

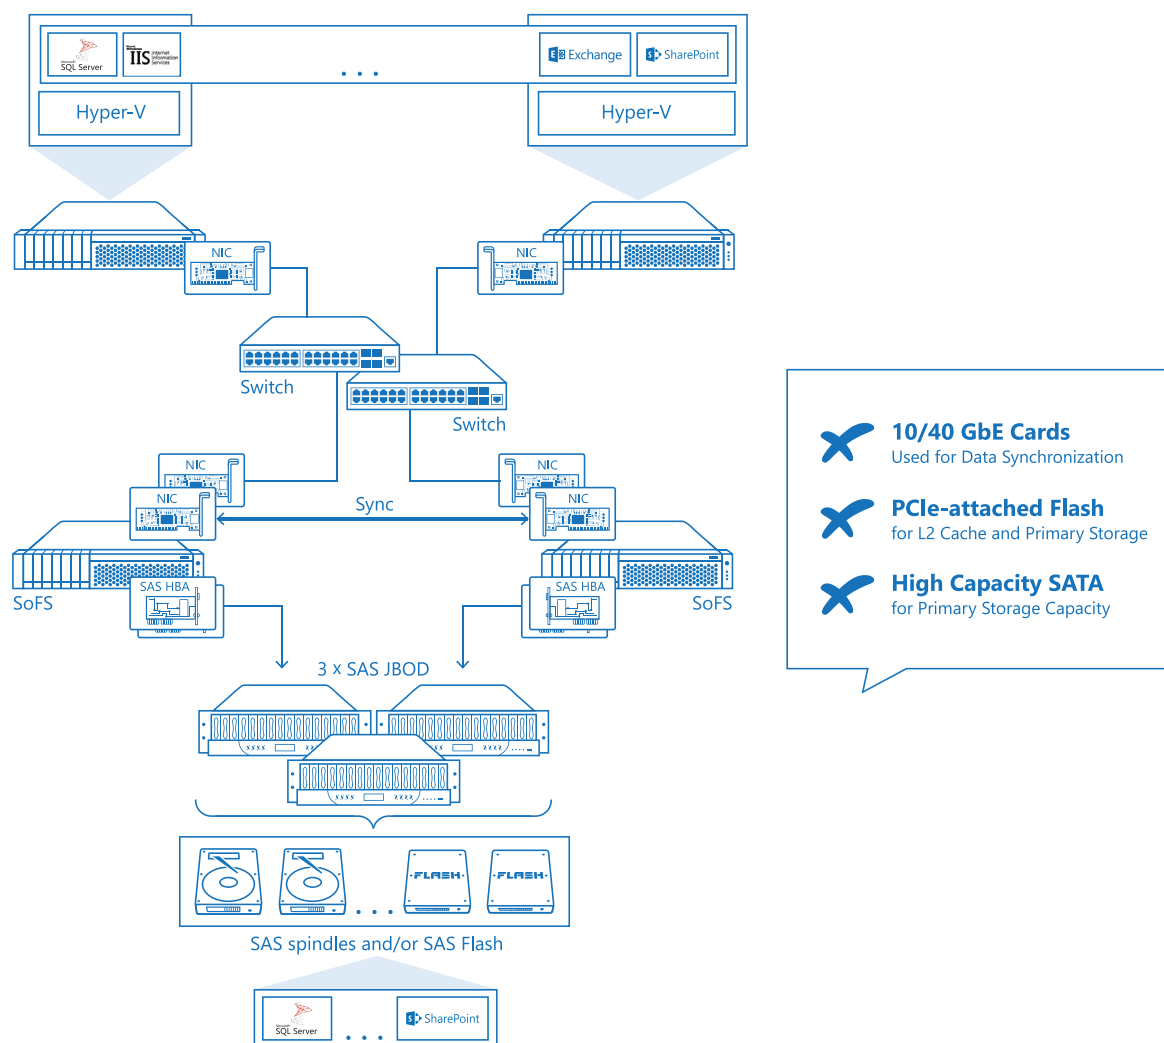
# Scale-Out File Servers (SoFS)

## INTRO

Microsoft now can be considered a serious virtualization player after release of Windows Server 2012 R2 and Hyper-V. They see the future of VM storage as SMB3, running virtualization and storage layers separately from each other. The most important reason is that using SMB protocol, Microsoft controls the whole specification, so it is possible to apply changes quickly. Most surely, upcoming storage initiatives will be built on top of SMB protocol, not iSCSI, which is considered to be legacy technology by Microsoft.

## PROBLEM

**High implementation cost.** Properly configured SMB3 based storage configuration is quite the same as high-end SAN or NAS appliances in terms of price. Also, there are two additional Windows Server licenses required. It is fine for test and development tasks, “quick and dirty” prototyping, but having SPoF inside of the cluster is a bad idea.



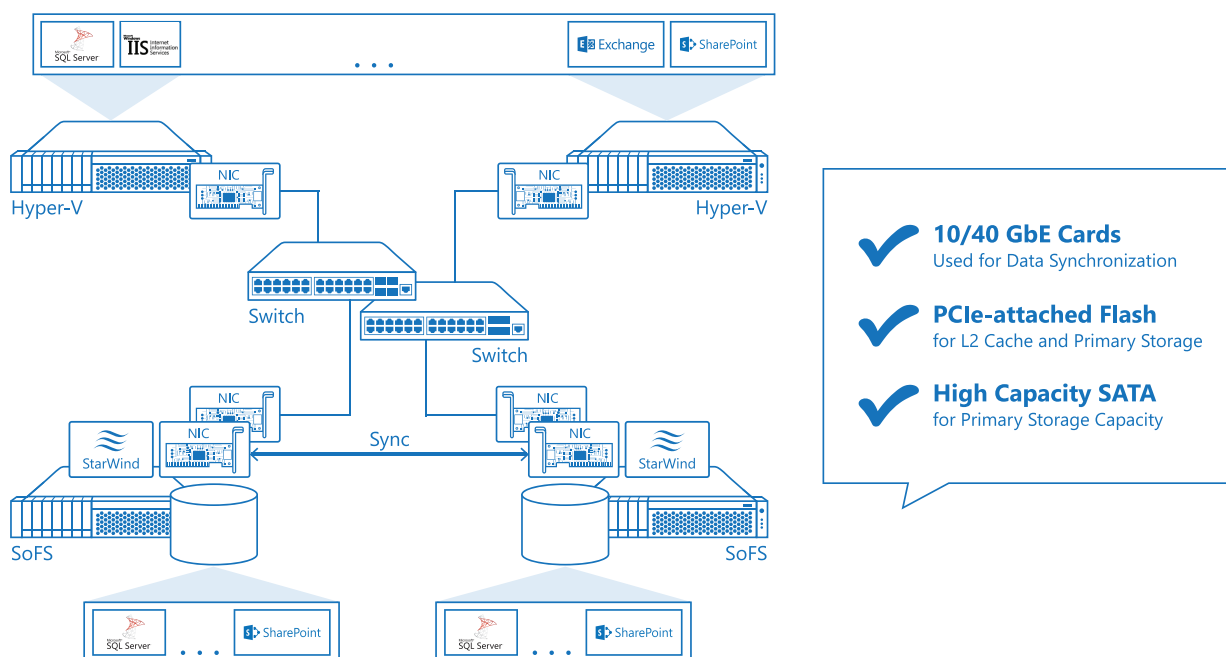
Microsoft's idea is Scale-Out File Servers, which is basically two or more Windows Server 2012 R2 machines, exporting the Continuously Available (CA) SMB3 shares to the clients. The problem here is that SoFS is built on top of the CSV, which requires block-based shared storage beneath it. This is when good old SAN or Clustered Storage Spaces come to play. There is no

point to layer SoFS as iSCSI SAN, since iSCSI can be passed through to client applications without additional layers. Using Fiber Channel SAN is considered only if there is already an FC based system behind; otherwise new hardware raises the price. Thus, Microsoft came to the idea of using Clustered Storage Spaces as the block shared backend for SoFS.

Typical single JBOD installation is the single point of failure, if connected through multiple paths. This is unobvious since all of those components are hidden inside the box. Microsoft's approach raises the implementation costs due to requirement for multiple SAS JBODs and lots of SAS equipment. The storage becomes too expensive and performs slower than high-end SAN.

## SOLUTION

StarWind replaces block back end for Scale-Out File Servers resulting in **significant CapEx and OpEx decrease**. That's achieved through storing the data on cheap SATA spindles. StarWind increases the performance by using uber-fast PCIe Flash for L2 or data storing and 10-40/56 Gbps PCIe-Ethernet NICs to run data synchronization.



**StarWind Virtual SAN** also brings the beneficial features complementing the Microsoft technologies:

- In-line deduplication increases utilization of usable space, which is especially the issue with Flash. Also it increases the performance for the write-intensive VMs, which perfectly complements the Microsoft Off-Line Deduplication optimized for reads and storing backups;
- RAM Write-Back caching replaces Microsoft CSV cache and Flash Write-Back cache for maximum possible performance and reliability without breaking the bank. Additionally it prolongs Flash life by absorbing the write I/Os;
- Log-Structuring File system (LSFS) eliminates I/O blender, which boosts the performance of VM-specific workload and prolonging the Flash life, while improves ability to use Flash, since there are no spot burns.

## CONCLUSION

**StarWind Virtual SAN** takes Microsoft SMB3 and SoFS initiatives and makes them much less expensive while boosting the performance at the same time, as the result of removing having the SAS-only based storage backend. As the result it is possible to use cheap high-capacity SATA spindles, uber-fast PCIe-attached Flash for storage and high-speed 10 and 40/56 Ethernet NICs for data synchronization.