

Commentary

May 10, 2010

EMC VPLEX™ Launches the Virtual Storage Era

Magicians make objects appear, disappear, or change appearance. EMC is doing the same thing with information and long-held data center beliefs by introducing its new VPLEX product family. The first two products, now available, are VPLEX Local™ and VPLEX Metro™, which deliver new, key Virtual Storage capabilities. VPLEX Local makes information in heterogeneous storage systems appear similar to access, and VPLEX Metro makes metropolitan-scale distance disappear. Information can now change its appearance seamlessly, from one location to another. Storage provisioning and data mobility have just gotten a lot easier.

VPLEX Starts the Virtual Storage Era That Leads to the Private Cloud

EMC's introduction of VPLEX officially heralds the arrival of the age of Virtual Storage. Virtual storage is the logical separation of information from physical storage media *and* location. This is at the core of what makes virtual storage unique — the added ingredient of location independence. Local and distributed federation, EMC's names for the ability of resource pools to work cooperatively at a single site or over distance, deals with the location issue and is also mandatory for virtual storage.

Virtual storage along with server virtualization creates the agile infrastructure to enable the journey to the private cloud.

IT organizations cannot wait for the future to happen; they have to *make* it happen. And with EMC Virtual Storage and VPLEX, the future starts *now*.

The VPLEX Platform Is Built on Storage Federation Engines

VPLEX physically manifests itself as an appliance. An appliance is dedicated to a single special-purpose function, and for VPLEX that function is storage federation.

A VPLEX engine takes a holistic view of all the storage resources that it federates. That holistic view sees the data as independent of the physical data storage devices themselves and their geographic location.

Physically, an individual engine resides in an approximately 3U high rack-mountable chassis. The VPLEX engine leverages other EMC storage engine technologies. That means that there is no new product hardware risk, as the existing technology is already well-proven and reliable.

VPLEX engines logically sit in front of storage systems, controlling the I/O flow that gives them the location and

media information that is necessary for them to deliver federation.

VPLEX software enables an I/O entering the appliance from anywhere to be serviced by any engine in the cluster. With the ability to have one, two or four engines in a single cluster, customers can start small and grow larger with predictable service levels. Leveraging distributed cache coherence, VPLEX provides for automatic sharing, balancing and failover of I/O across the storage infrastructure.

VPLEX Current Releases

The first two products in the VPLEX family are VPLEX Local, which operates within a data center, and VPLEX Metro, which works within distances that have been traditionally acceptable for synchronous remote mirroring -- roughly 100 kilometers (60 miles). EMC has stated that VPLEX Geo™, which will operate at asynchronous distances (i.e. thousands of miles), will be available in 2011. And later, VPLEX Global will enable data mobility and access anywhere on the planet, including multiple locations and service providers. But that is then; this is now.

VPLEX Local

A VPLEX Local cluster can virtualize up to 8,000 logical unit numbers (LUNs). That LUN count maximum holds true whether a single VPLEX engine is used or the maximum of four engines. Adding additional engines on top of an initial engine to create a cluster increases performance and resiliency. Scaling is on a true linear basis, with N+1 performance scaling and N-1 resiliency (where N is the number of additional engines).

VPLEX, as an enabler of virtual storage, must be vendor agnostic, because an IT infrastructure functions as a composite of heterogeneous products. That is why VPLEX can support a range of host environments, including VMware, Vblocks, and Cisco's UCS, as well as a number of storage arrays, including Symmetrix and CLARiiON (naturally!) and some arrays from Hitachi and IBM (HP and NetApp are also planned).

VPLEX Metro

VPLEX Metro delivers distributed federation over distance. A total of 2 VPLEX Metro clusters can be connected together, providing up to 8 engines and a maximum of 16,000 LUNs. Each cluster is physically separate and distinct (for example, 8,000 LUNs is the maximum for any one cluster). However, both clusters may be in the same data center (and even the same rack), or they may be separated by 100 km (the maximum synchronous distance).

One use of shared LUNs is to be able to transparently move or relocate applications between host clusters or locations within synchronous distances. This is done in conjunction with host-based clustering technologies. You should check to see what technologies are supported -- for example, VPLEX Metro offers support for VMware, Microsoft MSCS and CSV; and support for a number of clustered file systems. In a VMware VMotion move, VPLEX can make the data available to the new location in seconds.

Putting VPLEX Local to Work

VPLEX Local enables local federation of heterogeneous storage arrays. This

capability alone is not unique to EMC, but what is unique is VPLEX's implementation and architecture. To understand further differences, it is important to know not only what VPLEX does, but also where it is going and how it will evolve within the data center.

To fully comprehend this evolution, you should note that the choice of a storage federation platform is a key decision for any IT organization. Selection implies that the vendor will also supply the storage management software (which discovers, monitors, and controls physical storage assets) and data management software (such as controlling migration and replication software) as well. Thus, implementation of VPLEX Local means that the physical arrays will remain heterogeneous, but everything else will be homogeneous.

This is not a bad thing; in fact, it's quite good. It is a necessary foundation for the next generation of the data center, the private cloud, or whatever your future IT infrastructure will be called. The principle might be called *fungibility*, which means the ability to be exchanged in whole or part for another of like nature or kind. Note, for example, that even though two heterogeneous arrays use the very same brand of hard disk drive, they are really not interchangeable. That rigidity is costly in terms of people, time to accomplish tasks, and money, and so will be unacceptable in the future. Isn't fungibility, then, one of the key benefits of federation? The answer is yes, and VPLEX enables fungibility.

Standardization

VPLEX Local simplifies multi-array management via centralizing and standardizing storage operations, from the standardization of LUN presentation to having one way to provision storage across mixed generations and mixed platforms. Turnaround of requests to provision storage to applications thus becomes much, much faster.

Data mobility

Data mobility gets into a second basic principle of virtual storage, *fluidity*. Fluidity is about the ability of a substance (in this case, data) to flow from one storage unit container to another. Data mobility is about moving data to the right place at the right time, according to real business needs, such as when VMotion moves a VM, and rather than having to move all of the data to the new location before using it, VPLEX enables the new location to access the data within seconds and to move only as much data as necessary for this situation. VPLEX virtualizes the storage, freeing the data from the physical constraints of the system.

Resiliency

VPLEX enables mirroring across mixed platforms without requiring host resources to improve resiliency. This improves the availability of critical applications while at the same time making better use of storage resources.

VPLEX Metro Does Its Thing

VPLEX Metro introduces what EMC calls distributed federation. EMC defines distributed federation as "cooperation between storage elements to enable data to be shared, accessed, and relocated transparently across

resources over distance." Let's see how VPLEX Metro instantiates distributed federation.

Multi-Site Application and Data Mobility

VPLEX Metro enables the transparent sharing and balancing of resources between data centers. For example, VPLEX Metro works in combination with VMware to enable VMotion over distance, transparently moving and relocating virtual machines, along with their corresponding applications and data access, over metropolitan distances. The ability to relocate, share, and balance infrastructure resources between data centers, without a lot of hard work, time, and cost, should be very welcome.

Multi-Site Application and Data Mobility

With VPLEX Metro, multiple users across two locations can access information from a single copy of data in real time, without the operational overhead and time that would otherwise have been required to copy and distribute data across the locations. This ability to support concurrent read and write access to data by multiple hosts across two locations is the reason that EMC devoted considerable resources solve the distributed cache coherence problem. Distributed cache coherence enables the automatic sharing, balancing, and failover of storage domains within and across VPLEX engines and clusters.

Improved Multi-Site Resiliency

VPLEX Metro has its own synchronous mirroring capability that enables mirroring across legacy multi-vendor arrays without host impact. This is true active-active mirroring: no restart of the application is required. Overall, this increases availability for critical

applications and eases the planning burden, as there is no requirement that the storage arrays in the two sites have to be identical.

So Why Does VPLEX (and Virtual Storage) Matter?

VPLEX matters because it enables much greater *flexibility* in an IT infrastructure. The existing IT infrastructure is the antithesis of flexible. Rigid is a more apt term. Major change requires a contortionist act.

That cannot continue. Not only is rigidity too expensive, but the time-bounded constraints on changing quickly enough may leave an enterprise in a competitive straightjacket.

In the case of the data center, VPLEX's increased flexibility translates into increased:

- Ease/speed of *adaptation*,
- *Fungibility* (using resources interchangeably), and
- *Fluidity* (data mobility).

Federation Leads to Greater Flexibility

EMC likes to point out that virtual servers and virtual storage are taking similar paths of evolution. The first phase of virtualization, via virtual machines and storage virtualization, has been a consolidation of resources to enable efficiencies that were not possible before. But traditional storage virtualization technologies do not go far enough; that is to say they do not allow for data movement over distance. With VPLEX, a new stage is underway with federation — call it the “pools of cooperation” phase — which enables a new usage model that was not possible up to now, a model in which heterogeneous resources can be provisioned

Commentary

instantly, dynamically, and on-demand. Those cooperating pools of resources lead to greater flexibility in the IT infrastructure, such as dynamic application and data movement, independent of geographical constraints.

Business As Usual Is Not Enough

Each organization has its own unique mix of applications and data, but one thing seems certain: at any organization, the new solutions will not be kept within the confines of one physical data center. It is how the organization manages the blend between local and long-distance that will be unique to each organization.

Logically, a data center may appear to be one place; but physically, the data center will seem to be everywhere. Even though applications and data have to reside physically on a piece of media at a physical geographical location at any moment in time, those applications and data have to be able to be moved non-disruptively and dynamically to where they can best serve the needs of the business. For the full realization of these capabilities, VPLEX Geo and VPLEX Global will help lead the way.

VPLEX Geo

VPLEX Geo will enable the balancing, relocating, and failover of applications and data among data centers over extended distances. VPLEX Geo will also enable instant access to data without the latency and data consistency issues that have been typically associated with distance (thanks to VPLEX's unique implementation of distributed cache coherence). One example of why this is important is software development. The software development process can be simplified

and speeded up, because developers in different geographies can have simultaneous access to development code and data.

VPLEX Global

VPLEX Global will remove two big constraints on data center design, (1) that no more than two data centers can be included and, (2) that external service provider sites cannot be part of the mix. Expansion to three or more sites and support for external providers enable a new operational business model for IT, where the resources of multiple cross-organizational sites can be aggregated into a single "virtual" data center. Think of the efficiencies, and the flexibility, that can ensue.

VPLEX Global also removes any remaining barriers on size or distance, taking the final step to fulfill EMC's vision of the private cloud. Virtual storage with no barriers over long distances will enable a number of things, including aggregation of big data centers from separate ones and 24 x forever-running applications without required restarts. This will yield even greater flexibility, leading to greater efficiency, reduced costs, improved productivity, and outstanding responsiveness to changing business requirements.

Conclusion

IT has always been a service in the sense that it is a service *function*; IT now has to become a service in its true sense, that of being a service *provider*. Above all, flexibility is needed for this transformation to occur, and that is what VPLEX promises to deliver.

Organizations will move inevitably toward the private cloud and, in doing

so, they'll likely find VPLEX a powerful enabler.

With this journey – as with many others, key steps should not be missed. Similarly, opportunities are missed by waiting for the “perfect solution”. Organizations must start at the beginning and complete one level before moving on to the next. All that waiting accomplishes is deferring the time to achieve both the incremental and eventually full value of the private cloud.

So, for the savvy IT organization, VPLEX and the virtual storage future should start today, with numerous benefits and competitive advantages to be reaped along the way.

David Hill

Analyst Name: David Hill
Topic Area: Virtual Storage

Mesabi Group LLC
26 Country Lane
Westwood, MA 02090
www.mesabigroup.com

Mesabi Group LLC is an affiliate of Valley View Ventures that aims to provide thought leadership and sound advice to both vendors and users of information technology. .

Phone: (781) 326-0038
email the author: davidhill@mesabigroup.com

The information contained in this publication has been obtained from sources Mesabi Group LLC believes to be reliable, but is not warranted by Mesabi Group LLC. Commentary opinions reflect the analyst's judgment at the time and are subject to change without notice. Unless otherwise noted, the entire contents of this publication are copyrighted by Mesabi Group LLC, and may not be reproduced, stored in a retrieval system, or transmitted in any form or by any means without prior written consent by Mesabi Group LLC