint replication to function normally. By following this methodology, the RecoverPoint local copy (replicas) volumes are able to be utilized following applicable standard RecoverPoint best practices. Refer to the RecoverPoint Administrators Guide 3.4.1 for more information on using the remote copies provided by RecoverPoint CDP. Operations like RecoverPoint "image mode" and "failover" all work per the normal processes from the remote site perspective. It is important to note, however, when running in failover mode, doing a failback, or any time writes are occurring from the replica devices (local copies) back to the production devices (VPLEX device) the VPLEX virtual volume *must be taken offline prior* to starting the failover or restore operation. The following steps provide more detail on this.

These steps assume that each RecoverPoint production volume has a one-to-one pass-through configuration (device capacity = extent capacity = storage volume capacity) to VPLEX and has a *RAID 0 or RAID 1 (single extent only) VPLEX device*

geometry. Figure shows the logical relationships between the host, virtual volumes, and corresponding VPLEX and RecoverPoint components.

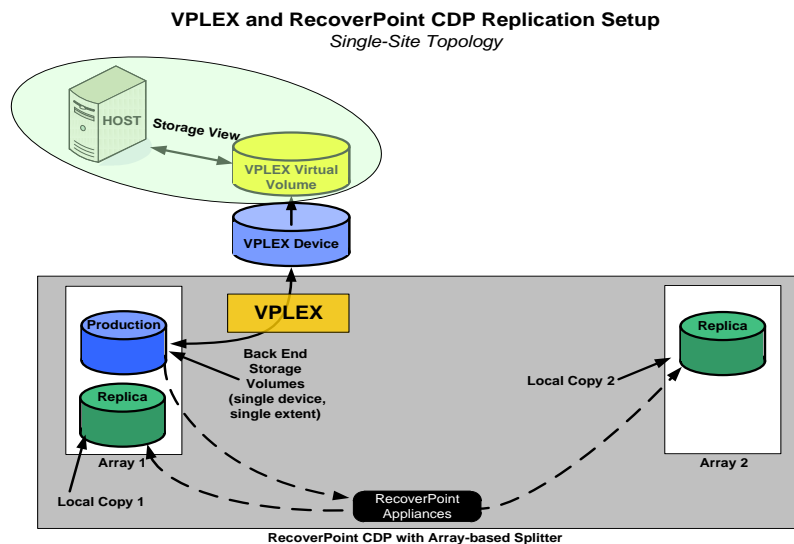


Figure 8. VPLEX and RecoverPoint CDP

Initial setup

To perform RecoverPoint local replication of VPLEX virtual volumes some initial setup work is required:

1. (Array-specific step) Within the array CLI or GUI, identify and record the source (production) devices on within the array you wish to replicate with RecoverPoint CDP. Reference the array specific online help documentation for further details on how to create and identify LUNs within the GUI and CLI.
2. (RecoverPoint-specific step) Setup RecoverPoint CDP relationship between production (source) and local replica (target) devices. The detailed RecoverPoint CDP replication procedure(s) can be found in the RecoverPoint Administrator's Guide 3.4.1 on EMC.Powerlink.com. This EMC web portal is available to customers and employees only.
3. (VPLEX-specific steps) Confirm that the production (source) volumes from the underlying array are visible to VPLEX. As necessary, perform array LUN masking and SAN zoning for array LUNs to VPLEX back-end ports.
4. Perform one-to-one encapsulation through VPLEX:
 - a. Claim storage volumes from array
 - b. Create extents
 - c. Create devices (single extent RAID 0 Geometry)
 - d. Create virtual volumes for each RecoverPoint source volume

5. Present virtual volumes based on RecoverPoint production (source) volumes to host(s).
 - a. If necessary, create storage view(s).
 - b. Add virtual volumes built from RecoverPoint production (source) volumes to storage view(s).
 - c. If necessary, perform zoning of virtual volumes to hosts following standard zoning procedures.

Restoring VPLEX virtual volumes with RecoverPoint CDP

When using RecoverPoint to restore storage volumes that are virtualized by VPLEX the primary consideration is the VPLEX read cache. During the restore process RecoverPoint is writing to the back-end storage volume(s) outside of the VPLEX I/O path. When writes occur outside of the VPLEX I/O path, VPLEX read cache may not match the data on the storage volume. *To avoid this situation VPLEX read cache invalidation for each restored storage volume is required.* Per-volume read cache invalidation is accomplished by removing the corresponding VPLEX virtual volumes from any storage views they are members of. The act of removing a VPLEX virtual volume from a storage view causes the read cache for that volume to be discarded (not written to disk). The VPLEX virtual volumes being restored must remain removed from all storage views until the RecoverPoint writes (non-VPLEX I/O path writes) are completed. Once the restore activities have finished and the underlying back-end storage volumes are in the desired state, then the virtual volumes can be added back into a storage view and be accessed normally. Figure (below) shows the logical relationships between the host, the local array containing the restore source, and the array containing the restore target during a RecoverPoint restore process. Notice that the host storage view has been removed as this is a requirement during the restore process.

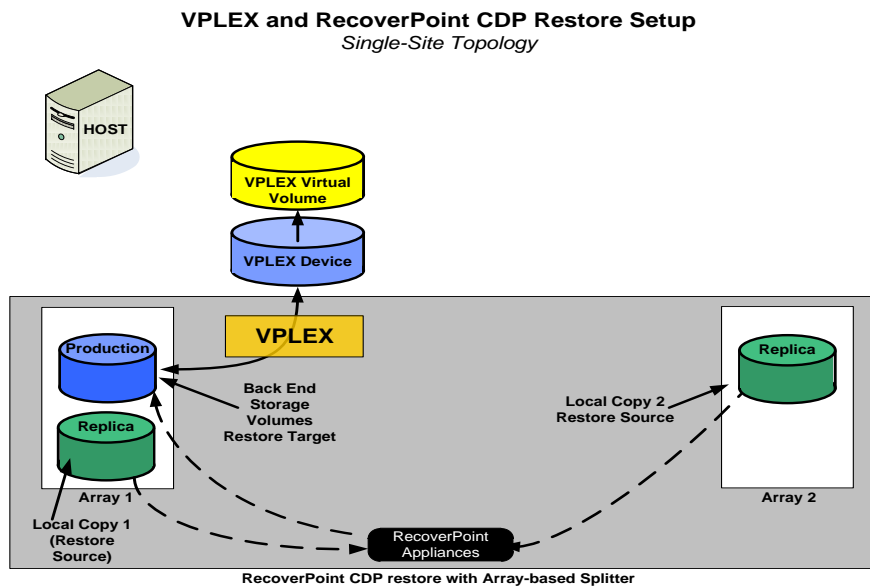


Figure 9. VPLEX and RecoverPoint CDP Restore setup

As mentioned earlier, it is assumed that each production device (restore target in this case) has a one-to-one storage volume pass-through configuration (device capacity = extent capacity = storage volume capacity) to VPLEX *and a RAID 0 or RAID 1 (single extent only) device geometry*. Please note that the RAID 1 steps are only applicable for those virtual volumes that are built on top of VPLEX RAID 1 devices.

To perform a restore to a VPLEX virtual volume using RecoverPoint CRR follow these steps:

1. Identify the VPLEX virtual volumes you wish to restore using RecoverPoint CRR.
2. Shut down any applications using the VPLEX virtual volumes you wish to restore.
3. As necessary, unmount the virtual volumes to be restored.
4. Remove host access to corresponding VPLEX volumes you wish to restore by removing them from *all* storage views.
 - a. If the virtual volume has local visibility and/or is a member of a single storage view, using the VPLEX CLI run:


```
/clusters/<cluster name>/exports/storage-views>removevirtualvolume -v storage_view_name -o virtual_volume_name -f
```
 - b. If the virtual volume is built from a distributed RAID 1 device and is a member of storage views in both clusters, using the VPLEX CLI run:

```

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

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

```

Note: It is vital to keep track of the VPLEX-assigned LUN number for each virtual volume you plan to restore. This is because when you remove the virtual volumes from a storage-view and then add them back they may not get numbered in the same sequence. You can obtain this information by executing a long listing (ll) in the corresponding storage view context from the VPLEX CLI or by clicking on the host storage view from the VPLEX Management Console.

5. For RAID 1 VPLEX devices detach the second mirror leg (mirror leg that will not be written to as part of the RecoverPoint restore process):

a. Detach the mirror leg. From the VPLEX CLI run:

```

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

```

6. Wait 30 seconds to ensure that the VPLEX read cache has been invalidated for each virtual volume. This can be done concurrently with step 7.

7. Perform a standard RecoverPoint restore procedure for the source volumes you identified in step 1. See the *RecoverPoint Administrator's Guide 3.4.1* available on EMC Powerlink for more information on the RecoverPoint steps.

8. Confirm the IO Status of storage volumes that correspond with the RecoverPoint source volumes (restore targets) are “alive” by doing a long listing against the storage-volumes context for your cluster.

For example:

```

vplexcli11> ll /clusters/cluster-1/storage-elements/storage-volumes/

```

Name	VPD3 ID	Capacity	Use	Vendor	IO Status	Type
Symm1554Tdev_061D	VPD03T3:60000970000192601554533030363144	100G	used	EMC	alive	normal
Symm1554Tdev_061E	VPD03T3:60000970000192601554533030363145	100G	claimed	EMC	alive	normal
Symm1554Tdev_061F	VPD03T3:60000970000192601554533030363146	100G	claimed	EMC	alive	normal
Symm1554Tdev_0620	VPD03T3:60000970000192601554533030363230	100G	claimed	EMC	alive	normal
Symm1554Tdev_0621	VPD03T3:60000970000192601554533030363231	100G	claimed	EMC	alive	normal
Symm1554Tdev_0622	VPD03T3:60000970000192601554533030363232	100G	used	EMC	alive	normal
Symm1554Tdev_0623	VPD03T3:60000970000192601554533030363233	100G	used	EMC	alive	normal
Symm1554Tdev_0624	VPD03T3:60000970000192601554533030363234	100G	used	EMC	alive	normal
Symm1554Tdev_0625	VPD03T3:60000970000192601554533030363235	100G	used	EMC	alive	normal
Symm1554Tdev_0626	VPD03T3:60000970000192601554533030363236	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. The VPLEX ‘resurrect’ command can be used to revive dead devices resulting from the restore process. See the VPLEX 5.0 CLI guide for more details on the resurrect command.

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
```

9. For RAID 1 devices re-attach the second mirror leg (mirror leg that was detached during restore in step 5):

a. To re-attach the second mirror using the CLI run:

```
device attach-mirror -m <2nd_mirror_leg_to_attach> -d
/clusters/local_cluster_name/devices/existing_raid_1_device
```

b. Turn logging back on for distributed RAID 1 devices:

```
/distributed-storage/distributed-devices> set-log -d
/distributed-storage/distributed-devices/distributed_raid_1_device_1
```

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

10. Restore host access to virtual volumes that were restored. Add the virtual volume back to the view, specifying the original LUN number (noted in step 2) using the VPLEX CLI:

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

11. Rescan devices and restore paths (for example, powermt restore) on hosts.

12. If necessary, mount devices.

13. Restart applications.

Replicating VPLEX virtual volumes with RecoverPoint CRR

EMC VPLEX Local and Metro's write-through caching model ensures host write order and consistency to back-end storage arrays and, when using array-based splitting, to RecoverPoint appliances. By mapping storage volumes provided by the source array in their entirety (1:1 mapping) through the VPLEX virtualization stack EMC RecoverPoint can be leveraged to provide remote replicas of production VPLEX virtual volumes. From a RecoverPoint perspective the virtual volumes are simply source

volumes and the one-to-one encapsulation ensures that the host view of the SCSI LUN it is writing to matches what is provided by the array. This configuration method ensures the underlying storage devices are left unchanged by VPLEX, enabling RecoverPoint replication to function normally. By following this methodology, the RecoverPoint remote copy volumes are able to be utilized following applicable standard RecoverPoint best practices. Refer to the RecoverPoint Administrators Guide 3.4.1 for more information on using the remote copies provided by RecoverPoint CRR. Operations like RecoverPoint “image mode” and “failover” all work per the normal processes from the remote site perspective. It is important to note, however, when running in failover mode, doing a failback, or any time writes are occurring from the target site (remote copy site) back to the source site (VPLEX site) the virtual volume at the source site *must be taken offline prior* to starting the failover or restore operation. The following steps provide more detail on this.

These steps assume that each RecoverPoint source volume has a one-to-one pass-through configuration (device capacity = extent capacity = storage volume capacity) to VPLEX and has a *RAID 0 or RAID 1 (single extent only) VPLEX device geometry*. Figure 10 shows the logical relationships between the host, virtual volumes, and corresponding VPLEX and RecoverPoint components.

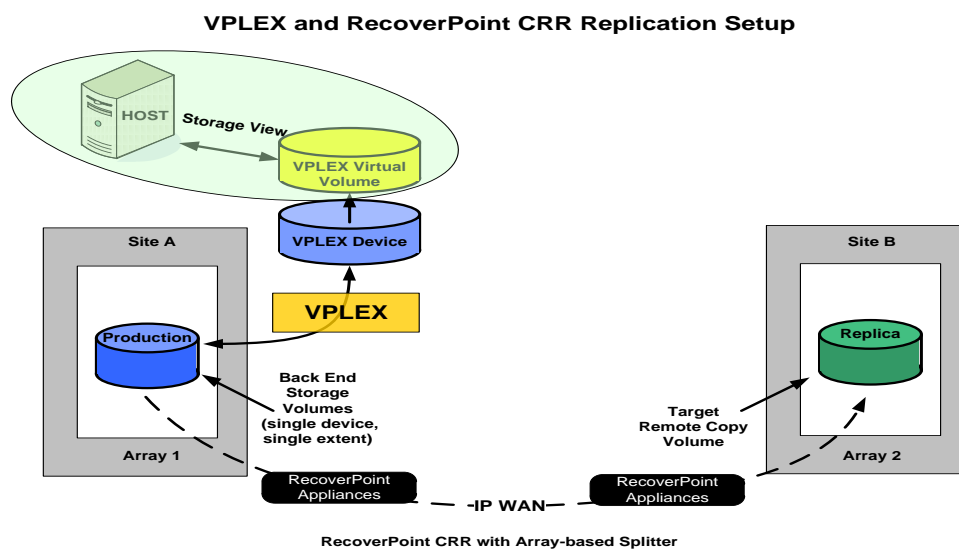


Figure 10. VPLEX and RecoverPoint CRR

Initial setup

To perform RecoverPoint replication of VPLEX virtual volumes:

1. (Array-specific step) Within the array CLI or GUI, identify and record the source volumes (LUNs) on the array you wish to replicate with RecoverPoint CRR. The native array online help documentation will typically provide details on how to create and identify LUNs within the GUI and CLI.

2. (RecoverPoint-specific step) The RecoverPoint CRR replication procedure(s) to configure the desired RecoverPoint target remote copy volumes is in the RecoverPoint Administrator's Guide 3.4.1 on Powerlink.EMC.com. This EMC web portal is available to customers and employees only.
3. (VPLEX-specific steps) Confirm that the production (source) volumes from the array are visible to VPLEX. As necessary, perform array LUN masking and SAN zoning for array LUNs to VPLEX back-end ports.
4. Perform one-to-one encapsulation through VPLEX:
 - a. Claim storage volumes.
 - b. Create extents
 - c. Create devices (single extent RAID 0 Geometry)
 - d. Create virtual volumes for each RecoverPoint production (source) volume
5. Present virtual volumes based on RecoverPoint source volumes to host(s).
 - a. If necessary, create storage view(s).
 - b. Add virtual volumes built from RecoverPoint production (source) volumes to storage view(s).
 - c. If necessary, perform zoning of virtual volumes to hosts following normal zoning procedures.

Restoring VPLEX virtual volumes with RecoverPoint CRR

When using RecoverPoint to restore storage volumes that are virtualized by VPLEX the primary consideration is the VPLEX read cache. During a restore process RecoverPoint is writing to the back-end storage volume(s) outside of the VPLEX I/O path. When writes occur outside of the VPLEX I/O path, VPLEX read cache may not match the data on the storage volume. To avoid this situation VPLEX read cache invalidation for each restored storage volume is required. Per-volume read cache invalidation is accomplished by removing the corresponding VPLEX virtual volumes from any storage views they are members of. The act of removing a VPLEX virtual volume from a storage view causes the read cache for that volume to be discarded (not written to disk). The VPLEX virtual volumes being restored must remain removed from all storage views until the RecoverPoint writes (non-VPLEX I/O path writes) are completed. Once the restore activities have finished and the underlying back-end storage volumes are in the desired state, then the virtual volumes can be added back into a storage view and be accessed normally. Figure 11 shows the logical relationships between the host, the remote array (restore source), and the local array (restore target) during a RecoverPoint restore process. Notice that the host storage view has been removed as this is a requirement during the restore process.

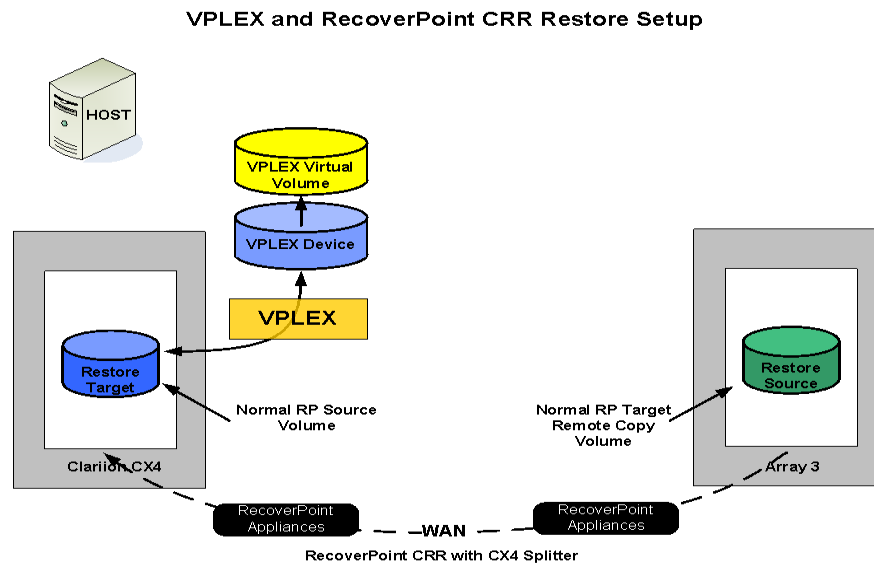


Figure 11. VPLEX and RecoverPoint restore setup

As mentioned earlier, it is assumed that each source device (restore target in this case) has a one-to-one storage volume pass-through configuration (device capacity = extent capacity = storage volume capacity) to VPLEX *and a RAID 0 or RAID 1 (single extent only) device geometry*. Please note that the RAID 1 steps are only applicable for those virtual volumes that are made from RAID 1 devices.

To perform a restore to a VPLEX virtual volume using RecoverPoint CRR follow these steps:

1. Identify the VPLEX virtual volumes you wish to restore using RecoverPoint CRR.
2. Shut down any applications using the VPLEX virtual volumes you wish to restore.
3. As necessary, unmount the virtual volumes to be restored.
4. Remove host access to corresponding VPLEX volumes you wish to restore by removing them from *all* storage views.
 - a. If the virtual volume is built from a local device and/or is a member of a single storage view, using the VPLEX CLI run:

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

- b. If the virtual volume is built from a distributed RAID 1 device and is a member of storage views in both clusters, using the VPLEX CLI run:

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

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

Note: It is vital to keep track of the VPLEX-assigned LUN number for each virtual volume you plan to restore. This is because when you remove the virtual volumes from the storage view and then add them back, they will not necessarily be assigned the same lun numbers. You can obtain this information by executing a long listing (ll) in the corresponding storage view context from the VPLEX CLI or by clicking on the host storage view from the VPLEX Management Console.

5. For RAID 1 VPLEX devices detach the second mirror leg (mirror leg that is not being replicated by RecoverPoint):
 - a. Detach the mirror leg. From the VPLEX CLI run:


```
device detach-mirror -m <device_mirror_to_detach> -d <distributed_device_name> -i -f
```
6. Wait 30 seconds to ensure that the VPLEX read cache has been invalidated for each virtual volume. This can be done concurrently with step 7.
7. Perform a standard RecoverPoint restore procedure for the source volumes you identified in step 1. See the *RecoverPoint Administrator's Guide* available on EMC Powerlink for more information on the RecoverPoint steps.
8. Confirm that the IO Status of the array storage volumes that correspond with the RecoverPoint production (source) volumes is “alive”. This can be done using the ll command (long listing) against the storage-volumes context for your cluster.

For example:

```
Vplexcli1>> ll /clusters/cluster-1/storage-elements/storage-volumes/
```

Name	VPD03 ID	Capacity	Use	Vendor	IO Status	Type
Symm1554Tdev_061D	VPD03T3:60000970000192601554533030363144	100G	used	EMC	alive	normal
Symm1554Tdev_061E	VPD03T3:60000970000192601554533030363145	100G	claimed	EMC	alive	normal
Symm1554Tdev_061F	VPD03T3:60000970000192601554533030363146	100G	claimed	EMC	alive	normal
Symm1554Tdev_0620	VPD03T3:60000970000192601554533030363230	100G	claimed	EMC	alive	normal
Symm1554Tdev_0621	VPD03T3:60000970000192601554533030363231	100G	claimed	EMC	alive	normal
Symm1554Tdev_0622	VPD03T3:60000970000192601554533030363232	100G	used	EMC	alive	normal
Symm1554Tdev_0623	VPD03T3:60000970000192601554533030363233	100G	used	EMC	alive	normal
Symm1554Tdev_0624	VPD03T3:60000970000192601554533030363234	100G	used	EMC	alive	normal
Symm1554Tdev_0625	VPD03T3:60000970000192601554533030363235	100G	used	EMC	alive	normal
Symm1554Tdev_0626	VPD03T3:60000970000192601554533030363236	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

```

9. For RAID 1 devices re attach the second mirror leg (mirror leg that was detached during restore in step 5):

a. To re-attach the second mirror using the CLI run:

```

device attach-mirror -m <2nd_mirror_leg_to_attach> -d
/cclusters/local_cluster_name/devices/existing_raid_1_de
vice

```

b. Turn logging back on for distributed RAID 1 devices:

```

/distributed-storage/distributed-devices> set-log -d
/distributed-storage/distributed-
devices/distributed_raid_1_device_1

```

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

10. Restore host access to virtual volumes that were restored. Add the virtual volume back to the view, specifying the original LUN number (noted in step 2) using the VPLEX CLI:

```

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

```

11. Rescan devices and restore paths (for example, powermt restore) on hosts.

12. If necessary, mount devices.

13. Restart applications.

Replicating and Restoring VPLEX virtual volumes with RecoverPoint CLR

EMC VPLEX Local and Metro deployed with RecoverPoint Concurrent Local and Remote data protection (CLR) combines the operational recovery capabilities of CDP and CRR solutions. No procedural changes are required relative to the CDP and CRR steps describe in the earlier sections of this document. The RecoverPoint replicas at each location can be used independently and have the same properties as if they had been

deployed into individual CDP/CRR solutions. Reference the earlier procedures for the necessary operational recovery steps for each site in a CLR configuration.

VPLEX data mobility considerations with RecoverPoint

In addition to providing access anywhere, improved resiliency, and a single provisioning interface across multiple heterogeneous storage arrays, VPLEX provides non-disruptive data mobility within and between storage arrays. When using VPLEX with RecoverPoint it is especially important to be aware of this fact. As the RecoverPoint array-based splitter is based on writes being sent to a specific array source LUN, if VPLEX does a mobility job from that source LUN to another LUN then RecoverPoint will no longer be performing replication for that source volume. Figure 12 shows the situation where VPLEX has used its data mobility feature to move from a source volume being replicated by RecoverPoint CRR to a new volume with the same array.

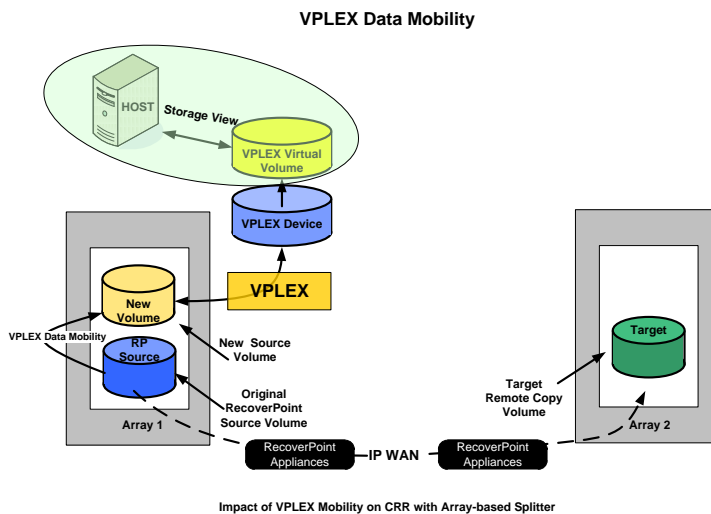


Figure 12. VPLEX Data Mobility and RecoverPoint

In order for replication to be restarted, the RecoverPoint relationship between the original source volume and a target must be removed and then the new source volume appropriately configured and then a “full sweep” will need to be performed on the device. The *RecoverPoint Administrator’s Guide 3.4.1* explains the necessary steps and the impact of this process in detail. When mobility jobs are planned for VPLEX virtual volumes, it is vital that the planning also takes into account any corresponding RecoverPoint CDP and CRR relationships.

Conclusion

To protect vital data customers are frequently turning to new technology such as RecoverPoint continuous remote replication. In this white paper we examined the technical impact of VPLEX on RecoverPoint remote replication practices. The necessary steps to replicate and restore VPLEX storage volumes using RecoverPoint with the various array-based splitters were explained. As was shown, RecoverPoint with array-based splitting was able to provide all its normal operational recovery capabilities with a few procedural adjustments. By planning for the impact of VPLEX read cache and mobility operations, RecoverPoint can be successfully deployed with array-based splitting technology.

References

- [EMC VPLEX page](#) on EMC.com
- [White Paper: Workload Resiliency with EMC VPLEX – Best Practices Planning](#)
- [TechBook: EMC VPLEX Metro Witness Technology and High Availability](#)
- [RecoverPoint Administrator's Guide 3.4.1](#)
- Navisphere Online Help Documentation