

Virtualization Reality: Why StarWind Virtual SAN Solutions **Deliver Value When Compared to Microsoft**







1. Shared Nothing Live Migration

Shared Nothing Live Migration makes it possible to move VMs from one running physical Hyper-V host to another without downtime. At that, you need neither a shared storage nor clustered Hyper-V hosts. Shared Nothing Life Migration helps manage scheduled downtime; however, it does not protect from unplanned downtime (hardware or software failure).

The big issue with this approach is the inability to build a guest VM cluster because it requires shared storage (FC or iSCSI). In addition, migration takes too much time if there are many virtual machines, since this process implies the movement of an entire VM on-disk image and VM memory content from one physical Hyper-V host to another.



StarWind utilizes the same hardware and software and solves all the issues mentioned above. It provides assistance with both planned and unplanned downtime. With StarWind, it is easy to build a guest VM cluster because it is a fault-tolerant iSCSI SAN itself. The VM data store can be any size because all the data is synchronized between hosts.



Solution capabilities	Microsoft	StarWind
Assistance with planned downtime	\bigcirc	\bigcirc
Protection from unplanned downtime	8	\bigotimes
Guest VM cluster	\bigotimes	\bigotimes
VM data store size	The bigger data store size is, the longer VM migration takes.	StarWind does not depend on the VM data store size.



2. Hyper-V Replica

Hyper-V Replica ensures effective VM level disaster recovery. It enables VM replication and manual failover in case of a sudden server breakdown; however, it does not deliver automatic failover. Data is asynchronously replicated with a minimum 5 minute interval, so that after the failover there is data loss for the referenced period.



Along with manual failover, StarWind ensures automatic failover, thereby providing continuous availability, zero data loss, and the best RTO and RPO values. Since StarWind makes synchronous mirroring, there is no difference between the data kept on both storage servers.



Solution capabilities	Microsoft	StarWind
Manual failover	©	\bigcirc
Automatic failover	8	
Data difference	Minimum 5 minutes	Data is synchronously mirrored (no difference)
Utilization of slow WAN connections	Good	Fair



3. SMB 3.0

SMB 3.0 provides a very simple way to cluster Hyper-V hosts. At least three hosts and all the associated network infrastructure are required to build a solution with a SMB 3.0 shared storage. As a result, the price of this kind of an installation is rather high. In addition, such a configuration makes the storage server holding VM images a single point of failure providing no fault-tolerance. Moreover, performance is not great because all read and write requests are sent over a comparably slow network.



From the other side, StarWind requires only 2 already existing Hyper-V hosts to build a completely fault-tolerant Hyper-V cluster. All read requests are processed locally (DAS) and only writes require an acknowledgement over LAN. Therefore, StarWind provides a considerably higher performance compared to the SMB 3.0 model. The StarWind architecture also eliminates a single point of failure by building a fully redundant configuration.



Solution capabilities	Microsoft	StarWind
Number of nodes	3	2
Performance	LAN	LAN/DAS
Uptime	Single point of failure	Full fault tolerance



4. Hyper-V cluster with SAS JBOD as shared storage

SAS JBOD is a very easy-to-use solution for the creation of high performance Hyper-V clusters. However, it requires dedicated hardware (actual SAS JBOD enclosures, SAS controllers, and SAS disks). It is impossible to use the existing SAS or SATA disks already mounted on physical Hyper-V servers. A single JBOD does not guarantee 100% fault-tolerance because not all the components inside it are redundant. Only dual JBODs provide an adequate level of failure resistance.



With StarWind, there is no need for extra disks, dedicated controllers, or chassis because it builds highly available shared storage by using the existing commodity hardware, network, and cheap internal disks SATA that are already present on Hyper-V servers. StarWind solutions reduce the hardware requirements at least twofold, thereby saving on costs, and they provide improved performance at the same time.



Solution capabilities	Microsoft	StarWind
SAS JBOD enclosures, SAS disks, SAS controllers	Required	Not required
Disk interface	SAS	SATA or any type of disk
Single point of failure	Not all components are redundant	All components are at least duplicated



5. Scale-Out File Server

Scale-Out File Servers allows for clustering two or more SMB 3.0 servers (up to four) for fault-tolerance, high performance, and load balancing. External shared storage – FC, iSCSI, or SAS JBOD – is also necessary to cluster Scale-Out File Servers. Such a complex architecture has a long I/O path (the reads are first processed over the network, and then they are addressed to the local storage – LAN, DAS) which degrades the performance.

SMB 3.0 can only scale with multiple clients. Single client requests are processed by just one node and, therefore, there should be many clients to increase the load and attain better utilization.





StarWind requires only two physical servers to provide fault-tolerance and high performance. That reduces the hardware expenditures at least by half compared to the Microsoft reference configuration. The reads are performed locally, which shortens the I/O path and reduces latency. StarWind utilizes iSCSI MPIO, which helps distribute client requests between the nodes. ISCSI MPIO provides automatic load balancing and performance increase.



Solution capabilities	Microsoft	StarWind
Number of servers	Minimum 4	2
External shared storage: FC, iSCSI, SAS JBOD	Required	Not required
I/O path	Long	Short
Client request processing	SMB can scale with multiple clients only	iSCSI MPIO with automatic load balancing, and performance increase