THE CLIPPER GROUP
Navigator™

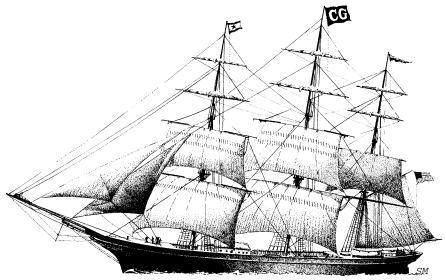

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2010023 May 10, 2010

EMC Bends the Speed-of-Light Problem — VPLEX Delivers a New Storage Vision

Analyst: Jim Baker

Management Summary

It is called the *Speed-of-Light Problem* and it has been challenging the computer industry from its very beginnings. It is variously described but for this paper, it will be defined as how long it takes signals running down the path to get from point *here* to point *there*. The inherent delay may not be that long, but it is nonetheless real, and it is long enough to introduce problems with consistency of data and getting it delivered where needed.

We learned about the speed of light in school, when our teachers made the preposterous suggestion that it took some time between the time a beam of light left the Sun and the time the light actually reached us here on Earth at our favorite beach. It was not as instantaneous as most of us thought. Instead, there was a built-in transit time delay. This concept was made even clearer when later considering what happens when a star burns out somewhere in deep space. If the star is far enough away, we could be seeing it (or more precisely the light from it) in the night sky even though the star no longer existed. The light was still inbound and could have been so for millions of years, because of the time the light had spent traveling great distances to reach us.

The late United States Navy Rear Admiral Grace Hopper (inventor of the COBOL programming language) used to carry a nanosecond demonstration around in her briefcase. She would pull out a bundle of copper wires each 11.8 inches in length and hand them out to her audience. She explained that we each were holding a nanosecond in our hand because that was the distance that an electron traveled down a copper wire in one billionth of a second in a non-vacuum. Her comparison for a microsecond was a full spool of wire nearly a thousand feet long. Even when using fiber optics, the speed of light limits how fast data can be moved from one location to another.

EMC (and every other large IT vendor) routinely has made the joke that they were "working on the Speed-of-Light problem" and likely would solve it. With the introduction of EMC's new prepackaged, distributed cache solution called *VPLEX*, certain aspects of the Speed-of-Light problem are now being addressed. To learn how EMC is making this happen and why it is very interesting, please read on.

Storage Limited by the Speed of Light

Consider, if you will, the problems associated with data located on the other side of the world from where it is needed, in what might be a worst-case scenario for Earthly communications. Data that might be residing on a very slow-spinning disk and ultimately must be transferred across a long distance using an infrastructure with varying speeds on various telecommunication legs until it is presented to the server or user requesting it. All too often, however, that delay in delivery is unacceptable.

IN THIS ISSUE

- Storage Limited by the Speed of Light 1
- EMC VPLEX — Something New in the Hardware/Software Stack 2
- VPLEX Delivers 4
- VPLEX Product Family 5
- Other Considerations 8
- Conclusion 9

On the surface, it would seem oh-so easy to address this distance issue. Just make a second copy of the data and keep it closer to those who might want to use it. However, this introduces the possibility that the two copies of data might not be the same (i.e., out-of-sync with each other) and, if so, which copy is the most current, assuming that this is the one that is desired? The storage dilemma surrounding *wanting it now* but (a) not wanting it to be the wrong version, (b) not wanting to have to plan too far in advance, and (c) not wanting to pay too much to solve this problem presents a classic example of *wanting to have your cake and eat it, too*. Something always had to give – at least until now.

EMC VPLEX — Something New in the Hardware/Software Stack

With its announcement at EMC World 2010, EMC has stepped up to the challenge of how to make data more fluid, with its own vehicle for *having your cake and eating it, too*. EMC's approach is one of local and distributed federation, combining hardware and software in a new product family called *VPLEX*.

- *Local federation* is the transparent cooperation between storage elements to enable data to be shared, accessed, and relocated transparently within a single site.
- *Distributed federation* is similar cooperation between storage systems that enable data to be shared, accessed, and relocated transparently across resources, over a distance.

VPLEX creates this new federation layer in front of heterogeneous block storage systems, like EMC's *CLARiiON* and *VMAX*, as well as arrays from other, non-EMC vendors. VPLEX employs a scale-out clustering hardware architecture, which lets you start small and grow easily as the amount of data behind it grows. This is important as enterprises store more and more data across more and more storage systems.

VPLEX uses *advanced data caching* to improve I/O performance and to address storage array contention. VPLEX introduces a new point of control, so that data can be managed in the sea of storage units required to hold all of the enterprise's expanding volumes of data. Access to data from anywhere in the storage network is achieved by inserting a VPLEX cluster into the storage hardware/software stack. VPLEX clusters can consist of 1, 2 or 4 *engines*.

Exhibit 1—
VPLEX's Place in the Architecture

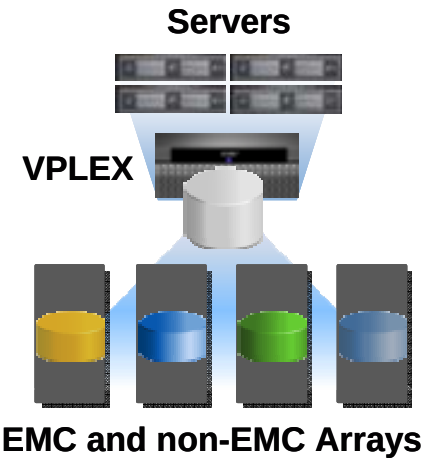

Source: EMC

VPLEX clusters transparently sit between the servers and storage systems (a.k.a. "arrays"). (See Exhibit 1, above.) By intelligently caching data in one or more VPLEX engines, the servers think that they are talking to storage systems but a local VPLEX cluster captures those requests, manages all interactions with other VPLEX clusters (if any) and the underlying storage (arrays), and controls what is stored where, including in its own cache. All VPLEX engines, within a cluster or between two connected clusters, contain the same index to what is stored, so that each engine knows what is stored where, including what is stored uniquely, partially, or redundantly in or behind one or more clusters. (See Exhibit 2, on the next page.) The index also contains metadata (data about data), including usage data and patterns of association with other data. It uses the location(s) of the data and the metadata to optimize reading and writing, by its placement (where it is stored) and by the related data that is pre-fetched (in anticipation), when associated data is retrieved. All of this is done to improve local performance (reading and writing) and to increase availability. This is what makes VPLEX truly special. In addition, a common graphical user interface (GUI) makes usage intuitive and the same, regardless of the involved brands and models of storage systems.

Read and Write Performance with Fidelity

The "prime directive" for storage is to accept whatever is written by the controlling application, store it without altering it, and return

Exhibit 2 — VPLEX Maintains Distributed Cache Coherency

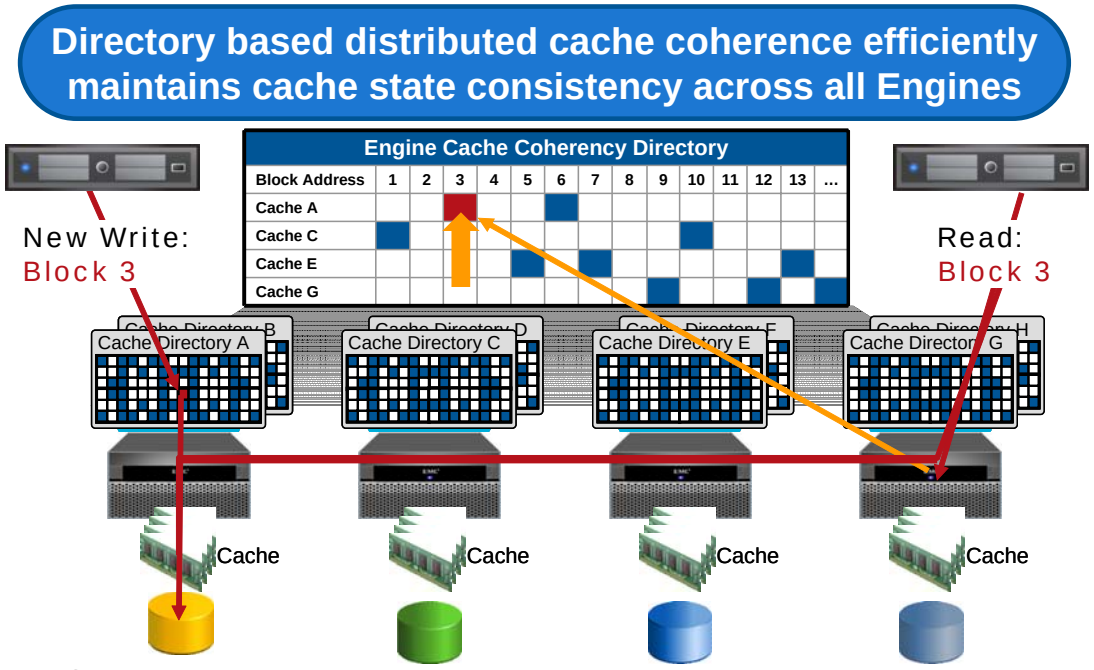

Source: EMC

it when requested exactly as it was originally written – with no changes allowed. Additions (such as metadata) may be attached or associated, but the payload must remain untouched. The VPLEX design adds new intelligence in front of the storage subsystem. This intelligence decides what is to be stored where, with the choices being in its own cache (for the "hottest" data) or in traditional storage systems that sit behind it, which also can be seen as tiers of storage with different performance characteristics. In the one-location sense, this seems simple enough. VPLEX is an inline controller that federates heterogeneous storage and offers new opportunities for increasing data mobility and availability. However, VPLEX was not conceived just for this limited case.

Think of multiple VPLEX clusters deployed in each of your enterprise's data centers, where they sit in front of a collection of storage systems (from EMC and others). Thus, each intelligently controls what is placed where, based upon policies defined and collected usage metadata. VPLEX becomes your storage resource optimization and execution engine.

Now, consider that the VPLEX clusters in your enterprise data centers are automatically aware of each other. Not only do they know what data is where, they can work together to pre-stage the delivery of data, transparently. Some of the data can be duplicated (or stored, even more than once on different engines), to satisfy availability, performance, reliability, and/or compliance requirements. It can be stored in parts and moved between locations, all transparently and without regard to the peculiarities of the behind-the-scenes storage systems being used. Because each VPLEX engine knows where everything is (because each contains an up-to-date index), any of the individual engines can service a request, either from its own cache or the storage systems behind it, or from the cache of another engine or from the storage systems behind that VPLEX cluster.

And there is more. Multiple copies might be stored behind multiple VPLEX clusters, increasing the speed of storing and retrieving in large organizations (i.e., those with many readers and writers of data). Of course, this presents the challenge of knowing what has changed, so that the most current data is or can be made available everywhere within the enterprise. VPLEX manages this through its *universal index*, which is maintained in an identical state on all connected VPLEX engines. VPLEX's method for keeping all this in sync is elegant, and it is made simple because the data in the index is tiny when compared to the large blocks of data that it catalogs.

In effect, each read or write that comes through the VPLEX cluster triggers a look-up in the index to determine whether anything has changed, anywhere. If something has changed,

Exhibit 3 — The VPLEX Engine

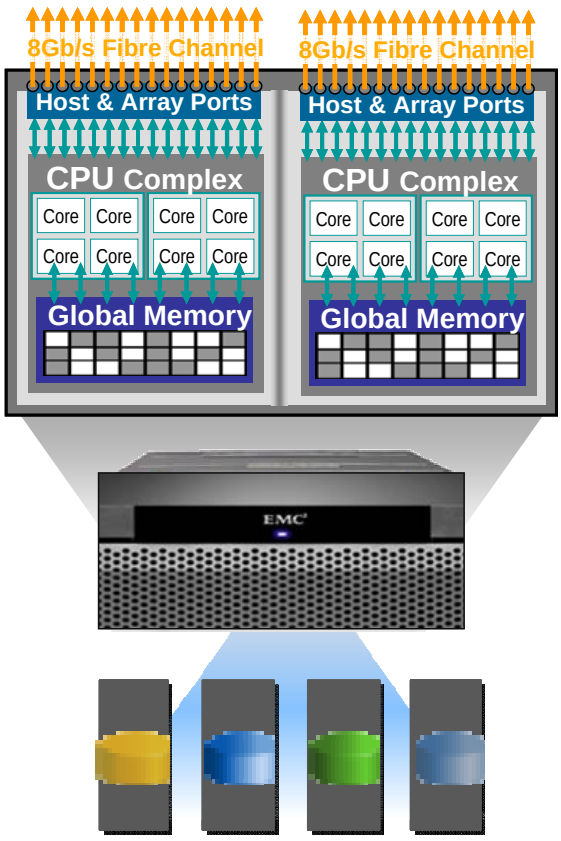

Source: EMC

and the engine controlling the latest data needs only to share what has changed with the requesting engine, the volume of data to be transported has been reduced significantly, especially in comparison with a full transfer of updated block(s)

All of this is done automatically, without human handling and without the involved servers (making the requests on behalf of applications, virtual machines, or users) or storage systems (holding the data) being aware that VPLEX has become the intelligent, controlling intermediary. This automation is what defines VPLEX. Of course, storage administrators also can use VPLEX to move data non-disruptively from storage arrays that are being taken out of service, either to be retired or to be repaired or updated, or to specific locations, say to respond to a planned outage. VPLEX provides the ultimate in ongoing data mobility.

In addition, you might look at VPLEX as an important enabler of private storage clouds. VPLEX provides the transparency required so that the users and applications don't really care about what is stored behind VPLEX or even the location(s) at which the data is being stored. As long as it is being stored, managed, and delivered to the qualities of service specified, the user will not care about the details. The magic is in VPLEX's federation and optimization of the entire storage interaction.

VPLEX Delivers

VPLEX is a prepackaged hardware and software solution[1] based on unique technology that combines scale-out clustering and advanced data caching with innovative distributed cache coherence intelligence.[2] Its goal is to take data location out of the storage equation or at least minimize its mostly negative influence, hence addressing the Speed-of-Light challenges. Up to four VPLEX engines share each other's cache resource. (As shown in Exhibit 3, at the left) The engines themselves include built-in, redundant switching technology. VPLEX configurations are designed for high availability, with dual controllers, redundant power supplies, and integrated battery backup.

For partially solving the age-old Speed-of-Light problem while keeping all of the data coherent, VPLEX is delivered in a deceptively small package. It comes in a single cabinet approximately 75 inches high by 24 inches wide by almost 40 inches deep. This packaging allows one, two, or four VPLEX engines with Intel multi-core processors. Read cache memory is shared and can reach up to 64 GB raw per engine. Each engine has a call-home feature for remote support and diagnostics. Arrays can interconnect with VPLEX using standard Fiber Channel SANs using host fan-in or array fan-out configurations. Its IOPS rating is up to 935,000 I/Os per second for a fully configured quad VPLEX. This figure is, of course, application dependent, as is the maximum transfer rate of 10.8 GB/s. VPLEX cluster capacities are shown in Exhibit 4, on the next page.

[1] Some might call this "an appliance", but that is more applicable to the packaging than the functionality. In many ways, what VPLEX does is more like being an "intelligent virtual switchboard", knowing where data is at all times, connecting the right data to those who want it, and intelligently optimizing all of the associated activities.

[2] For more on data coherency, see *The Clipper Group Captain's Log* dated April 22, 2010, entitled *In Search of Worry-Free Data Coherency Across the Enterprise*, available at http://www.clipper.com/research/TCG2010020.pdf.

Exhibit 4 — Systems Capacities in Each VPLEX Cluster

Max. Virtualized Capacity	No known limit
Max. Virtual Volumes	8,000
Max. Storage Elements	8,000
Min./Max. Virtual Volume Size	100MB/96TB
Min./Max. Storage Extent Size	No VPLEX Limit/96TB
Number of Initiators	400

Source: EMC

Includes Support for Non-EMC Arrays

VPLEX adds cache resources to any storage system including competitors IBM and HDS in May 2010 and NetApp, HP *EVA* series, and IBM/LSI *DS4000* and *DS5000* soon thereafter. The benefits of placing a large cache ahead of round, brown spinning disks has long been a distinguishing characteristic of EMC's top-of-the-line *Symmetrix*, *DMX*, and *VMAX* storage systems. If requests for data can be satisfied by data already in VPLEX's cache, there is no need to incur another disk I/O and perhaps incur rotational latency of the drive itself. Time is saved implicitly as disk accesses are avoided. The cache itself is very fast (approximately, thirty times faster than a 15K RPM disk drive, according to EMC).[3] EMC is placing the VPLEX prepackaged solution in the data path ahead of its and competitors' storage systems, thereby increasing VPLEX's knowledge of and its ability to control placement on these devices. This maximizes interoperability and data sharing. A companion benefit is the easy intermixing of legacy storage systems with new technology storage systems, as they work their way onto the floor.

VPLEX Product Family

The VPLEX family will eventually consist of four products:

- ***VPLEX Local*** – for managing data mobility and access within a single data center
- ***VPLEX Metro*** – for mobility and access between data centers within synchronous distances of each other
- ***VPLEX Geo*** – between two distantly-separated data centers (say, one or more thousands of kilometers apart)
- ***VPLEX Global*** – among three or more distantly-separated data centers

The VPLEX Local and VPLEX Metro offerings will be available this month. The VPLEX Geo offering is planned for 2011, with VPLEX Global being released sometime after that.

VPLEX Local

VPLEX Local use cases are those where only one cluster sits in a single data center. (See Exhibit 5, on the next page, for an illustration of a fully-configured VPLEX Local.) It provides common management to the storage systems that are connected to it. There are three categories of benefits for using VPLEX locally. Each is discussed below.

(1) Enabling mobility between EMC and non-EMC arrays

Support for non-EMC arrays means that users can implement the benefits of VPLEX without having to sweep the floor of older storage, whether from EMC or others. Older arrays (those supported by VPLEX) can remain in use for as long as the customer desires. Data can be moved from and to these storage systems without affecting users, thus reducing unwanted downtime. This is made easy by VPLEX.

If non-EMC storage systems have a particularly desirable feature not yet implemented on EMC offerings, there is no need to forego that feature or function. In this way, EMC is honoring the customer's prior buying decisions, even when the deal did not go in EMC's favor originally.

(2) Simplifying multi-array storage management

By standardizing LUN presentation and management, the differences between the arrays are no longer relevant, except from a performance perspective. VPLEX masks the differences between storage vendors. Users and applications will see no difference in what they read and write to storage.

Simple tools are provided to provision and allocate storage devices for standardized LUN presentation and management. Users easily can pool and aggregate capacity across multiple storage systems even if they are from different vendors and, in the process, increase the productivity of the storage administrators. The result is improved efficiency for each involved storage system.

[3] The comparison point is a read hit on the VPLEX cache at approximately 200 microseconds versus a read miss against a spinning disk of approximately 6 milliseconds.

Exhibit 5 —
VPLEX Local with 4 Engines

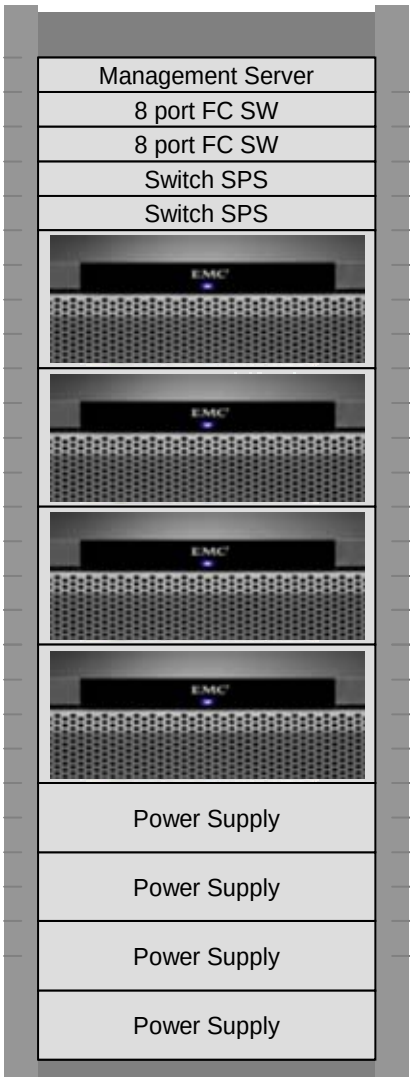

Source: EMC

(3) Increasing resiliency and extending life for existing arrays

EMC is careful to advise that it is important to keep in mind that VPLEX Local is not intended to be used as a one-time migration tool. Due to the number of restrictions at the time of general availability, including the need to disrupt application availability to introduce and remove a VPLEX cluster, existing EMC migration technologies and solutions should be considered when planning a technology refresh. Also, host-level mirroring is not supported. However, VPLEX is a viable solution for ongoing non-disruptive data mobility that can be leveraged for continuing storage system refreshes or moves. This is a better way to use VPLEX (as opposed to just a one-time event).

Adding VPLEX to installed storage farms increases the high-availability quotient for all storage subsystems under its management and control, since it allows data to be mirrored on heterogeneous storage systems. With VPLEX, this is an easy solution.

VPLEX Metro

VPLEX Metro involves two clusters connected within synchronous distance of one another. (See Exhibit 6, on the next page.) This is usually a two-datacenter scenario, whether they are on the same campus or in adjacent metro areas. The classical metro scenario is cross-town from New York City Manhattan to Brooklyn or across the river into suburban New Jersey. When two VPLEX Metro clusters are connected together, they form a *Metro-Plex*. There are four main VPLEX Metro use cases.

(1) Enabling data mobility and workload relocations between sites

The primary use case for VPLEX Metro is relocating workloads and data over synchronous distances (less than 100 kilometers). The benefit to the enterprise is transparently sharing and balancing server and other computing resources between data centers using self-activating policy definitions.

In combination with *VMware*, VPLEX helps enable *VMotion over distance*, allowing users to transparently move and relocate virtual machines (VMs) and their corresponding applications and data over supported distances. This is a huge advantage, makes a compelling argument in favor of VPLEX, and exemplifies the fluidity of storage federation.

(2) Distributing and sharing data across sites

VPLEX Metro also allows data to be distributed and shared across sites. A single copy of data can be accessed by multiple servers in two different data centers. This allows immediate access to information in real time and eliminates the operational overhead and time required to copy and distribute data across locations. As multiple users are sharing the same copy of data, it still remains synchronized. This innovation likely will eliminate or at least drastically reduce overnight batch transmissions of copies (and more copies) of data to remote stores.

(3) Increasing workload resiliency across sites

This feature allows applications (and users) to mirror volumes within, across, and between data centers. This does not require a user to re-

Exhibit 6 — VPLEX Metro with 4 Engines at Two Locations

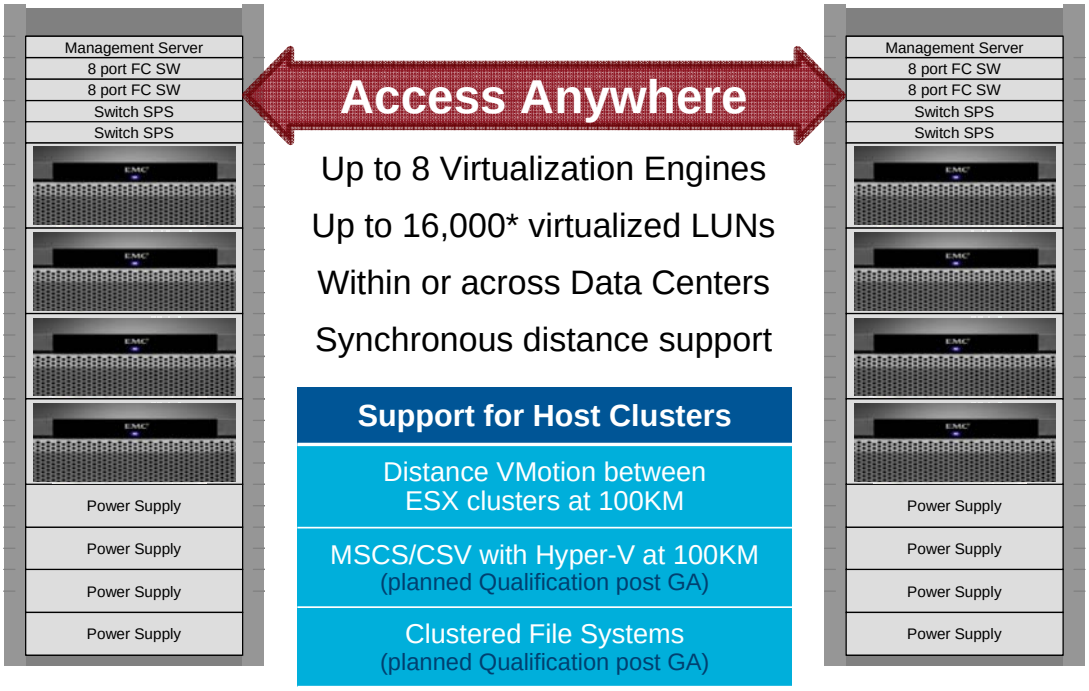

*Total LUNs depend on the specific configuration

Source: EMC

place their current replication technology, and, in fact, VPLEX can provide a level of disaster recovery via continuous operations. VPLEX Metro maintains a consistent, crash-recoverable image of all LUNs shared between the two clusters of a Metro-Plex and it can be configured to operate upon failure of either site or VPLEX Metro cluster. VPLEX eliminates the complexity of swapping target and source volumes at the remote (DR) site by presenting the shared LUNs as always-active at both clusters. However, VPLEX Metro will block I/Os at the remote site, should it lose contact with the primary site (due to failure or link loss). A simple operator intervention is required to re-enable I/O at the remote site in a DR situation. With this configuration, the loss of the DR site will not cause the primary site to stop I/Os. The benefit is increased data availability to the servers and users near each data center by offering a way to mitigate unplanned application outages.

(4) Enabling VMotion over distance

EMC Global Solutions has created working prototypes to show how *SQL Server*, *SharePoint*, and *SAP* applications benefit from *VMotion* over distance, which offers the ability to move an application from one virtual partition to another across a significant distance. Non-disruptive migrations and relocations are typical, even over 100 kilometers.

The importance of showing real-world examples and integrations from other vendors, such as SAP and Microsoft, serves two distinct purposes: business-relevant evidence of the technology's applicability and proof points that entice users to consider how these technologies might be used successfully in their own enterprises.

VPLEX Geo (Planned Future Release)

VPLEX Geo also involves two clusters connected together over distance – but with this product, asynchronous distances will be supported. Asynchronous means that the data flows more slowly between the two clusters and that one may be lagging behind the other (for recently changed data). This is not a trivial difference and will be provide a substantial benefit to customers when it's available in 2011. There are two main VPLEX Geo use cases.

(1) Enabling migrations and relocations between sites over asynchronous distance

VPLEX Geo's asynchronous support allows migrations and relocations over extended distances. Importantly, this is still a two-site scenario. In combination with VMware and VMotion, applications can be moved between locations, geographies, and even service providers

dynamically. IT organizations could implement new operational models leveraging separate data centers or data centers that have a favorable electrical utility contract. Key to success is the ability to move, balance, relocate, and fail over applications transparently between sites, at a distance. The aggregation of all these storage resources across huge distances opens up heretofore-unconsidered business models and unmatched flexibility.

(2) Distributing shared data access across sites over asynchronous distances

By taking limited Fiber Channel distances out of the equation, VPLEX Geo allows cross-continental and transoceanic collaboration. For example, software developers from different geographies can all have shared, quick access to the same software code or test data, not only simplifying and speeding the development process, but also leveraging engineering resources in real time without geographical limits.

Do not underestimate the potential savings that might come from sharing software licenses. There likely is a justifiable quantitative payback from using fewer copies of software while simultaneously increasing the effectiveness of software developers. VPLEX gives them access to servers that might be sitting idle in a distant enterprise data center.

VPLEX Global (Planned Future Release)

VPLEX Global use cases are similar to those of VPLEX Geo, but they have an added level of functionality, allowing for three or more VPLEX clusters to be connected together simultaneously across geographic distances. Obviously, this scenario is the most complex to develop and test, which explains why it is at least two years away. There are two main VPLEX Global use cases.

(1) Global data access and on-demand workload relocation across locations and service providers

This is almost the same as the first VPLEX Geo use case for cooperation and sharing of resources, except now three or more data centers can be involved, with at least one VPLEX Geo cluster in each. Without too great a stretch of the imagination, one can envision an amalgamation of enterprise storage assets in its several-to-many data centers with leased offsite and remote cloud storage assets.

(2) Fully Enabling the Private Cloud

With VPLEX Global, you could truly move thousands of VMs over thousands of miles, fully building out your private cloud. You can move batch-processing applications to locations with lower energy cost. Achieving boundary-less workload balancing and relocation finally is achievable. The future looks bright as you aggregate big data centers from separate ones and increase your duty cycle to 24-by-7-by-forever. Just like other VPLEX scenarios, federated applications can run without ever having to suffer the pains of a restart.

Certain Applications Are Not Ideal for VPLEX

Some applications, by their very nature, may not be a good fit for VPLEX. The most notable ones are mega-transaction processing systems, where the data just cannot be out of synchronization and where data from a previous transaction affects the next one: Stock market systems, credit authorization applications, and airline reservation systems come to mind. Fortunately, as a percentage of enterprise computing and storage, these applications represent a small minority of the total operational environment. The vast majority of IT projects and scenarios will tolerate the few seconds or minutes that it takes to update the VPLEX engines to achieve cache coherency. VPLEX won't move data unless absolutely necessary. Rather, it will send the most used parts to the cache, and then request and extract outliers on demand.

Additional VPLEX Solutions and Services

EMC Global Solutions, in cooperation with VMware, has produced a reference architecture for moving partitioned applications and vBlock data across synchronous distances non-disruptively. This is done using VPLEX Metro in combination with VMware and VMotion over distance, and is currently available from EMC.

Other Considerations

New Opportunities for Storage Administrators

What could Storage Admins be doing if they were not concentrating on the minutiae of today's storage micro-operations, such as LUN definitions, initial and subsequent provisioning, monitoring hot spots, load balancing and integrating, updating and decommissioning storage capacities? VPLEX is a "next-generation" wake-up call. As quickly as the shortage of storage administrators became the new bottleneck potentially limiting IT success, the decade-long

crisis will subside, due to the high level of automation in storage solutions, as exemplified by VPLEX.

If you are a storage administrator, it is time to rethink how you add continuing value to the enterprise IT equation. Less intimate and detailed caretaking will allow increased focus on the data (content), the business requirements for that data, the policies under which it needs to operate, and, still, lowering the total cost of storage, all while following the Four Storage Commandments: *Make storage (1) easier, (2) more transparent, (3) more efficient, and (4) more effective*[4]. For storage administrators, this is an opportunity to upgrade your career.

Architecture First, Products Next, Confidence Always

It is not unusual for an IT vendor to describe its new seemingly-always-significant and breakthrough architecture to the potential buying community. As analysts, we hear these proclamations weekly. Most users, appropriately, are skeptical upon first hearing another "new architecture" story for the first time. However, potential buyers tend to lower their skepticism dramatically when reference architectures have been documented and tested, when successful beta testimonials from credible customers are available, and when vendor's plans, complete with roadmap timelines, can be viewed. EMC has provided all three with this announcement. VPLEX Local and VPLEX Metro have been announced. VPLEX Geo is expected during the first half of 2011 while VPLEX Global is due out "later", presumably within a few years of VPLEX Geo's release. Nonetheless, you may have to try VPLEX to believe it is real, which is common in large enterprises. No doubt, VPLEX will create compelling new value in the right environments.

Conclusion

VPLEX constitutes a bold vision – a *game changer*, so to speak. It jostles with the Speed-of-Light problem and cleverly reduces its effect. It forces us to rethink many of the generally accepted principles of how to store and manage data. It requires the customer not only to "see the light" but also to place trust in EMC to make it all happen.

Most of us have used FedEx or UPS delivery services at one time or another. We vaguely understand that trucks pick up our packages, send them through local consolidation centers, and put them onto aircraft bound for a centralized airline hub, where they are sorted and transferred to outbound aircraft and then set for delivery by truck the next day to remote locations all over the world. We buy the service of getting our packages from here to there without understanding every step in the process or even having to think about it. VPLEX has the potential to do the same for storage, raising it from a collection of discrete pieces and processes to fully-implemented, easy-to-use Storage-as-a-Service. And just as we can verify that our packages have indeed arrived at their intended destinations, so too will VPLEX create its own proof points and credibility.

EMC has made a strategic investment in VPLEX. Not only must VPLEX deliver on all of its operational promises, it also must convince users that this is all worthy of adopting a new architectural approach. EMC has not done this without much forethought, research, and testing. EMC's announcement of VPLEX means that it believes that this product is ready for prime time use, and additional innovations and uses are to be expected. Its disclosure of the VPLEX Geo and VPLEX Global offerings means that they already know that these future products will work effectively (at least in theory). These all are strong indicators that VPLEX will be important and provide substantial value to EMC and its customers. You need to check this out!

[4] As discussed in *In Search of Worry-Free Data Coherency Across the Enterprise. See Footnote #1.*

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

***Jim Baker* is a Senior Contributing Analyst for The Clipper Group.** Mr. Baker specializes in storage hardware and software, the relationship between technology and its application in business, and emerging technological trends affecting users and vendors in the Information Technology community. He recently joined The Clipper Group after three decades in hardware and software product marketing and product management at companies as large as EMC and Honeywell Information Systems and as small as IPL Systems and Datapoint Corporation. He most recently served as Research Manager for Storage Software at IDC in Framingham, MA. Mr. Baker earned a Bachelor of Science in Marketing degree from the University of Illinois at Urbana-Champaign and an MBA, also from UIUC.

- ***Reach Jim Baker via e-mail at jim.baker@clipper.com or at 781-235-0085 Ext. 212. (Please dial "212" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group *Captain's Log***, **The Clipper Group Voyager**, ***Clipper Notes***, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.