



# **Enterprise Storage Solution for Hyper-V Private Cloud and VDI Deployments using Sanbolic's Melio Cloud Software Suite**

**April 2011**



# Executive Summary

Large enterprise Hyper-V™ deployments with a large number of virtual machines (VMs) can be demanding on the storage infrastructure. In these environments, there is a lot of concurrent input/output (I/O) that contends for network and disk resources. This is especially true for example, in virtual desktop infrastructure (VDI) deployments, when many machines might boot concurrently at the beginning of the work day. That activity, commonly referred to as a “boot storm,” requires a specialized storage platform with the scalable performance and capacity to meet the needs of the company. In addition, ease of deploying images, identifying and eliminating I/O bottlenecks, effectively protecting corporate data, managing quality of service, and providing sufficient storage performance to limit user latency perceptions are all important administrative considerations for this type of deployment.

This white paper explains how the new capabilities of Microsoft® Windows Server® 2008 R2 and Hyper-V R2 combined with the flexible, scalable shared storage provided by the Sanbolic® Melio Cloud™ software suite offers organizations the ideal storage solution for achieving the greatest return on enterprise-class Hyper-V deployments.

The key benefits that an organization using this solution can experience include:

- **Active-active access to data on SAN volume:** All Windows workloads, including Hyper-V and Windows server applications running on virtual machines, have concurrent read-write access to files on shared SAN volumes.
- **Scalable performance:** Capable of supporting large number of VMs via:
  - **Striping across storage controllers:** Melio can aggregate performance across multiple storage controllers.
  - **Fully-symmetrical architecture:** Eliminates storage bottlenecks associated with a virtualized server environment. All servers are equal members of the cluster.
  - **No lock overhead:** Full block storage performance is maintained while scaling to large numbers of VMs (i.e., Melio has no lock overhead when used with exclusive lock objects such as Hyper-V .vhd files).
  - **Enterprise-scale host support:** Melio supports more than 200 hosts per

cluster, allowing multiple Hyper-V clusters to be deployed on a single storage volume.

- **Disaster recovery:** Synchronous and asynchronous mirroring provides protection from storage hardware failure or site failure.
- **Provide flexible data protection including:**
  - **Flexible snapshot granularity:** Any host can take a snapshot of the entire environment or a specific VM.
  - **Flexible backup support:** Snapshots can be handed off to standard VSS-based backup software or copied to second tier storage.
- **Centrally-managed unified storage platform:** Management is simplified with:
  - **Storage Live Migration:** Logical volumes can be moved between storage arrays while applications remain active.
  - **“Software SAN” capability:** Internal disk drives of Hyper-V hosts can be aggregated and made available as a highly available shared block storage resource for proof of concepts or smaller deployments.
  - **Microsoft toolset support:** Administrators can centrally assign storage to VMs using Active Directory® and Windows native security tools.
  - **Dynamic volume expansion:** Storage volumes can be dynamically expanded.
  - **Granular performance management:** Performance can be managed on a per server, process, virtual machine, file, or directory basis.
- **Extending Shared Storage Capability to Virtual Machines:** Active-active access to data on SAN volumes and all other storage features are available to virtual machines for application data—enabling application scale-out, improved HA options, and industry-leading storage management features.

This paper is intended for those who have a working knowledge of Microsoft Windows Server 2008 R2 and virtualized Windows Server environments. This paper also assumes that the reader has a basic understanding of SAN environments, storage virtualization, and clustering terminology.

# Introduction

The benefits of virtualization are well known in the IT community. By breaking the link between software and hardware, companies are creating a foundation for a more dynamic, flexible, and efficient data center. However, there are a few important storage related considerations that need to be addressed before transitioning to a virtualized data center.

First, the selection of a storage platform that organizations deploy to support virtualized computing environment is very important and should be carefully considered. The right storage system platform should ensure low latency for maximum performance, provide scalable capacity and bandwidth as application demands change, and automatically balance application workloads as needed. The storage platform should also provide organizations with the tools necessary to simplify their virtualized storage environments. In addition, the right platform should provide the pooling of disk resources to optimize and protect the application environment as well as to eliminate storage silos.

Second, the scale of the deployment must be matched to the storage platform performance capabilities. Large enterprise deployments of Hyper-V with a large number of VMs can be demanding on the storage infrastructure. These environments can generate a significant amount of input/output (I/O) requests. The selected storage platform needs to include capabilities that can satisfy all conditions and scale seamlessly to meet the performance needs of the company.

This white paper provides a solution that meets each of the above requirements. Microsoft Windows Server 2008 R2 Failover Clustering and Hyper-V, coupled with an appropriate SAN storage array and Sanbolic's Melio Cloud Storage software Suite, can be used to build a highly scalable and highly available storage solution that is perfect for supporting Hyper-V private cloud deployments or VDI deployments comprised of a large number of VMs.

# Solution Components

## Microsoft Virtualization Technologies

Microsoft virtualization provides an architecture that can profoundly affect nearly every aspect of the IT infrastructure management lifecycle. It can drive greater efficiencies, flexibility, and cost effectiveness throughout the organization. A standard Microsoft virtualization implementation is typically structured using the Windows Server 2008 R2 Hyper-V role to enable virtualization and Windows Server Failover Cluster to handle high availability and disaster recovery requirements. System Center Virtual Machine Manager (SCVMM) is typically used to simplify virtualization.

### **Windows Server 2008 R2 Hyper-V SP1**

Hyper-V is the hypervisor-based virtualization technology from Microsoft that is integrated into all Windows Server 2008 R2 x64 Edition operating systems. As a virtualization solution, Hyper-V enables users to take maximum advantage of the server hardware by providing the capability to run multiple operating systems (on VMs) on a single physical server. The availability of Hyper-V as a role in a mainstream Windows operating system provides several key advantages (see chart on following page):

Features	Benefits
Built-in Windows Server 2008 R2 technology	Enables enterprises to easily leverage the benefits of virtualization without adopting a new technology.
Broad device driver support	The new 64-bit micro-kernelized hypervisor architecture leverages the broad device driver support in the Windows Server 2008 R2 parent partition to extend support to a broad array of servers, storage types, and devices.
Symmetric Multiprocessor (SMP) support	Hyper-V supports SMP on VMs.
Host high availability	Windows Server 2008 R2 clustering provides high availability to VMs to minimize unplanned downtime.
Dynamic Memory Allocation	Dynamic memory balancing supports increasing virtual machine density, as well as adjusting virtual machine memory on-the-fly based on actual workload requirements
Shared storage high availability	Microsoft MPIO dynamically routes I/O to the best path and protects against connection failures at any point between a Hyper-V host and shared storage including NICs/adapters, switches or array ports.
Easy VM migration	Live migration capability is available to support business continuity during planned and unplanned downtime and over distances.
Volume Shadow Copy support (VSS)	Provides a robust host-based backup of VMs by leveraging the existing Windows VSS-based infrastructure.
Easy extensibility	Easy extensibility is available by using the standards-based Windows Management Instrumentation (WMI) interfaces and APIs.

## Windows Server Failover Clustering

Sanbolic, Inc.  
[www.sanbolic.com](http://www.sanbolic.com)  
[www.twitter.com/sanbolic](http://www.twitter.com/sanbolic)

sales@sanbolic.com  
 Phone: 617 833 4242  
 304 Pleasant St., Watertown, MA 02472

---

Failover clustering in Windows Server 2008 R2 helps to ensure that mission-critical applications and services such as e-mail and line-of-business applications are available when needed.

### **System Center Virtual Machine Manager**

Microsoft SCVMM 2008 R2 is enterprise-class management software that enables administrators to easily and effectively manage both the physical and virtual environments from a single management console and thus avoiding the complexity of using multiple consoles typically associated with managing an IT infrastructure. The key capabilities of SCVMM 2008 R2 include:

<b>Features/Capabilities</b>	<b>Benefits</b>
Enterprise-class management suite	Manages both Hyper-V and VMware ESX virtualization environments.
Intelligent VM placement	Provides support for the intelligent placement of VMs.
System Center Operations Manager 2007 integration	Integrates with System Center Operations Manager 2007 to provide proactive management of both virtual and physical environments through a single console by leveraging PRO.
Native P2V/V2V migration	Provides native capability for physical-to-virtual migrations and virtual-to-virtual migrations.
Failover integration	Provides integration with failover clustering to support high availability and the live migration of VMs.
Automation	Provides easy automation capabilities leveraging Windows PowerShell.

## Sanbolic Melio Cloud Storage Software Suite

Sanbolic, Inc.  
[www.sanbolic.com](http://www.sanbolic.com)  
[www.twitter.com/sanbolic](http://www.twitter.com/sanbolic)

sales@sanbolic.com  
Phone: 617 833 4242  
304 Pleasant St., Watertown, MA 02472

Sanbolic's Melio Cloud software suite provides an advanced, host-based storage fabric solution for enterprise scale Hyper-V deployments using high performance SAN storage. It also provides a shared application data LUN for Hyper-V VMs.

**Note:** *When configuring Hyper-V clusters within SCVMM 2008 R2, the shared storage resource is selected by highlighting the "Clustered Shared Volumes" storage option. It is important to note, however, that Melio FS is a standalone file system and a separate product from Microsoft's Clustered Shared Volumes.*

Melio Cloud storage software suite is built around a 64 bit cluster file system that provides multiple servers with simultaneous access to data stored on the same storage partition or LUN on a SAN. 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. Employing advanced transaction managers, symmetrical architecture, full journaling, and dynamic clustering capabilities, Melio FS ensures the integrity of the data stored on the shared LUN by controlling read/write access by multiple servers at any given time. It also allows granular quality of service management as well as cluster wide snapshots.

LaScala Cluster Volume Manager is a host-based volume manager that is a component of the software suite and incorporates advanced transaction management, locking, and clustering technology based on the Melio file system 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 Window ACLs and Active Directory. LaScala also provides synchronous or asynchronous mirroring that provides protection from storage array or site failure. LaScala supports "storage Live Migration"—the ability to move an logical volume from one storage array to another, while the application using it remains active. Used in conjunction with Sanbolic's Melio clustered file system, LaScala provides high performance shared access to data on volumes striped across many disks residing on multiple storage controllers. The key benefits of using Melio FS and LaScala as a shared resource for Windows Server 2008 R2 Hyper-V SP1 and SCVMM 2008 R2 deployments include (see chart on following page):

Features	Benefits
Live Migration support	Provides a shared LUN solution for Windows Server 2008 R2 Hyper-V. Windows Server 2008 R2 Hyper-V supports Quick and Live Migration.
Concurrent shared access	Concurrent shared access to application data on SAN storage from VMs improves flexibility and availability, as well as application scale-out for high I/O applications.
Centralized configuration	Central configuration and dynamic storage assignment using LaScala Clustered Volume Manager.
Supports over 200 hosts per LUN	Supports large host clusters. Is able to provide shared access to a LUN containing VHDs to more than 200 hosts. This allows multiple Hyper-V clusters to be deployed on a single storage volume.
Unlimited volume size	There are no volume size limitations.
Synchronous and asynchronous mirroring	Provides protection against storage array or site failure
Software SAN capability	Create a highly available shared block storage resource using internal drives of Hyper-V hosts
Storage live migration	Move an active volume between storage arrays while the application using it remains active
Monitoring and quality of service (QoS) capabilities	Provides storage performance monitoring and QoS assignment.
Distributed snapshots	Snapshots of the entire environment or of any individual VM can be made from any Hyper-V host in the cluster.
Easy installation	Melio FS and LaScala install easily using any vendor's SAN storage.
Microsoft virtualization integration	Designed for easy integration with Windows Server 2008, Windows Sever 2008 R2, and SCVMM 2008 R2.

**Melio Cloud software suite functionality**

## Technology Integration

Sanbolic's Melio Cloud Suite is designed to fully utilize the capabilities of any industry-standard SAN storage while integrating easily with Hyper-V and SCVMM, as well as with Windows application servers running on VMs. To provide enterprise level scalability, Melio uses a fully symmetrical architecture, so all Hyper-V hosts in the storage cluster have concurrent read-and-write access to a shared LUN over the SAN fabric. In addition, Melio is designed for several hundred hosts accessing a shared LUN concurrently and for a volume and file size of up to 18 exabyte's. Therefore, the full performance of SAN storage hardware is passed through with essentially no lock overhead when Melio is used to provide access to VHD files on a shared LUN from multiple Hyper-V hosts.

Furthermore, QoS management allows specific hosts or VMs to have priority access to storage bandwidth in the event that a constraint does occur. 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 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. The integrated management benefits of using Melio Cloud storage suite when deploying large numbers of Hyper-V VMs are discussed in the following three sections.

### **Volume Management**

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. Melio DataCenter Suite uses Windows ACLs to define access rights at the volume level in addition to the disk and the file level. Logical volumes can be remotely configured and assigned to physical or virtual servers using the standard Windows security interface. Sanbolic's volume manager uses only basic disks for all operations. Sanbolic's basic disks support stripe sets, volume sets, and mirror sets using fiber channel, SAS, or iSCSI LUNs. These disks can be dynamically expanded. Therefore, a Hyper-V deployment using 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. A logical volume can also be moved from one array to another while the application using the volume remains active.

### **VSS-based Snapshots**

In a large enterprise environment, data protection is a key consideration. 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 entire Hyper-V environment or of individual VMs from any of the Hyper-V hosts. This snapshot capability can be used 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. This can provide flexible options for recovering individual VMs while utilizing existing storage capacity. Melio Data Center Suite also allows snapshots to be scheduled and then automatically copied to a designated location.

### **Performance Monitoring and Quality of Service Management**

Performance monitoring and QoS management can be a key task 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 host or on a per VM basis.

---

## Conclusion

Large enterprise deployments of Hyper-V with a high VM count benefit from an optimized SAN storage solution that can effectively support the intensive demands imposed on network and disk resources. The solution outlined in this document is unique because Sanbolic's fully symmetrical architecture enables the infrastructure to scale to enterprise performance levels. The symmetrical architecture maintains the full performance of the storage platform as the deployment is scaled due to its lack of overhead or performance bottlenecks.

Organizations with enterprise license environments that provide access to Windows Server 2008 R2 Hyper-V should consider evaluating Hyper-V virtualization for large deployments and demanding workloads. In addition, the use of enterprise SAN storage hardware and the Sanbolic Melio Cloud storage software suite should be utilized to enhance the performance, scalability, data availability, and management features of Hyper-V environments.

## Additional Information

For those interested in implementing the blueprint provided by this paper, please contact Sanbolic or a Microsoft certified partner to discuss specific needs. To learn more about the technologies involved, please see the related links below.

**For more information, see:**

Microsoft Virtualization solutions:

<http://www.microsoft.com/virtualization/solutions>

Sanbolic Clustered File System for Hyper-V

<http://www.sanbolic.com/Hyper-V.htm>