

HIGHLY AVAILABLE MULTI-DATA CENTER WINDOWS SERVER SOLUTIONS USING EMC VPLEX METRO AND SANBOLIC MELIO 2010

Abstract

This white paper demonstrates key functionality demonstrated in a lab environment using EMC® VPLEX™ Metro, supporting geographic clusters and Sanbolic Melio 2010 software.

November 2010



Copyright © 2010 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is”. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Part Number h8114

Table of Contents

Executive summary	4
Key value propositions	4
VM mobility: Transparently share and balance resources between data centers	4
Federated file sharing over distance to provide instant data access to support remote users	5
Increase resiliency of applications against unplanned and planned outages.....	5
EMC VPLEX overview	6
VPLEX Local.....	6
VPLEX Metro	6
Application availability during data center outages when using VPLEX	8
Sanbolic Melio 2010 overview	8
Clustered file system using a distributed lock manager	8
Volume management	8
VSS-based snapshots and mirroring.....	9
Performance monitoring and quality-of-service management	9
SQL Server clustering	10
Technology integration	10
Lab configuration	10
Lab demonstration	11
Provisioning storage from the VPLEX	12
Storage configuration and management.....	12
Granular storage performance monitoring	13
Granular storage QoS management.....	13
Consolidation of virtual machine images on a single LUN.....	14
Support for Live Migration across VPLEX Metro.....	15
Conclusion	17
References	19

Executive summary

The EMC® VPLEX™ storage federation solution together with Sanbolic Melio software enables simultaneous read and write access to a block storage volume for servers running Microsoft Server 2008 R2 with Hyper-V technology at separate geographic locations. This architecture provides new options for the following:

- Disaster recovery (DR) and disaster avoidance
- Moving workloads across locations
- Supporting organizations in remote offices
- Enterprise-wide virtual desktop solutions
- Deploying public or private cloud computing environments that span multiple data centers

Unlike typical active-passive clustering, active-active clusters provided by the VPLEX-Melio solution enable workloads to be running simultaneously on both sides of the cluster. This allows better utilization as well as shared access for collaboration, and load balancing of workloads across locations, while providing DR capability in the event that one facility becomes unavailable.

Key value propositions

VM mobility: Transparently share and balance resources between data centers

Large clusters of Hyper-V hosts and Live Migration of Hyper-V VMs are fully supported, enabling dynamic migration of virtual server or virtual desktop workloads between office locations and/or cloud data centers.

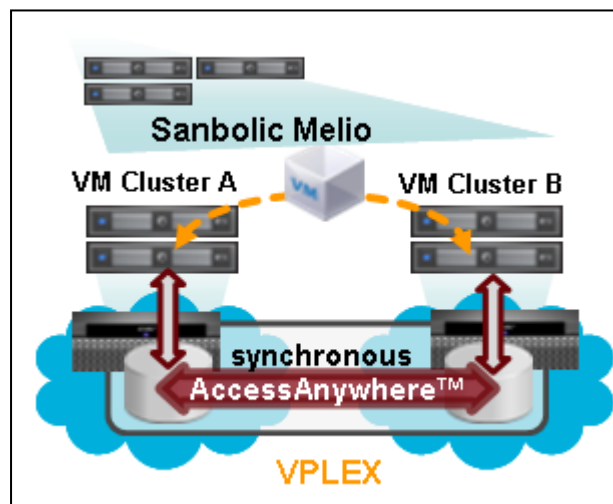


Figure 1. Supports Live Migration of VMs across VPLEX nodes

Federated file sharing over distance to provide instant data access to support remote users

Virtual and/or physical servers can concurrently read and write to a common block storage volume across locations, supporting remote offices and collaboration by making shared data available locally on each site.

Shared data scale-out clusters of physical or virtual servers provide capability to dynamically increase capacity for key workloads.

The VPLEX solution provides immediate benefits such as increased resiliency for unplanned outages and centralized storage management with the ability to dynamically scale out and load balance your applications, servers, and other business needs between local and geographically dispersed sites.




Figure 2. Supports read-write access to a file share across both nodes

Increase resiliency of applications against unplanned and planned outages

Active-active access to files on a block storage volume across geographic locations provides DR functionality without the need for “passive” standby servers that are only utilized in the event of a failure of the primary site.

Dynamic provisioning of Windows server workloads and storage SLA management can be managed on a per-server or per-application basis across locations.

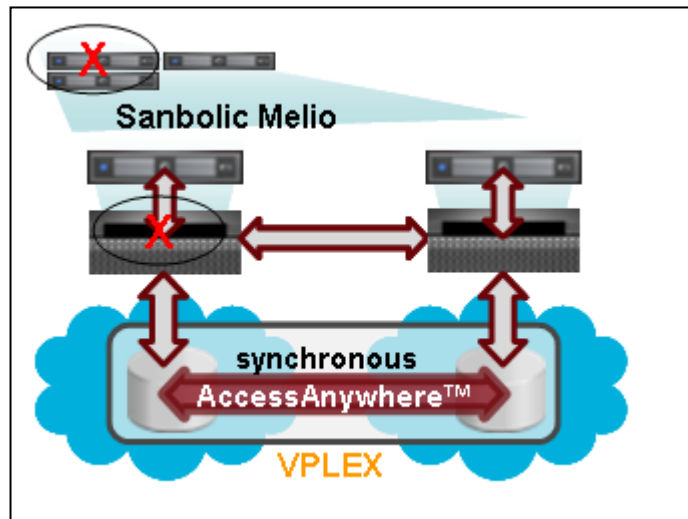


Figure 3. Applications remain available when one node fails

For this white paper we have demonstrated key functionality in a lab environment using VPLEX Metro, which supports geographic clusters up to 5 ms apart¹. VPLEX Metro is generally available, and Sanbolic Melio 2010 software has been fully qualified on the platform.

EMC VPLEX overview

EMC VPLEX is an enterprise-class storage area network-based federation solution that aggregates and manages pools of Fibre Channel-attached storage arrays that can be either collocated in a single data center or multiple data centers that are geographically separated by metropolitan area network (MAN) distances. EMC VPLEX Metro provides non-disruptive, heterogeneous data movement and volume management functionality within synchronous distances both within and between data centers. With a unique scale-up and scale-out architecture, advanced data caching, and distributed cache coherency, VPLEX provides workload resiliency, automatic sharing, and balancing and failover of storage domains, and enables both local and remote data access with predictable service levels.

EMC AccessAnywhere™, available with VPLEX, is a breakthrough technology from EMC that enables a single copy of data to be shared, accessed, and relocated over distance. EMC GeoSynchrony™ is the VPLEX operating system.

VPLEX Local

VPLEX Local provides simplified management and non-disruptive data mobility across heterogeneous arrays.

VPLEX Metro

VPLEX Metro provides data access and mobility between two VPLEX Clusters within synchronous distances. VPLEX Metro supports two forms of distributed devices:

¹ This is EMC's current Metro distance limitation. Applications may perform at higher latencies but would require EMC RPQ.

metro-distributed virtual volumes and remote virtual volumes. Metro-distributed virtual volumes provide synchronized copies (mirrors) of the volume's data in each cluster. The mirrored volume appears and behaves as a single volume and acts in a similar manner to a virtual volume that uses a RAID 1 device. Remote virtual volumes provide access to a virtual volume whose data resides in one cluster. Remote virtual volumes, like metro-distributed volumes, are able to take advantage of VPLEX distributed coherent cache and its prefetch algorithms to provide better performance than a SAN extension solution.

VPLEX uses a unique clustering architecture to help customers break the boundaries of the data center and allow servers at multiple data centers to have concurrent read and write access to shared block storage devices. A VPLEX Cluster, shown in Figure 4, can scale up through the addition of more engines, and scale out by connecting multiple clusters to form a VPLEX Metro configuration. In the initial release, a VPLEX Metro system supports up to two clusters, which can be in the same data center or at two different sites within synchronous distances (within 5 milliseconds (ms) – approximately up to 60 miles or 100 kilometers apart). VPLEX Metro configurations help users to transparently move and share workloads, consolidate data centers, and optimize resource utilization across data centers. In addition, VPLEX Clusters provide non-disruptive data mobility, heterogeneous storage management, and improved application availability.

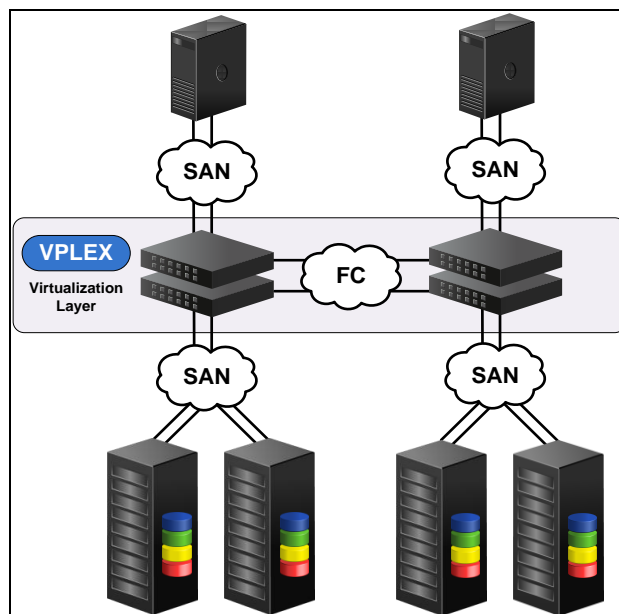


Figure 4. VPLEX Cluster

Sanbolic's file system has been designed to support cluster latency of up to 3,000 ms, and thus will support active-active VPLEX clusters in the currently supported Metro offering.

Application availability during data center outages when using VPLEX

A VPLEX Metro distributed between two data centers can be used to protect against data unavailability in the presence of a data center outage by mirroring data between the two data centers. Data access will remain available for metro-distributed virtual volumes whose preferred cluster is the surviving data center. For those volumes whose preferred cluster is the data center with the outage, access to the data for these volumes can be restored in the surviving cluster with the invocation of a manual command to resume the suspended I/O. When combined with failover logic for host clusters or active-active environments using Melio 2010, this provides infrastructure that is able to restore service operations quickly even in the presence of an unplanned data center outage.

Sanbolic Melio 2010 overview

The combination of Sanbolic's Melio FS[®] and LaScala[®] software provides the ability to centrally configure and assign logical volumes to the VPLEX LUN across geographic locations, and to have physical or virtual servers in all locations concurrently read and write from these volumes. Without a clustered file system VPLEX can support active-passive failover across locations — if a server fails in one location, the application can be restarted in the second location. Using Melio clustered file system servers in both locations allows you to run a workload off the same dataset simultaneously — supporting scale-out and collaboration, as well as providing inherent DR capability. Melio can also support large clusters of Hyper-V hosts across locations, enabling Live Migration of virtual machines across locations.

Clustered file system using a distributed lock manager

Melio FS is a 64-bit clustered file system that provides multiple servers with simultaneous access to data stored on the same storage partition or LUN on a SAN using a multi-layer distributed lock manager. As a general-purpose cluster file system, Melio FS supports access to all data types stored on the LUN and can support LUNs as large as 18 million terabytes. Melio employs an advanced transaction manager, a symmetrical architecture, full journaling, and dynamic clustering capabilities. It ensures the integrity of the data stored on the shared LUN by controlling read/write access by using multiple servers at any given time.

In addition to providing Hyper-V clusters that enable virtual machine migration across geographic locations, and supporting active-active clusters of physical Windows servers, the joint VPLEX and Sanbolic solution will support active-active clusters of virtual machines or of mixed physical and virtual machine clusters across geographic locations.

Volume management

LaScala Cluster Volume Manager is a host-based volume manager that incorporates advanced transaction management, locking, and clustering technology based on the Melio FS technology. Using LaScala, multiple hosts can share access to and

administer storage volumes spanning multiple storage controllers. Access to volumes for each host can be defined using native Windows ACLs and Active Directory. Used in conjunction with Sanbolic's Melio FS, LaScala provides high-performance shared access to data on volumes striped across many disks residing on multiple storage controllers.

Sanbolic's software is designed for a shared storage environment. Therefore, all servers in the shared environment can observe the entire volume structure and automatically see any changes made to it. All operations are fully transactional so multiple changes can be made to the disk structure and then committed simultaneously. Logical volumes can be remotely configured and assigned to physical or virtual servers using the standard Windows security interface. This can help simplify deployments across locations.

Sanbolic's Data Center Suite uses only basic disks for all operations. Sanbolic's basic disks support stripe sets, volume sets, and mirror sets using Fibre Channel. These disks can be dynamically expanded. Therefore, a Hyper-V deployment using VPLEX and Sanbolic can expand the storage of an active environment either by expanding the LUN on the storage array, or by adding another LUN to the logical volume on the host. These changes can be made while an active I/O is present.

VSS-based snapshots and mirroring

Melio provides a flexible distributed snapshot capability that uses a VSS interface to provide a more granular and flexible data protection architecture and compatibility with existing VSS-based backup products. Melio is able to create a snapshot of the volume from any host in the cluster. Since the snapshot is coordinated with the distributed lock manager, all servers accessing the volume are quiesced and cache data flushed to disk before the snapshot is taken, ensuring data coherency. This snapshot capability can be accessed by any industry-standard VSS backup tool or through the Melio interface. The snapshot can also be mounted as a file share on any of the hosts in the cluster and automatically copied to another location. LaScala also provides asynchronous mirroring capability.

These tools provide additional flexibility in protecting data, and for moving data onto a VPLEX volume or across VPLEX Clusters.

Performance monitoring and quality-of-service management

Performance monitoring and quality-of-service (QoS) management can be key tasks in a large environment where performance bottlenecks can occur at the storage layer. Storage layer bottlenecks are often difficult to identify. Melio natively supports full perfmon and sysmon statistics that can be used to identify bottlenecks or exported to other reporting tools. If storage-related bottlenecks are discovered such as with a runaway VM, Melio QoS management can be configured to prevent the problem server or process from limiting storage bandwidth performance available to critical servers. QoS can be assigned on a per-server or per-process basis.

SQL Server clustering

Sanbolic's Melio 2010 software now includes an application clustering service. This gives users the ability to consolidate all their SQL databases onto a single data volume. It also provides the ability to cluster the SQL instances utilizing this shared volume in an active-active mode. Melio 2010's AppCluster service provides ease of management, high availability, and very fast failover or load balancing of SQL databases across instances running on multiple physical or virtual SQL servers. AppCluster removes the complexity and the need for expensive underutilized hardware associated with typical clustering tools. Shared access to storage across n+1 nodes in a cluster running Melio ensures application recovery within seconds on any other node. In a VPLEX environment physical and virtual SQL servers can share a volume containing the databases across locations, providing the ability to load balance or fail over SQL databases across SQL instances located in more than one data center.

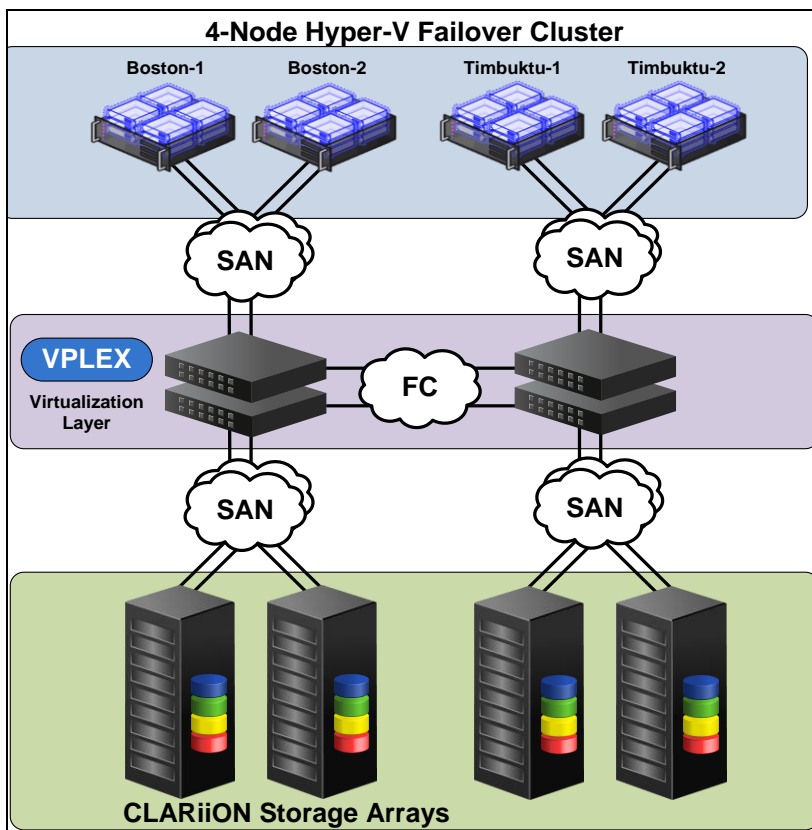
Technology integration

EMC's enterprise-level storage capabilities have been proven over time to be extremely effective at supporting large Microsoft virtualized infrastructures. Sanbolic's Melio Data Center Suite is designed to fully utilize the capabilities of the EMC CLARiiON® or Symmetrix® storage and VPLEX, while integrating easily with Microsoft Hyper-V and System Center Virtual Machine Manager, as well as with Windows application servers running on physical or virtual servers. Melio will support file and volume size environments of up to 18 exabytes, and efficiently passes through the performance of CLARiiON or Symmetrix storage, while supporting key software applications such as EMC PowerPath®.

Sanbolic has worked closely with Microsoft to ensure that Melio works seamlessly with Hyper-V. Both Windows Server 2008 R2 Hyper-V and SCVMM R2 recognize Sanbolic storage volumes as active-active storage resources. This simplifies the configuration of Hyper-V host clusters. Live Migration and Quick Migration are both fully supported. After initially creating the VMs using Melio, Hyper-V host clusters can be configured and managed from within SCVMM R2.

Lab configuration

The demonstration configuration includes two VPLEX locations, named Boston and Timbuktu, respectively, each with CLARiiON CX4™ storage. VPLEX presents a common LUN in both locations and each is managed using Sanbolic's Cluster Volume Manager and formatted with the Melio file system, both of which are components of Melio 2010 Data Center Suite. This LUN is used to support four clustered Windows 2008 R2 servers. Boston1 and Boston2 are attached to one side of the VPLEX. Timbuktu1 and Timbuktu2 are attached to the other side of the VPLEX. All servers have concurrent read-write access to the LUN. Twenty Hyper-V virtual machines were configured on the cluster. A file share was also configured on each physical server. All file shares were clustered into a common volume on the shared LUN.



Lab configuration

Microsoft environment

- Windows Server 2008 R2 Hyper-V
- Windows Server Failover Clustering
- Systems Center Virtual Machine Manager

Sanbolic

- Melio 2010 Data Center software

EMC

- CLARiiON CX4 storage
- VPLEX Metro

Figure 5. Diagram of lab configuration

An ANUE distance simulator is installed between the two VPLEX locations and hosts. Each virtual machine is a Windows 2008 R2 Server Core instance, and each has a boot disk and a second disk used as a file share. All VHDs for the VMs reside on the X: drive, which is formatted with Sanbolic’s Melio shared file system. This allows any node in the cluster to access the X: drive simultaneously.

The control server is connected to the file share disk of each VM. IOmeter was used to drive I/O to the second disk of each VM to generate I/O traffic on the storage. A single I/O profile used in IOmeter testing was a 64k block size, 100% sequential, and 100% read.

Lab demonstration

The lab configuration was used to demonstrate the following functionality:

- Creation of the Sanbolic Melio Clustered File System over VPLEX-based federated storage in a Microsoft Windows environment
- Accessibility to both virtual and physical hosts , with concurrent read and write to a common SAN volume across locations
- Dynamic provisioning of Windows Server workloads and Server SLA management on a per-server or per-application basis across locations
- Demonstrate LaScala performance monitoring and QoS setting

- Demonstrate active-active access to files on a SAN volume across metro distances
- Demonstrate DR functionality without need for “passive” standby servers
- Demonstrate Microsoft Live Migration support for clusters of hypervisor hosts, enabling dynamic migration of virtual Microsoft workloads between metro-distant locations

Provisioning storage from the VPLEX

To begin using a VPLEX Cluster, you must provision and export storage so that hosts and applications can use the storage. Provisioning and exporting storage refers to the tasks required to take a storage volume from a storage array and make it visible to a host. This process consists of the following tasks:

- Discovering available storage
- Claiming and naming storage volumes
- Creating extents from storage volumes
- Creating devices from the extents
- Creating virtual volumes on the devices
- Registering initiators
- Creating a storage view
- Adding initiators, ports, and volumes to storage view

Storage configuration and management

The LaScala volume manager was used to bring the shared LUN under management and to format it with the Melio file system. LaScala is designed for a clustered environment, so the Melio disk was automatically visible in the volume manager interface of the two Boston servers and the two Timbuktu servers. Each server automatically has shared access to it, from both sides of the VPLEX.

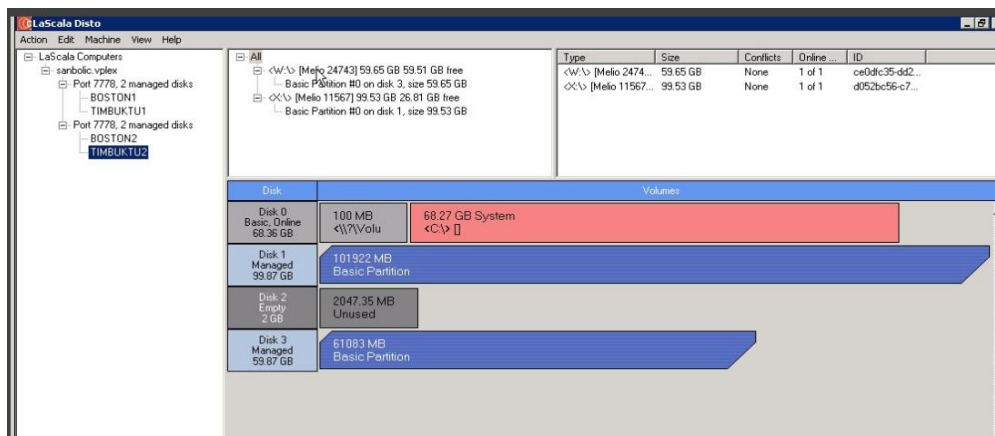


Figure 6. LaScala volume manager interface

Granular storage performance monitoring

Melio provides extensive Perfmon performance counters, which are available on each of the servers. Several examples of the counters were demonstrated. The reporting capability provides the capability to monitor in detail the performance on the shared storage resources on the VPLEX.

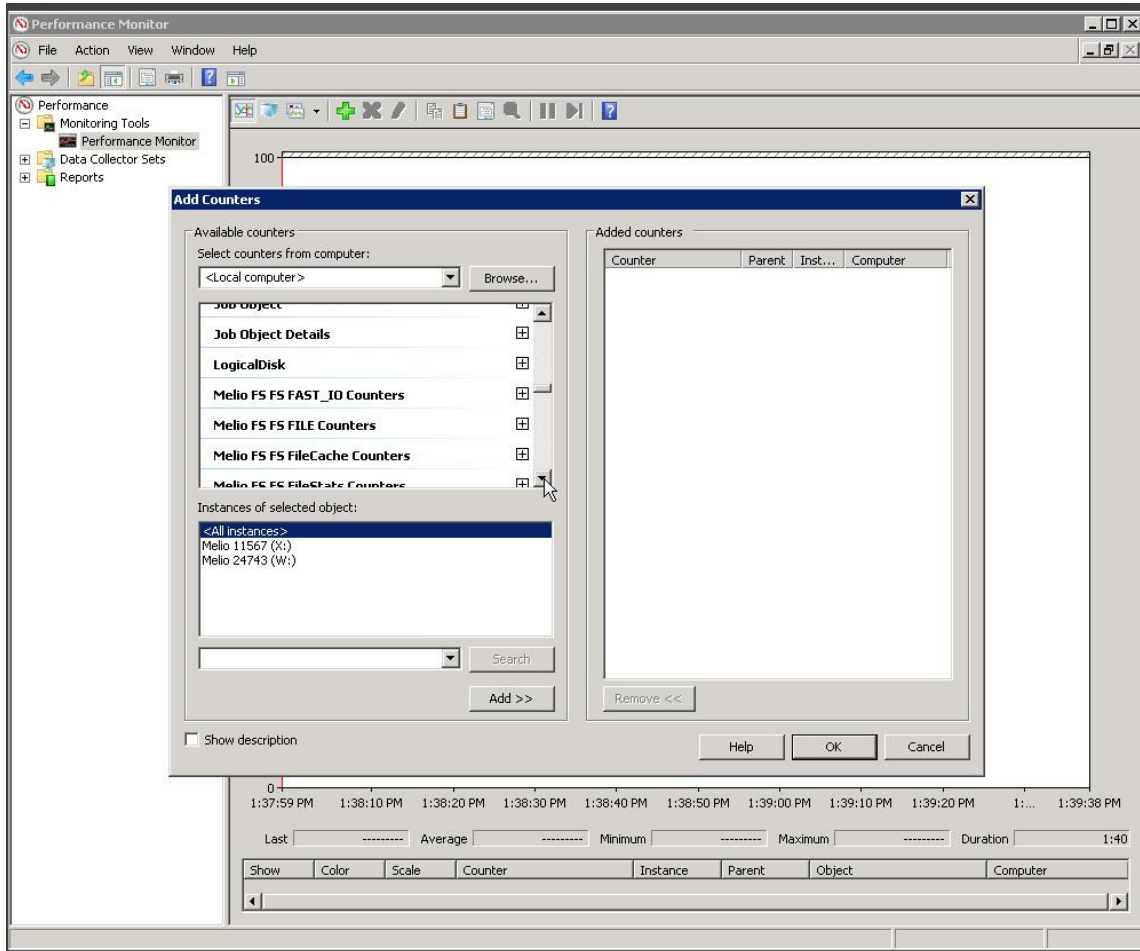


Figure 7. Examples of performance available within Melio

Granular storage QoS management

Melio was then used to assign storage QoS both for specific applications and for specific files and directories. The QoS capability allows storage resources to be prioritized in a shared storage environment across both sides of the VPLEX, either to guarantee access to critical processes or virtual machines, or to limit the bandwidth available to workloads (such as a virus scan) so as not to affect the SLAs of other workloads running on the cluster.

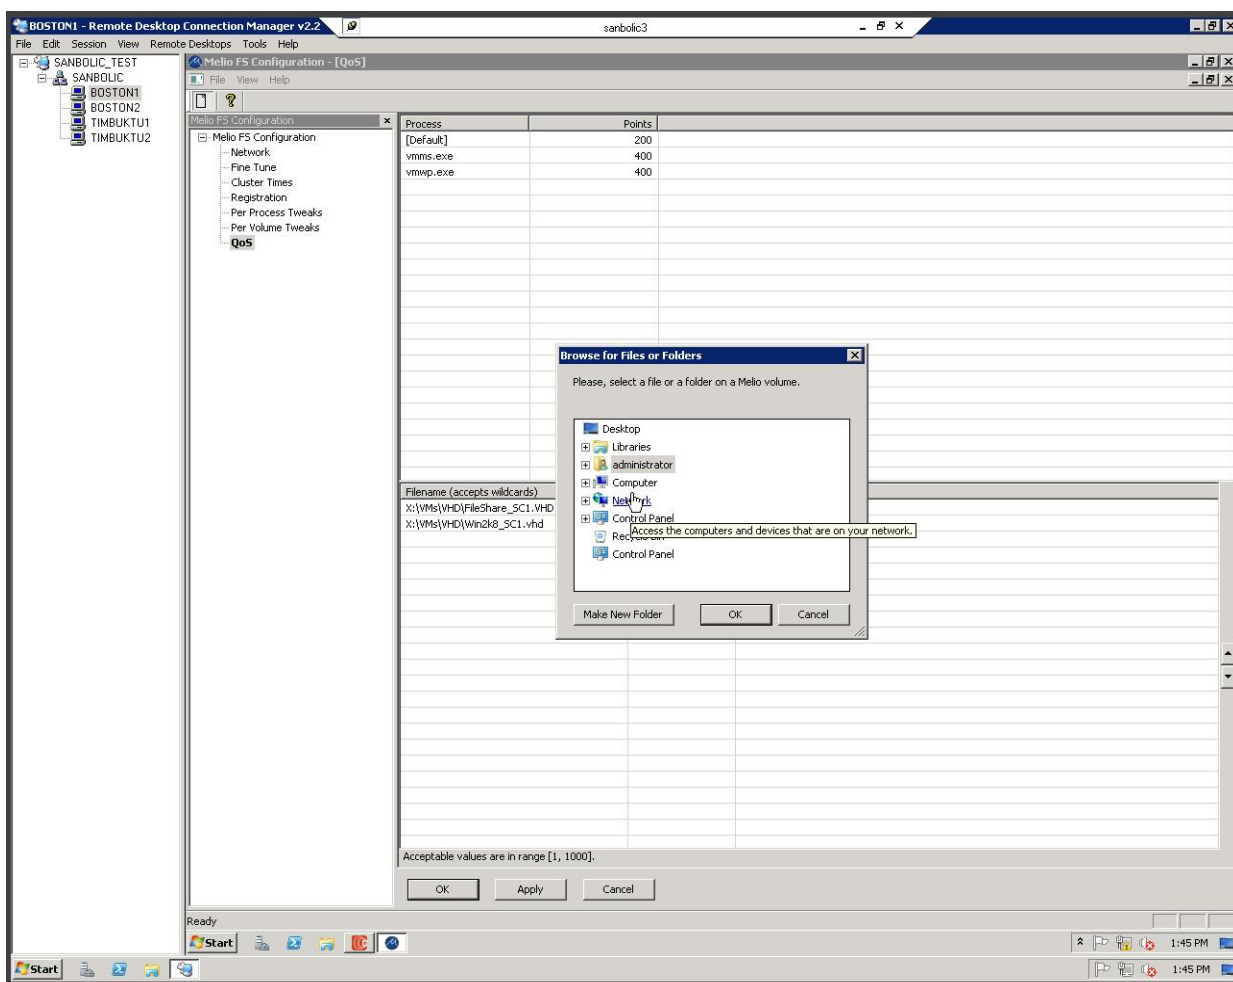


Figure 8. Melio interface for assigning storage quality of service

Consolidation of virtual machine images on a single LUN

Since Melio is recognized by Hyper-V and the Microsoft System Center virtual machine manager, ease of configuration of Hyper-V environments is also demonstrated. All of the virtual machines are stored on the shared Melio volume and accessible from all servers in the cluster across both VPLEX locations.

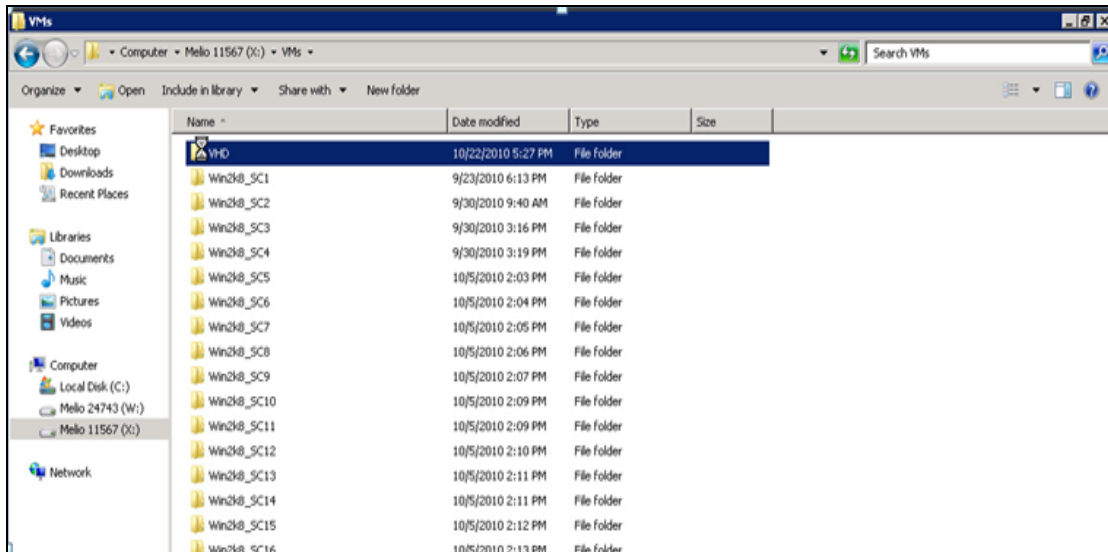


Figure 9. Virtual machine files stored on shared Melio volume

Support for Live Migration across VPLEX Metro

Live Migration across VPLEX locations was demonstrated. The ANUE was used to simulate Fibre Channel latency, and IOmeter was used to measure I/O performance.

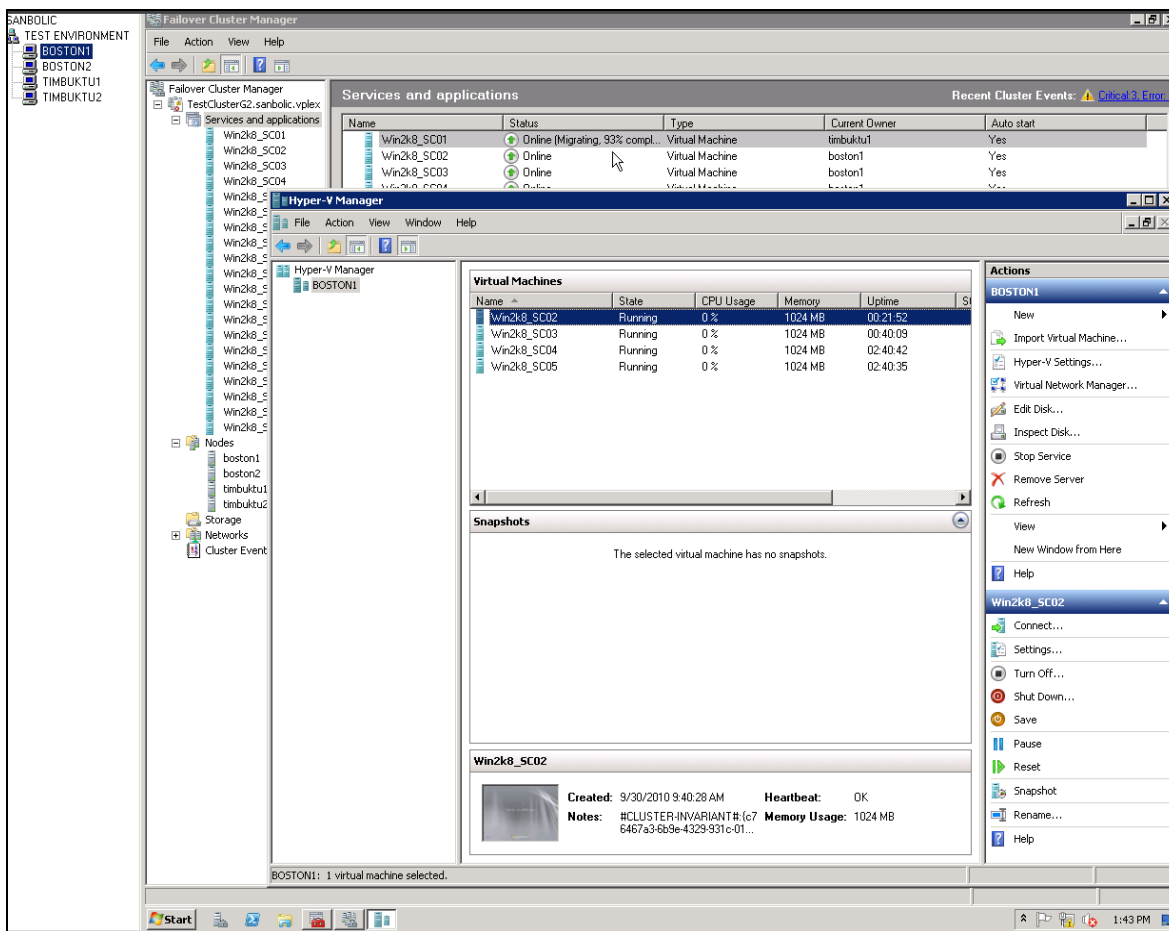


Figure 10. Migration of virtual machines across the VPLEX nodes

Finally, a file share was created on each of the four physical servers and clustered into a common SAN volume. A Microsoft DFS namespace was used on each side of the VPLEX to create a single address for the Boston cluster, and a second for the Timbuktu cluster. Either server in each cluster could respond to a file request, accessing a single copy of files on the Melio volume on VPLEX. The lab work demonstrated all four servers in the geographic cluster simultaneously writing to a shared VPLEX-Melio volume. This enables a highly available CIFS serving environment, with the VPLEX eliminating latency in file response times, while supporting collaboration across the locations.

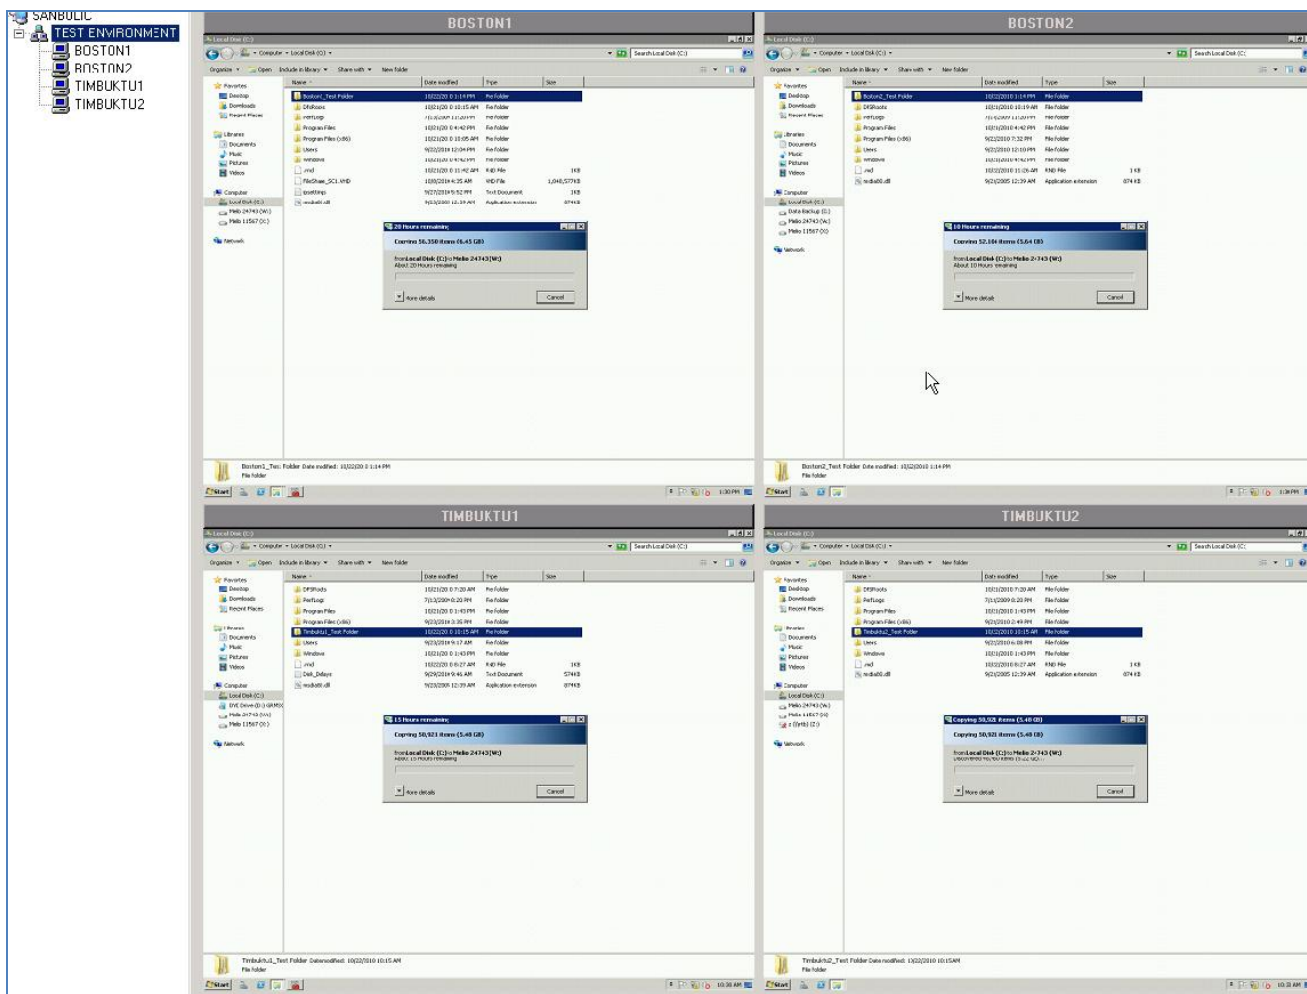


Figure 11. All servers in the VPLEX cluster simultaneously write to the Melio volume

Conclusion

VPLEX storage federation technology provides new capability to distribute workloads across data centers, improve collaboration, and efficiently provide DR protection using active-active rather than active-passive secondary data center architecture. A clustered file system enables customers to fully utilize this capability and deploy active-active workloads, providing concurrent shared read-write access to shared data in two or more geographic locations

Sanbolic's Melio 2010 software integrates easily with VPLEX, EMC storage, and Windows Server technologies, and provides a fully featured clustered file system together with central logical volume management, granular QoS management, and extensive performance reporting. Melio also provides application clustering for SQL Server and IIS.

The joint solution provides a unified environment across physical and virtual servers and across geographic locations. The VPLEX-Melio solution:

- Supports any Microsoft workload (with appropriate extensions/configurations)
- Provides a CFS that spans both physical and virtual environments
- Provides a CFS that will leverage a federated volume across distance
- Provides accessibility to users, again over distance, who may use the shared file system for a variety of needs concurrently (reducing TCO)
 - Ensures continual updating of volume data and metadata on multiple discrete servers, thereby providing consistency
 - Provides a highly scalable Windows file share infrastructure
 - Provides HA and resiliency at the host level
- Provides file share and application mobility to meet or exceed SLAs during planned or unplanned failures (Live Migration)
- Provides a high degree of HA at the storage level via VPLEX hardware redundancy and Metro-based extensibility. Active-active access to files on a SAN volume across geographic locations provides DR functionality without the need for “passive” standby servers that are only utilized in the event of a failure of the primary site.
- Provides file share administrators a robust, easy-to-use management interface for volume configuration, QoS management, and host- and array-based snapshots.
- Supports Microsoft Hyper-V workload mobility over distance. This enables dynamic migration of virtual server or virtual desktop workloads between office locations and/or cloud data centers.

The demonstration used EMC CLARiiON storage, but VPLEX can also be used with other Fibre Channel storage that customers may already have in their data center. These capabilities provide compelling use cases for enterprise, hosting, and cloud deployments:

- A corporate environment can use these capabilities to support remote offices and multi-site VDI locations, balance workloads across locations, move workloads across time zones, improve collaboration by having key data locally available, and provide DR capability without the cost of passive standby servers.
- Cloud hosting environments now have the option not only to dynamically provision capacity for customers at a single data center, but also across multiple locations, while providing and enforcing storage SLAs on a shared virtualized platform. Workloads like websites can be balanced across locations and scale out across multiple locations. Customers can be guaranteed DR capability without the capital cost of passive standby data center capacity.

VPLEX Metro and Sanbolic Melio 2010 are available today.

References

For further information, read the EMC-Microsoft-Sanbolic Proof of Concept:

- [Optimized Storage Solution for Enterprise Scale Hyper-V Deployments](#)

The Sanbolic website provides information on the following:

- Storage Management for Enterprise Hyper-V:
<http://www.sanbolic.com/cloudstorage.htm>
- Data Availability and Application Scalability:
<http://www.sanbolic.com/virtualization.htm>
- Melio Data Center Product Suite:
http://www.sanbolic.com/Melio_data.htm

You can contact Sanbolic at 617-833-4242 or sales@sanbolic.com.

Register for evaluation software at <http://clients.sanbolic.com/?action=client-registerform>.

Visit the [VPLEX page](#) on EMC.com for more information.