

EMC VPLEX™: Solutions for Local and Distributed Federation

**Silverton Consulting, Inc.
StorInt™ Briefing**

Introduction

Today almost everything is increasingly more mobile. Acknowledging this direction, EMC recently announced its new VPLEX family and its first two generally available products– VPLEX Local™ and VPLEX Metro™. Before EMC VPLEX, information mobility was only available through the use of special purpose technologies, specifically designed applications and systems that tolerated long I/O latencies and other uncommon requirements. Such special accommodations were simply not available to an overwhelming majority of IT operations, large and small.

Leveraging a new and unique cache coherent architecture, EMC VPLEX is now offering these same mobility benefits to customers without the requisite expensive and single use application and system redesign. Specifically, these noteworthy benefits include:

- **Balanced data center workloads and improved resiliency** – with multiple data centers, optimal workload distributions can be difficult to maintain. Workloads at any particular data center change frequently. Without data processing mobility among the data centers in a group, each site can be over- or underutilized. However, with effective data processing mobility, peak workloads at one data center can be quickly moved to an alternate, less busy site. In addition, scheduled outages are less problematic as a data center’s workload could be easily redirected to another member of the federated group.
- **Improved data center economics/energy savings** – with energy emerging as a more significant expenditure, operating cost differences between data centers in a group are becoming more pronounced. In fact, because of demand pricing, energy costs vary significantly not only according to time-of-day but also to data center location. With data processing mobility among the federated data centers of the group, more energy intensive workloads could be moved to an optimum energy site and result in noticeable savings.
- **Automated disaster avoidance** – with many companies today, adequate disaster plans for business operation continuity is an absolute necessity. In slow-forming disaster situations, data processing mobility could be used to move workloads to another, unaffected data center and thus, sustain ongoing business operations.

EMC VPLEX™: Solutions for Local and Distributed Federation

In today's global business environment, these advantages are indeed significant and maybe even vital to a global federated group of data centers. When large, worldwide companies have adopted this trend toward rapid data processing mobility, the cost effectiveness and resiliency of this facet of their overall business will certainly be improved.

Components of the federated data center

What's needed to bring about data processing mobility? This capability is enabled by virtual machine (VM) mobility, network mobility and data mobility. Of these three components, in the first two, VM and network mobility, the pivotal technology already exists to make a straightforward move to the federated data center group. For example, both VMware's VMotion and Microsoft Hyper-V's Quick Migration can already move VMs from server to server within a single, local data center; network mobility requires a system address reconfiguration and does not require technology advances.

In contrast, the third component, data mobility, requires new, emerging technology advances. Using older existing technology, terabytes (TBs) of data can be migrated across a network but the process is oftentimes unacceptably lengthy. The emerging technology avoids the slow migration process by leaving the data placement intact and adds advanced caching technology to allow concurrent access to the data by multiple data centers. At an extreme, a data center could in fact be a "storage-less" site and still process data efficiently using this advanced caching technology.


EMC VPLEX technologies

EMC has made data mobility a reality by introducing its VPLEX system architecture, which enables a federated data center group. Each VPLEX system is arranged in a cluster of engines. VPLEX Local provides local federation across heterogeneous arrays whereas VPLEX Metro permits two clusters within limited distances to access the same physical LUN. The VPLEX system architecture uses distributed cache coherence, active-active data sharing and scale-out performance, all new technology innovations developed by EMC.

Distributed cache coherence

Historically, accessing the same data from two locations would entail long I/O latencies. EMC's VPLEX architecture minimizes these latencies by introducing a distributed caching system in front of data center storage. Using VPLEX Metro in a multi-site setup, a portion of highly used data is quickly copied from one data center's storage to another data center's cache. This advanced caching technology thus provides rapid access to the highly used data in the cache while retaining the bulk of the data at its original location.

EMC VPLEX™: Solutions for Local and Distributed Federation



EMC's VPLEX system architecture also advances caching technology by introducing cache coherence at a more granular block level. This allows the engines in a VPLEX cluster to simultaneously update the same data by creating a global cache directory for all data blocks being managed by VPLEX. In turn, the global directory maintains the current state of data held in all VPLEX cache engines, even between sites. Thus, any update to data is simultaneously reflected in the global directory of all VPLEX systems and can be retrieved from an updated location by any federated data center.

Active-active data sharing

EMC's new VPLEX architecture allows bi-directional active-active data sharing where a federated group of data centers could simultaneously access and update data copied to both locations. However, to maintain data coherency under these conditions, the global directory records updates to copies of the data on storage as well as in cache.

Note, active-active data sharing cannot support simultaneous writes to the same data block without incurring synchronization delays. For example, applications like travel reservation programs where many users are all vying for the last available seat, may be impacted. But for the vast majority of cases where write updates are distributed throughout a set of data

with many readers, VPLEX active-active data sharing involves no additional synchronization delays.

Scale out performance

Traditionally, many caching solutions were limited to one or at most two units and thus limited in performance. But enterprise storage demands that any caching system must scale in performance along with capacity. EMC has accomplished this necessary scalability in its VPLEX architecture by using a cluster of engines that operates as one system to service the I/O workload, providing coordinated, coherent access to data. A VPLEX cluster could therefore increase performance noticeably, whether adding more engines to the VPLEX cluster itself or installing additional clusters in the data center.

EMC VPLEX solutions

Using distributed cache coherence, active-active data sharing and scale out performance, EMC's VPLEX solutions provide enterprise class operations with demanding availability, reliability and serviceability specifications. Currently, EMC

EMC VPLEX™: Solutions for Local and Distributed Federation

offers two versions of its VPLEX system architecture, one providing federation within a single data center and another for distributed federation across a group of data centers.

EMC VPLEX Local

For the single data center, VPLEX Local provides non-disruptive capabilities across heterogeneous arrays, including:

- **Local data mobility** – being able to move data, non-disruptively, from one physical storage subsystem to another
- **Centralized administration** - being able to define and configure storage LUNs for multiple physical storage subsystems within a single unified management console
- **Resource pooling** – being able to create and allocate storage across different storage subsystems within a data center

The VPLEX environment also provides ongoing technology refreshes by moving data from an older subsystem to a newer one without disrupting normal IT operations. Additionally, customers may benefit from storage tiering with VPLEX Local as data can be moved to higher or lower performing storage transparently.

VPLEX Local supports one, two or four engines in a cluster. Each VPLEX engine consists of two redundant directors, with each director supplying 16-8GFC ports (8-storage subsystem and 8-host connection ports). Having two directors in each engine supports high availability with a combined total of 16-host and 16-storage links. Each VPLEX engine is connected to the other engines in the cluster by FC interconnects. VPLEX Local also supports native backend array functionality, including EMC Symmetrix™, EMC CLARiiON™, HDS 9970/9980 and USP-V™, and IBM DS8000™. For any size cluster (1, 2 or 4 engines), VPLEX Local can import up to 8,000 physical storage elements and create up to 8,000 virtual volumes.

EMC VPLEX Metro

To support federated data centers located within synchronous distances (same building, same campus, or between two different data centers or sites within 100km) of each other, EMC introduces VPLEX Metro. This offering enables two VPLEX Metro clusters to interconnect two data centers at synchronous distances and supports up to 16,000 LUNs. With the distributed data access and cache coherency provided by this system, federated data center groups can now access all data across both locations without having to move the data. Additionally, VPLEX Metro's unique capabilities can provide support for disaster avoidance and business continuity concerns of federated data centers.

With VPLEX Metro, LUNs can be located in separate locations and, if necessary, data can be active-active shared across both of them, providing maximum flexibility. VPLEX Metro also supports storage-less sites that may be ideal for use in remote office and central hub configurations.

EMC VPLEX future enhancements

In addition to the Metro and Local configurations described above, EMC has also announced plans for two future versions of its VPLEX system architecture:

- **VPLEX Geo™**, which will support distributed data access for federated data centers at asynchronous distances (up to thousands of kilometers apart). Such capabilities could be used to move work across continents combining two data centers separated by vast distances into one federated environment.
- **VPLEX Global**, which will support multi-cluster federations (more than two clusters connected at one time) across asynchronous distances, capable of being deployed across the same building, campus or multiple geographically disparate sites. Such capabilities would supply follow-the-sun data center operations where data processing would continue from one continent to another.

Thus, the planned EMC enhancements should further expand the distance and configuration capabilities present in the current, already impressive VPLEX Metro and Local products. Indeed, future VPLEX offerings should provide first-rate, global distributed data access, cache coherency and shared data to benefit the federated data centers of tomorrow.

Summary

Federated data centers with mobile processing are now a reality. EMC has successfully developed a new technology to create the missing link, data mobility, along with its requisite coherent caching and desired bi-directional active-active data sharing. As such, new federated data centers will be more efficient, economical and resilient under ever increasing workload demands and economic complexities.

While EMC's VPLEX Local and Metro systems available today are already impressive, the promised future EMC VPLEX Geo and Global systems will further enable even larger, continent and worldwide federated data centers. Thus, EMC's future systems will eliminate the physical data processing constraints of today's enterprise class system customer and allow true 24 X 7 global operations.

Silverton Consulting, Inc. is a Storage, Strategy & Systems consulting services company, based in the USA offering products and services to the data storage community.