

Product Brief

EMC VPLEX: Virtual Storage Beyond Real Walls

Date: May 2010 Author: Mark Peters, Senior Analyst

Abstract: *EMC's new VPLEX provides data accessibility across distance (some now, more to come) that enables enormous flexibility and, potentially, entirely new IT operational models. In the same way that VMware can be seen as a "software mainframe" in terms of applications and the cloud, so EMC is adding the "virtual storage mainframe" that embraces distance, rather than having to overcome it. This new approach holds promise to turn data centers into application and service hubs for flexible, geo-dispersed IT.*

What VPLEX is and Does

Let's start with what VPLEX is *not*. First, it is not simply storage virtualization in terms of abstracting physical elements into logical views to provide non-disruptive mobility and management across arrays; to such functions, it adds the dimensions of scale and distance in a federated manner. Secondly, it's not about replicating between systems. It is *one integrated system, with one namespace*. The traditional idea of a "master" and "slave" that dominates storage (where one storage entity owns a particular volume) needs to be forgotten. Instead, "cache-coherency" works with a shared cache directory that amounts to a map of where all the data is; data is not replicated, map coordinates are transferred.

This is the essence and secret sauce of VPLEX since the role of VPLEX Metro in a SAN environment is both as a *target* and an *initiator*. From the host perspective, VPLEX Metro is a target and from the back-end storage array perspective, it is also an initiator. What this boils down to is seamless and predictable local or remote data access and mobility between two VPLEX Metro clusters within synchronous distances. Data center walls are no longer the barriers they have been. This new concept is what EMC calls "AccessAnywhere" storage (and what might more prosaically be called a stretched LUN for storage-focused readers) and it allows data to be moved, accessed, and mirrored transparently between data centers, effectively allowing storage and applications to work as though those physical boundaries were not there.

Functional Capabilities: EMC refers to this as a "next generation private cloud infrastructure." Whatever label it ends up with, these new capabilities offer both valuable and intriguing potential. Here are some, in no particular order:

- **Non-disruptive data mobility** of virtualized storage arrays between EMC and non-EMC platforms, and sharing LUNs between VPLEX clusters—perhaps for migration or technology refreshes.
- **Pooling and aggregating capacity**, which can improve storage utilization and change the way "virtual-super-data-centers" are managed.
- **Distributed mirroring** across mixed platforms without requiring host resources, offering increased resiliency and availability including enabling applications to run continuously without restarts or interruptions.
- **Transparent sharing and balancing of resources within and across physical data centers**, where virtual machines can now be moved in toto, storage included, *while* applications are kept running. This can be done for workload balancing or to move work to low energy cost locations (as the geographic ability of VPLEX grows, this follow the moon approach for operations could be augmented by a follow the sun approach for knowledge workers, since physical separation of the data center and the work will no longer matter). Support for VMware to enable VMotion over distance between ESX clusters is also supported at general availability.
- **Instant, real-time data access for remote users**, where data is distributed and access is shared across sites, allowing concurrent read and write access to data by multiple hosts across two locations, providing real-time data access to remote physical data centers without local storage.

While the actual available products will take center stage, EMC's VPLEX announcement also covered its general architectural approach and future strategy. It is about bringing federation to bear in order to relieve the geographic limitations that have prevented storage from playing a full, flexible role in the growing wave of resource-pool-based

computing. As EMC itself states, “Federation enables IT to quickly and efficiently support the business through pools of resources that can be dynamically allocated. This flexibility elevates the value [that] IT offers within the business, as application and data movement is possible for better support of services. Together, cooperating pools of server applications and storage enable a new model of computing—dynamic IT as a service.”¹

The Actual VPLEX Announcement: This is not the place for an exhaustive description of VPLEX. However, a quick summary of the highlights is important to put the later analysis in context:

Architecture: VPLEX combines scale out clustering, advanced data caching, and distributed cache coherence to allow data on distributed federated storage resources to be accessed and shared between locations.


Platform: An appliance based on a field-proven, redundant-everything EMC storage engine (for the technically focused, this has Intel multi-core processors using 8-Gbps Fibre Channel ports to provide 16 host and 16 back-end connections, and 64 GB of read cache memory).

Products: The two initially released products are VPLEX Local (provides local federation across heterogeneous arrays) and VPLEX Metro (provides distributed federation where two clusters can be connected within 100km synchronous distance). Both products support a range of hosts (VMware, Vblocks, and [Cisco’s UCS](#)) and storage arrays (EMC Symmetrix /CLARiiON, plus high-end [Hitachi](#) and [IBM](#) on day one with [HP](#) and [NetApp](#) to follow). The asynchronous VPLEX Geo is slated for 2011, and VPLEX Global is also planned.

Amidst all the market chatter and debate that will be generated around VPLEX, it is also worth looking at the market that the products enter. Of course, the areas where users will be investing (shown in Figure 1) were investigated before VPLEX was commonly known; however, it is telling that the product can play a role in addressing the top planned storage investment areas cited by users from SANs and replication to virtualization and power-efficiency.

Figure 1. Top Areas for Storage Investments in 2010

With regards to specific spending plans for data storage, in which of the following areas will your organization make the most significant investments over the next 12-18 months? (Percent of respondents, N=286, five responses accepted)



Source: Enterprise Strategy Group, 2010.

¹ Source: VMotion over Distance for Microsoft, Oracle, and SAP Enabled by VCE Vblock 1, EMC Symmetrix VMAX, EMC CLARiiON, and EMC VPLEX Metro—An Architectural Overview, EMC April 2010.


The User and Market Value of VPLEX

ESG Research Shows Practical Value: Until now, data access over distance has usually required replication (regular mirroring or copies of some kind) often accompanied by pain—whether operational, financial, or both. The elegance of VPLEX in addressing this issue should not be allowed to get in the way of appreciating the pragmatic appeal of its capabilities. After all, as the data in Figure 2 reminds us, the most important considerations for justifying IT investments remain very much about *financial and business value*. In this regard, VPLEX looks well placed for success as it can address many of the top considerations affecting IT investment, including the top two considerations cited by respondents: reducing operational expenditure and business process improvement.

The potential sales news is even better for EMC when the data in Figure 2 is ranked by “spending mode”—those users that describe themselves as being in “growth mode” have slightly different considerations leading their investment choices. For these organizations in spending mode, their number one consideration is business process improvement. This is an area where VPLEX clearly delivers value, not just by doing “known tasks” better, but further by opening up the opportunity for hitherto impossible tasks. Furthermore, the same research also shows that “leading edge” users (those that stay on top of technical trends and purchase as soon as new tools are available—in other words, perfect VPLEX target purchasers) are far *less* focused on cost reduction than most (17% less than the average) and far *more* focused on improved business intelligence and delivery of real time business information (50% more than average). This is an area where VPLEX can deliver enormous value.

Figure 2. Most Important Considerations for Justifying 2010 IT Investments

Which of the following considerations do you believe will be most important in justifying IT investments to your organization’s business management team over the next 12-18 months? (Percent of respondents, N=515, three responses accepted)



Source: Enterprise Strategy Group, 2010.

Even when the research is viewed from the perspectives of broader market segments, the results augur well for EMC’s new baby; when asked about their most important IT priorities, enterprise users—obviously a stalwart target segment for VPLEX—responded highly above average in favor of certain priorities that will help drive interest in VPLEX, most notably the increased use of server virtualization (the number one priority), regulatory compliance initiatives, the implementation of IT governance frameworks, and data center consolidation.

ESG Lab Testing Shows Operational Value: ESG Lab has also tested VPLEX and found it both excellent and very easy to use. Not surprisingly the headline story is VMotion over distance, where ESG's Lab experience was very positive. As already covered, instead of failing over to another site using SRM or a stretched host cluster, users can move VMs at will between data centers (up to 100km apart with VPLEX Metro). What cannot be fully appreciated from EMC's materials is how easy and transparent this is to do. From a VMware management perspective, the simple drag and drop process used to move VMs between physical servers in a data center is exactly the same with VPLEX. It's fast (minutes) since the data is on both sides of the mirror. Another way to do this is with VMware's Storage VMotion, but it can take hours to move the data and finish the migration compared to minutes with VPLEX. During testing, ESG Lab moved a VM running a heavy Iometer write workload 100Km in just one minute and twelve seconds with no interruption at the application level.

Unexpected surprises are good if they are also pleasant and EMC has delivered a couple with VPLEX. For instance, there's what could be called the "easy test drive" ability. It's something that isn't in the EMC materials: the claim process used to virtualize an existing VM with VPLEX can be reversed. That's a good thing—users could perhaps try it with their existing applications and VMs over a weekend and go back to their traditional tested infrastructure on Monday. This was a notable and missing capability in many early storage virtualization solutions. The Lab team was also impressed by the heavy investment in exhaustive real-world solutions-based testing with Oracle, SAP, SharePoint, and SQL Server, resulting in well documented solutions-based proof-points available at G.A.

The Bigger Truth

There are two "mega trends" driving the storage market: value and simplicity. These jointly deal with the issues of continually increasing data volumes and operational complexity.

If a vendor or product does not have a useful contribution in either one of these areas—and ideally, they should contribute to both—then that vendor or product will probably die. Products that are overly specialized, are techno-geek 'crack,' or that require an army of specialists to tweak all the knobs will be increasingly irrelevant as storage continues to become more of a flexible and automated servant of IT and less of a rigid and frustrating shackle restraining it. We must remember that storage is no more and no less than a tool—no one makes money by having storage in grand isolation. It is a means to an end, and not an end in itself.²

The bottom line for VPLEX, and hence the good news for EMC, is that measured against such criteria, it appears to have market applicability and relevance in two main ways:

1. **The practical and immediate business impacts** of delivering financial and operation value, which are a prerequisite for any market success. This includes management ease; better resource utilization; and power, cooling, and space advantages.
2. **A more conceptual and longer term "play"** (strongly tying back to value and simplicity) is that VPLEX (with its flexibility and heterogeneity) puts EMC squarely in the current and growing rush to collapsed-stacks and resource-pools. Ignoring this and continuing to simply produce standalone systems would soon have placed EMC in the role of King Canute. Instead, EMC is now part of the advance.

Like many good things, VPLEX isn't conceptually complex. Indeed, its underlying sophistication and engineering wizardry belies its straightforward value proposition. Removing distance as a storage barrier brings us one step nearer to fully flexible IT, which is more about applications and delivering business value than it is about implementing and running a bunch of technologies. Server virtualization was one piece of the solution, VPLEX represents another. Yes, it still has to be field-proven, and yes, there are bound to be a few minor initial implementation limitations, and certainly the impact of truly massive distances on performance is some time from even being tested, but the supposed rule of distance as a virtual brick wall is broken as VPLEX Metro breaks through the real walls of data centers. Of course, other vendors are likely to cry humbug about such "unnecessary" functions until they, too, offer similar "vital" functions! After all, who needs radios in cars or computers in homes!?! Since VPLEX is not a full storage virtualization solution (such as Hitachi's USP-V, IBM's SVC, NetApp's V-Series, or even EMC's own Invista), EMC would do well to work to make it at least aware of traditional virtualization tools (such as thin provisioning for instance); for now, many users will choose to run VPLEX in

² Source: ESG Brief, [What's on the Block for Block Storage in 2010?](#), January 2010.

association with their existing traditional storage virtualization and we would not be surprised to see EMC integrate these functions over time. VPLEX federation is something new and different. Moreover, we absolutely know the business need exists because of the limited and specialist capabilities that have sprung up in this space before—but never with the scale, distance, ease, and affordability of VPLEX. Detractors might also say that the world already has remote mirroring using tools such as EMC’s own SRDF, but that is to overlook the massive rise of server virtualization and the new use case of VMotion over distance. The new capability will also open up new operational opportunities—and other ideas that have yet to be imagined.

Indeed, whatever else happens, we will all have to start thinking differently. After all, according to conventional wisdom in 1491, the world was flat. In 1492, it was not. In 2005, Pluto was a planet. In 2006, it was not. Similarly, in 2009, data access and management across distance was a major anchor. In 2010, that rule changed. Perhaps we need a new name for data centers—or at least a new conceptual understanding of what constitutes a wall?