

# EMC<sup>®</sup> VPLEX<sup>™</sup> METRO WITH VERITAS STORAGE FOUNDATION HIGH AVAILABILITY SOLUTIONS

## ABSTRACT

This white paper describes the implementation of Veritas High Availability Extended Clusters using EMC VPLEX Metro configuration.

January 2016



EMC WHITE PAPER

To learn more about how EMC products, services, and solutions can help solve your business and IT challenges, [contact](#) your local representative or authorized reseller, visit [www.emc.com](http://www.emc.com), or explore and compare products in the [EMC Store](#)

Copyright © 2016 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as 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.

Part Number H11975.2

## Table of Contents

|   |           |
|---|-----------|
| <b>Executive summary</b> .....  | <b>4</b>  |
| Audience .....  | 4         |
| <b>Introduction</b> .....   | <b>5</b>  |
| <b>Products and features overview</b> .....   | <b>5</b>  |
| VPLEX.....  | 5         |
| VPLEX product offerings.....  | 6         |
| VPLEX architecture highlights .....   | 7         |
| Veritas High Availability solutions .....   | 7         |
| Veritas Storage Foundation .....  | 8         |
| Veritas Cluster Server .....  | 8         |
| Use cases.....  | 9         |
| <b>Implementing VPLEX Metro with Veritas clustering technologies</b> .....            | <b>10</b> |
| Extended Veritas VCS deployment topologies in a VPLEX Metro environment.....          | 10        |
| Extended Veritas SFRAC/SFCFS deployment topologies in a VPLEX Metro environment ..... | 11        |
| Failure scenarios and expected impacts.....   | 12        |
| <b>VPLEX Metro with extended Veritas HA lab configurations</b> .....                  | <b>15</b> |
| Lab configuration and setup .....   | 15        |
| VPLEX Metro setup .....   | 17        |
| Hosts and Veritas High Availability setup.....  | 20        |
| Veritas Veritas Cluster validation test.....  | 21        |
| <b>Conclusion</b> .....   | <b>22</b> |
| <b>References</b> .....   | <b>22</b> |

## Executive summary

EMC® VPLEX™ is an enterprise-class solution that breaks the physical barrier by enabling users to access to a single copy of the data at different geographical locations. It enables to geographically stretch virtual and physical host clusters. This enables transparent load sharing between multiple sites while providing the flexibility of relocating workload between sites in anticipation of planned events. In the stretched clusters the cluster nodes are spread across the sites. In the event of unplanned failure at a site, the cluster nodes on other sites are not impacted providing high availability to the running applications. This solution is made possible by the VPLEX AccessAnywhere clustering technology which allows read/write access to distributed volumes across distance where the volumes have the exact same SCSI LUN identity.

The VPLEX product line includes VPLEX Local (single site SAN federation), VPLEX Metro supporting synchronous distributed volumes with round-trip latency up to 5 ms, and VPLEX Geo support asynchronous distributed volumes with round-trip time up to 50 ms. This paper highlights VPLEX Metro working in conjunction with Veritas Storage Foundation High Availability (HA) Solutions on extended distance. The solution mentioned in this white paper focus on extending the distance of the two supported Veritas Storage Foundation High Availability Solutions using VPLEX Metro:

- Asymmetric or Active/Passive Cluster configuration
- Symmetric or Active/Active Cluster configuration

Main benefits of VPLEX Metro in an extended Active/Active Veritas HA solution:

- Continuous application availability through network, server, storage, and site failures
- Scale-out architecture and full read-write access to the same data at both sites (no idle hardware)
- Increased resiliency to entire site failure scenarios or data center maintenance operations without application downtime
- Host cluster nodes are distributed across two sites for maximum availability

EMC VMAX™ and VNX™ Series is used in this paper as the storage array behind VPLEX clusters. Refer the best practices guide on implementing the storage array for better performance and data protection.

## Audience

This white paper is intended for Systems administrators, Storage administrators, and IT architects responsible for architecting, creating, managing, and using IT environments that focus on high availability solutions with Veritas , VPLEX, and EMC Block storage technologies. The white paper assumes readers are familiar with Veritas Clustering technology, EMC VPLEX, and the EMC Storage array technologies.

## Introduction

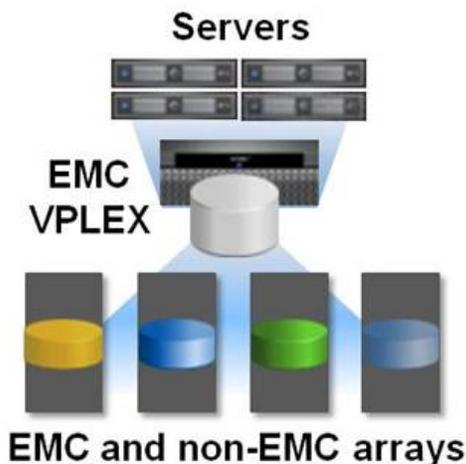
Extended Veritas Clusters provides a way to scale out, utilize storage and server resources at multiple sites, and provide an increased resiliency to failure scenarios or maintenance operations without application downtime - allowing organizations to eliminate database downtime, and continue business processing uninterrupted, even in the case of complete site failures.

This white paper introduces readers to the EMC VPLEX family, VPLEX Metro cluster architecture, and features and functionality that are relevant to an extended Veritas High Availability deployment. This paper also discusses the resiliency of the extended Veritas Clusters against different failure conditions.

## Products and features overview

### VPLEX

EMC VPLEX is a storage virtualization solution for both EMC and non-EMC storage, as shown in [Figure 1](#). The storage behind VPLEX can be heterogeneous, supporting both EMC storage and common arrays from other storage vendors, such as NetApp, HDS, HP, and IBM.



**Figure 1. EMC VPLEX Local**

VPLEX can be extended across geographically dispersed data centers to provide simultaneous access to storage devices through the creation of VPLEX Distributed Virtual Volumes.

VPLEX technology provides non-disruptive, heterogeneous data movement and volume management functionality. Because of these capabilities, VPLEX delivers unique and differentiated value to address four distinct requirements:

- The ability to dynamically move applications and data across different compute and storage infrastructures, either within or across data centers.

- The ability to create high-availability storage and compute infrastructure geographically dispersed with unmatched resiliency.
- The ability to provide efficient real-time data collaboration over distance.
- The ability to seamlessly migrate data between storage arrays.

### VPLEX product offerings

EMC offers VPLEX in three configurations to address customer needs for high-availability and data mobility as seen in Figure 2:

- VPLEX Local
- VPLEX Metro
- VPLEX Geo (Not currently supported with Veritas HA products)

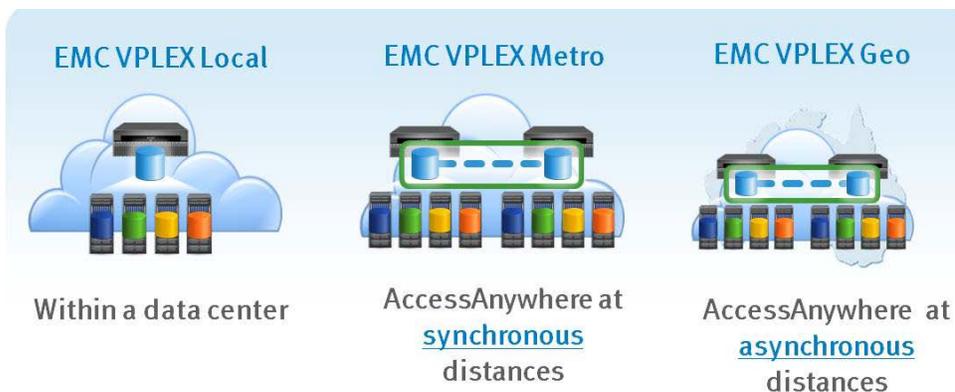


Figure 2. VPLEX topologies

#### VPLEX Local

VPLEX Local provides seamless, non-disruptive data mobility and ability to manage multiple heterogeneous arrays from a single interface within a data center.

VPLEX Local allows increased availability, simplified management, and improved utilization across multiple arrays.

#### VPLEX Metro with AccessAnywhere

VPLEX Metro with AccessAnywhere enables active-active, block level access to data between two sites within synchronous distances up to 5 ms round-trip time (RTT).

#### VPLEX Geo with AccessAnywhere

VPLEX Geo with AccessAnywhere enables active-active, block level access to data between two sites within asynchronous distances up to 50ms RTT latency. *VPLEX Geo is not currently supported with Veritas HA solutions and will not be detailed in this document.* For more information about VPLEX Geo visit <http://support.emc.com>, search by product: VPLEX Geo.

## VPLEX architecture highlights<sup>1</sup>

The VPLEX family uses a unique clustering architecture to help customers break the boundaries of the data center and allow servers at multiple data centers to have read/write access to shared block storage devices. VPLEX Local includes a single cluster while VPLEX Metro includes two. A VPLEX cluster consists of one, two, or four engines as seen in Table 1. Each VPLEX engine provides SAN/WAN connectivity, cache and processing power with two redundant directors.

**Table 1. VPLEX hardware components**

| Feature        | Description  |
|----------------|--|
| VPLEX Cluster  | Contains one, two, or four engines   |
| VPLEX Engine   | Contains two directors, management modules, power supplies, battery power and fans |
| VPLEX Director | Contains some of the I/O modules, SSD, CPU and RAM                                 |

VPLEX Local uses write-through caching and lets writes pass directly and get acknowledged first at the storage behind the VPLEX volumes before acknowledging them back to the host. With EMC storage, such as VMAX Series and VNX Series, where writes only need to register with the storage persistent cache, application write response time is optimal. VPLEX Metro also uses write-through cache but will acknowledge the writes to the application only once they were registered with both local and remote storage. In all VPLEX deployments, reads can benefit from the VPLEX cache, and in VPLEX Metro, read-hits are served from the local VPLEX cluster cache.

Refer to the VPLEX product documentation, located at <https://support.emc.com>, for further details.

## Veritas High Availability solutions

Veritas High Availability solutions protects and keeps the applications, data and infrastructure resilient and available for the most critical services running in the data center. Veritas supports mainstream business applications for high availability providing 24x7 data availability.

---

<sup>1</sup> The details in this section are based on VPLEX release 5.2 and may be different in other releases. The *VPLEX Product Guide* provides exact version details.

The Veritas Storage Foundation High Availability (SFHA) consists of the following products:

- Veritas Cluster Server (VCS) for Active/Passive Cluster configuration
- Storage Foundation for Oracle RAC (SFRAC) using Cluster Volume Manager (CVM) for Active/Active Cluster configuration
- Storage Foundation for Cluster File System (SFCFS) for Active/Active Cluster configuration

Storage Foundation High Availability includes Veritas Storage Foundation and Veritas Cluster Server. Veritas Cluster Server adds high availability functionality to Storage Foundation products.

### Veritas Storage Foundation

Veritas Storage Foundation includes Veritas File System (VxFS), Veritas Volume Manager (VxVM), and Veritas Dynamic Multipath (VxDMP).

- Veritas File System (VxFS) is a high performance journaling file system that provides easy management and quick-recovery for applications. Veritas File System delivers scalable performance, continuous availability, increased I/O throughput, and structural integrity.
- Veritas Volume Manager (VxVM) removes the physical limitations of disk storage. Users can configure, share, manage, and optimize storage I/O performance online without interrupting data availability. Veritas Volume Manager also provides easy-to-use, online storage management tools to reduce downtime.
- Veritas Dynamic Multi-Pathing (DMP) provides multipathing functionality for the operating system native devices configured on the system. The product creates DMP metadevices (also known as DMP nodes) to represent all the device paths to the same physical LUN.

### Veritas Cluster Server

Veritas Cluster Server (VCS), which provides clustering solution, provides the following benefits:

- Minimize planned or unplanned downtime
- Facilitates the consolidation and the failover of servers
- Supports wide range of applications in heterogeneous environments
- Built-in cluster simulator tool allows cluster configurations to be configured and tested without requiring a cluster testing environment or impacting the production environment.

Veritas Cluster Server is included in all available SFHA product suites.

## Use cases

The following are some uses cases of high availability infrastructure using VPLEX Metro and Veritas High Availability solutions tested by EMC E-Lab™ for shared data access and high-availability.

### Active/Passive cluster configuration

- Veritas Cluster Server (VCS)
- Storage Foundation and High Availability Solutions

### Active/Active cluster configuration

- Storage Foundation for Cluster File System (CFS) – for general purpose clustered file system
- Storage Foundation for Oracle RAC – using Cluster Volume Manager (CVM) by configuring Oracle ASM diskgroup over CVM

Refer to Oracle *Real Application Clusters (RAC) on Extended Distance clusters with EMC VPLEX Metro for Best Practices Planning*.

## Implementing VPLEX Metro with Veritas clustering technologies

This section provides information for implementing VPLEX Metro with Veritas clustering technologies.

### Extended Veritas VCS deployment topologies in a VPLEX Metro environment

EMC VPLEX breaks the physical barriers of data centers and allows users to access data at different geographical locations concurrently. In an extended Veritas VCS clusters with VPLEX Metro, the active cluster nodes are running on a dedicated VPLEX cluster. For example, in a 4-node cluster implementation, Nodes 1 and 2 are active nodes and Nodes 3 and 4 are passive cluster nodes. In case of an unplanned failure on site with active nodes, the passive nodes at the surviving site take the active role and the application continues with minimal downtime.

Figure 3 shows an example of Extended Veritas VCS clusters with VPLEX Metro.

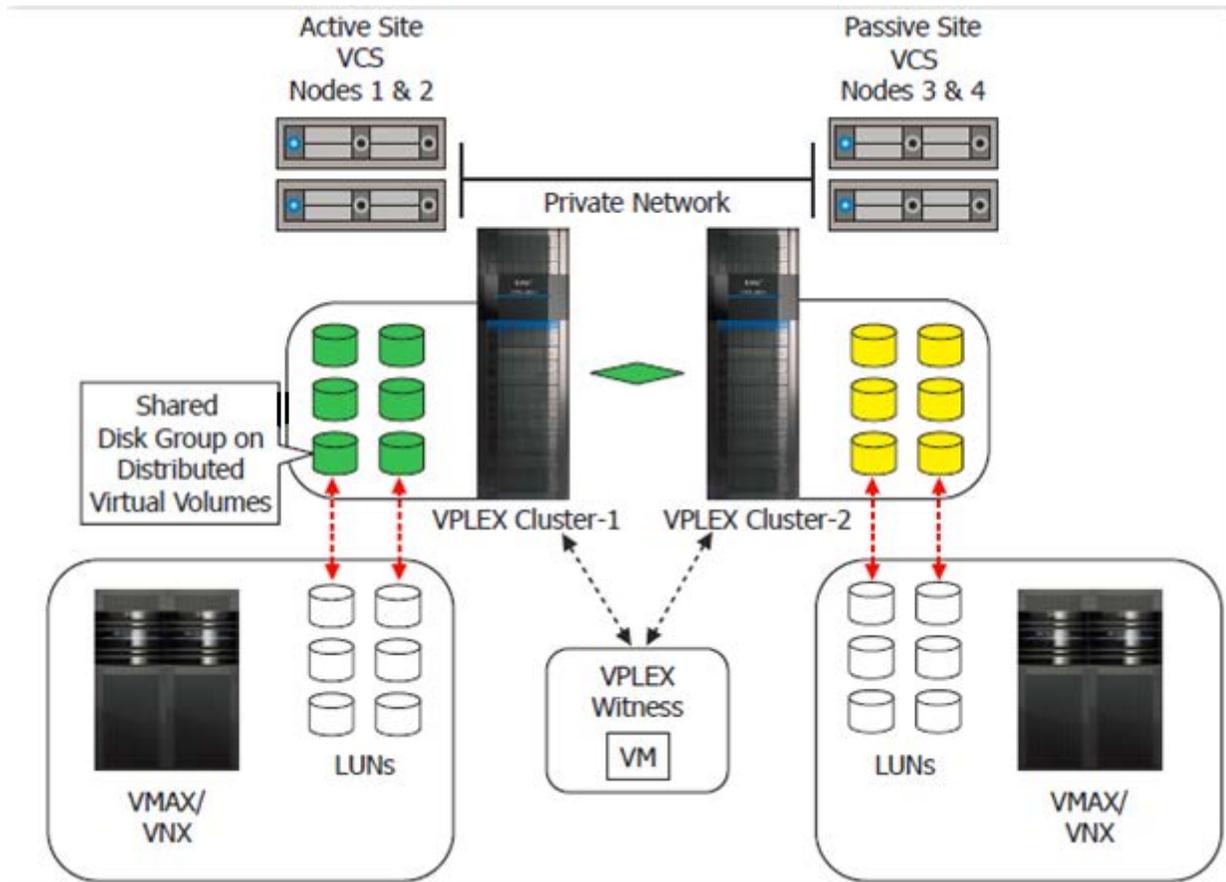


Figure 3. Extended Veritas VCS clusters with VPLEX Metro

## Extended Veritas SFRAC/SFCFS deployment topologies in a VPLEX Metro environment

With an active/active clusters such as Veritas SFRAC or SFCFS, the application runs without any downtime as all the nodes are active at both the sites. The Veritas clusters can be physically extended up to 100km distance between the sites with a maximum of 5 ms round trip latency. For example, in a 4-node cluster implementation, each site has 2 nodes zoned to the VPLEX cluster. In case of unplanned failure on one of the site, the surviving site keeps the data available to the cluster volume manager / file system.

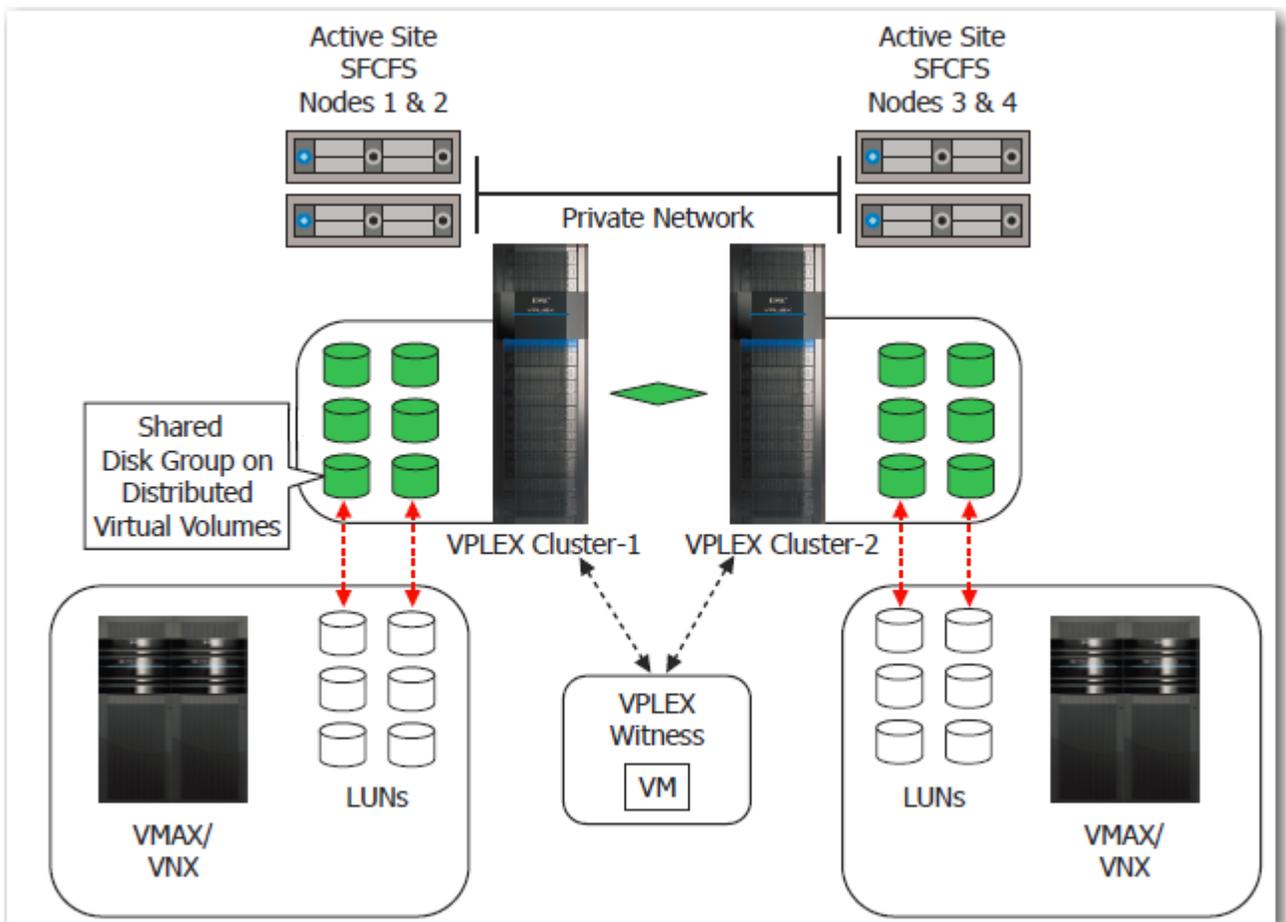


Figure 4. Extended Veritas SFRAC/SFCFS clusters with VPLEX Metro

## Failure scenarios and expected impacts

The combination of extended Veritas HA solutions and VPLEX Metro provides improved availability and resiliency to many failure conditions and therefore increased the availability of mission applications.

Refer to the *EMC VPLEX with GeoSynchrony Product Guides* for more information regarding VPLEX connectivity best practices.

Table 2 and Table 3 summarize the list of failure scenarios and the best practices that will make the host cluster able to continue operations in each of them.

**Table 2. Host failure scenarios in VPLEX-Metro and impacts**

| Scenario                       | Veritas VCS (Active/Passive)  | Veritas SFHA with SFCFS/SFRAC (Active/Active)   |
|--------------------------------|---|---|
| Host HBA port failure          | Multiple host initiators are required to avoid single point of failure. With EMC PowerPath® and Veritas DMP (VxDMP) multipathing software, it does automatic path failover and load balancing. In Active/Passive cluster configuration, HBA failure can cause slight performance degradation since the hot stand-by nodes do not participate in load balancing. | In an active/active cluster configuration the HBA failure on a single node will not cause a noticeable performance impact as the workload is shared on other active nodes. In the event of a path failure to the storage array, EMC PowerPath and Veritas DMP software automatically selects the next available path for I/O requests without any intervention for the administrator. |
| Host hardware failure or crash | Surviving nodes on the Active site handle the I/O traffic without any downtime.   | Since the active nodes are distributed across the sites, the application keeps running even in the impact of complete host failure in a given site.   |
| Site failure                   | Minimal downtime implied until host and application can resume operations. The cluster diskgroup is automatically failed over to hot stand-by passive site and application need to be restarted at the failover site. In case of a failback to the primary site another restart of application is required.   | By installing VPLEX clusters and Witness in independent failure domains (such as another building or site) it becomes resilient to a lab, building, or site failures.<br><br>The VPLEX cluster in the failure-domain not affected by the disaster will continue to serve I/Os to the application.   |

Table 3. VPLEX failure scenarios and impacts

| Scenario                                  | VPLEX behavior   | Extended Veritas SFHA cluster impact   |
|---|--|--|
| Single path failure – Front-end (FE) port | VPLEX will continue to provide access to the distributed volume and cluster host multipathing will be expected to failover to the alternate paths.   | None   |
| Single path failure – Back-end (BE) port  | VPLEX will switch to provide access to the Metro distributed virtual volume via alternate paths to the same alternate paths to the same BE array and continue to provide access to the Metro distributed volumes exposed to the hosts.                       | None   |
| VPLEX hardware component failure          | VPLEX has fully redundant directors and power supplies.  | None   |
| VPLEX Interconnect partition              | If both sites are still available, the VPLEX Preferred Cluster detach rules will determine which cluster resumes I/Os and which suspends, without downtime for host connected to the surviving Cluster.  | Cluster nodes connected to Preferred VPLEX cluster will not be impacted without any application downtime.<br><br><b>Note:</b> In an Active/Passive cluster it is very important to configure active cluster in the Preferred VPLEX Cluster to avoid host downtime.       |
| VPLEX cluster unavailable                 | VPLEX Witness will allow I/O to resume at the surviving VPLEX cluster. <i>Refer VPLEX Metro Witness Technology and High Availability TechBook – VPLEX Witness failure schemantics section for various possible VPLEX Witness failure handling scenarios.</i> | Cluster nodes connected to the surviving VPLEX cluster will continue operations with 0 RTO (Recovery Time Objective).<br><br><b>Note:</b> In an Active/Passive cluster it depends if the active site goes down an application restart is required with minimal downtime. |

Although the architecture of VPLEX Metro is designed to support concurrent access at multiple locations, the current version of the product supports a two-site configuration separated by synchronous distance with a maximum round trip latency of 5 ms between the two sites. In addition, extended Veritas HA solutions with VPLEX Metro may require extension of a VLAN to different physical data centers to facilitate client connectivity and Cluster node interconnect. Technologies such as Brocade's VPLS and Cisco's Overlay Transport Virtualization (OTV) can be leveraged to provide the service.

# VPLEX Metro with extended Veritas HA lab configurations

## Lab configuration and setup

The following sections describe the technologies and components used in the test cases documented in this paper.

### Physical environment

Figure 5 illustrates the overall physical architecture of the extended SFHA configuration deployment that was used for the tests shown in this paper. The extended SFHA consisted of four cluster nodes, two at each simulated data center (Site 1 and Site 2).

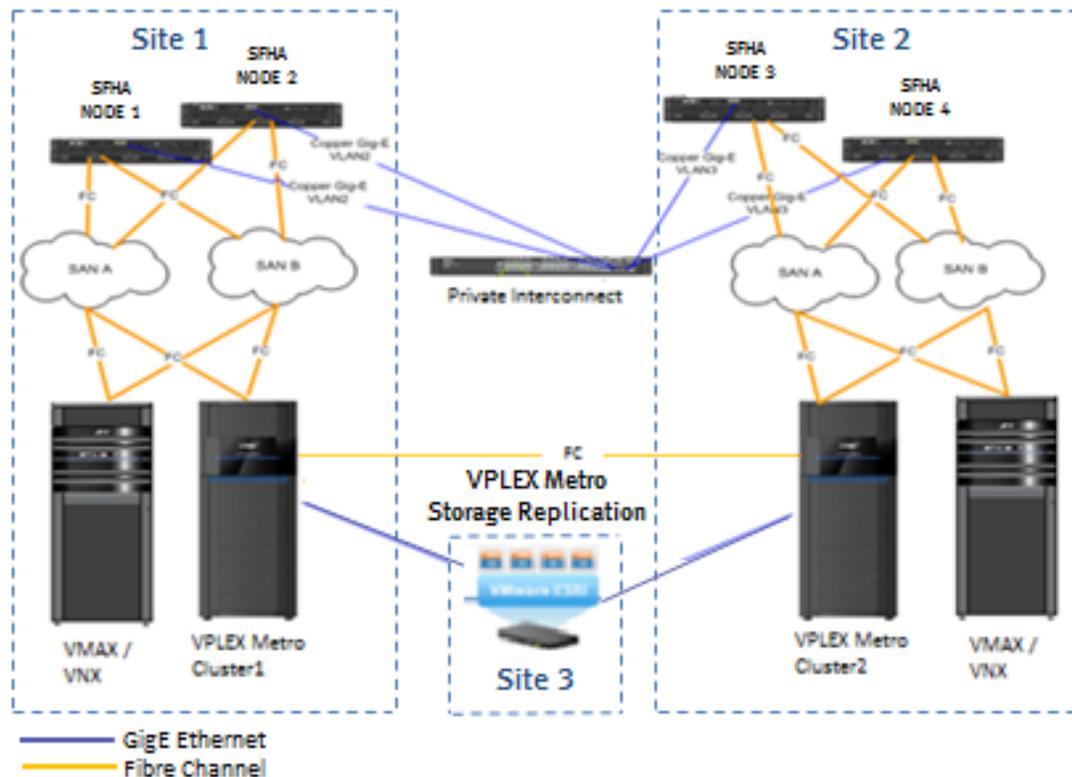


Figure 5. VPLEX Metro configuration for extended Veritas HA solutions

The hardware setup included setting up Fibre Channel connectivity among the hosts, VPLEX clusters, and EMC storage arrays (VMAX/VNX) to the redundant fabric switches. In each site, the hosts are zoned to their local VPLEX cluster front-end ports and the storage to the VPLEX cluster back-end ports. The VPLEX clusters are zoned to each other for synchronous storage replication. The Veritas cluster Private Interconnect between sites uses GigE Ethernet for IP connectivity.

Table 4 shows the VPLEX, hosts, and VMAX storage array hardware details. They included 2 x 2-engine VPLEX clusters setup in Metro configuration, and 2 x 1-engine VMAX Series technology providing local storage to each VPLEX cluster. The 4 x RAC nodes consisted of Dell 2850 dual core servers with 8 GB of cache. VPLEX Witness was set and connected to each VPLEX cluster.

**Table 4. Extended Veritas HA hardware environment**

| Hardware                                     | Quantity | Release and configuration  |
|--|----------|--|
| EMC VPLEX Metro                              | 2        | VPLEX Metro with GeoSynchrony 5.2<br>Two engine and four directors on each cluster               |
| VPLEX Witness                                | 1        | Dell R900 running VPLEX Witness virtual machine (VM)   |
| VMAX Series                                  | 2        | Single engine VMAX with Enginuity 5876,<br>112 x 450 GB/15k FC drives using Virtual Provisioning |
| VNX Series                                   | 2        | VNX7500 with 60 x 300 GB/10k SAS drives  |
| Standard x86_64 Server (cluster nodes)       | 4        | 2 x dual core, 8 GB RAM  |
| QLogic QLA2462 (2 HBA ports used per server) | 4        | Dual-port, 4Gbps Fibre Channel-to-PCI-x 2.0<br>266-Mhz Adapter                                   |

The main guidelines during the connectivity step are to maximize hardware redundancy such as using two switches, more than a single HBA port, multipath for dynamic path failover, and load balancing.

Table 5 describes the host software used.

**Table 5. Host software**

| Software   | Release                                   |
|--|---|
| Server OS  | Red Hat Linux Release 6 Update 3 [x86_64] |
| EMC PowerPath  | Version 5.7.1 for Linux x86_64            |
| Veritas Storage Foundation High Availability for Linux | VCS 6.0.1 and SFCFS 6.0.1                 |

The storage setup and device allocation planning are not outlined in this white paper. It is recommended to provision same size storage volumes on storage arrays connected to both sides of the VPLEX cluster for maximum space efficiency.

## VPLEX Metro setup

This section discusses how to set up VPLEX Metro cluster.

### VPLEX Metro cluster setup steps

Figure 6 lists the main tasks that are required for VPLEX Metro setup.

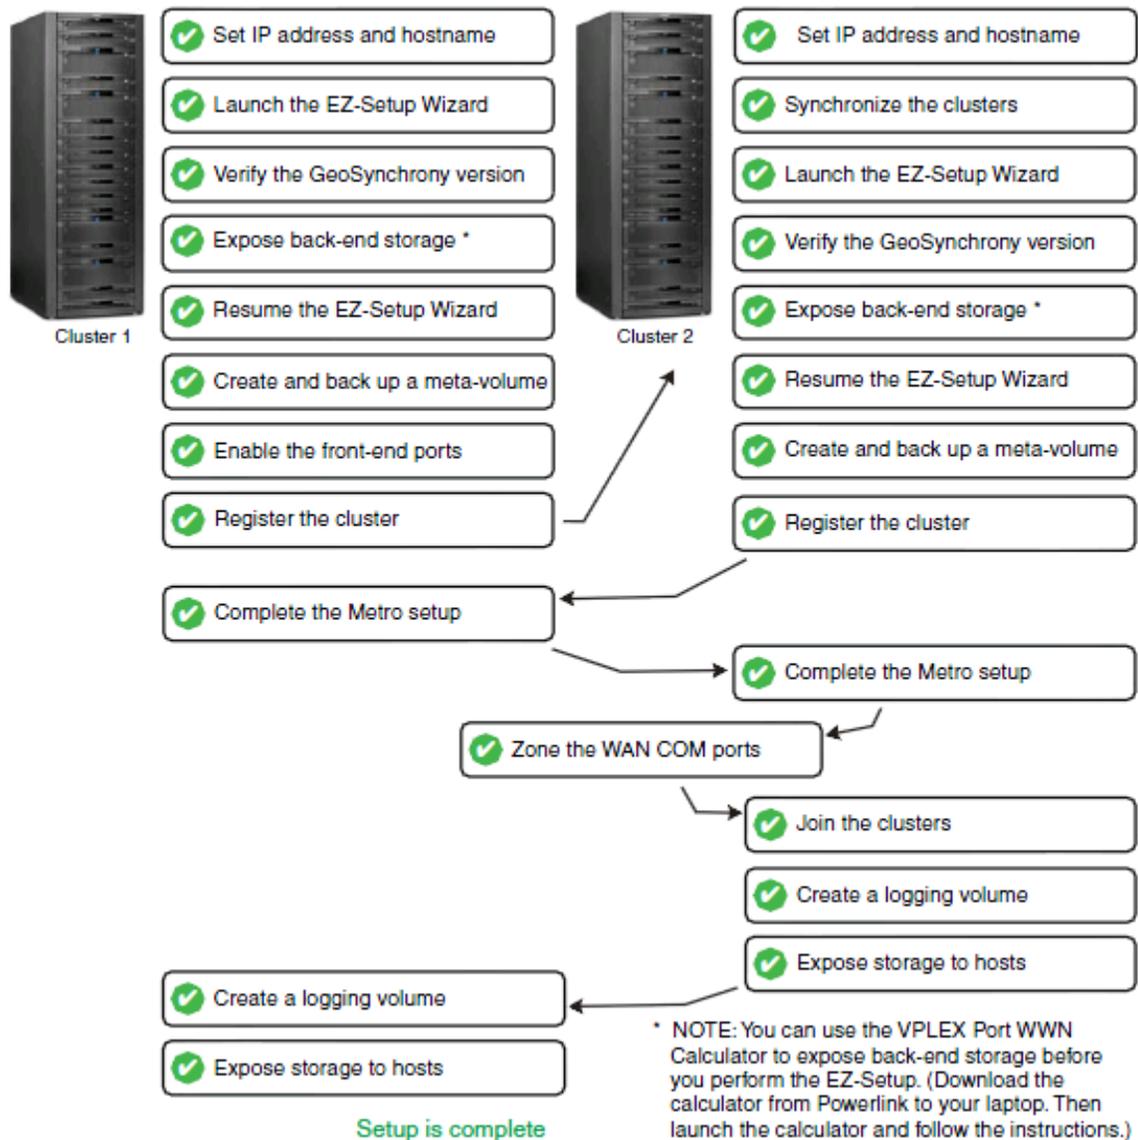


Figure 6. Overview of VPLEX Metro setup tasks

**Note:** You must set up both VPLEX Metro clusters as described. You cannot set each cluster up individually and then join them later.

## Set up VPLEX Metro cluster connectivity

The two key components of VPLEX Metro inter-site communication are FC and IP. Metro can use either Fibre Channel or 10 Gigabit Ethernet for connectivity between directors of each cluster. Connectivity between VPLEX management servers at each cluster are over a secure VPN tunnel. A VPLEX Metro should be set up with redundant (dual fabrics) and completely independent Fibre Channel or 10 Gigabit Ethernet networks between clusters for inter-director communication. This provides maximum performance, fault isolation, fault tolerance, and availability.

To check VPLEX Metro Cluster status, run **cluster summary** command. For example:

```
VPlexcli:/> cluster summary
Clusters:
Name      Cluster ID   TLA           Connected   Expelled   Operational   Health State
-----
cluster-1  1           PMTELABA000010 true        false      ok            ok
cluster-2  2           PMTELABA000009 true        false      ok            ok
Islands:
  Island ID  Clusters
  -----
  1         cluster-1, cluster-2
```

## VPLEX Metro host connectivity

To ensure the highest level of connectivity and availability to extended Veritas High Availability cluster nodes even during abnormal operations, each node in the extended Veritas HA deployment model should have at least two physical HBAs, and each HBA should be connected to front-end ports on different directors on EMC VPLEX. This configuration ensures continued availability of the cluster nodes even if one of the front-end ports of the VPLEX goes offline for either planned maintenance events or unplanned disruptions.

When a single VPLEX engine configuration is connected to an extended Veritas cluster node each HBA should be connected to the front-end ports provided on both the A and B directors within the VPLEX engine. Connectivity to the VPLEX front-end ports should consist of first connecting unique hosts to port 0 of each I/O module emulating the front-end directors before connecting additional hosts to the remaining ports on the I/O module. If multiple VPLEX engines are available, the HBAs from the cluster nodes should be connected to different engines.

The connectivity from the VPLEX engines to the storage arrays should follow the best practices recommendation for the array. A detailed discussion of the best practices for connecting the back-end storage is beyond the scope of this paper.

## VPLEX Metro administration

Administration of VPLEX Metro running GeoSynchrony 5.2 can be done through the EMC Unisphere™ for VPLEX Management Console. Some additional advanced functionality is provided by with VPLEX CLI. On authenticating to the secure web-based GUI, the user is presented with a set of on-screen configuration options, listed in the order of completion. The EMC Unisphere for VPLEX Management Console online help provides more information about each step in the workflow.

Figure 7 shows the operational status of VPLEX-Metro cluster.



Figure 7. EMC Unisphere for VPLEX management interface

In the Provision Storage tab select Provisioning Overview to launch the EZ-Provisioning wizard to claim storage and create virtual volumes. The first step in the process of provisioning storage from EMC VPLEX is the discovery of the storage arrays connected to it. This step rarely needs to be executed since the EMC VPLEX proactively monitors for changes to the storage environment. The second step in the process is the “claiming” of storage that has been exposed to EMC VPLEX. The process of claiming the storage creates the object’s “Storage Volume” that is shown in Figure 8. The Create Storage View wizard enables you to create a storage view and add initiators, ports, and virtual volumes to the view. Once all the components are added to the view, it automatically becomes active. When a storage view is active, hosts can see the storage and begin I/O to the virtual volumes. After creating a storage view, you can only add or remove virtual volumes through the GUI. To add or remove ports and initiators, use the CLI. The *EMC VPLEX CLI Guide* provides comprehensive information above VPLEX Metro commands.

## VPLEX Metro with VPLEX Witness

VPLEX Witness is installed as a closed virtual machine deployed in a failure domain separate from either of the VPLEX clusters (to eliminate the possibility of a single fault affecting both the cluster and the VPLEX Witness). VPLEX Witness connects to both VPLEX clusters over the management IP network. By reconciling its own observations with the information reported periodically by the clusters, the VPLEX Witness enables the cluster(s) to distinguish between inter-cluster network partition failures and cluster failures and automatically resume I/O in these situations.

### Hosts and Veritas High Availability setup

This section details the steps for running Extended Veritas Storage Foundation Cluster File System (SFCFS) on Red Hat Linux 6.3 in an active/active cluster configuration. It is noted that the Extended Veritas Cluster implementation using VPLEX-Metro is the same for all supported OS by Veritas Storage Foundation and High Availability (SFHA) solutions. Refer to the *EMC Support Matrix*, available at <https://elabnavigator.emc.com>, for supported configurations.

The following table summarizes the steps to be taken to configure Storage Foundation Cluster File System (SFCFS). Detailed installation guides are available at [https://sort.Veritas.com/documents/doc\\_details/sfha/6.0.1/Linux/ProductGuides/](https://sort.Veritas.com/documents/doc_details/sfha/6.0.1/Linux/ProductGuides/)

#### High-level steps:

| Step | Action  |
|------|---|
| 1    | Make sure the Veritas recommend OS packages are installed prior to SFCFS package install  |
| 2    | Configure server nodes private network, OS /etc/hosts file, OS kernel parameters and passwordless ssh connectivity between the cluster nodes                |
| 3    | VPLEX-Metro is SCSI-3 PR compliant, uses Coordinator disks for I/O fencing. Refer ESM for sizing and disks consideration.                                   |
| 4    | Run “installer –precheck” script available in the SFHA installation media.  |
| 5    | Run “installer” from the installation media and select Storage Foundation Cluster File System High Availability.  |
| 6    | Run “installsfcfsa<version> -configure to complete the configuration.   |
| 7    | Verify Array Support Library (ASL) for the VPLEX is installed, if not download from <a href="https://sort.Veritas.com/asl">https://sort.Veritas.com/asl</a> |
| 8    | Setup disk-based I/O fencing using /opt/VRTS/install/installsfcfsa<version> - fencing   |
| 9    | Verify I/O fencing configuration “vxfenadm –d” and “vxfenconfig –l”   |
| 10   | Verify all the configured Ilt links and gab ports are started by running “Iltstat” and “gabconfig –a”   |
| 11   | Initialize disks and create a shared disk group. Refer <i>Veritas SFCFS Admin guide to create a cluster resource group</i> .                                |
| 12   | Validate cluster status by running “hastatus –summary”  |

## Veritas Veritas Cluster validation test

Check the status of Veritas cluster and make sure all the resource status are online before proceeding with the validation test.

```
[root@durxha239056]# hastatus -sum
```

| -- SYSTEM STATE |         |        |
|-----------------|---------|--------|
| -- System       | State   | Frozen |
| A durxha239056  | RUNNING | 0      |
| A durxha239057  | RUNNING | 0      |
| A durxha239058  | RUNNING | 0      |
| A durxha239059  | RUNNING | 0      |

| -- GROUP STATE |              |        |              |        |
|----------------|--------------|--------|--------------|--------|
| -- Group       | System       | Probed | AutoDisabled | State  |
| B cvm          | durxha239056 | Y      | N            | ONLINE |
| B cvm          | durxha239057 | Y      | N            | ONLINE |
| B cvm          | durxha239058 | Y      | N            | ONLINE |
| B cvm          | durxha239059 | Y      | N            | ONLINE |
| B vnxdg        | durxha239056 | Y      | N            | ONLINE |
| B vnxdg        | durxha239057 | Y      | N            | ONLINE |
| B vnxdg        | durxha239058 | Y      | N            | ONLINE |
| B vnxdg        | durxha239059 | Y      | N            | ONLINE |
| B vmaxdg       | durxha239056 | Y      | N            | ONLINE |
| B vmaxdg       | durxha239057 | Y      | N            | ONLINE |
| B vmaxdg       | durxha239058 | Y      | N            | ONLINE |
| B vmaxdg       | durxha239059 | Y      | N            | ONLINE |
| B vxfen        | durxha239056 | Y      | N            | ONLINE |
| B vxfen        | durxha239057 | Y      | N            | ONLINE |
| B vxfen        | durxha239058 | Y      | N            | ONLINE |
| B vxfen        | durxha239059 | Y      | N            | ONLINE |

Before running any cluster tests it is important to choose a good I/O generator tool to generate continuous workload. The testing in E-lab was done using in-house cluster aware I/O tool which runs file system I/O on the cluster nodes. The test scenario mentioned before was executed manually and results shows the solution worked as expected.

## Conclusion

EMC VPLEX Metro running the GeoSynchrony operating system is an enterprise-class SAN-based federation technology that aggregates and manages pools of Fibre Channel attached storage arrays that can be either co-located in a single data center or across multiple data centers that are geographically separated by metro distances. Furthermore, with a unique scale-up and scale-out architecture, EMC 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.

Extended Veritas High Availability cluster dispersed in two data centers within metro distance, backed by the capabilities of EMC VPLEX Metro, provides simplified deployment topologies and storage administration, non-disruptive storage scalability and tech refresh.

In addition, the capability of EMC VPLEX to provide non-disruptive, heterogeneous data movement and volume management functionality within synchronous distances enables customers to offer nimble, efficient, and cost-effective cloud services spanning multiple physical locations.

## References

The following documents include more information on VPLEX Metro can be found on <https://support.emc.com>.

- *Implementation and Planning Best Practices for EMC VPLEX Technical Notes*
- *Using VPLEX Metro and RecoverPoint with VMware SRM in an Oracle Extended RAC E-Business Suite Environment White Paper*
- *EMC VPLEX with GeoSynchrony Product Guides*
- *Conditions for Stretched Hosts Cluster Support on EMC VPLEX Metro White Paper*

For product-specific information, refer to:

- VPLEX at <http://www.emc.com/storage/vplex/vplex.htm>
- Veritas High Availability at <http://www.Veritas.com/products-solutions/families/?fid=storage-fdn-ha>