

#SANLess Clusters:

High Availability for SQL Server in VMware Environments



SIOS

Server Virtualization Delivers Substantial Benefits

The benefits of server virtualization have become widely recognized in enterprise data centers. As a result, more and more data center operations are being moved to virtual server and cloud environments. A detailed Market Pulse study sponsored by VMware and Intel explored the reasons for the widespread and rapid adoption of server virtualization. According to the study, “server consolidation and disaster recovery/business continuity are the two top drivers and major applications behind IT virtualization initiatives. Other drivers include improving systems reliability/availability, more flexible development and testing environments, and lower data center operating costs.”

Server virtualization is part of a broader trend toward software-defined and cloud-based infrastructures that also involve storage and networks. But as with most major trends, there can be some downside associated with the benefits. For companies moving important applications that require high availability protection, the benefits of moving to a virtual environment may be offset by the limitations imposed by shared-storage clustering in vSphere. Before moving business critical applications to VMware environments, IT managers should understand these limitations and the simple steps they can take to avoid them.



SIOS

Limitations of Shared Storage Clusters in VMware Environments


To move important applications from physical servers to a VMware environment, IT managers need to provide high availability and disaster protection. However, creating a traditional shared storage cluster in a VMware environment is challenging in several ways.

In a traditional failover cluster, two or more physical servers (cluster nodes) are connected to a shared storage system. The application runs on one server, and in the event of a failure, clustering software such as Windows Server Failover Clustering (WSFC) moves the application to a standby node. Similar clustering is also possible with virtualized servers in a vSphere environment,

but this requires a special configuration for the storage.

Whether the servers are physical or virtual, the requirement of shared storage can create a single point of failure, as shown in the diagram. A storage area network (SAN) can have a high availability configuration, of course, but that increases its complexity and cost, and can adversely affect performance, especially for transaction-intensive applications like SQL Server.

Shared storage also imposes limitations for high availability and disaster recovery provisions, and these limitations can, in turn, create a barrier to migrating business-critical applications to vSphere.



In a traditional failover cluster, two or more physical servers (cluster nodes) are connected to a shared storage system. The application runs on one server, and in the event of a failure, clustering software such as Windows Server Failover Clustering (WSFC) moves the application to a standby node. Similar clustering is also possible with virtualized servers in a vSphere environment, but this requires a special configuration for the storage using Raw Device Mapping.

Sacrificing Important VMware Features

Creating a shared-storage cluster in vSphere requires the use of Raw Device Mapping (RDM) technology, which adds significant complexity that can make it difficult or impossible to use popular VMware features such as vMotion and VM Cloning.

In a traditional cluster created with WSFC in vSphere, RDM is used to provide virtual machines (VMs) with direct access to the underlying storage (SAN).

Significant Limitations Imposed by RDMs

However, a failover cluster that uses RDM complicates or prevents the use of several important VMware features that employ virtual machine disk (VMDK) files. For example, RDM prevents the use of VMware

snapshots and any feature that requires snapshots, such as Virtual Consolidated Backups (VCBs).

RDM also complicates data mobility, an important attribute of virtual environments. For example, with RDM, **you cannot:**

- Convert VMs into templates to simplify deployment (VM Cloning).
- Use vMotion with certain configurations to migrate VMs dynamically among hosts.

These pains associated with the use of RDM can completely undermine the potential gains that most IT departments hope to achieve with vSphere.

HIGH AVAILABILITY WITHOUT SACRIFICING FEATURES OR FLEXIBILITY

Raw Device Mapping Adds Complexity & Prevents the Use of Key VMware Features

- No VMware snapshots.
- No Virtual Consolidated Backups (VCBs).
- No VM Cloning.
- No vMotion with certain configurations.



HA Clusters without RDMs for Maximum Flexibility

An easy way to provide high availability protection for SQL Server without RDM or the limitations RDM imposes is to build a SANless cluster.

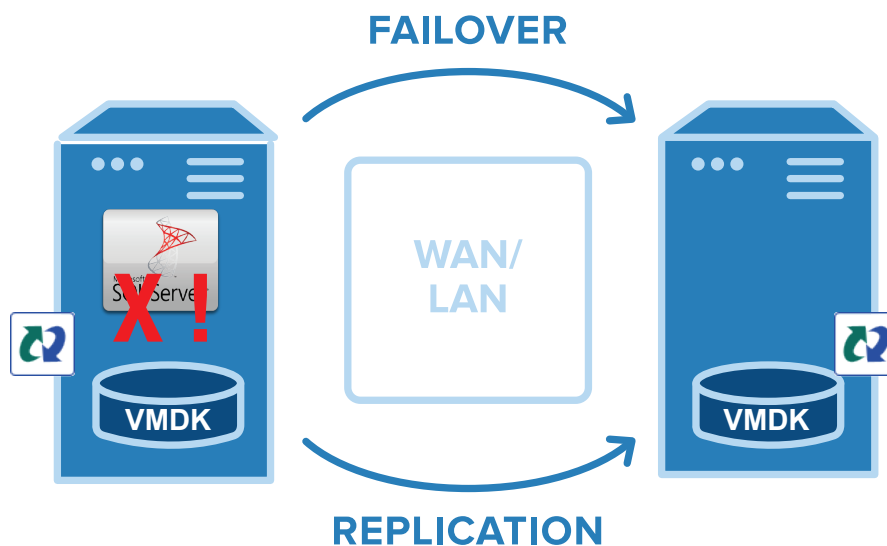
SANless clustering software, such as SIOS DataKeeper Cluster Edition, can be added as an ingredient in a Windows Server Failover Clustering (WSFC) environment. The software uses real-time block-level replication to synchronize local storage in each cluster node, making the nodes appear to WSFC like a single virtual SAN. This synchronization enables you to create a “shared nothing” SANless cluster using VMs as nodes in a vSphere environment without using RDM.

SIOS SANless clusters provide intuitive configuration wizards, simple management and reduced costs—all of which help

maximize the benefits of moving applications to vSphere. SANless clusters are also hardware agnostic. Any storage device or subsystem that is presented to the Windows operating system as a block device and appears in Windows Disk Management with a drive letter can be used as part of a valid configuration.

Maximum Flexibility

As shown in the diagram, each VMDK can be attached to different VMs in an N-node cluster using each VM’s independent storage. The software provides LAN/WAN-optimized, real-time block-level replication in either a synchronous or asynchronous mode. In effect, it creates a RAID 1 mirror across the network. It automatically changes the direction of the data replication (source and target) as needed after failover and failback.



In a SANless cluster, each VMDK can be attached to different VMs in an N-node cluster using each VM’s independent storage. The software provides LAN/WAN-optimized, real-time block-level replication in either a synchronous or asynchronous mode. In effect, it creates a RAID 1 mirror across the network. It automatically changes the direction of the data replication (source and target) as needed after failover and failback.

SQL AlwaysOn Availability Groups vs SANless Clusters

SANless clusters have additional advantages over traditional clusters for protecting SQL Server clusters running on virtual servers. The table below compares three options for providing high availability protection for SQL Server: SQL Server AlwaysOn Availability Groups; SQL Server AlwaysOn Failover Clustering alone; and SANless clusters created with SQL Server AlwaysOn Failover Clustering and SIOS DataKeeper Cluster Edition.

Note that the AlwaysOn Failover Cluster without SIOS software has a dependency on SAN and, therefore, has all of the limitations imposed by RDM. These clusters cannot provide disaster recovery across data centers.

The SQL Server AlwaysOn Availability Groups option introduces other limitations. For example, this option does not support more than two nodes for automatic failover and it requires the more expensive SQL Server Enterprise Edition.

Because clusters built with SIOS DataKeeper use efficient block replication technology and are fully integrated with WSFC, they protect the entire SQL instance, including the database, logons and agent jobs—all in an integrated fashion. Contrast this approach with AlwaysOn Availability Groups, which protect only user-defined databases, and require IT staff to manage every cluster node manually and individually.

	AlwaysOn Availability Groups	AlwaysOn Failover Cluster	SIOS SANless Cluster with AlwaysOn Failover Cluster
Supports SQL Server Standard Edition	○	●	●
Protects distributed transactions	○	●	●
Protects unlimited no. of databases	○	●	●
Protects system databases (master, MSDB, etc.)	○	●	●
Automates database administration	○	●	●
Supports more than two nodes in failover configuration	○	●	●
Replicates data other than SQL	○	○	●
Supports multisite clusters	●	○	●



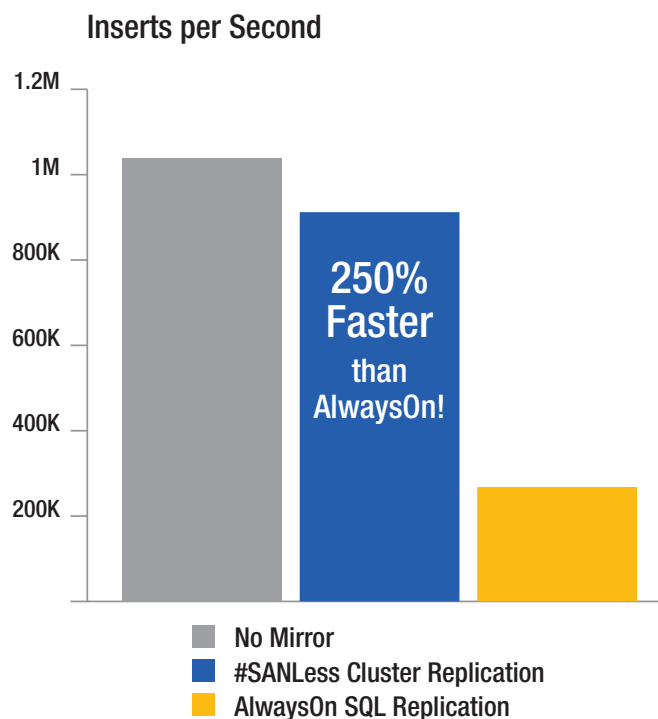
Performance Advantages of SANless Clusters

Factors like server CPU or processor contention and the use of SAN storage can significantly reduce application performance in VMware environments. Highly transactional applications like SQL Server are particularly vulnerable to these performance-related factors. Administrators often improve application performance with host-based caching by using high-performance solid state storage to maximize throughput.

The diagram below shows the 60-70 percent performance penalty associated with using AlwaysOn Availability Groups. These test results also show how a high-availability SIOS SANless cluster performs nearly as well as configurations not protected with any data replication or mirroring.

The test configuration used here used SanDisk Fusion-io PCIe flash cards to facilitate insertion of rows into a database as fast as possible; in this case, over 1 million rows per second with no mirroring. The block-level replication engine used by SIOS DataKeeper imposed at most a 10-15 percent overhead when configured for synchronous replication between two nodes over a 10 Gbps LAN.

SANless clusters built with SIOS DataKeeper are able to achieve such impressive performance because the SIOS software driver sits immediately below NTFS. As writes occur on the primary server, the driver writes one copy of the block to the local VMDK and another copy simultaneously across the network to the VMDK on the remote secondary server.



Maximizing vSphere Benefits with SANless Clusters

Server virtualization brings a number of compelling benefits that increase efficiency and productivity while reducing both capital and operational expenditures. But providing high availability and disaster recovery for virtual machines using shared, virtual storage comes with a serious limitation: the creation of a single point of failure.

VMware's vSphere is normally configured for high availability and rapid recovery from disasters, but the Raw Device Mapping technology necessitated by shared storage makes it difficult or impossible to use popular VMware features like vMotion and VM Cloning.

SIOS DataKeeper Cluster Edition is purpose-built to overcome the limitations associated

with shared storage, in general, and the requirement to use RDM with Windows Server Failover Clusters, in particular. As shown in the diagram, SIOS software creates a SANLess shared-nothing, hardware-agnostic N-node cluster capable of meeting the most stringent recovery point and recovery time objectives.

Just as important, a SIOS SANLess cluster is easy to implement and operate with both physical and virtual servers running either Windows or Linux. And because it utilizes local storage, a SIOS SANLess cluster delivers superior performance, especially for transaction-intensive SQL Server applications.

SANLESS CLUSTERS FOR VMWARE ENVIRONMENTS

- **High Availability Protection**
- **Disaster Recovery**
- **Easy Implementation and Configuration**
- **Full Integration with Windows Server Failover Clusters**
- **No VMware feature Limitations**



Ready to Learn More

About getting high availability for vSphere without sacrificing key
VMware features visit our website: <http://us.sios.com>

or call +1.650.645.7000