Ptak / Noel

EMC VPLEX and the Road to Global Data Access

Executive summary:

Using virtual machines is a proven way for enterprise data centers to reduce IT expense and increase application availability. However, keeping applications online as they move around the infrastructure means that the link between application and data remains unbroken each time the VM moves. This document looks at virtualized data centers and at their need for agile storage. It then describes the EMC VPLEX, clustered storage that provides necessary data mobility across a single data center or between two synchronously connected sites. It presents a number of likely use cases for both the near term and long term, and suggests the changed computing models that will result from local, regional, and worldwide data federation.

Contents

Introduction

IT procedures continue to change as data centers adapt to meet new sets of business needs. During the past several years, the use of virtual machines has removed the need for applications to be physically associated with specific servers. The most widely used example of this is VMware's VMotion, which allows the transparent migration of virtual machines between hardware hosts that share the same storage. Because applications are associated with virtualized servers, rather than physical servers, they can be moved across the data center with relative ease whenever a need for application mobility arises. This technology has proven to be so effective that it is now nearly ubiquitous. It is hard to find any enterprise data center that does not use it extensively.

Virtual Storage

An evolution from storage virtualization technologies is now coming online that makes data as accessible and available as applications. Just as virtual machines remove the link between applications and a particular server, storage virtualization today inserts a virtual layer between hosts and storage, freeing data from the need to be associated with a specific physical storage array. Once the data is associated with a virtual, rather than a physical device, it may be easily moved across the infrastructure as user needs or resource availabilities change. Virtualization provides transparent mobility and has become the preferred technology for data migration between storage arrays.

Because the process is automated, always on, and secure, virtualizing storage simplifies the management of multi-server environments. Properly managed, it maximizes utilization rates, performance, and reliability, and in some cases even extends the useful life of older arrays. Data becomes available instantly, dynamically, and on-demand.

Still, to date it has not been possible to include all corporate data in this storage pool. Data distributed over distant locations or locally held data on arrays that cannot be managed with the rest of the storage, must be dealt with separately.

If it were possible to include all corporate data within a virtualized storage pool, irrespective of where the data physically resides, the result would be to improve storage performance, agility, and mobility to the point where virtualized storage would be on a par with virtualized servers and applications. Every server would be able to point to every data set in the storage pool, greatly enhancing the flexibility and resilience of those applications that have already been virtualized and opening up a new set of possibilities for new computing models. Those IT departments wishing to take advantage of storage as a service (sometimes referred to as "utility" storage) would find that the opportunity to do this has become immediately available. By dissolving the tie between data and physical storage, data centers would potentially have the best of both worlds: they can provide all data to all applications, and they could do so economically. Automated, dynamically allocated data could then be provisioned when applications need it, and de-provisioned (and thus freeing up resources) when it is no longer required by the application. Virtual storage's automation and storage federation capabilities will move today’s consolidated storage infrastructures to a dynamic set of resources that work in concert with virtual servers.

Storage in the Virtualized Data Center

Data center managers appreciate the flexibility that virtualization provides, but while enterprise data centers now count their VMs in the thousands, storage technology remains a constraining factor. Data mobility beyond a single array is typically labor-intensive and often ponderous, limiting the ability to move data in a way that ensures virtualized applications can easily access the information they need. The challenge is to give the data the same mobility as the applications and to make sure that managing that mobility is an easy – and ideally, automated – task.

Two technologies combine to deliver this now: federated storage management and private cloud computing.

Federated storage management refers to a system of autonomous, often heterogeneous, centrally managed storage devices that can share data because of their common management structure. Federated management extends the limits of a virtualized storage pool. Systems are often physically distributed across the data center, the entire system will likely include arrays from multiple vendors and the centralized management of all storage devices reduces the complexities associated with storage provisioning, which because of automation can be done dynamically and on-demand.

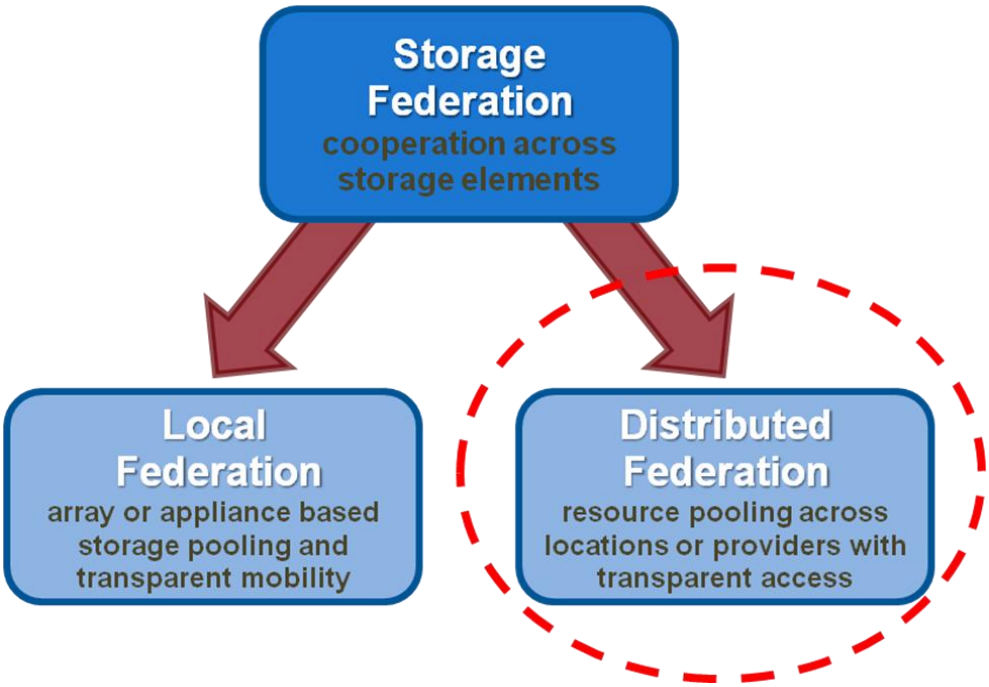

Figure 1 Federated storage

Distributed federation is now available for the first time, providing cooperation between storage elements to enable data to be shared, accessed, and relocated transparently across resources over distance. This extends the virtualized storage pool between distant data centers, allowing them to share, access, and move data between geographically distant resources. Data available at each site would be accessible to the other site, making the entire set of federated storage available to the entire set of virtualized servers.

The private cloud is a controlled environment in which a company exercises its own control over security, availability, performance, IT issues that concern corporate governance, and so forth. Logically behind a corporate firewall, a private cloud is a mix of corporate data centers that provide on-demand workload relocation and other services.

EMC VPLEX

VPLEX is a storage federation platform that removes the necessity of assigning data to physical storage, making data as agile as the virtualized applications that use it.

What It Does

EMC's VPLEX delivers federated storage management to data centers with virtualized server environments so they can have a single, coherently managed IT system. The architecture moves data transparently, both within a single data center (local federation) and between two separate data centers (distributed federation)[1]. It pools storage capacity within the federated environment, providing multi-array and multi-site information access and delivers simple, efficient scale out across one or both data centers. Data movement is automated, requires no staff intervention, and is transparent to users.

VPLEX specifications

16 8 GB/s front-end FibreChannel ports
16 8 GB/s back-end FibreChannel ports
FibreChannel interconnect between directors
Intel quad-core CPUs
64 GB read cache

2 high availability directors per engine
Up to 4 engines per cluster
Redundant power supplies
Integrated battery backup
"Call home" and remote support

8,000 virtualized LUNs per system
N+1 performance scaling
N-1 fault tolerance
4KB data blocks

Platform support: ESX, UCS, Windows, Solaris, HP-UX, Linux, AIX
Fabric support: Brocade, Cisco, McData

Note: VPLEX Metro requires 1 rack per site

What It Is

VPLEX is a federation platform with advanced data caching techniques that deliver improved I/O performance and reduce contention for physical storage resources. Because cache coherency is maintained across all caches within the distributed VPLEX environment, VPLEX can provide automatic data sharing, load balancing, and failover of all storage domains under management. VPLEX systems may be installed individually or in clusters of up to four engines, allowing sites to start small, but still scale out when necessary with predictable performance and service levels. VPLEX Local™ provides a single cluster at one site and VPLEX Metro™ supports up to two clusters that can be connected to one another within synchronous distances.

How It Works

Because it provides concurrent synchronous read/write access for multiple hosts that may be as much as 100KM apart, VPLEX must maintain cache coherency across widely distributed caches and must keep latencies to a minimum. EMC does this with a

[1]VPLEX is a family of products. In its first release, VPLEX manages synchronously, both across a single data center (VPLEX Local) and between two data centers separated by a Metropolitan Area Network of up to 100KM (VPLEX Metro). Later releases will provide continent-wide asynchronous communications between two sites (VPLEX Geo™, scheduled for release in 2011) and, later, multi-site global connectivity (VPLEX Global™).

metadata directory. The directory identifies which engine within the cluster owns each block of data, shares this information across all engines within the cluster, and updates the information for each I/O. Because only a small amount of metadata is actually passed between the engines within a cluster (much less than the 4K data blocks that are being updated), network traffic is minimized, enabling simple, synchronous data communications. For example, on a read request, VPLEX automatically identifies the engine that "owns" the data block and sends the read request to that engine, which accesses the actual data. When a write operation modifies the table, the directory rechecks the table on subsequent read requests coming from another engine. If the newly written data is still in cache, that cache can be read directly rather than going to disk. Because VPLEX distributes this directory between clusters (and therefore, between sites), VPLEX clusters can extend across metro distances.

Conclusion: Ptak/Noel on Using VPLEX

Storage virtualization has been around for a while now and has proven its value. What EMC brings to the table with VPLEX is an ability to make storage federation as agile as the virtualized applications that the data serves. How useful will VPLEX be in adding functionality to the data center? Ptak/Noel has looked at likely use cases for both the present VPLEX release and for the future releases that are now in-plan. Our conclusions follow.

Present Use Cases

At its first release, VPLEX synchronously manages a heterogeneous mix of federated storage systems across both a single data center (VPLEX Local) and between two data centers connected by a MAN within synchronous distance (VPLEX Metro). Ptak/Noel expects managers at enterprise data centers to see value in several areas:

- Because management extends over both EMC storage (Symmetrix and CLARiiON) and high-end systems from HDS and IBM, likely use scenarios include both providing relief for ongoing data migrations between unlike systems (HDS to EMC, for example) and data movement that occurs in support of technology refreshes to systems throughout the infrastructure. This will be even more useful once VPLEX can manage devices from additional vendors[2].
- VPLEX's ability to mirror storage across mixed platforms (while using no host resources) should extend the life of existing arrays. Rather than decommissioning older hardware, older arrays can be leveraged for reuse as mirror targets to provide greater resiliency to business-critical data.
- Using VPLEX Metro to link distributed data centers synchronously over a MAN will surely be advantageous for companies that maintain more than one data center within a geographical region. Because a VPLEX Metro storage pool extends to MAN distances, users at each location get real-time access to identical data sets. It is easy to imagine such sites operating as peers, simultaneously accessing active data from both sides of the extended cluster, while at the same time maintaining online archives across metro distances. While it would not make good sense to use VPLEX solely as a migration tool, it is useful to keep in mind that data migration has historically been both labor-intensive and time-consuming. Thus, VPLEX's ability to do rapid and transparent migrations in addition to its other functions should be seen as another value-add.
- Data centers that create a well-managed data pool aggregating data across multiple arrays (and in

[2] An obvious next step here would be to add top-line devices from HP and Sun, as both are built on a platform supplied by HDS.

case of Metro, across two sites) will see noticeable improvements in storage utilization rates.

- Perhaps most significantly, data centers that presently use VMotion and DRS to pool server resources locally will now be able to aggregate and move data over distance in a like fashion. By relocating data when they move VMs and applications across the infrastructure and between sites, IT managers can keep applications active throughout the process.

Future Use Cases

Future releases will deliver VPLEX functionality asynchronously over much greater distances. VPLEX Geo will enable asynchronous cross-continent data communication between two sites in 2011. Sometime after that, EMC will ship VPLEX Global, which will offer worldwide many-to-many data center connectivity.

Both offerings will provide widely separated data centers with the capability for long distance data mobility and data relocation. Because this will be low latency data distribution and data sharing between globally separate access points and because it will be accomplished with none of the data consistency issues that have always been associated with such long-haul traffic, Ptak/Noel sees data centers in the future taking advantage of this asynchronous access to:

- Aggregate all data in their worldwide data centers within a single, centrally managed storage pool.
- Provide users worldwide with access to multiple hosts over unlimited distances.
- Do global load balancing, moving processing and/or data around the world to take advantage of available unused capacity.
- Enable efficient power management by constantly moving batch processing jobs to locations where electrical power is less expensive.

Futurethink: New Computing Models

VPLEX frees data from its association with physical storage. Ptak/Noel is of the opinion that when VPLEX Geo and VPLEX Global are introduced, they will provide the groundwork for new ways to operate data centers. Once physical location becomes irrelevant, moving or accessing data across the continent or across the globe will be no more of a challenge than using data that sits within the walls of today's data center. Resources -- eventually federated across the globe – will be consumed on demand, with applications and data automatically moving to wherever they will be most efficiently consumed. Using resources when they are needed (and then releasing them back to the general pool when they are no longer in use) is often referred to as a *utility* model. With such a model, few need to own specific storage resources; rather, they need only make sure they have access to storage services. VPLEX Geo and VPLEX Global will enable accessing those services when they are outside the data center walls, moving utility services into the cloud.

In the future, it is easy to envision data centers that get many – and in some cases, perhaps all – services from the cloud. Those data centers will likely operate quite differently than do those of today. Increasingly, automation will let them build cooperating pools of resources that will be consumed on demand. They will then be able to:

- Move thousands of VMs over thousands of miles;

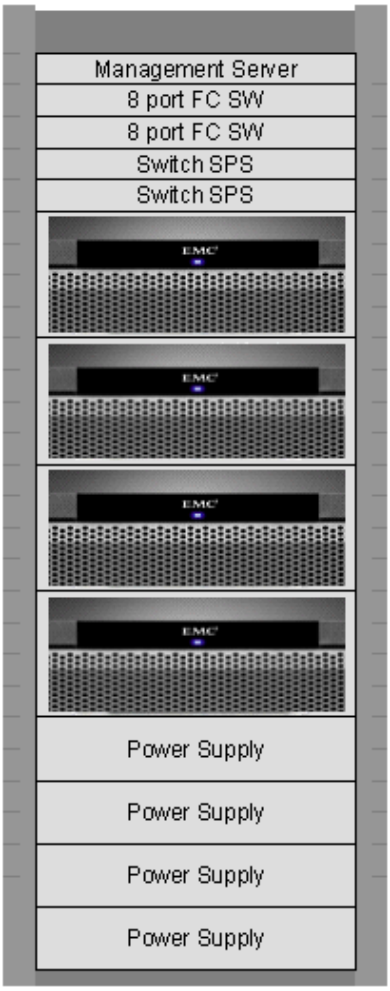

Figure 2 EMC VPLEX

- Dynamically move applications and data across geographies, and even across service providers;
- Move workloads to batch process in locations where energy is cheap;
- Balance workloads by deploying processing and data to wherever resources are available;
- Aggregate a global virtual data center from globally distributed smaller ones, creating a centrally managed global data pool that delivers much more efficient use than was ever possible before;
- Eliminate downtime because they now can run or recover applications without interruption or restart.

At that point, the term *virtual storage* will take on a larger meaning than *storage virtualization* has today. Virtual and federated storage from the cloud, completely abstracted from physical hardware and geographic location, will deliver services in a way that is transparent to the user. If performance, agility, and mobility can be delivered anywhere, securely, transparently, and on-demand, physical location will become irrelevant. As far as processes and users are concerned, data will reside in the cloud, and no one will care where the hardware is.

This white paper was sponsored by EMC Corporation.

Publication Date: April 30, 2010

About Ptak, Noel & Associates LLC

We help IT organizations become "solution initiators" in using IT management technology to business problems. We do that by translating vendor strategy & deliverables into a business context that is communicable and actionable by the IT manager, and by helping our clients understand how other IT organizations are effectively implementing solutions with their business counterparts. Our customers recognize the meaningful breadth and objectively of our research in IT management technology and process.
www.ptaknoel.com

About the Author

Mike Karp (mkarp@ptaknoel.com) is Vice President and Principal Analyst at Ptak Noel Associates. Much of his interest, writing and research focuses on driving efficiency into data center operations, on digital archiving, and on reducing the agonies associated with managing complex systems. Mike also has been known to look at the impact of technology on both technologists and on society at large, and often considers technology's relationship with the social infrastructure. His writing on storage and how storage delivers value to the business process is widely available on the web. Prior to joining Ptak | Noel, Mike led the storage practice at Enterprise Management Associates and was Director of Storage Services for the Hurwitz Group. He holds degrees from Franklin & Marshall College and Brandeis University.