Dell EMC VPLEX

Note: Evaluator Group research of Storage Virtualization has been discontinued. The following Product Analysis has been archived and is no longer being maintained for updates as of January 2023.

The Dell EMC VPLEX is a network-based, virtualized storage product designed to create high availability (HA) storage between datacenters and provide the ability to move applications and data between storage environments. The Dell EMC VPLEX will compete with other HA options and storage area network (SAN) virtualization, including Hitachi Virtual Server Processor (VSP), along with IBM Spectrum Virtualize Software (SVC) and NetApp Fabric Attached Storage (FAS) series.

Highlights

- VPLEX Controllers
  - In-band virtualization appliance
  - Uses VMAX controller hardware
  - GeoSynchrony firmware operates VPLEX and is supported only on VPLEX systems
  - MetroPoint for three site replication
  - VS2 and VS6 models (VS6 targeted at all-flash storage)
- Offered in Three Configurations
  - Small: One VPLEX engine
  - Medium: Two VPLEX engines
  - Large: Four VPLEX engines
- VPLEX Storage Virtualization
  - Support for Dell EMC PowerMax, VMAX, DMX, and CLARiiON/VNX, and other third-party vendors' back-end storage, including ALUA storage arrays
  - Logical unit number (LUN) aggregation and concatenation with back-end storage
    - Provides pass-through LUNs
    - LUN concatenation, with RAID 0 and RAID 10 support
  - Support for thin provisioned volumes to be moved between arrays
  - VPLEX Virtual Edition as vAPP for lower cost implementation
- Single Cluster from One to Four VPLEX Engines
  - Management server (with secure virtual private network (VPN) in Metro-PLEX)
  - Internal Fibre Channel (FC) switch pair for inter-director communication
  - Each node has five FC modules, and one Gigabit Ethernet (GbE) module
- Management and Support
  - Management provided by management server
  - Four GbE interfaces on management server (dual connect to private VPLEX local area network (LAN))

  - Support for VPLEXcli (command line interface (CLI)) and VPLEX console (graphical user interface (GUI))
  - Call home support via Dell EMC ESRS gateway
  - VMware: vCenter Plug-in, vCOPS integration
- VPLEX Witness
  - Monitors connectivity and validates communication for VPLEX clusters
  - Operates in a separate failure domain
  - Automates and validates failovers and restarts
- Integration with Dell EMC RecoverPoint for continuous local and remote replication
- Support for VMware vStorage API for Array Integration (VAAI) and vStorage APIs for Storage Awareness (VASA)
- Support for up to 1,600 host initiators per local VPLEX configuration
- Support for large block inputs / outputs (I/Os) over 10 GbE interface to increase performance and bandwidth utilization
- Unisphere for VPLEX management interface
- Site Recover Manager (SRM) Suite for VPLEX for health reporting, capacity planning, and performance metrics
- Tiering of data to cloud providers

Overview

VPLEX leverages Dell EMC's common hardware design now employed in the VNX and the Symmetrix VMAX. The hardware has the identical form factor, with I/O connection differences from the other products. The embedded software (firmware) running on VPLEX is known as GeoSynchrony. The distributed caching design of VPLEX originated with startup YottaYotta, which was acquired by Dell EMC in September 2008.

VPLEX systems may be joined together into an architecture that joins or federates multiple systems together. Federation is also known as loose clustering or a shared nothing cluster. This type of architecture allows a single copy of data to be shared, accessed, or relocated anywhere within the federation, regardless of location. Each node in the VPLEX system maintains a map of where data is located. No data or resources are shared between the clustered nodes; instead, the mapping table provides a reference to where specific data is located.

Evaluator Group Comments: Dell EMC's VPLEX is designed for open system environments, specifically VMware workloads. Although it does provide virtualization capabilities, it is not marketed as a general-purpose storage virtualization product.

The VPLEX system is targeted to address several specific issues in datacenters that are difficult to address with any single system design. These include:

- Non-disruptive workload relocation of data between systems for maintenance or performance
- Non-disruptive workload resiliency during unplanned events
- Pooled (storage) resources without physical location limitations

VPLEX Local

A VPLEX Local deployment provides data migration and failover capabilities for volumes within a datacenter. This capability is supported with storage from Dell EMC and from a limited list of competing vendors' storage systems.

VPLEX Metro

A VPLEX Metro system provides a single LUN to be actively shared across metro distances without failover. This capability does not exist in competing storage virtualization technologies today.

MetroPoint provides three-site replication with two sites for local continuous replication and a third site as a remote disaster recovery (DR) copy.

VPLEX Geo

A VPLEX Geo system provides a single LUN to be actively shared across geographic distances using asynchronous technology without failover. This capability does not exist in competing storage virtualization technologies today.

VPLEX Virtual Edition

VPLEX Virtual Edition is a vAPP intended as a low-cost, entry-level replication solution that supports Internet Small Computer Systems Interface (iSCSI). It can be used for replicating data in and out of cloud service providers, in addition to between two sites.

Product Use Cases

Dell EMC VPLEX does not easily fit into any existing product category. Although its features and design most closely resemble other network-based SAN virtualization products, VPLEX is not supported by Dell EMC for general SAN virtualization uses. The specific use cases VPLEX is designed and supported are provided below.

VPLEX Local Uses

- Non-disruptive data migration between Dell EMC and non-Dell EMC storage within a datacenter
- Support for VMware vMotion migration between two systems

VPLEX Metro Uses

- Non-disruptive data migration between Dell EMC and non-Dell EMC storage within a datacenter
- Create a volume that is presented to two active locations simultaneously (as shown below)
- Synchronize and replicate data between two sites, up to 100 kilometers (km)
- Support for VMware vMotion migration between two datacenters
- Support for Microsoft Hyper-V Live Migration to move a virtual machine (VM) between datacenters

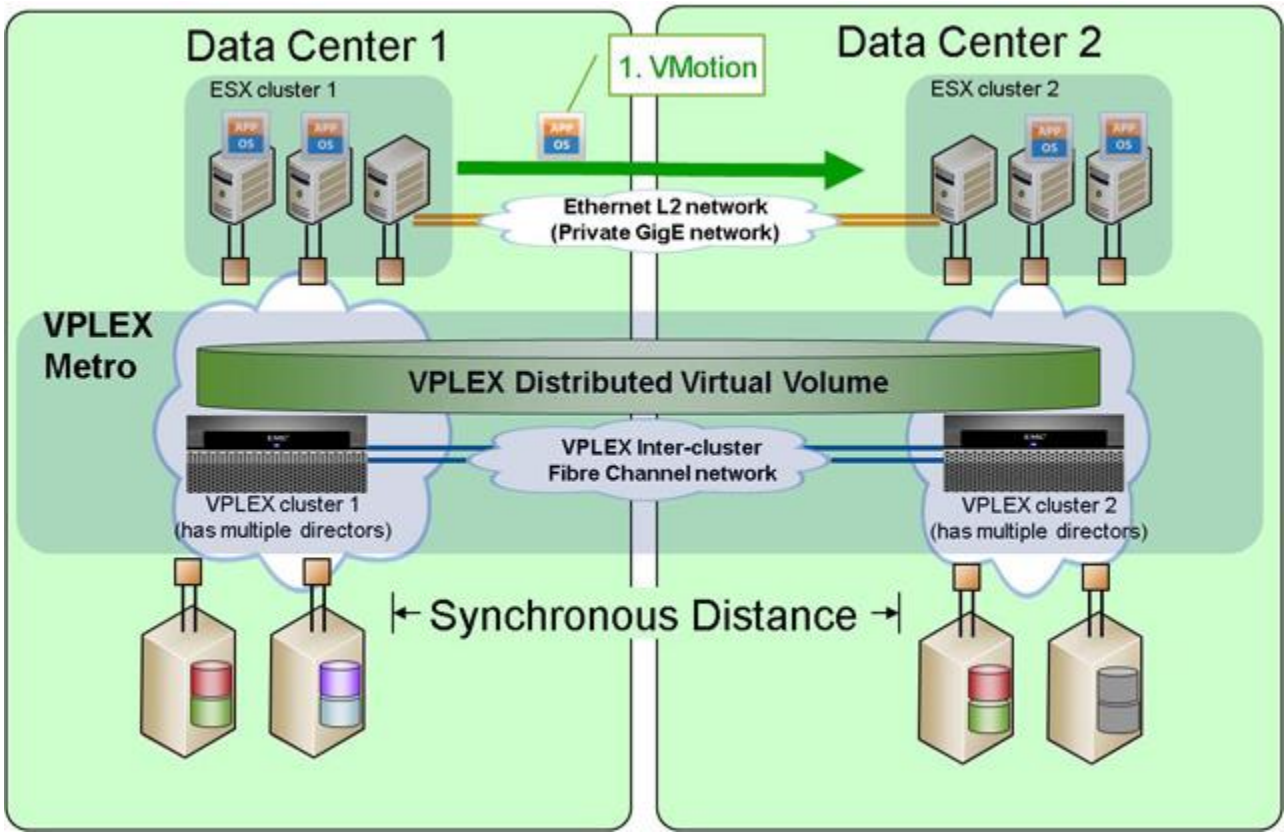

Figure 1: Dell EMC Metro, VPLEX with a Mirrored Volume

VPLEX Geo Uses

- Non-disruptive data migration between Dell EMC and non-Dell EMC storage within a datacenter
- Create a volume that is presented to two active locations simultaneously (a stretched cluster)
- Synchronize and replicate data between two sites, using asynchronous technology where the latency is 50 milliseconds (ms) or less
- Support for VMware vMotion migration between two datacenters
- Support for Microsoft Hyper-V Live Migration to move a VM between datacenters

Product Architecture

The VPLEX system has design aspects similar to a shared disk cluster in that each VPLEX controller has access to all data within the system. In this respect, VPLEX aggregates multiple physically distributed caches into one logical cache, through a high-speed interconnect and a cache-mapping table. However, **each node is assigned ownership of a data space, or LUN**, which makes a VPLEX most similar to a federated cluster.

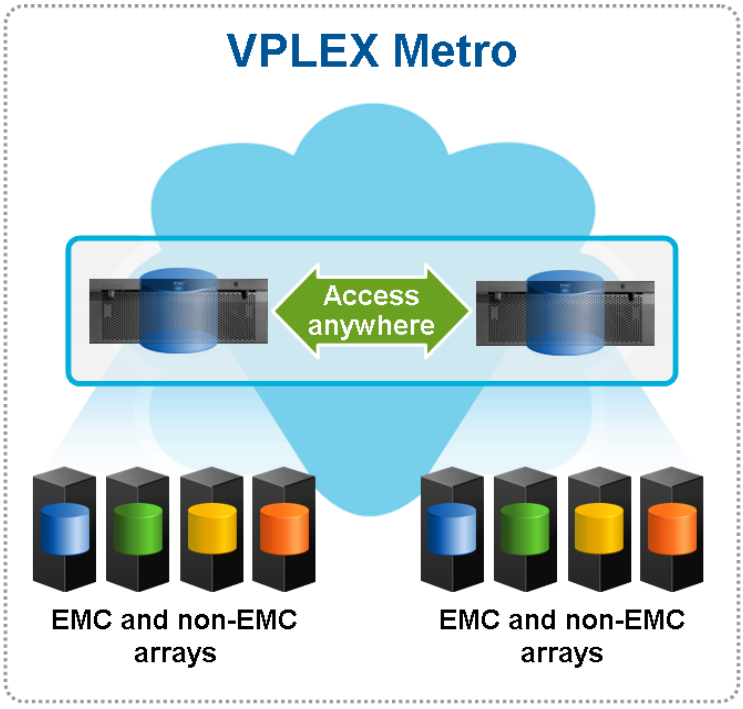

Figure 2: Dell EMC VPLEX (Source: Dell EMC)

Dell EMC accomplishes distributed cache coherence using techniques that compensate for latency and bandwidth issues that can cause problems when physical distance separates federated nodes. The VPLEX's cache coherency technology was acquired with Dell EMC's purchase of YottaYotta in 2008. This technology is trademarked by Dell EMC as AccessAnywhere, and is planned for inclusion within more of Dell EMC's products in the future according to Dell EMC statements.

Evaluator Group Comments: A federated architecture is well suited to scaling up data that may be portioned. Workloads that maintain local reference of data are best suited to a loosely coupled or federated approach. Additionally, the VPLEX's support for distributed mirroring can support workloads that access shared data between sites.

The caching tables are also maintained on non-volatile storage in a dedicated meta-data volume. VPLEX meta-data volumes reside on back-end connected storage and require at least 78 Gigabytes (GB). This meta-data is also mirrored between multiple back-end storage systems for redundancy. Dell EMC additionally recommends that meta-data volumes are backed up or replicated periodically to maintain resiliency.

Model Comparison

	VPLEX Local VS2	VPLEX Metro VS2	VPLEX Local VS6	VPLEX Metro VS6
Basics				
Max Initiators	1,600	3,200	1,600	3,200
Max Capacity	No limit	No limit	No limit	No limit
RAID Levels Supported	0,1, 10, Concatenate	0,1, 10, Concatenate	0,1, 10, Concatenate	0,1, 10, Concatenate
Engines per Cluster	1-4	1-4	1-4	1-4
LUN's				
Max Back-End LUNs	8,000	16,000	12,000	24,000
Max Front-End LUNs	8,000	16,000	12,000	24,000
Distributed LUNs	N/A	5,000	5,000	5,000
Front-End LUN size	100 MB - 32 TB	100 MB - 32 TB	100 MB - 32 TB	100 MB - 32 TB
Back-End LUN size	No limit - 32 TB	No limit - 32 TB	No limit - 32 TB	No limit - 32 TB
Replication / Migration				
Max Migrations	25 (local)	25 (remote)	25 (remote)	25 (remote)
Transfer Size	4 KB - 32 MB	4 KB - 32 MB	4 KB - 32 MB	4 KB - 32 MB
Connectivity				
FC—Front-End	8 - 32 @ 8 Gb/s	8 - 32 @ 8 Gb/s	8 - 32 @ 16 Gb/s	8 - 32 @ 16 Gb/s
FC—Back-End	8 - 32 @ 8 Gb/s	8 - 32 @ 8 Gb/s	8 - 32 @ 16 Gb/s	8 - 32 @ 16 Gb/s
Inter-Engine Communication	8 Gb/s FC	8 Gb/s FC	40 Gb/s InfiniBand	40 Gb/s InfiniBand
Physical				
Controller Size	4 U	4 U	4 U	4 U
Mgmt Node	1U	1 U	In base engine	In base engine

Table 1: VPLEX Model Comparison

The figure below shows a logical view of the functions within a director above the physical enclosure.

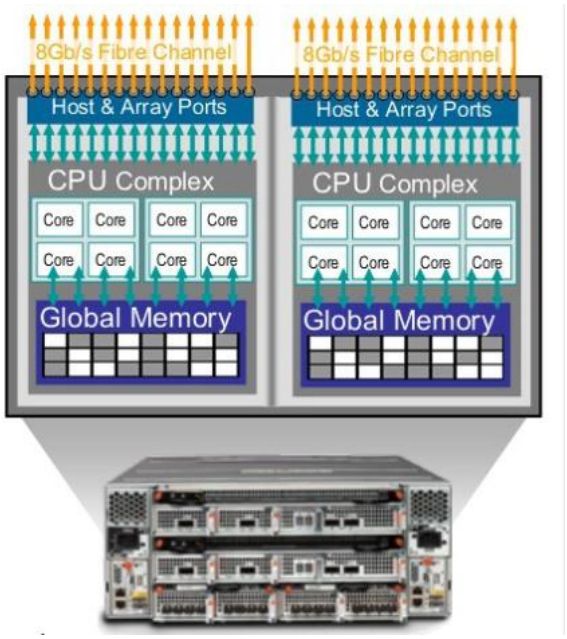

Figure 3: Dell EMC VPLEX Director (Source: Dell EMC)

Federated clusters are most suitable to solving data access and storage problems that are partition-able, or those that are easily divided into workloads assigned to specific nodes.

In order to accomplish transparent network virtualization between systems, Dell EMC VPLEX requires a specific Cisco technology available on Cisco Nexus switches. The Cisco Overlay Transport Virtualization (OTV) technology uses encapsulation within an Ethernet frame to redirect data transparently. Cisco OTV provides the following features:

- Provides media access control (MAC) address routing, to extend a virtual LAN (VLAN) over any distance
- Extends layer 2 across distances by encapsulating an Ethernet frame within an Internet Protocol (IP) packet
  - This provides the ability to route the frame to another location
  - Once encapsulated, any IP-capable device can route the frame
- Supports multi-pathing, multi-homing, and multicast replication
- Automatic load balancing

Data access in a federation differs significantly from traditional replication. With traditional replication, there is a primary site and a secondary site for failover or recovery of data. Only one site may access or update information at a time. In contrast, with VPLEX, data is active at both sites at the same time.

Evaluator Group Comments: Dell EMC's VPLEX is an in-band, virtualization appliance, designed as a loosely coupled cluster or federated cluster.

All virtualization appliances move the primary point of control from a storage system to the point of virtualization. Moreover, the primary point of control and value will become VPLEX devices when they are deployed. Back-end storage devices will necessarily become less important. Dell EMC VPLEX may be used locally to enable non-disruptive technology refreshes of back-end storage. Non-disruptive VPLEX updates are supported by moving data to a remote system.

VPLEX is well suited to enable VMware and Hyper-V migrations, because it allows an application to run at either site without the data needing to be physically relocated. Although there is a performance penalty while the data is being moved, this capability allows access to the same volume simultaneously at two different locations. This is a unique feature of VPLEX.

The VPLEX does not currently support point-in-time data protection points, commonly known as snapshots, or TimeFinder clones using the VMAX/Symmetrix name. The ability to transparently access and move data within a VPLEX cluster removes much of the need for replication products, such as Symmetrix Remote Data Facility (SRDF). However, without snapshot support, customers will be forced to rely on other, much less elegant or efficient means of protecting specific data points in time. This could be a high-value enhancement to VPLEX.

VPLEX supports LUN aggregation and partitioning, known as volume concatenation and slicing. RAID levels supported on VPLEX local include a pass-through LUN (aka volume encapsulation), along with RAID C (concatenation), RAID 0, and RAID 1. Additionally, VPLEX Metro supports distributed RAID 1 (similar to a stretched RAID 10 configuration). A feature known as "Remote Export" is supported that allows a volume from a VPLEX cluster to be presented at a remote VPLEX cluster. Write-through caching is used on VPLEX, so that data is written to back-end storage before writes are acknowledged back to the host.

Remote Device

A remote device, or distributed LUN, requires the use of a logging volume on each node. The logging volume maintains a record of blocks written during a link failure. Without a logging volume, a full resynchronization of every distributed device is required after a link failure. The logging volume requires 10 GB on each node for every 320 terabytes (TB) of distributed LUN's in a Metro VPLEX.

Distributed Cache

Dell EMC VPLEX uses a distributed cache coherency model that is similar to other federation models. A single global cache-mapping table resides within each node. The built-in cache manager is responsible for maintaining consistency, and updating the cache table. A cache map is similar to a bitmap table used for other purposes, such as point in time snapshots and asynchronous replication. A bitmap contains

exactly one bit to indicate status for an entire block. The bitmap represents whether a particular block has changed between two nodes, as shown in the figure below.

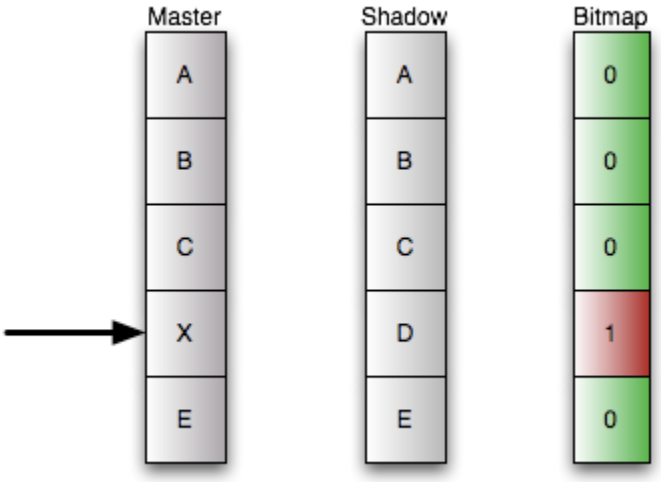

Figure 4: Bitmap Example

In this example, a block is different between the master and shadow copy, which is represented by exactly one bit. As a result, bitmaps can represent large amounts of capacity. The cache coherency ensures that the bitmap is synchronized between nodes. As a result, only a relatively small amount of data must be transferred between VPLEX systems.

Cache Coherency

Cache coherency is maintained through two different functions. A local node cache manager maintains local cache and communicates with back-end storage. A distributed node cache manager is responsible for cache coherence, protection of cache updates, invalidating cache after updates, and node failure recovery.

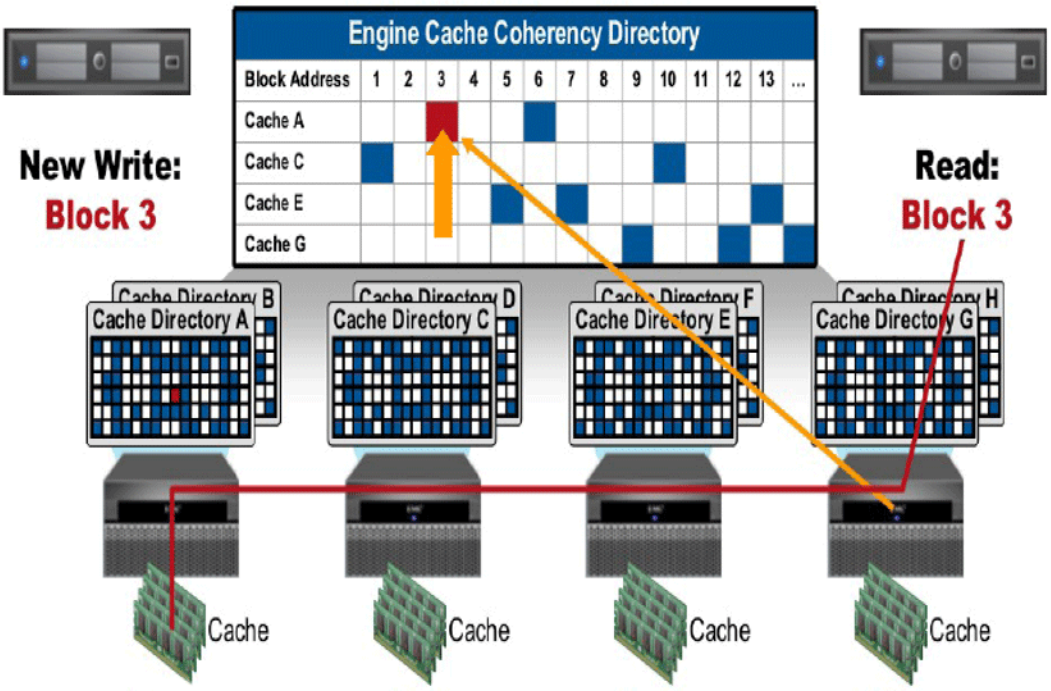

Figure 6: Dell EMC VPLEX Distributed Bitmap Cache

The distributed manager uses a directory with one directory per host volume. The directory is split into chunks with 4,096 entries per chunk. Each directory uses a bitmap as previously described, with one directory entry per global cache page. The directory entries list which VPLEX node owns a LUN, which node last wrote to that LUN, and any maintaining access locks.

The directory chunk is managed by the meta-directory, which maintains chunk ownership. Directory chunks can migrate between nodes. Each cache page entry tracks the following on a per VPLEX node basis:

- Page ownership
- Last node to perform a write
- Bits to provide LUN access locking and queuing to locks

Logging Volumes

The VPLEX system includes two 50 GB solid-state drive (SSD) devices (mirrored for HA) on each node for use as a boot and logging device.

A logging volume is required in order to create a distributed device, or a remote device. A logging volume must exist within each node of a VPLEX cluster. The logging volume tracks data written during an inter-cluster link failure. This logged data is used to re-synch a volume by sending only changed blocks.

Logs are also a bitmap, using one bit for every page of distributed storage space. Thus, 10 GB of log space handles 320 TB of distributed capacity.

Cache Reads

When a read request is received, the VPLEX cache manager searches the local cache. If the data is present, it is returned. If not in the local cache, the global cache is searched. If data is present in the global cache, the local nodes cache is updated. If data is not in either cache, it is serviced directly from the back-end storage, with both the local and global caches on the requesting nodes being updated.

Cache Writes

All writes use a write-through mode, so that data is written to back-end storage before the write is acknowledged to the host. Both local and global caches are updated with the new data (if it exists) prior to data being written to the back-end storage.

VPLEX Metro/Geo—Distributed

When a LUN is configured as a distributed LUN, this model is most similar to an active-active configuration, with access available via either controller. Read operations will occur with limited delays, although writes will incur some penalty to update the remote side. If the link fails, a rule can stipulate which cluster will continue processing I/O, while the other side of the mirror remains suspended. The Small Computer System Interface (SCSI) reservations are used to fence access when communication issues exist.

Evaluator Group Comments: Due to the overhead with searching cache, invalidating and updating cache, and then writing data to the back-end system, delays can occur when multiple VPLEX nodes attempt to access the same LUN simultaneously in a VPLEX Metro or Geo configuration. This is similar to the "ping-pong" effect of controller failover between controllers when multi-pathing round-robin software is used on active / passive systems. The expected delay would be at least one round-trip delay for cache lookup and invalidation (up to 10 ms), plus the delay for one roundtrip data transfer (10 ms plus data transfer time), plus processing overhead. This operation is likely to take at least 30 ms, which can be significant in I/O-intensive operations that are susceptible to latency. Therefore, Dell EMC does not recommend using a remote shared LUN in this manner, but only for data movement, migration, or failover scenarios. This is best described as an "active-passive" mode with primary access occurring via one path, or one node.

However, Dell EMC has internal benchmark results of an application without VPLEX, and after inserting a VPLEX appliance into the data stream. In the case of local I/O, the impact is less than 5% in most cases. In the case of a VPLEX metro solution, the impact ranges from 5% to 15% depending on the workload. In read-intensive environments, the result can be a performance increase due to the positive effect of caching.

In all cases, WAN bandwidth is a critical factor in the performance of VPLEX Metro configurations.

Controllers

The Dell EMC VPLEX consists of one director, which houses a pair of highly available controllers. This was shown previously in Figure 4. VPLEX directors are physically and logically very similar to those used in the Dell EMC VNX and the Dell EMC VMAX systems. For more information and analysis of those systems, see the product analysis by Evaluator Group of these individual products.

RAID Groups

The RAID group size and the number of disks in a RAID group can vary within the limitations for support. Table 2 below lists the officially supported RAID configurations.

Volume Type	Comments
Encapsulated	Volume encapsulated from back-end storage (no translation)
RAID 0	Striped aggregate of LUNs on local VPLEX
RAID C	Concatenation of LUN's on local VPLEX
RAID 1	Mirror of LUN's on local VPLEX
Distributed RAID 1	Mirror between two VPLEX systems (Local or Metro or Geo)
Remote Export	Presentation of a LUN by one VPLEX, controlled by another VPLEX

Table 2: VPLEX Local and Metro—Supported RAID Groups

Drive Support and Sparing

Dell EMC VPLEX systems do not directly use or manage devices for use by hosts. All device models, sparing, and other features are managed by the underlying storage systems providing back-end capacity to VPLEX systems.

LUN Management

LUNs, or volumes, are managed both on the back-end storage systems providing storage, and within a VPLEX cluster. The types of volumes supported are listed above in Table 2. Expansion and concatenation features are supported by VPLEX.

Virtual / Thin Provisioning

Thin provisioned volumes may be migrated by a VPLEX.

RAID Migration

Up to 25 migration sessions can run concurrently. A RAID migration includes:

- Convert source from RAID 0 to RAID 1
- Convert a source from RAID 1 to RAID 0

Data Encryption

No support for encryption.

Storage Virtualization within a VPLEX

The levels of storage virtualization within a VPLEX from bottom (storage) to top (host) are:

1. Back-End LUN: Storage volumes from back-end storage
2. Extent: Any subset of the capacity on a storage volume
3. Device: Configured from one or more extent in RAID 1, or RAID 0, or concatenated (RAID-C) configuration
4. Virtual Volumes (vVols): Provisioned to the host
5. VPLEX Front-End Ports
6. Registered Initiators: An HBA identified by “VPLX”
7. Storage View: Grouping of ports, initiators, and vVols for LUN mapping and masking

Reliability, Availability, and Serviceability Features

VPLEX uses a "no single point of failure" hardware design. All major components are deployed in pairs. The VPLEX director contains a pair of redundant controllers. Each director contains redundant power supplies, processor and memory boards, fans, I/O connectivity, and other features.

Connectivity to front-end hosts should be configured via redundant FC SAN connectivity. Connectivity to back-end storage should also be configured via redundant FC SAN connectivity.

The VPLEX GeoSynchrony operating system provides software fault tolerance, along with cache consistency and mirroring as described previously.

Mirrored Cache

Cache is mirrored and distributed between all controllers within a VPLEX federation or cluster.

Concurrent Maintenance

Maintenance during operations are supported, including firmware, upgrades are possible on each controller within a director.

Remote Call Home (Limited Availability)

Dell EMC VPLEX supports Dell EMC's common call home platform, known as Dell EMC Secure Remote Support (ESRS) gateway. This is the same call home system used by Symmetrix VMAX, VNX, and other Dell EMC systems.

VPLEX Witness

VPLEX Witness is an optional feature to provide a third-party verification of the operation and control of failover and restart of a VPLEX node. Operating as a VM deployed on a customer's ESX server in a different failure domain than VPLEX, VPLEX Witness can monitor communications and connectivity to determine the health of a VPLEX cluster. The VPLEX Witness is connected over IP and will coordinate seamless failover and restart and solve the "split-brain" problem that might potentially occur with two communicating systems with a partial failure occurring.

Evaluator Group Comment: The overall RAS characteristics of VPLEX are good from both a hardware perspective, as well as a software or firmware design. By leveraging Dell EMC's existing fault-tolerant hardware platform, the VPLEX is able to provide similar HA features.

Performance

Currently, VPLEX has no independent performance benchmark results available.

Dell EMC internal benchmarks show the total I/Os for a four-director cluster can scale up to 3,000,000 I/Os. The bandwidth of a four-director cluster is 23.2 GB/s with the latest VS2 generation hardware.

Large block I/Os over a 10 GbE network interface are supported to increase performance and bandwidth utilization.

Connectivity

Back-End Storage

Dell EMC maintains an extensive support matrix. Dell EMC currently claims support for products including IBM DS8000, Hitachi, HPE, NetApp, and others.

Multi-Path I/O to Hosts

Multi-pathing support from a VPLEX to hosts is provided through Microsoft MPIO for Microsoft platforms. Additional multi-path support is provided using Dell EMC PowerPath.

Multi-Path I/O to Back-End Storage

Dell EMC includes multi-pathing and failover capabilities within VPLEX to back-end storage, when redundant path and port configurations are used.

Evaluator Group Comment: Always check with vendor for the latest supported configurations.

Advanced Features

Common advanced features include management and data protection options.

Management:

Unisphere for VPLEX is the management tool that provides a common look and feel to the management for VMAX and VNX. A separate VPLEX management server is used for Unisphere. One port is connected to a dedicated management LAN. The other connections are provided for redundant connections in-rack private networks.

The management server hardware has the following features:

- The management server supports, Secure Shell (SSH), Hypertext Transfer Protocol Secure (HTTPS), and Internet Protocol Security (IPSec) VPN connections
- For a remote configuration, two management servers must be connected to communicate, over a VPN
- VPLEX management has three roles: service, administrator, and user.

VPLEX has both Unisphere for VPLEX (a management GUI) and a CLI known as VPLEXcli.

Data Migration

Up to 25 migration sessions can run concurrently. This process is essentially moving from a local RAID to a distributed RAID 1.

Migration occurs as:

- Convert source from RAID 0 to RAID 1
- Add target as a mirror
- Synch the mirror
- Promote the target to be primary
- Remove source device convert back to RAID 0

Data Protection

RecoverPoint provides data protection for VPLEX with continuous local and remote replication. An appliance is used (the Recovery Point Appliance (RPA)) to perform the replication using either write splitting technology in the storage array (Dell EMC arrays) or intercepting write I/Os to arrays without write splitters (third-party arrays) in a virtual storage layer. Storage volumes at the primary and secondary sites are assigned to a consistency group to define the set of data to be replicated. RecoverPoint will ensure data consistency and write order fidelity for the set. A replication policy is established by the administrator in RecoverPoint for the replication actions for each set of data to be replicated. The replication behavior changes during operation based on the policy settings.

The RecoverPoint appliance is software running on a 64-bit Linux kernel. A single appliance can manage multiple storage groups, each with different policies. Two RecoverPoint appliances are required at a minimum for HA.

<u>Evaluator Group Comment:</u> Using RecoverPoint for data protection with VPLEX is a very wise move by Dell EMC. The RecoverPoint solution has great capabilities for continuous protection—far advanced over many other solutions. By providing a standard protection technology across multiple platforms, Dell EMC simplifies administrative efforts in many environments.

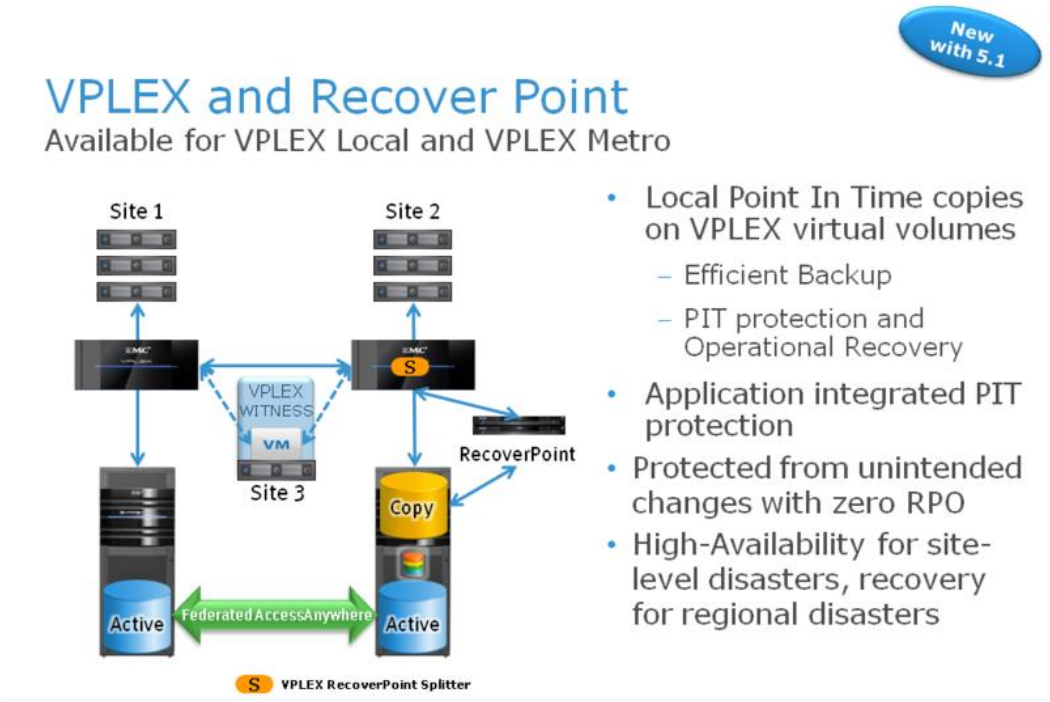

Figure 7: Dell EMC VPLEX and RecoverPoint (Source: Dell EMC)

Tiering

Through integration with CloudArray, VPLEX can tier data to cloud providers. Support clouds include VMware vCloud Air, Microsoft Azure, Amazon S3, and Google Cloud Platform.

Evaluator Group Comments

Dell EMC's VPLEX is a unique product. The local version of the product may be used to migrate data from one system to another, increasing the resiliency and performance of existing storage.

Strengths:

The VPLEX heavily leverages the Dell EMC VMAX hardware, which may help decrease costs and on-site sparing, while improving overall reliability through the user of proven hardware.

VPLEX provides interesting data sharing capabilities over metro distances. The in-band approach to storage virtualization is a straightforward and proven approach to storage virtualization compared to the earlier Dell EMC Invista. Invista was introduced as a general storage virtualization appliance, while VPLEX is positioned in specific deployment scenarios. The VPLEX announcement shows that Dell EMC has learned to create specific targeted deployment scenarios and show examples of specific problems that VPLEX solves.

Evaluator Group also notes that Invista represented a hybrid in-band / out-of-band approach to storage virtualization. It appears that after years of debate over the merits of in-band versus out-of-band virtualization, Dell EMC is now supporting the in-band approach, joining Hitachi, IBM, and NetApp, which are also in this camp.

With VPLEX, Dell EMC has taken the concept of using an in-band appliance and evolved it into a scale-out clustering architecture that provides N+1 performance, as well as N+1 resiliency (the VPLEX cluster increases both performance and resiliency as the cluster grows).

Potential Concerns:

Once a storage virtualization device is in use, it becomes the primary point for managing data access and protection. Managing data copies or access at any other point is not possible, due to the possibility for data inconsistencies. Snapshots or point-in-time copy features on back-end devices will provide limited use, unless tightly integrated with VPLEX device. This tight integration is likely to have advantages with Dell EMC arrays, such as VNX and Symmetrix. By using pass-through LUNs or "encapsulated" devices on a VPLEX, customers may use the underlying storage systems data protection features.

Evaluator Group advises customers to focus on the stated use cases: VM migration over distance, non-disruptive data migrations, and back-end storage technology refreshes. In particular, VMware or Hyper-V workloads deployed across a geographic span to enable non-disruptive HA capabilities are use cases for which VPLEX is well suited.

For customers looking to deploy one of the specific workloads highlighted by Dell EMC, VPLEX is a unique and effective tool.

More detailed information is available at http://evaluatorgroup.com