

High-Availability Data Center Design for Healthcare with VPLEX

By Brett Foy, Technical Architect—Ahead

One of my healthcare clients has the admirable goal—and board-driven mandate—to build and deploy an IT infrastructure capable of achieving 100% uptime. It's critical for systems with patient data to always be available to the doctors, nurses, and other medical professionals prescribing medications, performing surgeries, comparing images, and making diagnoses.

Starting the Journey

This 100% uptime goal was so important to the client that they decided to invest in a brand-new production data center with redundancy built into the physical data center (facilities, power, cooling) as well as the IT data center (network, server, storage). They decided to go all-in with virtualization: keeping consolidation ratios relatively low, while taking advantage of load balancing, Microsoft server clusters, and VMware affinity rules to limit the impact of any ESX host failure.

The client invested in an EMC VMAX storage array for production and implemented all of the tools and systems required to ensure data replication to an off-site recovery location. The client had also planned to redeploy other storage targets—both in their production data center and at their DR site for data replication. However, they realized that in a catastrophic failure—such as cabinet damage, dual-feed power loss, physical data center local impact, or total array failure—they really needed an automated failover solution that would not require any administrative interaction or interruption to host I/O.

Planning for the Worst Case

With the realization that a localized worst-case scenario (loss of single storage array) would severely impact a heavily virtualized and centralized infrastructure, the client determined that even relying

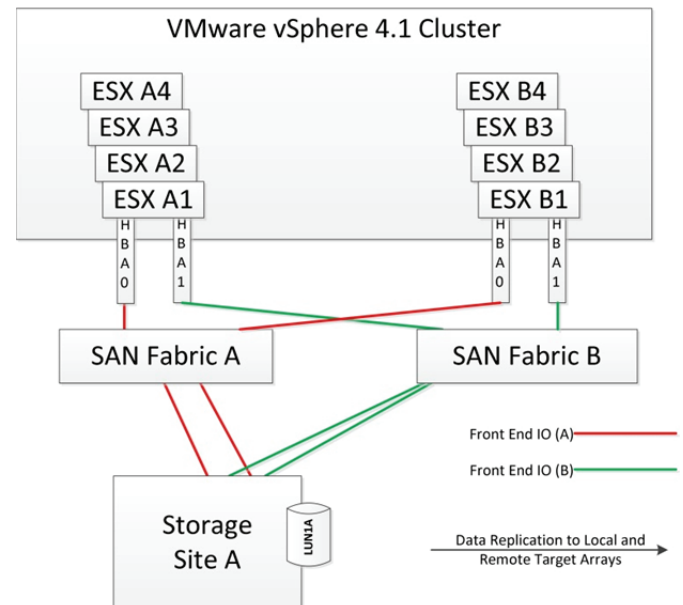


Figure 1: Original Data Center Design

on a single 5-9s array would not be enough. The client's revised requirement was to implement a no-data-loss active/active storage infrastructure with two completely distinct and separate storage arrays within their primary data center.

Our first meeting to review this new requirement was on the Wednesday before Thanksgiving. In that meeting we identified requirement. In the next meeting we whiteboarded the options to proceed using the EMC VPLEX technology.

New Technology, New Capabilities, New Issues

At EMC World 2010 in Boston, EMC announced VPLEX—a storage virtualization appliance that provides active/active data center capabilities for high-availability and multi-site active storage. For this healthcare client, VPLEX offered some unique capabilities within their production data center to give them what they wanted: an active/active

storage solution that provided protection against a catastrophic array failure.

Figure 2 shows the initial VPLEX Local configuration we proposed. Effectively, a LUN on the VMAX could be encapsulated, presented through the VPLEX to the host, and mirrored to the secondary array. In a local VPLEX configuration, VPLEX presents to the host a single virtualized LUN comprised of the mirrored members of both arrays. A key availability component is that the VPLEX does not handle any write-caching, since the write flows to both target arrays and is acknowledged by both arrays, at which time VPLEX acknowledges the committed write to the host.

In the event that either storage array (or all paths to either storage array) were eliminated, the VPLEX appliance would mark that array down and continue to serve host I/O in a non-disruptive fashion from the remaining array. (For simplicity the back-end paths through the SAN to the array are omitted in Figure 2.) When the “down array” comes back online, the VPLEX appliance manages the re-synchronization process, again, non-disruptively to the host until the mirrored status is recovered. *(This makes a **great** demo, which will be available soon in the Ahead Lab.)*

VPLEX Examined

The VPLEX architecture has two main components. The first is a fourth-generation software product acquired by EMC a few years ago that was engineered into the VPLEX platform as the GeoSynchrony operating system. The VPLEX technical platform is modeled on the VMAX high-availability multi-director/multi-engine hardware design. Built as a 5-9s solution in and of itself, we quickly identified one critical issue with this approach, specific to this client: all the solution did was move a catastrophic single point-of-failure up the stack.

For the same reason that the client was not comfortable with a single storage array, a single VPLEX cluster presented the same risk factor: a localized data center problem, cabinet damage, dual feed power loss, or an unlikely total system failure would leave no path to the virtualized mirrored storage.

This is when the meeting got really fun.

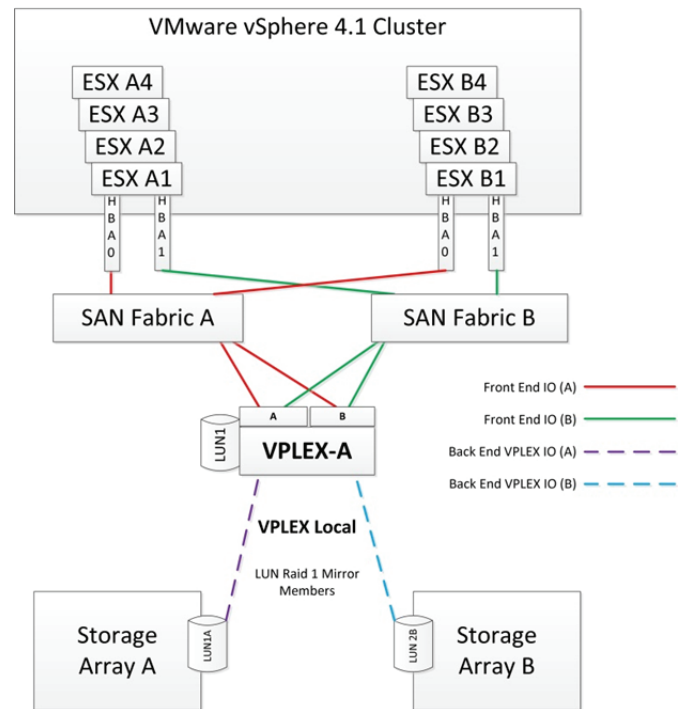


Figure 2: VPLEX Local for High Availability

In addition to announcing the VPLEX Local solution for storage virtualization and data migration, EMC had also announced and released VPLEX Metro at EMC World 2010. Initially, the primary use case for VPLEX Metro was building active/active data center solutions for synchronous-distance data centers for cluster aware applications. (EMC will follow this in future releases with VPLEX Geo and VPLEX Global for distant and multi-data center solutions.)

Revising the Approach

In the span of 20 minutes, we quickly changed our approach to solving the client problem. To begin with, the client wanted to build their high-availability solution in a single site—with the goal being never to have to fail over to a secondary site. Secondly, the client’s DR data center was outside the distance limitations for synchronous data access, so VPLEX Metro in a typical configuration between sites was not an option. Instead, we got to the whiteboard and developed a VPLEX Metro solution **contained in a single data center.**

VPLEX Metro introduces the concept of a stretched LUN. Instead of a mirrored active LUN in a single site, VPLEX

Metro brings a significant new technology called AccessAnywhere to the solution. Working with cluster-based storage-aware OS technologies such as VMware's vSphere, AccessAnywhere technology relies on advanced VPLEX Cache Coherency implementation to allow access to a shared LUN from multiple locations reading and writing from storage on different storage arrays. (*How VPLEX Metro works is the subject of many other widely available blog posts. Several solid references are provided at the end of this treatise.*)

The important thing for this story is how the configuration in Figure 3 was created and how we subsequently arrived at a final solution to meet the client's most critical requirement for on-site high-availability and non-disruptive fault recoverability.

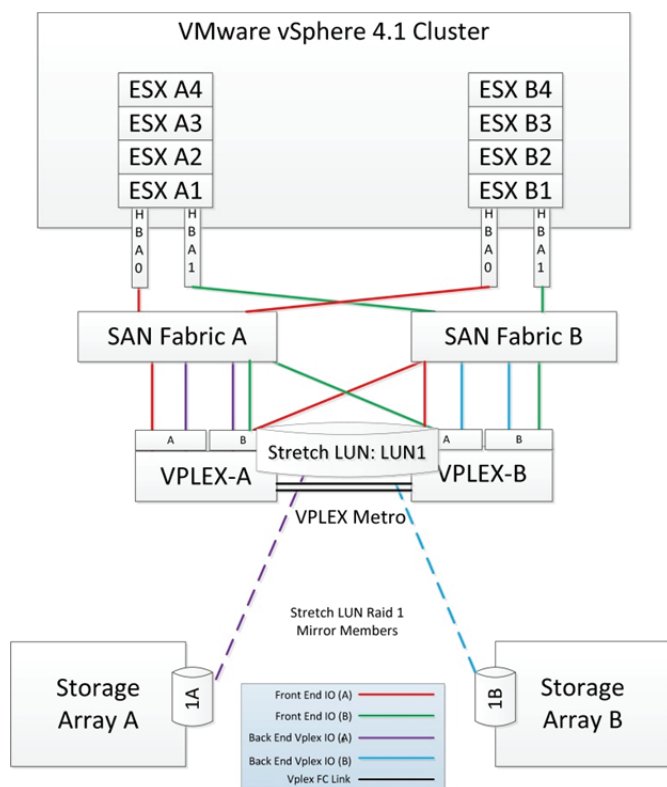


Figure 3: Initial VPLEX Metro Design

We developed the first draft of this configuration on the fly in the first meeting. In following meetings, we spent a lot of time analyzing how this solution should be implemented, and what impact this design solution would have on many other factors. How would this solution get the client closer to their goal of 100%

uptime? How would this solution impact DR and streamline DR recovery? What would happen in this solution if the link between VPLEX metro clusters were disabled? What would happen if a complete VPLEX cluster failure or catastrophic failure occurred? How does vSphere work in a VPLEX infrastructure (*and there are a lot of design considerations here*).

How's that for a 20-minute meeting taking an interesting turn and evolving into an entertaining and exciting technology conversation?

Decision-Making and End-State Design; Caveats and the Future

The overall driver for the technology selection was the requirement to build an active/active automated storage infrastructure to move closer to the goal of building a highly resilient data center... one that could handle the loss of any infrastructure component without impacting users. The solution we created on a whiteboard in a couple of hours looked like a great use case for EMC's VPLEX technology. And in fact, the client ultimately selected this technology and is in the process of deploying it today in their new and soon-to-be production data center. But nothing can be this easy—right? That is correct.

VPLEX Considerations

The VPLEX technology is managed through a feature-rich, but easy-to-use GUI. The technology implementation itself is fairly straightforward. But what are the design considerations and what tough decisions would be needed to roll this out in production?

Active/Active vs. DR

The client's primary and secondary DR data centers are about 120 miles apart. This is too far for a synchronous replication or data recovery solution. The client had already selected, evaluated, and deployed EMC RecoverPoint for data replication between existing CLARiiON storage arrays. They intended to use a combination of the SAN-based RecoverPoint splitter to manage data replication between a primary VMAX storage array and a DR CLARiiON array in the DR site.

Being highly virtualized and/or running Microsoft cluster technologies, the client had developed a fairly comprehensive DR restart plan, with the planned deployment of SRM for VMware and the use of Microsoft Cluster services to restart systems in the event of a true disaster. Central to both of these use cases was the ability of RecoverPoint to work with the underlying host OS to manage the recovery process. For VMware, SRM can leverage RecoverPoint and automation to restart guest OSs on remote hardware. For MS SQL Clusters, we planned to use RecoverPoint/Cluster Enabler (RP/CE) to coordinate the restart of MS Cluster members at the DR data center.

The original plan was for two replication streams from the production systems: a local CDP copy for on-site recovery (to protect against array failure) and a remote CRR copy for off-site DR protection. Figure 4 depicts how this was envisioned to work for Microsoft cluster services (and would have been similar for VMware SRM).

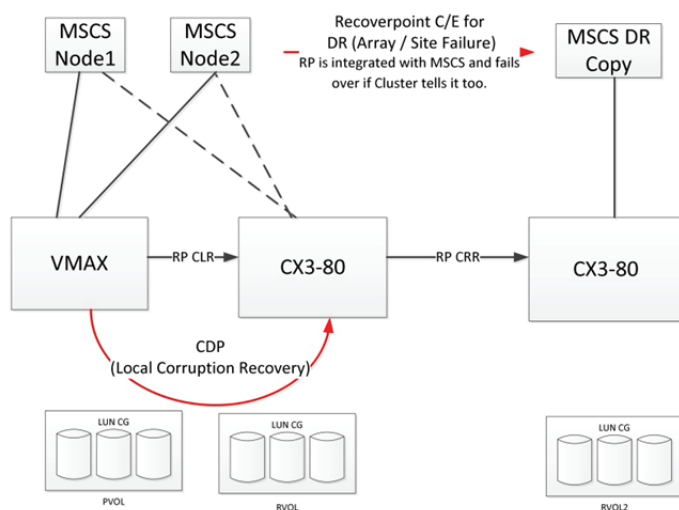


Figure 4: Original Local and DR Recovery Plan

The key item of note (and concern) to the client was this: in the event of a catastrophic array or data center failure affecting the VMAX, how would they decide to restart the infrastructure either locally or remotely?

The client wanted a solution that would detect an outage on the VMAX and non-disruptively mount all systems on the local CDP copy of the infrastructure prior to failing over to the DR site. However, all of the tools being referenced were used to create a DR/off-site capability.

Besides the tools and capability issue, there was a very simple question to be answered in the production data center: how and when do you decide to declare a VMAX array failure and either automate or manually bring the local CX3-80 online? The process of selecting a CDP recovery point and rezoning all hosts/mounting RecoverPoint LUNS locally is not a trivial task. If you set up the tools to work locally, you need to disable the tools from being able to work remotely. Either way, the solution would not be automated and no matter what, there would be a host impact during a failure event. On the plus side, DR itself could be largely automated in the event of a data center failure, but this client is really focused on local vs. distant recovery. The availability to achieve 100% local uptime was ultimately the driving force in the technology selection process.

So now the question became—how would introducing a “local” VPLEX Metro cluster in the production data center affect their ability to automate disaster recovery?

The first question to answer was how they would replicate data between data centers. The original data center design called for the use of RecoverPoint with the SAN-based Cisco splitter to mirror host writes destined to the VMAX to both CDP (local) and CRR (remote) copies of data. However, if you consider the VPLEX design, you can see a potential issue in that we were already altering the path of traffic in the SAN by introducing a storage virtualization layer (virtualized and encapsulated LUNS were being presented by VPLEX to the host). The idea of presenting two technologies in the data path (RP and VPLEX) seemed like a lot to consider and is, as of this writing, still an unsupported configuration.

What we did, however, was use array-based replication to get data off-site. Because we configured VPLEX to use encapsulated LUNS from both arrays, we had the full capability to do anything we wanted with the encapsulated VPLEX LUNS. Our client can implement local snaps/clones and can replicate specific LUNS (as opposed to entire pools), all without the VPLEX knowing or caring what is happening to its member LUNS. It was an easy decision to leverage the embedded CLARiiON RecoverPoint splitter to replicate data between CLARiiON SPs at both sites for all critical data being protected in the production site by the

VPLEX. Using the embedded RecoverPoint splitter allowed replication to take place outside of the data path—behind the VPLEX with no ramifications to the VPLEX storage virtualization.

While this design enabled the client to get data replicated off-site in a near-synchronous fashion, it *did* create two design constraints that were reviewed, weighed, and ultimately approved based on the primary functional design consideration of having an always-available local storage infrastructure. Issue number one was that while loss of the NS960 in the production data center may not cause disruption to production hosts, there would absolutely be a loss of data replication to the target site until the NS960 could be returned to production. This created the potential to affect DR RPOs in the event that the remainder of the data center or the production VMAX itself were to experience a failure or shutdown following the production NS960.

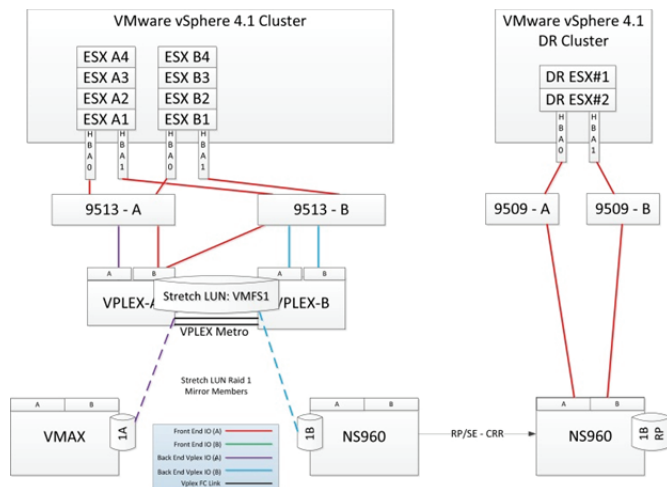


Figure 5: VPLEX Metro with Array-Based DR Replication

The second issue was that by using VPLEX with back-end array-based replication, the client lost the ability to leverage tools such as SRM and RP/CE. Essential to both SRM and RP/CE is that the hosts have knowledge of the datastore LUN structure and, by working with the underlying replication and array technology, can automate the failover and recovery of storage for the DR host environment. For now at least, there is no SRM SRA for VPLEX, so the client again decided to err on the side of

higher availability in the production site with a manual/scripted recovery instead of an SRM recovery for DR.

So in the end, the question of active/active vs. DR truly came back to the client's original functional design goal: keep the production site up at all costs. Active/active won—and DR will require more work because of it, but the technology is still flexible enough to do what the client needed to assure off-site data protection in a timely manner. It will just require more work to make recovery happen in a real DR event.

For many organizations, the question of active/active isn't typically a "same data center" question; it's a question of synchronous distance data centers for high availability in a multi-location infrastructure.

VMware and VPLEX Design Considerations for Failure Scenarios

There are trade-offs between building an active/active data center within one site and having automated disaster recovery. For this client, the goal is to never have to declare a disaster and shift production workload. Yes, the requirement and capability are there, but the business priority is to keep the production data center up and running with the absolute highest availability capabilities.

To do this, the client decided on a VPLEX Metro Cluster configuration in their production data center. This consisted of two storage arrays behind two VPLEX clusters. By presenting mirrored storage from two arrays to a VPLEX Metro Cluster stretched LUN, we've created an architecture where there are multiple paths to multiple access points to the same LUN being presented by multiple storage arrays.

This design presents a scenario by which storage is always available to the hosts. In a highly virtualized infrastructure, we have effectively created a very highly available SAN, VPLEX, and storage infrastructure—with the one real risk being loss of a physical ESX host. VMware provides coverage for that through HA (minimal disruption/automated recovery). For this client, VMware's Fault Tolerance wasn't (yet) a viable option. To further protect against host outages,

the client is still implementing SQL Clustering and other cluster technologies on top of VMware to protect their most critical HIS and EHR systems. At this point, we seem to have everything covered, meeting the business and functional goals of the organization.

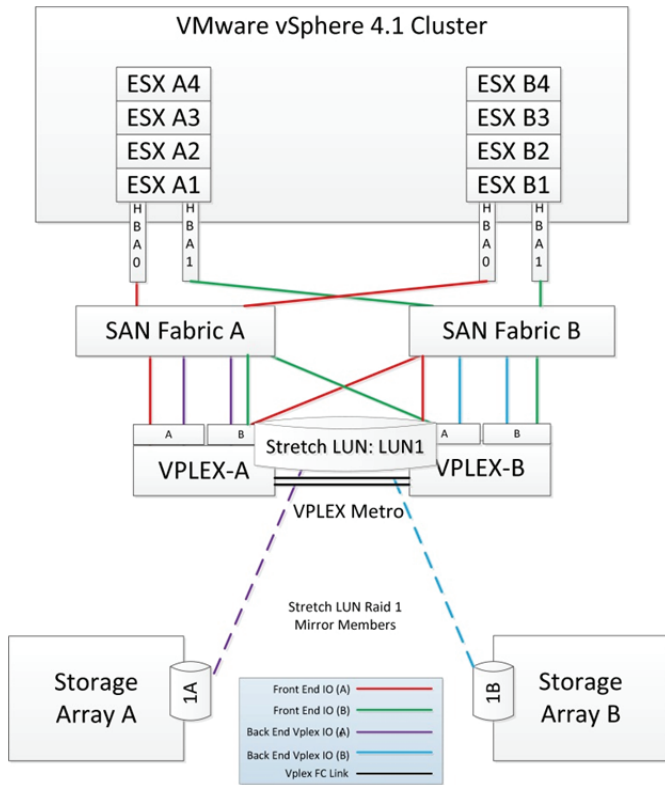


Figure 6: Intra-Data Center VPLEX Metro Design

That being said, implementation of VPLEX introduces some failure scenarios, however unlikely. The solution needs to be carefully architected to protect against failure and impact to host/storage access. Potential VPLEX failure scenarios include:

- loss of a component on either VPLEX Metro Cluster (FC Port, Engine, Cache, etc.)
- loss of connectivity between the two VPLEX Metro Cluster Members
- loss of a complete VPLEX Cluster

Loss of VPLEX Component

The VPLEX architecture is designed to be highly available from the perspective of a single instance. The VPLEX hardware architecture is based on the EMC

VMAX architecture built for 5-9s+ of uptime. A single-engine VPLEX has two directors that are redundant. Loss of either director or component of the director does not impact host access or storage connectivity since a properly configured SAN infrastructure allows active/active failover between the directors. This high-availability design is only strengthened by the addition of more engines/directors to the configuration. Simply stated, loss of any single (or in some cases, multiple) component on a single VPLEX cluster will not affect operations of host access to storage. The biggest consideration here would be impact to available I/O bandwidth—and the VPLEX throughput scalability exceeds most clients’ requirements in a single engine.




Figure 7: VPLEX Component Failure

Loss of Intra-Cluster Link Connectivity between Two VPLEX Metro Cluster Members

If the dedicated connectivity between VPLEX Metro Clusters is lost, but both Clusters are still up, the very real possibility for split brain exists. Without special error handling in this scenario, it is feasible to foresee that data could be continued to be written to the same LUN by hosts on each side of the Metro Cluster—creating a data corruption event because the separate copies would be out of synch. To prevent this split brain scenario and ensure that only one side of the Metro Cluster continues to allow writes to the stretched LUN, VPLEX introduces the concepts of preferred LUNS and sidedness. These concepts, explained in depth later in this document, play a crucial role in determining which hosts continue to have access to which LUNS. It is crucial to understand what these concepts mean from a design and planning perspective.



Figure 8: VPLEX Metro Intra-Cluster Link Failure

Loss of a VPLEX Cluster

In the unlikely event of a total cluster failure, the remaining VPLEX Metro Cluster must determine if the other cluster has failed or if there is a link failure. In the current release of VPLEX, there is no external monitor or witness to provide the cluster with this knowledge, so the remaining cluster member must assume a link failure condition exists to prevent potential split brain. Similar to the previous failure scenario, this means that access to the preferred LUNS of the “down” VPLEX will be marked unavailable for read/write I/O through the remaining VPLEX Metro Cluster member. To restore access to the remaining cluster’s non-preferred LUNS, an administrator must make the manual decision and allow VPLEX to enable those paths. Again, the key here is understanding what this means from a design and planning perspective!




Figure 9: VPLEX Metro Cluster Failure

Failure Scenarios—A Deeper Look

In this section we’ll dig a little more deeply into each of these failure scenarios—starting with some VMware constraints and requirements for deployment with VMware, how VPLEX actually manages a storage array failure, and followed a detailed look at by the three VPLEX failure scenarios.

The VPLEX solution is built to be highly available. While at first glance it may seem daunting on paper, the failure scenarios and associated impacts of those scenarios can be architected and designed to minimize impact. This is not a simple drop-in solution. It is an active/active data center solution—something new to the industry. And just like any other highly available infrastructure, it must be well understood and documented, and have all the caveats laid out in advance to plan a successful deployment.

Detailed Design and Considerations—vSphere

VMware vSphere 4.1 Infrastructure with VPLEX

Designing and deploying a VMware vSphere infrastructure without VPLEX requires significant planning and extensive design work to achieve desired results. Planning a vSphere deployment with VPLEX introduces new complexities, many of which are listed below:

- VMware clusters have an absolute maximum of eight hosts per cluster to conform to VMware HA Admission Control Policies.
- VMware hosts are all contained in a single vCenter instance and are a logical vSphere cluster, as opposed to multiple vCenters in a distributed SRM infrastructure.
- VMware is not aware of sidedness of the host layout, and therefore DRS and HA rules will not respect sidedness considerations.
- Each side of a typical VPLEX Metro Cluster solution should run < 50% total utilization to accommodate total failover of the other site.
- VMFS datastore design should take VPLEX distributed volume preferences into account, specific to placement of VMs on preferred LUN datastores.
- In a mixed infrastructure with VMware hosts connected to both VPLEX clusters, the VMware administrators should maintain sidedness of the virtual infrastructure manually.
- A-Side host guests should always be placed on VPLEX-A preferred distributed volumes (LUNs).
- B-Side host guests should always be placed on VPLEX-B preferred distributed volumes (LUNs).

VMware Recovery

- If any host fails, VMware HA will restart all eligible VMware guests on surviving cluster’s hosts.
- This is an availability event that will require a guest restart and will likely impact end users.
- In a multi-cluster VPLEX design, HA Admission Control Policy for host or site restart requires careful design and implementation, including potential for additional hosts to support admission control options.

Storage Infrastructure Failure/Recovery

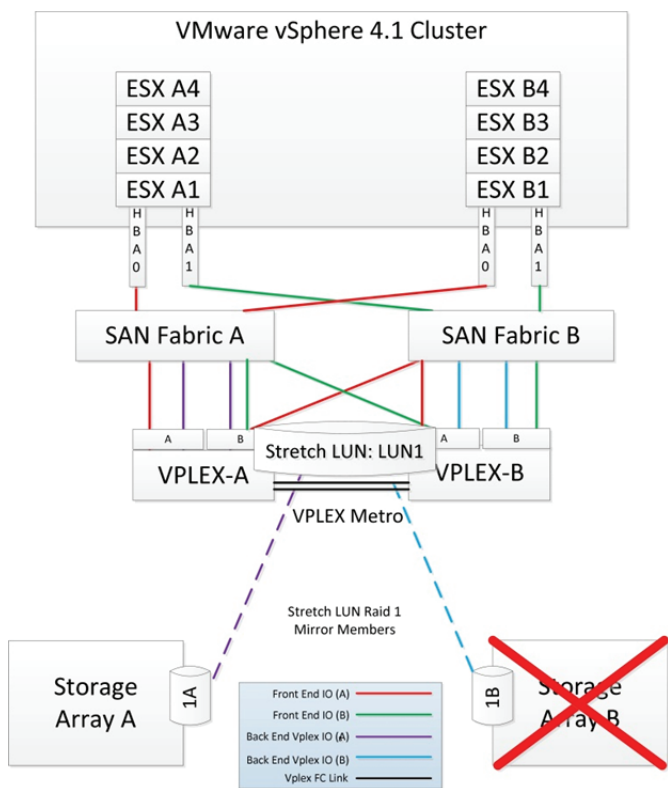


Figure 10: Failure Scenario, Loss of Storage Subsystem

Availability Considerations

When planning for the design of the storage infrastructure, the following considerations should be taken into account:

- LUN 1A from VMAX-A and LUN 1B from VMAX-B are, from the VPLEX perspective, Raid-1 mirrored LUNS, with mirroring controlled and managed by VPLEX.
- Each storage array is fully configured and highly available to accommodate any component failure without impact.
- In the unlikely event there is a total array failure (or loss of all data paths) to either storage array, the VPLEX will note that the storage path is unavailable and proceed with host read/write I/O from the remaining array with no host impact (sub-second delay).
- There is no host HA, restart, or clustering impact required, because the host never sees the outage, which has been made transparent by VPLEX.

- If required, it is possible to leverage snaps and clones for PIT recovery on either local mirrored storage array as required. (This can be mounted through or outside of VPLEX.)

Array Recovery

- A storage array failure of either leg is a non-disruptive event to the host.
- When the array is returned to service, VPLEX will manage re-synch of the primary active mirror volumes to the newly restored storage array LUNS. It will then re-add the restored array to participate in the mirrored read/write environment.
- In the event that the array that fails is the replication source, off-site replication will be suspended until the array is returned to service.

Detailed Design and Considerations—VPLEX: Component Failure, High Level

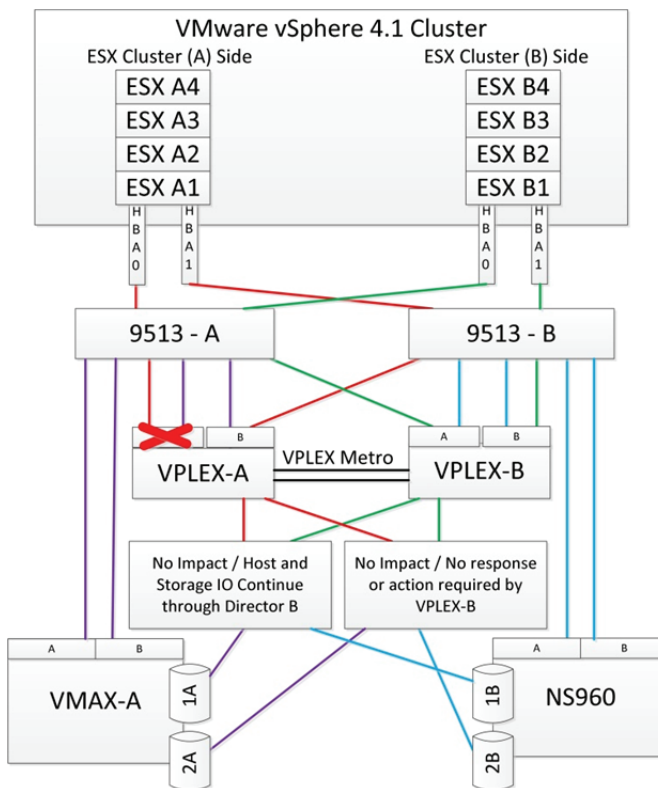


Figure 11: Failure Scenario, Loss of VPLEX Component

A component failure is a scenario in which a component on one of the VPLEX clusters fails. This could be a director, an FA or BE port, cache, or any other component. Because of the N+1 architecture inherent in the design of the VPLEX hardware, no director or engine component failure should affect availability.

VPLEX Infrastructure

- Each VPLEX cluster is a highly available implementation consisting of two engines and four total directors (128 GB cache, 32 FE/32 BE ports).
- Side-A hosts are connected to VPLEX-A. Side-A VMware Guests should be placed on VMFS distributed volumes that are VPLEX-A preferred.
- Side-B hosts are connected to VPLEX-B. Side-B VMware Guests should be placed on VMFS distributed volumes that are VPLEX-B preferred.
- All VPLEX LUNS are “encapsulated” array LUNS that are comprised of VPLEX-managed mirrored LUNS made up of members from both storage arrays.
- Each VPLEX cluster is fully configured and highly available to accommodate any component failure without impact.

VPLEX Failure/Impact—High Level

This assumes that guest sidedness is maintained.

- VPLEX intra-cluster link failure will have no impact to any hosts/guests.
- Total VPLEX cluster/cabinet failure will precipitate an outage for all hosts/guests connected to that VPLEX cluster, but hosts/guests on the other side will continue to run.
- By design, manual intervention is required to initiate HA or manual restart of failed guests on the other side cluster hosts.

Detailed Design and Considerations—VPLEX: Link Failure Impact/Recovery

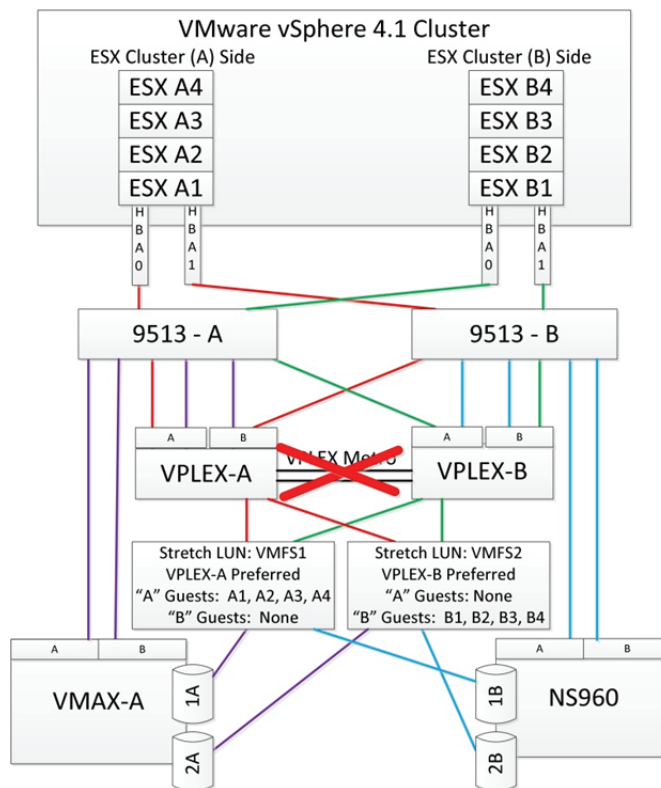


Figure 12: Failure Scenario, Loss Intra-Cluster Link

VPLEX Failure Mode 1: Link Failure

A link failure condition would take place if the physical dedicated link between the VPLEX clusters was severed or connectivity was lost in such fashion that both VPLEX Metro Clusters would still be running, but they could no longer communicate with each other.

VPLEX Impact

VPLEX will transition distributed virtual volumes on the non-preferred site to the I/O suspension state. On the preferred site, the distributed virtual volumes will continue to provide access.

This is why it is critical that guests of the A-Side hosts always run on VPLEX-A preferred distributed volumes. If A-Side guests run on B-Side preferred distributed volumes, they will become unavailable in this scenario.

VMware Impact

- **VMs running in preferred site:** No impact.
- **VMs running in non-preferred site:** These VMs will see all I/Os as suspended and the guest OS may hang during that time. If VM Monitoring is turned on, the VM would attempt to be reset on the same host, but the attempts will fail until the distributed virtual volumes are manually unsuspending. If the maximum reset limit, based on the VM Monitoring policy, is reached, the failed VMs will need to be restarted manually. The datastore will not be marked as unavailable, because the path is still active, but no I/Os will be processed.

Detailed Design and Considerations—VPLEX: Cluster Failure Impact/Recovery

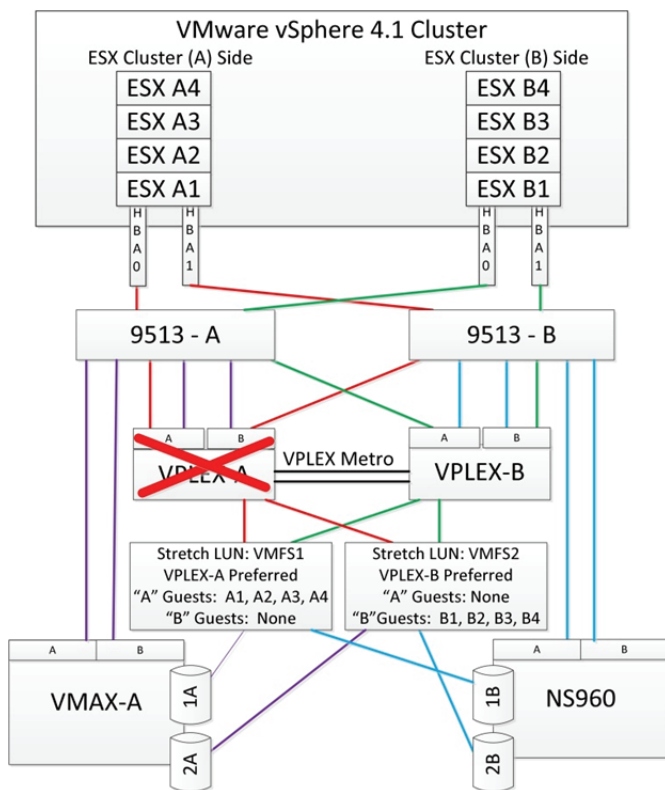


Figure 13: Failure Scenario, Loss of VPLEX Cluster

VPLEX Failure Mode 2: VPLEX Cluster Failure

A cluster failure event, by this definition, is an event that physically disables an entire cluster. This could most likely be caused by dual-feed power loss (EPO),

physical data center/rack damage, or some other type of catastrophic event.

VPLEX Impact

VPLEX will suspend I/O on the Metro-distributed virtual volume on the non-preferred site. Once it is determined by the administrator that the site has failed and it is not a case of inter-site communication failure, the volumes on the non-preferred site can be unsuspending (“resumed”) using the device resume-link-down command.

This is why it is critical that guests of the B-Side hosts not run on VPLEX-A preferred distributed volumes. If B-Side guests run on A-Side preferred distributed volumes, they will become unavailable in this scenario.

Note that this process is *intentionally manual*. While the automated resuming of I/O works in the site failure, it does not work in the VPLEX Cluster Partition case. Issuing the unsuspend command automatically on the non-preferred site would cause both sites to become simultaneously read-writeable, creating a potential split-brain condition.

VMware HA Impact—Preferred Site: VPLEX Cluster and ESX Host Loss

VMs running in preferred site: VMware HA will attempt (up to 5 times) to bring up the failed VMs on the ESX Servers supported by the non-preferred site for the Metro-distributed virtual volumes. These attempts will fail until the volumes are unsuspending on the non-preferred site. If the HA maximum restart limit is reached, the failed VMs have to be manually restarted on the non-preferred site. The datastore will not be marked as unavailable, because the path is still active, but no I/Os will be processed.

VMware HA Impact—Preferred Site: VPLEX Cluster Loss Only, ESX Hosts Still Up

Given a VPLEX cluster loss of the preferred site, running VMs will see a loss of storage connectivity. Different factors determine how VMs will respond to the loss of their storage paths. Depending on OS type, active storage I/O, and other factors, some guests may go down immediately, some may hang in a degraded state, and some may stay up for some time. If VM Monitoring is enabled, as each guest hangs due to the

loss of its storage, vCenter will attempt an HA restart on another eligible and surviving host that has an active storage path.

If VM monitoring is not enabled, or a faster restart time is required (rather than waiting for the guest to crash), recovery would require a cold migration through vCenter from the affected hosts to an unaffected host connected to the remaining cluster, after I/O access to stretched had been manually restored.

VMware HA Impact—Non-Preferred Site

VMs running in non-preferred site on the remaining cluster (accessing the failed clusters preferred LUNS) will see the I/O as being suspended and the guest OS may hang during this time. If VM monitoring is turned on, the guest OS would attempt to be reset, but the attempts will fail until volume is unsuspended. If the maximum reset limit, based on the VM Monitoring policy, is reached, the failed VMs will need to be restarted manually.

Conclusion

A single client meeting led to this treatise. In one design session, we threw around some ideas—and ultimately came to the design described here. The process of building this architecture to meet specific functional business requirements required a lot of learning and many questions asked and answered. As this treatise demonstrates, designing an active/active data center—in one data center or across sites—is not trivial. It should be reviewed with careful technical consideration in all aspects. With proper attention to detail, a high-availability data center as described in this treatise can be designed and implemented that meets both the business and technology requirements of an organization.

Resources

VMware Knowledge Base for VPLEX Implementation:
http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&docType=kc&externalId=1026692&licenseId=1&docTypeID=DT_KB_1_1&dialogID=124544475&stateId=0%200%20133613044

Chad Sakac's *Virtual Geek* Blog, "Understanding vSphere Disaster Recovery Avoidance Options," parts i, ii, and iii:

http://virtualgeek.typepad.com/virtual_geek/2010/11/understanding-vsphere-disaster-recoveryavoidance-options-part-i.html

http://virtualgeek.typepad.com/virtual_geek/2010/12/understanding-vsphere-disaster-recoveryavoidance-options-part-ii-understanding-emc-vplex-partition-and-vm-ha.html

http://virtualgeek.typepad.com/virtual_geek/2010/12/understanding-vsphere-disaster-recoveryavoidance-options-part-iii-areas-of-development.html

EMC's VPLEX Product Documentation:

<http://www.emc.com/collateral/hardware/technical-documentation/h7113-vplex-architecture-deployment.pdf>

For another Ahead perspective on the Active/Active data center, check out John Cole's Blog:

<http://www.thinkaheadit.com/2011/01/28/new-year%E2%80%99s-resolution-be-more-active-active/>

About the Author

Brett Foy is a data center technology consultant at Ahead, where he designs and architects advanced data center technologies around virtualization, computing, networking, storage, and backup/recovery. Prior to joining Ahead, Brett worked at a large IT manufacturer where he was responsible for the divisional Virtualization Data Center Practice as a practice manager. He also spent five years in a technology consulting and architecture sales support role. Prior to moving into the consulting and architecture role, Brett was an AVP for a multinational global insurance company, responsible for managing the North American data center infrastructure, focused on data center virtualization and consolidation around many technologies, including server, backup, desktop, email, security, web hosting, and SAP.

About Ahead

*Ahead is purpose built to enable companies to understand and take advantage of today's revolutionary technologies. These technologies drive the next-generation data center, which improves agility, availability, and responsiveness, and takes out cost. Our technologists, consultants, and sales team help clients evaluate the potential technical, operational, and financial benefits of these innovative solutions. Where the benefit is adequate, we work to properly architect, procure, and deploy these solutions in a timely and effective manner. Simply put, we help companies **THINK | LOOK | PLAN | MOVE AHEAD.***

©2011 Ahead, LLC. The information contained herein is subject to change without notice. Nothing herein should be construed as constituting a warranty or condition, express or implied, in fact or in law. All other marks and names mentioned herein may be trademarks of their respective companies.