

Lab Validation Report

EMC VPLEX Metro and VMware ESX

Enabling 100 km VMotion with New Distributed Storage Federation

By Brian Garrett

May 2010

Contents

| | |
|-------------------------------------|----|
| Introduction | 3 |
| Background..... | 3 |
| Introducing EMC VPLEX | 4 |
| How VPLEX is Used | 5 |
| ESG Lab Validation | 6 |
| The ESG Lab Test Bed | 6 |
| Getting Started | 8 |
| <i>ESG Lab Testing</i> | 8 |
| VMotion over Distance..... | 10 |
| <i>ESG Lab Tested</i> | 10 |
| Performance Analysis | 12 |
| <i>ESG Lab Testing</i> | 12 |
| <i>What the Numbers Mean</i> | 14 |
| <i>What the Numbers Mean</i> | 15 |
| ESG Lab Validation Highlights | 16 |
| Issues to Consider | 16 |
| The Bigger Truth | 17 |
| Appendix | 18 |

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This report was made possible through cooperative testing by Enterprise Strategy Group, EMC. This ESG Lab report was sponsored by EMC.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. Copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482.0188.

Introduction

This ESG Lab report explores how [EMC VPLEX](#) can be used to create a heterogeneous pool of storage resources that can be accessed and shared over distance. VPLEX-enabled VMotion over distance and the performance of SAP, [Oracle](#), and [Microsoft](#) SharePoint and SQL Server applications are examined as a mix of disk arrays spread over two data centers 100 km apart is turned into a centrally managed, distributed storage federation.

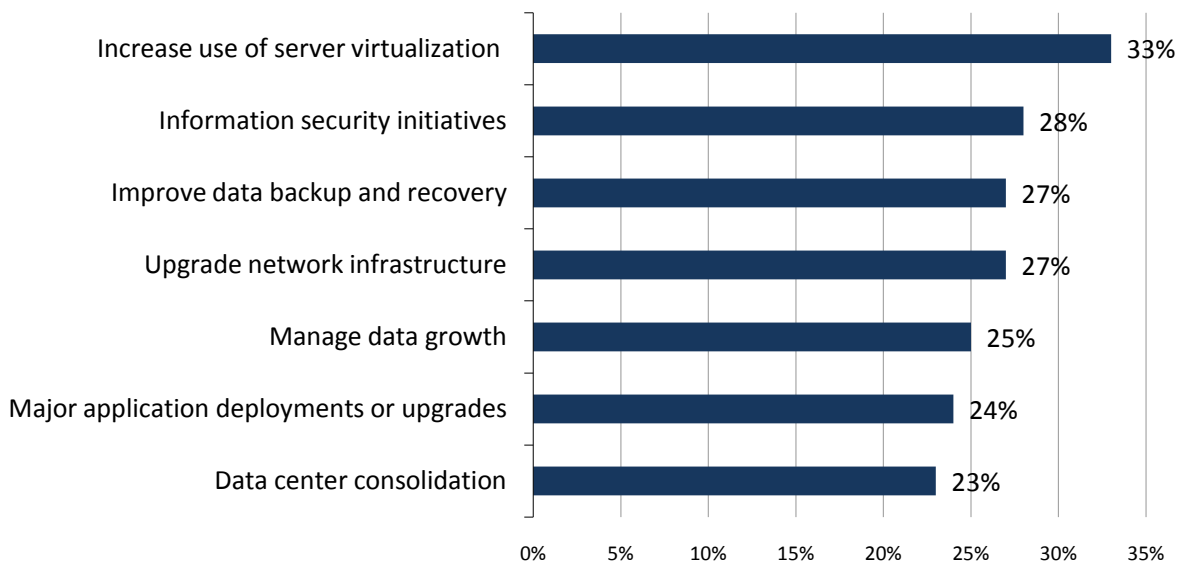
Background

ESG research indicates that increasing the use of server virtualization will top the list of IT priorities over the next 12 to 18 months. As shown in Figure 1, securing, protecting, and managing ever-growing stores of information assets are top priorities as well.¹

The research also indicates that the vast majority of organizations that have deployed server virtualization have done so in conjunction with networked storage (87%);² compared to islands of direct attached hard drives, utilization is greatly increased when applications share a pool of networked storage. Applications deployed on virtual machines sharing a networked pool of storage are more mobile and available than those deployed on direct attached hard drives. Put it all together and it's clear that forward-thinking IT managers are looking to combine the benefits of virtual server and virtual storage to increase the efficiency, flexibility, and security of their IT infrastructure.

Figure 1. Server Virtualization is a Top IT Priority

**What are your organization's most important IT priorities over the next 12-18 months?
(Percent of respondents, N=515, ten responses accepted)**



Source: Enterprise Strategy Group, 2010.

IT managers are looking to extend the benefits of a virtual infrastructure beyond the physical walls of the data center. The lofty strategic goal is to create a centrally managed pool of virtual resources spanning all data centers in an organization. A practical first step is achieving the ability to move a running application to another data center without having to declare a disaster.

¹ Source: ESG Research Report, [2010 IT Spending Intentions Survey](#), January 2010.

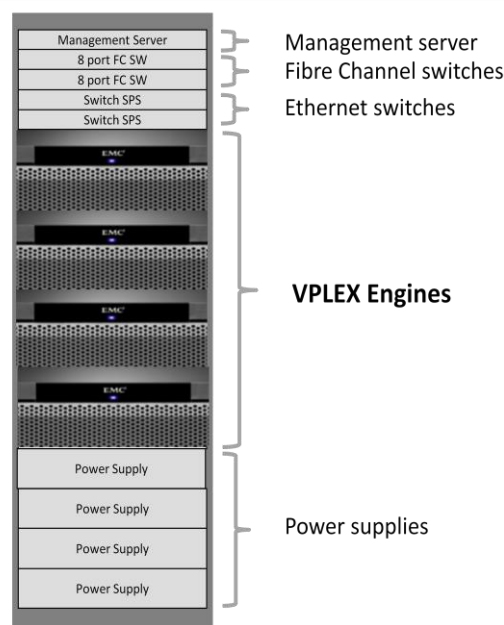
² Source: ESG Research Report, [The Impact of Server Virtualization on Storage](#), December 2007.

Introducing EMC VPLEX

VPLEX is a combination of hardware and software deployed between servers and Fibre Channel-attached disk arrays. It turns a heterogeneous mix of disk arrays into a distributed, federated pool of virtual storage that can be spread over multiple data centers. One of the key goals of the VPLEX architecture is to overcome the scalability and distance challenges that until recently have made deploying a large pool of centrally managed virtual resources that spans data centers prohibitively complicated and expensive.

VPLEX hardware is deployed as a rack-mounted solution (see Figure 2) with VPLEX engines at the core. The field-tested server hardware within enterprise-class EMC disk arrays is used for VPLEX engines. Each engine is comprised of a pair of highly available directors, each with a pair of quad core Intel CPUs. Directors have 32 FC ports, each with 8 Gbps of bandwidth, for connectivity to hosts and disk arrays. The solution is equipped with a management server, Fibre Channel switches, Ethernet switches, and power supplies.

Figure 2. EMC VPLEX Hardware



A cache coherent, active-active cluster spread over one to four VPLEX engines is supported in the first release. In other words, up to four VPLEX engines can be used to create single view of storage. Two VPLEX clusters can span two data centers at synchronous distances. All of the engines can be actively used and if any one of them fails, resources in other engines are used to ride through the failure without interruption. The first release of VPLEX has been qualified to work with a heterogeneous mix of disk arrays and IT infrastructure components. This list will grow as EMC continues to invest in VPLEX interoperability testing.

Table 1. VPLEX Interoperability

| | |
|-------------------|---|
| Disk Arrays | EMC Symmetrix, VMAX, DMX, CX, CX4; HDS 9970, 9980, VM, USP-V, BM DS8000 |
| Host Platforms | ESX, UCS, Windows, Solaris, AIX, HP-UX, Linux |
| Multipath Drivers | Power Path, VMware, NMP |
| Local Cluster | VCS, MSCS, Oracle RAC |
| Volume Managers | VxVM, AIX LVM, HPQ, HP LVM |
| SAN Fabrics | Brocade, McData, and Cisco |

VPLEX has been qualified to work with a wide variety of physical and virtual servers, operating systems, and applications. In the family of VPLEX products, VPLEX Local is intended for use within the confines of the data center while VPLEX Metro can be deployed within, across, and between data centers up to 100 kilometers apart. VPLEX Geo, which is planned for a future release, is designed to stretch the benefits of VPLEX between data centers separated by more than 100 km.

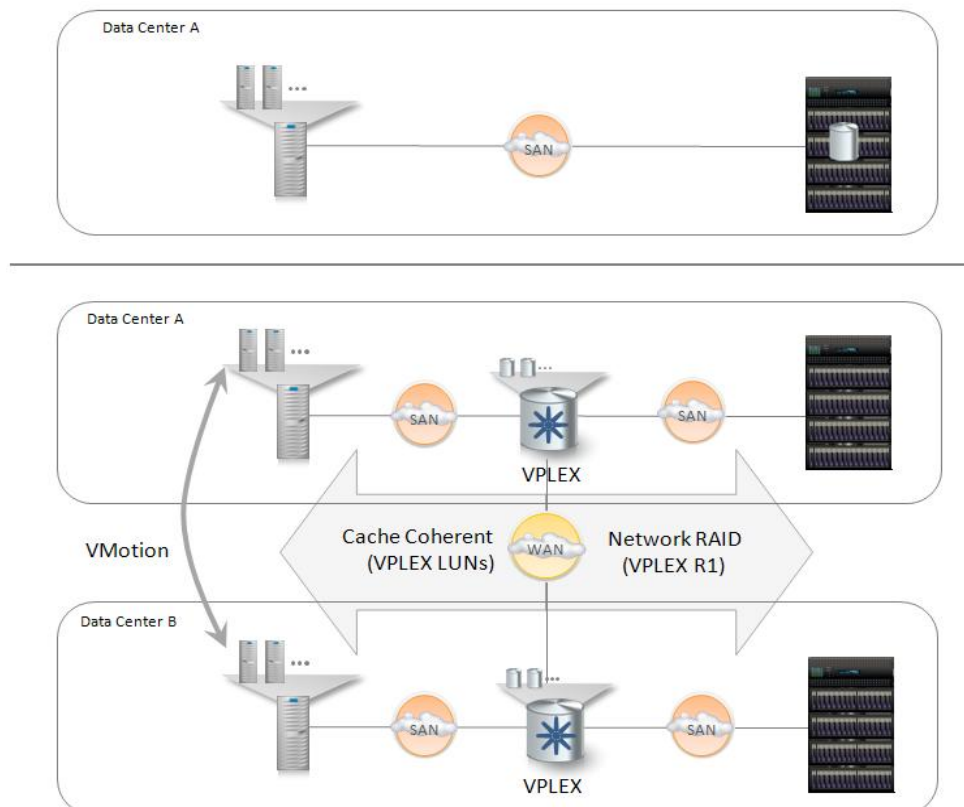
How VPLEX is Used

Some of the more powerful ways that the first release of VPLEX can be used include:

- Online mobility of applications between servers, storage arrays, and data centers.
- Online movement of [VMware](#) virtual machines between data centers up to 100 km apart.
- Online data mobility between EMC and non-EMC disk arrays located within and between data centers.
- Centralized online mobility of storage capacity between tiers of storage.
- Standardized processes for the configuration and deployment of storage capacity.

Now that we've seen what VPLEX is and how it can be used, let's take a look at how it fits into an existing virtual server infrastructure. As shown in Figure 3, VPLEX is configured between SAN attached hosts and storage. In the traditional configuration shown on top in Figure 3, a virtual server is accessing storage through a SAN using a LUN presented by a disk array. As shown towards the bottom, VPLEX engines have been inserted in the SAN as a cluster, between the servers and the disk arrays. In this example, two clusters have been deployed: one in data center A and another in data center B. The WAN-connected (up to 100 km) VPLEX clusters work together to present a cache coherent view of application data (VPLEX LUNs). Behind the scenes, network RAID is used to mirror LUN data to the second data center (VPLEX R1) using caching to make it seamless. This is what makes VMotion over distance a reality as it removes the need for a replication step. Instead of having to do a VMotion and a Storage VMotion, VPLEX is VMotion only. With VPLEX, Distance VMotion is as easy as VMotion within the data center.

Figure 3. VPLEX Metro Enabling VMotion Between Data Centers



ESG Lab Validation

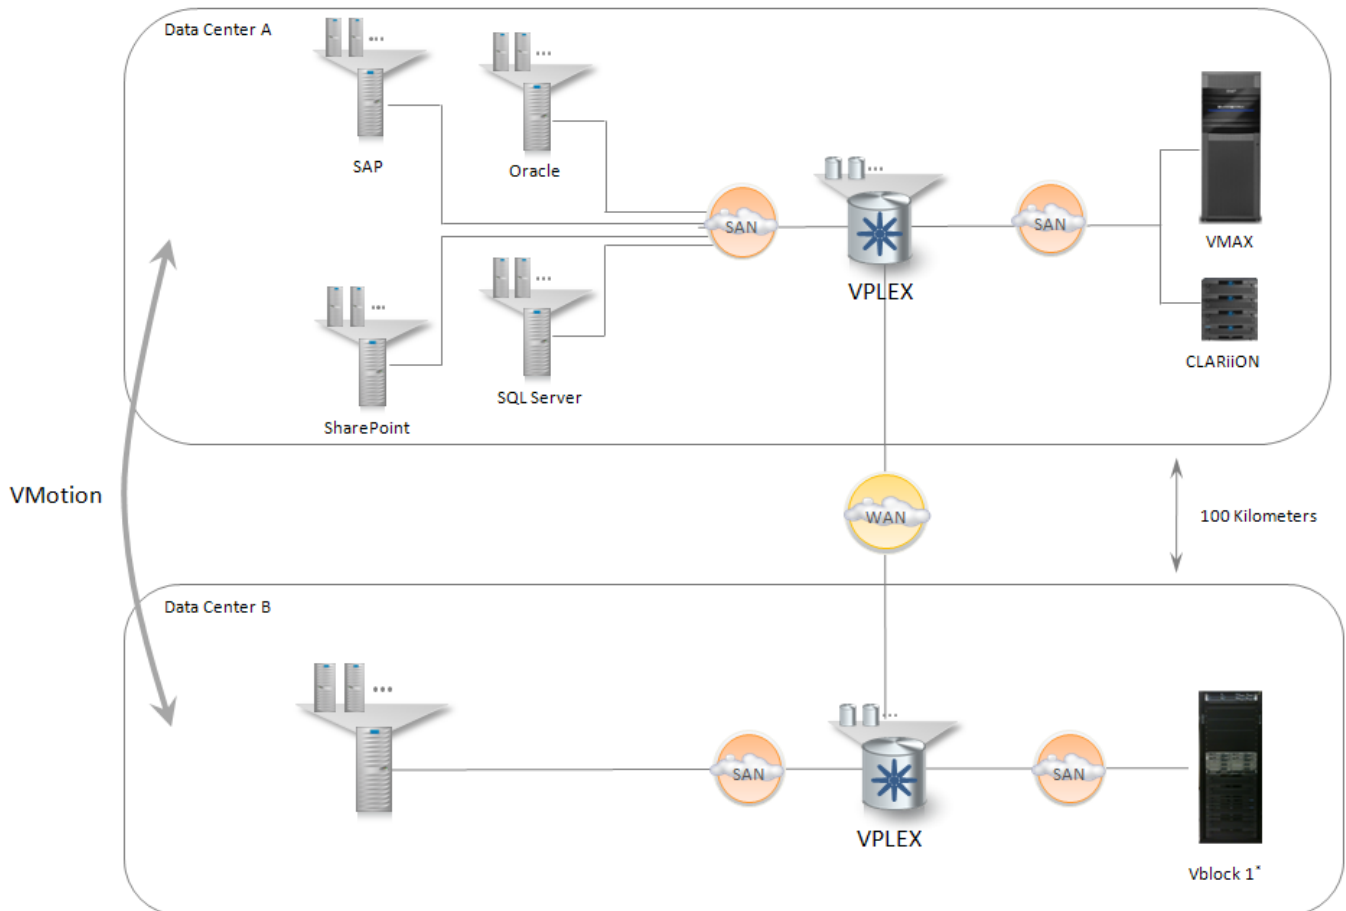
The balance of this report was designed to explore how VPLEX can be used in an existing VMware environment. VMotion between two data centers is demonstrated using a heterogeneous mix of disk arrays. The performance impact of VPLEX deployed within and between data centers is examined.

ESG Lab testing was performed in an EMC solutions lab in Hopkinton, Massachusetts. The test bed was designed to evaluate real-world application workloads running in a VPLEX enabled environment. Oracle, SAP, Microsoft SQL Server, and Microsoft SharePoint were tested using industry-standard workloads. ESG Lab’s analysis of the VPLEX was based on a combination of hands-on testing, an audit of EMC performance test results, and conversations with VPLEX Beta customers.

The ESG Lab Test Bed

The configuration used for ESG Lab testing is shown in Figure 4.³ SAP, Oracle, SharePoint, and SQL Server applications were deployed within VMware ESX-enabled virtual machines running on four physical servers in data center A.

Figure 4. The ESG Lab Test Bed



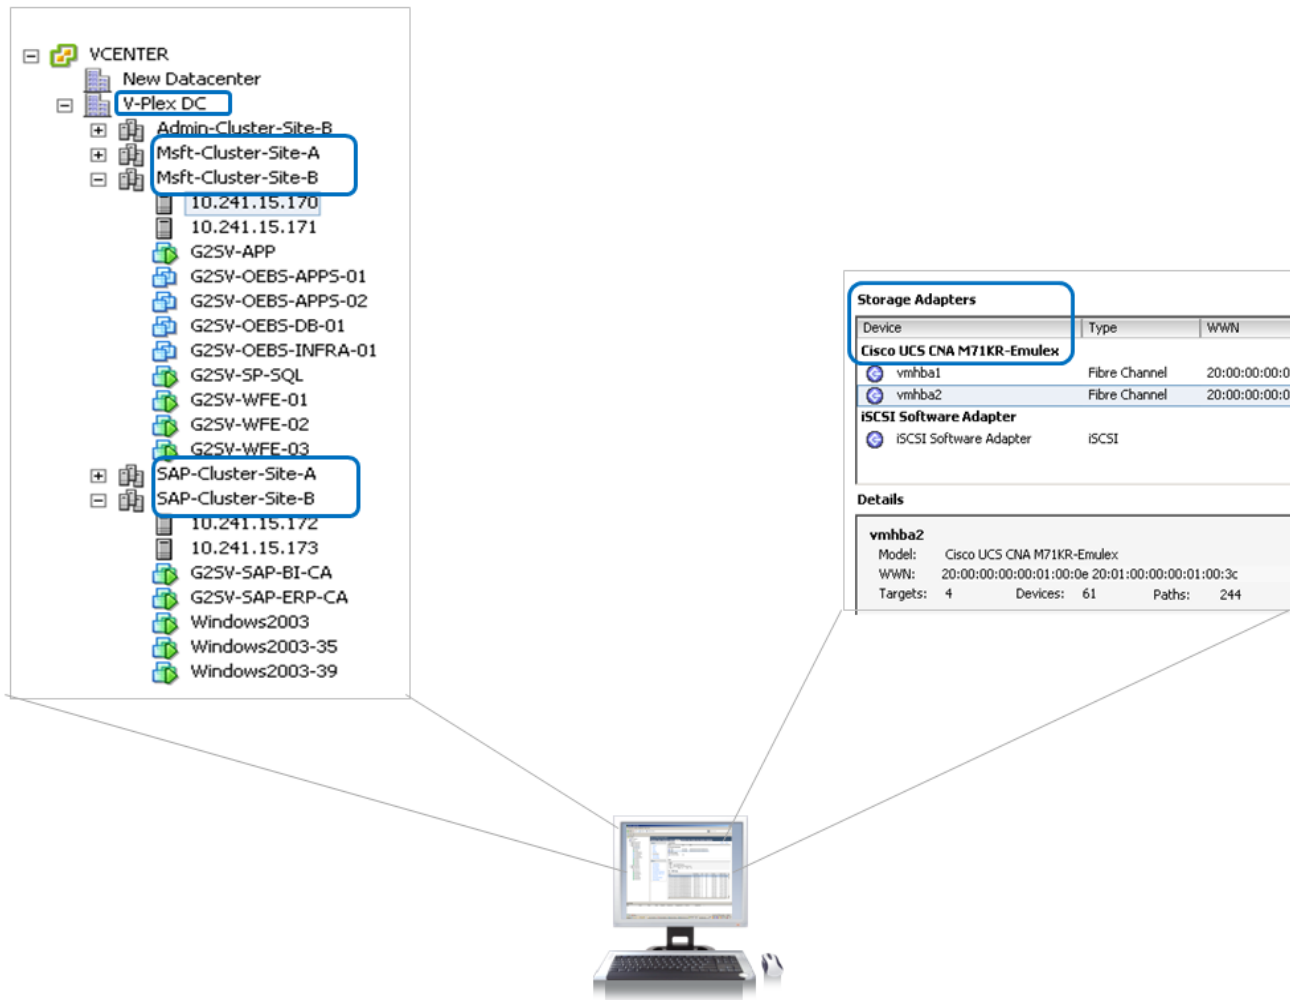
* Vblock 1 was used for EMC CLARiiON storage virtualized by VPLEX, Cisco servers virtualized by VMware and Cisco switching

³ Configuration details are documented in the Appendix and an EMC white paper entitled *Distance VMotion for Microsoft, Oracle, and SAP Enable by VCE vBlock, EMC Symmetrix VMAX, EMC CLARiiON, and EMC VPLEX Metro*, April 2010.

The servers in data center A were SAN connected to a pair of VPLEX engines, which was SAN connected to an EMC Symmetrix VMAX and an EMC CLARiiON disk array. The virtual infrastructure in data center A was WAN connected to a simulated data center B. WAN emulators (one for the LAN and another for the SAN) were used to simulate 100 km of distance between the two data centers. A Cisco UCS and an EMC CLARiiON within VCE vBlock 1 were used for the servers and storage in data center B.

A VMware vCenter view of the test bed is shown in Figure 5. Note that the two sites are not configured as separate data centers; instead, they are defined as two clustered pairs within one logical data center. It should also be noted that Converged Network Adapters (CNA) from [Emulex](#) and Converged Enhanced Ethernet (CEE) switches from [Cisco](#) were used to consolidate Ethernet and Fibre Channel networks onto a 10 Gigabit Ethernet network. This reduced the number of cables that needed to be connected to each server.

Figure 5. A VMware vCenter View of the ESG Lab Test Bed



Getting Started

ESG Lab testing began with the configuration of VPLEX between existing virtual machines and disk arrays. There are two ways that VPLEX can be configured with existing virtual machines: non-disruptive and disruptive. The non-disruptive method uses VMware Storage VMotion to migrate existing storage to a newly configured pool of storage. This requires additional storage and takes time, but can be done online. The disruptive method does not require new storage capacity, but does require that the virtual machines are not running during the configuration process. After VMware has detected the now VPLEX-managed LUN, the virtual machines can be powered on and resume normal operation. During ESG Lab testing, the disruptive method was used to configure VPLEX for use between an existing virtual machine and an EMC CLARiiON disk array.⁴ The process was later reversed to confirm that the VPLEX can be taken out of the configuration with no loss of data.

ESG Lab Testing

The home screen of the VPLEX Management console provides a good overview of the configuration process. As shown towards the bottom left in Figure 6, a storage volume within a disk array (shown in orange) is discovered and divided into extents. Extents are used to build VPLEX devices (shown in green). VPLEX devices are optionally mirrored using extents from a second disk array. The VPLEX device is used to create a virtual volume which is presented to hosts. Hosts that have been registered as initiators (shown in blue) access VPLEX volumes through a storage view (shown in red).

Figure 6. The VPLEX Management Console

⁴ ESG Lab confirmed that the non-disruptive mode is supported and routinely tested by EMC.

ESG Lab configured an existing virtual machine deployed on one of the physical servers in data center A. EMC refers to this process as encapsulation. An overview of the configuration process is listed below. Note that the VPLEX LUN encapsulation steps listed in step 4 can be performed using wizards listed towards the right on the VPLEX management console shown previously in Figure 6.⁵

- I. The virtual machine was turned on and a text file was created and saved to the desktop.
- II. The virtual machine was shut down.
- III. The LUN that was servicing the virtual machines was removed from the vSphere inventory.
- IV. The LUN was encapsulated:
 1. The LUN was discovered.
 2. The LUN was claimed.
 3. Extents were created.
 4. A VPLEX device was created.
 5. The VPLEX virtual volume was created.
 6. The virtual server's initiator was registered.
 7. A storage view was created and mapped to the initiator.
- V. The vCenter console was used to rescan and discover the VPLEX LUN.
- VI. The virtual machine was configured to use the VPLEX LUN.
- VII. The virtual machine was booted and the text file on the desktop was verified.

Thirty-three minutes after getting started on a pre-wired and -zoned test bed, an existing virtual machine was up and running in a VPLEX-enabled environment. The process was repeated in reverse to confirm that an existing virtual machine can not only be configured, but can optionally be de-configured without losing data.

Why This Matters

ESG Lab found that getting started with VPLEX was intuitive and straightforward. VPLEX can be seamlessly added to an existing SAN infrastructure with no interruption. When desired, an existing application can be configured to take advantage of a VPLEX distributed storage federation (and back again if needed) in a matter of minutes.

⁵ ESG Lab used a combination of VPLEX GUI wizards and VPLEX command line operations to complete the configuration.

VMotion over Distance

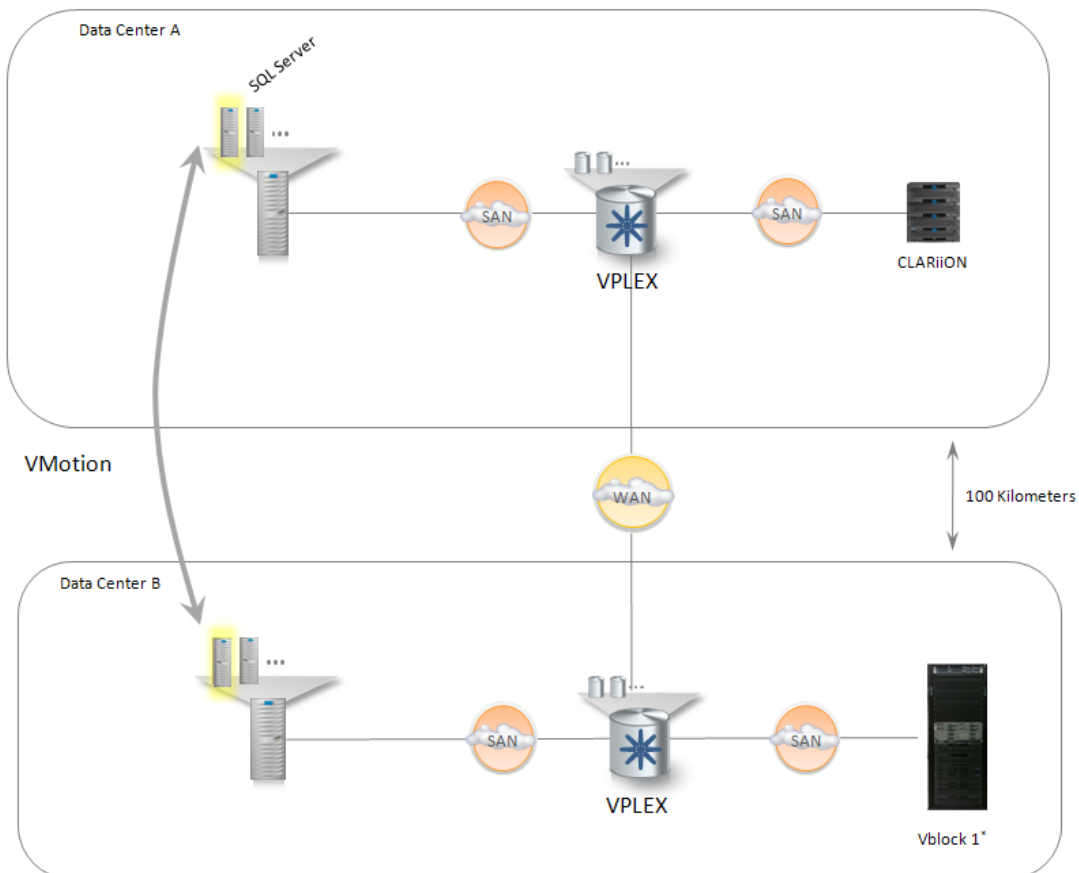
VMware VMotion is used to move a running application from one server to another within a data center. VPLEX Metro extends this capability as it enables the online migration of applications between data centers up to 100 km apart. VMotion over distance has not been practical until now due to the challenges associated with creating a seamless, latency free connectivity to storage at distance. VPLEX creates a decoupled layer between the application and storage as it maintains a virtual, mirrored, cache coherent copy of data in both data centers. As a result, VMotion works the same as it does within a single data center. There is no need to fail over the storage to a mirrored copy at a second data center (i.e., Site Recovery Manager is not needed). There is no need to pre-allocate storage and wait for the data to be mirrored by VMware (i.e., Storage VMotion is not needed). The combination of VPLEX Local and VMware VMotion provides similar value within a local data center. A storage engineer at a global financial institution explained the value of this when he said:

“We present large LUNs to ESX and carve them into a bunch of smaller VMware file systems for guests. The problem we have now is that if we want to move a guest to another server, then we have to disrupt all of those other guests that share the same LUN. Since VPLEX makes the LUN active/active in two locations, that means we can fail over without disrupting the other guests. That’s what got us interested in VPLEX. I know that you could argue that we could have used smaller LUNs so each guest has its own storage, but that would increase admin overhead and complexity.”

ESG Lab Tested

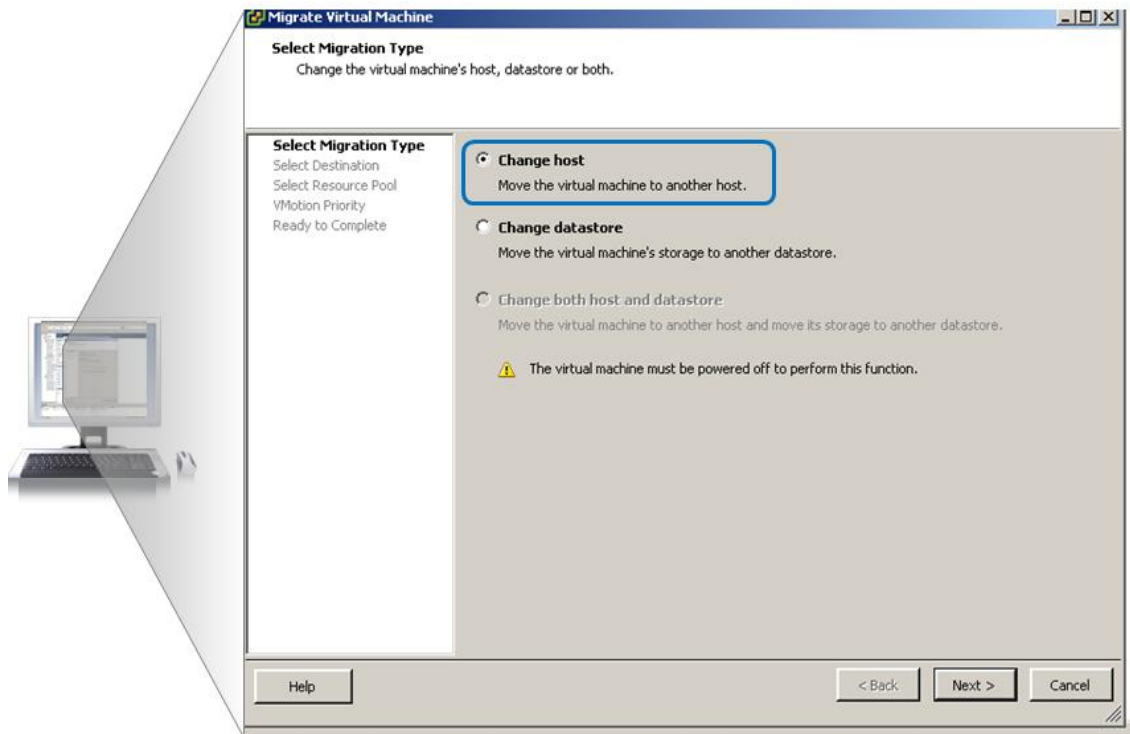
ESG Lab used VPLEX Metro and VMware VMotion to move a virtual machine between two data centers located 100 km apart. As shown in Figure 7, a SQL server virtual machine in data center A was moved online to data center B. VPLEX maintained a mirrored copy of the data on an EMC CLARiiON disk array at site A and a vBlock 1 at site B.

Figure 7. Testing Distance VMotion



The VMotion operation began with a drag and drop of the VM from cluster A to cluster B. After dragging and dropping the VM from a server in cluster A to a server in cluster B, the “Change host” option in the wizard panel shown in Figure 8 was used to kick off the VMotion operation.

Figure 8. Kicking Off a VPLEX-enabled VMotion



Note that the “Change datastore” option, which is used to start a VMware Storage VMotion, was not selected because VPLEX was maintaining a mirrored copy of the VM in data center B. In other words, the storage didn’t have to be moved to data center B because it was already there.

A write intensive IOmeter workload was running during and after the migration. The VMotion completed in one minute and twelve seconds. The workload continued throughout the migration and completed without error. The virtual machine remained up and operational throughout the move. The process was repeated in reverse as the VM was moved from cluster B to cluster A.

Why This Matters

IT organizations need to move applications from one data center to another for a variety of reasons including data center consolidations, planned maintenance, disaster avoidance, and load balancing. Traditional methods typically require planned downtime. If the application must remain online during the move, a costly and complex combination of host clustering software and disk array remote mirroring is needed.

ESG Lab has confirmed that the combination of VPLEX Metro and VMware ESX can be used to quickly and easily move a running virtual machine between data centers up to 100 km apart. The VMotion process looks and feels exactly as if VPLEX was not there. For administrators that are familiar with the drag and drop vCenter VMotion interface, there is nothing new to learn.

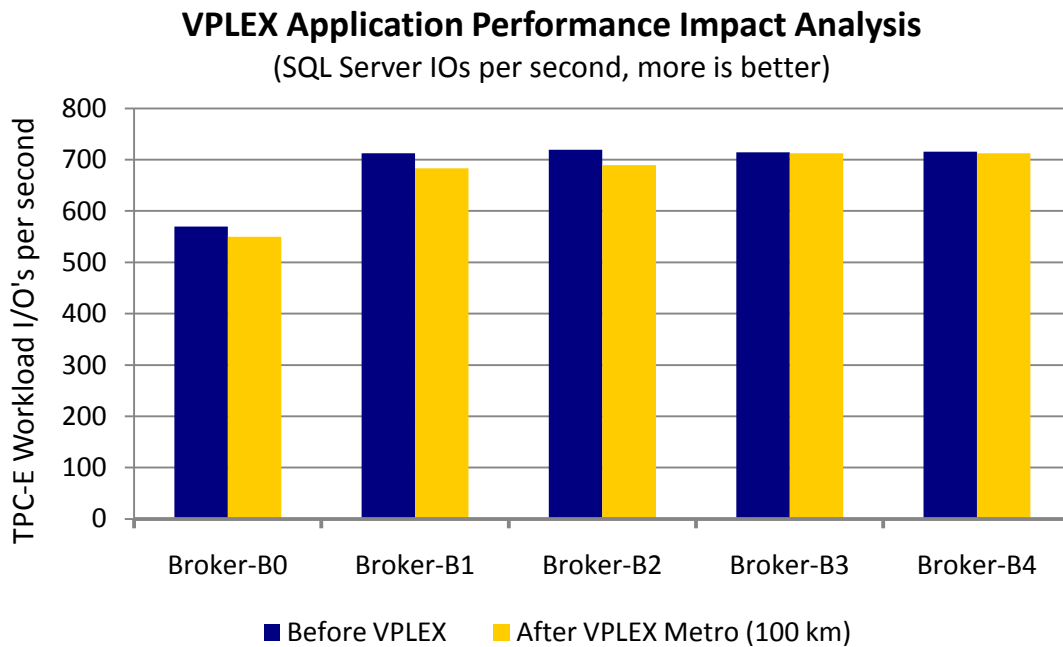
Performance Analysis

VPLEX uses advanced caching algorithms running on a scalable cluster of powerful hardware engines to overcome the performance challenges associated with deploying federated storage at scale and over distance. When a federated storage abstraction layer is inserted between hosts and storage, each I/O operation will take a bit longer. In other words, a federated storage abstraction layer adds latency. Scaling this abstraction layer to run over multiple nodes in a cache coherent cluster adds a bit more latency. And last but not least, stretching the abstraction layer between data centers adds even more latency. The efficiency and advanced caching algorithms at the core of the VPLEX architecture were designed to mitigate the effect that these latencies can have on application performance.

ESG Lab Testing

A one terabyte online transaction processing system (OLTP) database application was tested using the Microsoft TPCE Benchmark kit. The application was designed to emulate multiple users within a brokerage firm as they managed customer accounts, executed customer trade orders and tracked customer activity.⁶ Results were captured before and after VPLEX devices were configured with a goal of determining whether users would notice a difference in performance. The performance of the five busiest LUNs is depicted in Figure 9 and Table 2.

Figure 9. Measurably Low Microsoft SQL Server OLTP Application Overhead



⁶For more details on the performance benchmark methodologies and parameters, see the EMC white paper entitled *Distance VMotion for Microsoft, Oracle, and SAP Enabled by VCE vBlock, EMC Symmetrix VMAX, EMC CLARiiON, and EMC VPLEX Metro*, April 2010

Table 2: SQL Server Results

| LUN | Before VPLEX | | After VPLEX Metro | | Delta | |
|-----------|--------------|-------------|-------------------|-------------|------------|---------------|
| | IOPS | Latency(ms) | IOPS | Latency(ms) | IOPS | Latency |
| Broker B0 | 570 | 9 | 550 | 10 | 3.4% less | 11% more |
| Broker B1 | 713 | 8 | 684 | 9 | 4.0 % less | 12 % more |
| Broker B2 | 719 | 6 | 690 | 7 | 4.0 % less | 16 % more |
| Broker B3 | 715 | 5 | 713 | 5 | 0.2 % less | No difference |
| Broker B4 | 714 | 5 | 713 | 5 | 0.2% less | No difference |

What the Numbers Mean

- IOs per second (IOPS) is an indicator of the amount of work being done by the storage solution.
- As more IOPS are supported, more users can be supported. In other words, more IOPS is better.
- VPLEX Metro over 100 km reduced the number of IOPSs that the first Broker (B0) could perform by 3.4% (from 570 to 550 IOPS).
- Response times increased slightly, but not to an extent that users would notice. All response times were within Microsoft guidelines.

Why This Matters

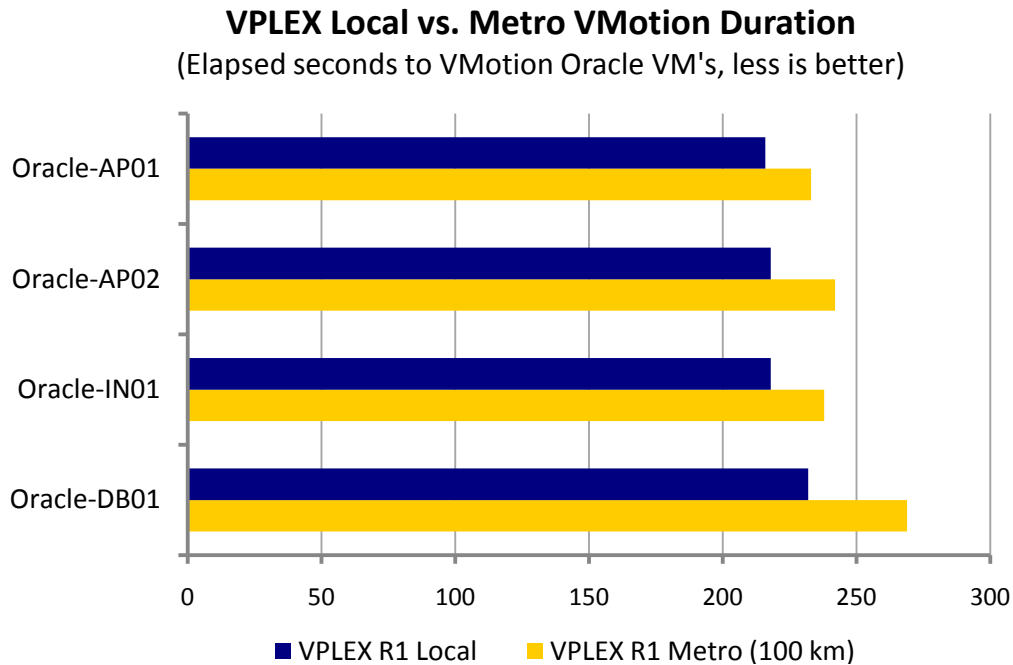
ESG research indicates that performance is a key concern when deploying applications in a highly consolidated environment. With multiple applications relying on a shared infrastructure, there is a concern that performance requirements can't be met predictably. As a matter of fact, 51% of ESG survey respondents that have deployed virtual servers connected to networked storage report that performance is their top concern.⁷ Replicating virtualized applications over a wide area network only increases the risk and therefore the concern as well. If you care enough to replicate an application between data centers, you probably care about its performance; multi-site flexibility at the expense of customer satisfaction is not a viable option.

ESG Lab has confirmed that the VPLEX Metro deployed in two data centers located 100 km apart introduced measurable, but manageably low, performance overhead for a Microsoft SQL Server 2008 OLTP application.

⁷ Source: ESG Research Report, [The Impact of Server Virtualization on Storage](#), December 2007.

The next test measured the time it took to VMotion a virtualized Oracle application as it was being load tested. The test was designed to ensure that the user experience remained the same while VPLEX-enabled VMotion was running. Using the Oracle Load Testing for Web application, each of two test scripts was allowed to ramp up to ten virtual users each. As expected, the average response time of transactions increased slightly (from 1.56 seconds to 1.62 seconds), but remained within acceptable limits. The time it took to complete the VMotion of the Oracle VMs from data center A to data center B is depicted in Figure 11.

Figure 10. VMotion Duration Analysis



What the Numbers Mean

- The graph depicts the elapsed time required to VMotion four virtual machines actively servicing 20 users accessing an industry-standard web server load testing tool from Oracle.
- As expected, the VMotion which moved the running virtual machines to a data center 100 km away (VPLEX R1 Metro) took a bit longer than the VMotion within a single data center (VPLEX R1 Local).
- In both cases, the VMotion completed in less than five minutes with no noticeable impact to users.

Why This Matters

Moving applications between servers and data centers can take hours, if not days, with traditional methods. ESG Lab has confirmed that the combination of EMC VPLEX and VMware VMotion can be used to quickly move running applications within and between data centers. As expected, VMotion within a data center completed quicker than a migration between data centers. In both cases, the VMotion completed in less than five minutes with no noticeable impact to users.

A long running Oracle operation was used to demonstrate how VPLEX caching can actually improve performance for read-intensive, cache-friendly applications. The time that it took to gather schema statistics for the Oracle web server databases was measured before and after VPLEX was configured.

Figure 11. Read Intensive Oracle Batch Job Analysis

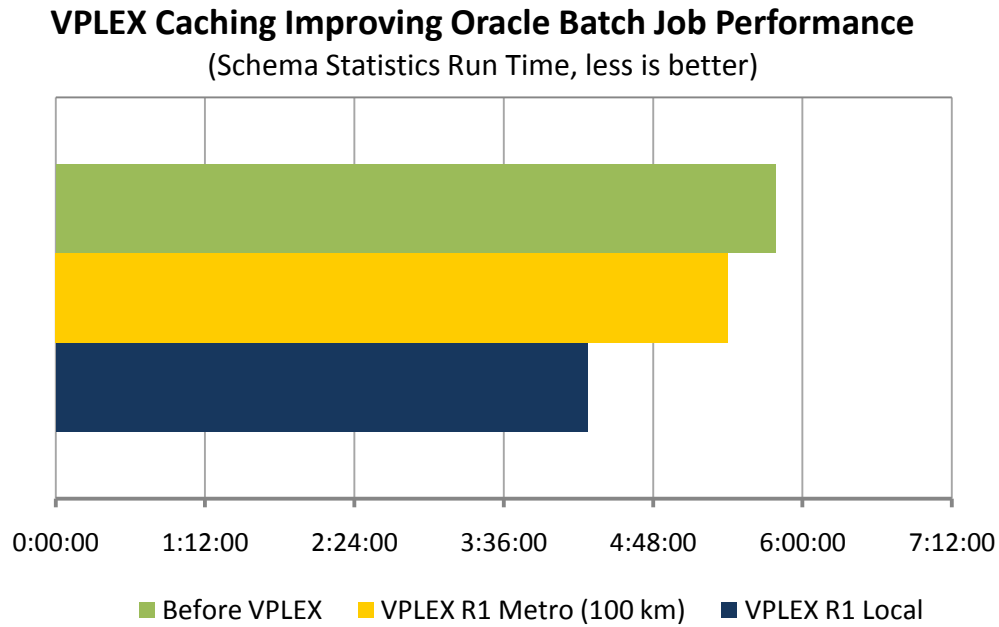


Table 3: Read-intensive Oracle Batch Job Analysis

| Configuration | Batch Job Duration | Delta |
|-------------------------|--------------------|------------|
| Before VPLEX | 5:46:55 | N/A |
| VPLEX R1 Metro (100 Km) | 5:24:12 | 6% faster |
| VPLEX R1 Local | 4:16:29 | 26% faster |

What the Numbers Mean

- An Oracle gather schematics operation is a long running, read-intensive batch job that is a good candidate for VPLEX caching.
- The Oracle batch job took five hours and forty-four minutes to complete before VPLEX was configured.
- VPLEX caching improved performance (6% with VPLEX Metro and 26% with VPLEX Local).

Why This Matters

Predictable performance is a critical concern when a mix of applications is consolidated within a virtual infrastructure. A performance bottleneck can lead to poor response times, lost productivity, and, in the worst case, lost revenue. ESG Lab has confirmed that VPLEX caching algorithms can improve performance for read-intensive applications deployed within a federated storage infrastructure deployed within—or between—data centers.

ESG Lab Validation Highlights

- ☑ The VPLEX management console and command line interface were used to configure an existing VMware-enabled virtual machine. Thirty-three minutes after getting started, the virtual machine was running with VPLEX configured between the virtual server and an EMC CLARiiON disk array.
- ☑ The VPLEX configuration process was executed in reverse to ensure that customers can optionally de-configure VPLEX without losing data.
- ☑ A running virtual machine was moved between two data centers 100 km apart in less than two minutes. The simple drag and drop VMotion interface that VMware administrators are familiar with was exactly the same with VPLEX.
- ☑ ESG Lab audited the results of exhaustive EMC application-level solution testing with SAP, Oracle, and Microsoft SharePoint and SQL Server.
- ☑ ESG Lab confirmed that the performance overhead introduced by VPLEX is manageably low. With four VPLEX engines working over two data centers located 100 km apart, the cached virtual storage abstraction layer has a measureable, but nominal, performance impact (0.2% to 4.0% less IOPS for a Microsoft SQL Server OLTP application).
- ☑ ESG Lab quantified the difference in the amount of time that it takes to execute a VPLEX-enabled VMotion within and between data centers. VMotion of four Oracle virtual machines took between 7.3% and 13.65% longer with VPLEX Metro. In both cases, VMotion completed in less than five minutes with no interruption to host applications.
- ☑ ESG Lab confirmed that for some application workloads, VPLEX caching can actually improve application performance. An Oracle batch job that completed in 346 seconds before VPLEX took 256 seconds to complete with VPLEX Local (26% faster) and 324 seconds with VPLEX Metro (6% faster).

Issues to Consider

- ☑ **WAN bandwidth.** Planning is needed to make sure that you can have a properly sized WAN connection between VPLEX Metro-enabled data centers. Like most organizations that have deployed synchronous remote mirroring for disaster recovery between two data centers up to 100 km away, a direct fiber connection (a.k.a., dark fiber) is recommended. If you have already installed Fibre Channel and Ethernet connections between buildings within a campus or metropolitan network, you should be all set. If not, you'll need to work with EMC and your WAN provider to make sure you can afford the WAN bandwidth needed to ensure that application performance is ensured during peak periods of activity.
- ☑ **GUI vs. CLI.** ESG Lab was impressed by the home screen of the VPLEX management GUI—especially the one-click access to wizards for each of the steps in the VPLEX configuration process. That said, the GUI could benefit from a number of enhancements with a focus on helping administrators manage an existing VPLEX deployment (e.g., inventory reporting, audit logging, topology viewers, and wizards for common tasks). As the GUI evolves into a more fully functional graphical user interface, administrators can use the well documented VPLEX command line interface (CLI).
- ☑ **More to Come.** There are a number of features and capabilities that will be needed before customers can realize EMC's ambitious goal of creating widely distributed storage federation unfettered by the limitations of scale and distance. Besides EMC's stated intention of supporting more scalability and distance in the future (e.g., VPLEX Geo), there are a number of incremental enhancements that ESG would like to see in future releases including Microsoft Active Directory integration, GUI enhancements, tighter integration with VMware vSphere, thin provisioning, and integration with EMC FAST.
- ☑ **Your Mileage may vary.** The performance results presented in this report are based on industry-standard benchmarks deployed in a controlled environment. Due to the many variables in each production data center environment, capacity planning and application level testing is recommended to assess the performance impact that VPLEX will have within your organization.

The Bigger Truth

VPLEX extends the benefits of virtual server technology to the storage infrastructure located within and between data centers. Like virtual server technology, VPLEX increases the flexibility and agility of consolidated infrastructure as it creates a single point of control for a consolidated pool of federated resources.

For those with experience in the storage industry, this should sound familiar. Storage virtualization solutions from EMC and others have been around for more than a decade. While it starts with the same goals of providing a flexible abstraction layer between servers and storage, VPLEX uses a radically new approach to tackle two fundamental challenges with traditional storage virtualization solutions: scale and distance. It uses advanced clustering and caching to extend the benefits of a virtual infrastructure beyond the walls of the data center. A cache consistent pool of storage capacity located within—and between—data centers turns a chaotic mix of disk arrays into a distributed federation of centrally managed storage capacity.

What are the practical implications of this new technology? The easiest to understand is the ability to non-disruptively move virtual server and storage resources to another data center. Instead of failing over to another site (e.g., with VMware Site Recovery Manager), ESG Lab moved a running virtual machine between two data centers in less than two minutes. The simple drag and drop VMotion interface familiar to VMware administrators was exactly the same with VPLEX. Compared to a VMware Storage VMotion, which could take hours to run in the background, the drag and drop approach was completed in minutes as VPLEX maintained a single cache consistent image of federated storage at both sites. Oracle SAP and Microsoft SQL Server and SharePoint virtual machines were moved between data centers with no interruption. An existing virtual machine was configured for use with VPLEX and the process was reversed.

ESG Lab confirmed that the performance overhead introduced by VPLEX is manageably low. VPLEX Local deployed within a single data center introduced measureable, but nominal, performance overhead (between 0.3% and 4.1% percent) during SQL Server testing. During Oracle load testing, VPLEX Metro configured between two data centers 100 km apart was only 3% slower than VPLEX Local in a single data center. As expected, VMotion between data centers 100 km apart was slightly slower than VMotion within a data center (7% to 14% longer)—but both were totally non-disruptive and completed in less than five minutes. Last, but not least, VPLEX caching actually improved performance for a long running Oracle batch job (7% faster).

VPLEX leverages the hardware engine at the core of EMC's enterprise-class disk arrays and field tested clustering software obtained through acquisition. Both have been deployed in extremely large mission critical environments. It was no surprise that the Beta customers ESG Lab spoke with were impressed with the stability of the solution during their evaluation of VPLEX with physical and virtual servers.

VPLEX is an innovation with the potential to change the way organizations deploy and manage storage infrastructure. Using clustering technology that's been field tested at scale in production environments, VPLEX has the potential to create a single view of storage that spans the globe. Existing EMC and VMware customers intrigued by the possibility of dynamically moving virtual servers between two locations within 100 km should view taking VPLEX for a test drive as a no-brainer. Customers shopping for a storage infrastructure upgrade that can maximize the value of the storage that they already own should consider adding VPLEX to their shopping list. Try it for a VMotion between data centers or an online migration between disk arrays and you can look forward to turning an out-of-control heterogeneous storage infrastructure into a centrally managed storage utility.

Appendix

Table 4. Test Bed Overview

| Storage | |
|--|--|
| EMC Symmetrix VMAX | Fibre Channel (FC) connectivity, 450 GB/15k FC drive |
| EMC CLARiiON CX4-480 | FC connectivity, 450 GB/15k FC drive |
| EMC VPLEX Metro | 2 VPLEX Metro storage clusters/4-director midsize configuration |
| Server | |
| SAP: Site A | Pair of 2 socket: 4 core: 2.7 GHz X5560 Intel Xeon , 96 GB RAM, Dual 10 GB Converged Network Adapters (CNAs) |
| Microsoft and Oracle: Site A | Pair of 4 socket: 6 core: 2.6GHz X7460 Intel Xeon , 128 GB, Two 10 Gb Emulex LightPulse LP21000 CNAs, Two Broadcom 5708 GigE adapters |
| VCE Block: Site B | Cisco Unified Computing System (UCS), Four blades of 2 socket: 4 core : 2.5GHz E5540 Intel Xeon, 48 GB, 2: Cisco UCS CAN M71KR-E-Emulex FCoE CNAs for Fibre Channel and Ethernet |
| Network | |
| Switches | Cisco Nexus 6120 switches |
| WAN Emulator | 1 GigE, 100 km distance |
| Fibre Channel SAN distance emulator | 1/2/4 Gbps FC, 100 km distance |
| Software – Virtualization | |
| VMware vSphere | 4.0 U1 Enterprise plus Build 208167 |
| VMware vCenter | 4.0 U1 Build 186498 |
| Software – SAN | |
| EMC PowerPath/VE | 5.4.1 Build 33 |
| Software – Guest OSs | |
| Red Hat Enterprise Linux | 5.3 |
| Microsoft Windows | 2008 R2 (Enterprise Edition) |
| Software – Middleware and Applications | |
| DB2 | 9.1 for Linux, UNIX, and Windows |
| Microsoft SQL Server | 2008 |
| Microsoft Office SharePoint Server | 2007 (SP1 and cumulative updates) |
| Microsoft Visual Studio Test Suite | 2008 |
| KnowledgeLake Document Loader | 1.1 |
| Microsoft TPCE BenchCraft kit | MSTPCE.1.9.0-1018 |
| SAP Enterprise Resource Planning | 6.0 |
| SAP Business Warehouse | 7.0 |
| Oracle E-Business Suite | 12.1.1 |
| Oracle RDBMS | 11GR1 11.1.0.7.0 |



Enterprise Strategy Group | **Getting to the bigger truth.**