



StarWind Virtual SAN and Microsoft SOFS

Cutting down SMB and ROBO virtualization cost by using less hardware with Microsoft Scale-Out File Server

StarWind Virtual SAN for Microsoft SOFS

*Cutting down SMB and ROBO virtualization cost by using less hardware
with Microsoft Scale-Out File Server (SOFS)*

By Greg Schulz

**Founder and Senior Advisory Analyst
Server and StorageIO Group @StorageIO**



November 18, 2014



StarWind Virtual SAN and Microsoft SOFS

Cutting down SMB and ROBO virtualization cost by using less hardware with Microsoft Scale-Out File Server

Introduction

This StorageIO Industry Trends Perspective (ITP) thought leadership white paper looks at addressing the shared storage needs of small medium businesses (SMBs) and remote office branch offices (ROBOs) leveraging Microsoft Scale-Out File Server (SOFS). The focus is on how Microsoft Windows Server 2012 [as well as 2012 R2], Server Message Block version 3.0 (SMB 3.0) and SOFS (Scale-Out File Server) combined with StarWind Virtual SAN management software unlock the full value of server-based storage. Unlocking the value of server-based storage is particularly important for ROBO and SMBs who have SQL Server databases and Hyper-V server virtualization. This approach provides a cost effective high-performance resilient storage option for small SMBs and ROBOs.

General Issues and Challenges

Organizations of all size, from SMB to ROBO to large enterprise environments have a reliance on information. This means supporting the aggregate workload performance needs of SQL Server, Hyper-V guest virtual machines and other applications. No one solution fits all needs or environments. Smaller environments typically do not have economies of scale to spread costs across many systems. In smaller SMB and ROBO environments there are not as many systems resulting in a higher cost overhead to share across fewer systems.

SMB and ROBO environments also may not have the specialty skills trained staff that are typically found in larger enterprise data center environments. ROBOs also have the need to replicate data for resiliency to a core data center location. Convergence or hyper-convergence should not cause hyper-complexity, hyper-compromise or hyper-costs for SMB and enterprise ROBO environments.

With many solutions designed for larger environments, extra hardware is absorbed and used as part of scaling up. However, for smaller environments the scale-down results in overhead. Of course, the same can hold true in that solutions designed and optimized for the SMB and ROBO can be used to scale-up, however, there can also be limits on how effectively and economically this can be done. To support sharing of direct attached storage (DAS), a solution should be hardware agnostic while removing complexity and support scale-up, scale-out and scale-down.

Aggregation causing Aggravation

Consolidation (aggregation of different applications, workloads and servers can cause server and storage I/O performance bottlenecks (Aggravation). This means server and storage I/O hardware and software need to be able to support aggregation without contributing too, or causing additional bottlenecks (aggravation). Instead, fast hardware needs fast software to support fast applications such as SQL Server and Hyper-V for virtual server consolidation.

Flexibility and Hardware Freedom

It is still impossible to have software and storage without some type of hardware. However, it is reality today to have freedom of choice pertaining to hardware including how much, what type and from whom, as well as leverage what you already own. This means using software that unlocks the value of your existing or new hardware requiring only what you need to support your current server and storage I/O performance as well as availability needs. Result is stretching IT budgets further supporting growth with resiliency and without increased complexity.



StarWind Virtual SAN and Microsoft SOFS

Cutting down SMB and ROBO virtualization cost by using less hardware with Microsoft Scale-Out File Server

Microsoft Server Storage Challenges

SQL Server and Hyper-V are common on Microsoft Windows Server systems including for SMB and enterprise ROBO environments. Some common server and storage I/O challenges pertaining to Microsoft Server systems supporting Hyper-V and SQL Server include:

- Storage space capacity – To support growth along with consolidation of servers and virtual desktops (VDI), along with databases holding more information more storage space capacity and performance are needed.
- Performance issues – Small random reads and writes (IOPs) for SQL Server databases as well as throughput (bandwidth) for large sequential reads and writes.
- Availability – Reliability as well as Resiliency including failover-over for High-Availability (HA), business continuance (BC), and disaster recovery (DR).
- Budget stretchers – Do more with available budgets to support growing applications in small environments
- Investment protection – Leverage existing server, storage and networking hardware you already have
- Simplify management – Remove complexity and costs associated with management of SQL Server, Hyper-V and other Microsoft Server environments.

Consolidation including server and VDI virtualization (aggregation) can cause bottlenecks (aggravation). These aggravations include the need for more performance, availability along with space capacity. This means as part of consolidating and virtualization applications and their workloads you also need to eliminate the aggravations caused by bottlenecks. However, this also means finding a balance between increasing complexities and cost, as well as enabling performance, availability and space capacity without compromise or additional overhead.

To address the above and other issues Microsoft has added new features and enhancements with Server 2012 [R2], along with SOFS and SMB 3.0. Some of these features are for general purpose, others are for Hyper-V and SQL Server including for smaller organizations as well as enterprise ROBO and larger environments

What is new with SMB 3.0?

SMB 3.0 provides enhancements to existing local and wide area file access including security encryption along with lower latency (e.g. better performance and security). Important new SMB 3.0 features include transparent failover of network paths, scale-out supporting Clustered Shared Volumes (CSV version 2) across all nodes for better resiliency and performance. Also new is multi-channel support enabling aggregation of several network paths ranging from 1GbE and 10GbE as well as higher-speeds including InfiniBand.

Another enhancement for Hyper-V Server virtualization and SQL Server databases is SMB Direct SMB Direct leverages network adapters that support Remote Direct Memory Access (RDMA) enabling remote file shares in a cluster to perform as though they are local direct accessed storage (faster performance).

Other SMB 3.0 improvements include Directory Leasing to cache commonly accessed meta-data to reduce network overhead improving performance via caching, as well as performance monitoring counters optimizing for small random reads and writes common to SQL Server environments along with enabling large Maximum Transmission Units (MTU) by default (also for performance). These enhancements address various issues and challenges for Microsoft Server, SQL Server and Hyper-V.



StarWind Virtual SAN and Microsoft SOFS

Cutting down SMB and ROBO virtualization cost by using less hardware with Microsoft Scale-Out File Server

Microsoft Scale-Out File Server (SOFS) and Windows 2012 R2

Microsoft Server Message Block (SMB) is an important technology for access and sharing data with different clients, servers and applications. With Server 2012 and 2012 R2, Microsoft has added enhancements for supporting scale-up, scale-out (and also scale down²) capabilities to support Hyper-V, SQL Server and general file serving among other application workloads. These capabilities can scale from a small SMB or enterprise ROBO environment to larger SMB and enterprise data center environments.

At the heart of Windows 2012 and 2012 R2 are SMB 3.0 enhancements which will enable a transition from its perceived role for Windows file serving (e.g. CIFS³) to that of a low-latency, high-performance server-to-server data and storage I/O access mechanism for database (e.g. SQL Server), VDI and server virtualization (Hyper-V). What this means is that while SMB still supports general file sharing among Windows and other servers over networks, there is low-latency, higher performance support for block-based storage applications such as SQL Server databases.

Scale-Out File Server (SOFS)

Designed to support the performance and resiliency scaling needs of Hyper-V and SQL Server. SOFS enables flexible and easy scaling using familiar Windows Server management tools. SOFS leverages SMB 3.0 (eliminates need for using shared iSCSI or SAN storage); Clustered Shared Volumes (CSV) and if needed can scale up to 64 nodes (e.g. Windows 2012 [R2] servers). Microsoft SOFS is an example of a converged solution enabling Clusters in a Box (CiB) with your choice of hardware; including using what you already own to protect your investment (reduces TCO¹).

Windows Server 2012[R2] enhancements include:

- Enable scaling of server and storage I/O capabilities
- Support Cluster in a Box (CiB)
- Tiering of shared storage using SSD's and HDD's⁴
- Basic read-cache using shared SSD⁵
- Data resiliency options (triple mirrors, dual parity)
- Server deduplication vs. Single Instance Storage (SIS)
- Clustering for load-balance (performance) and resiliency
- SMB 3.0 extending from legacy file serving to enabling low-latency high-performance storage for SQL Server and Hyper-V based applications (alternative to iSCSI and Fibre Channel SANs)
- Scale-Out File Server (SOFS)

¹ TCO = Total Cost of Ownership reduction by maximizing what you already own vs. spending budgets to replace

² Scale Down – Adapt to needs of smaller business and ROBOs with capabilities found in larger environments

³ CIFS = Common Internet File System or Windows File Shares also known as SMB (Server Message Block)

⁴ SSD = Solid State Devices, HDD = Hard Disk Drives

⁵ Shared flash based SSD read cache needs to be accessible by other nodes in the cluster (shared anything)

StarWind Virtual SAN and Microsoft SOFS

Cutting down SMB and ROBO virtualization cost by using less hardware with Microsoft Scale-Out File Server

Microsoft SOFS Reference Architecture

The Microsoft SOFS reference architecture is shown in figure-1 to the right. In the reference architecture shown external shared SAS JBOD disk enclosures are attached to the storage server nodes with switched SAS connectivity. Industry standard LAN technology including 10GbE connects storage server nodes and Hyper-V server nodes. Both shared SAS storage SAN and server-to-server LANs are configured for high-availability with redundant paths.

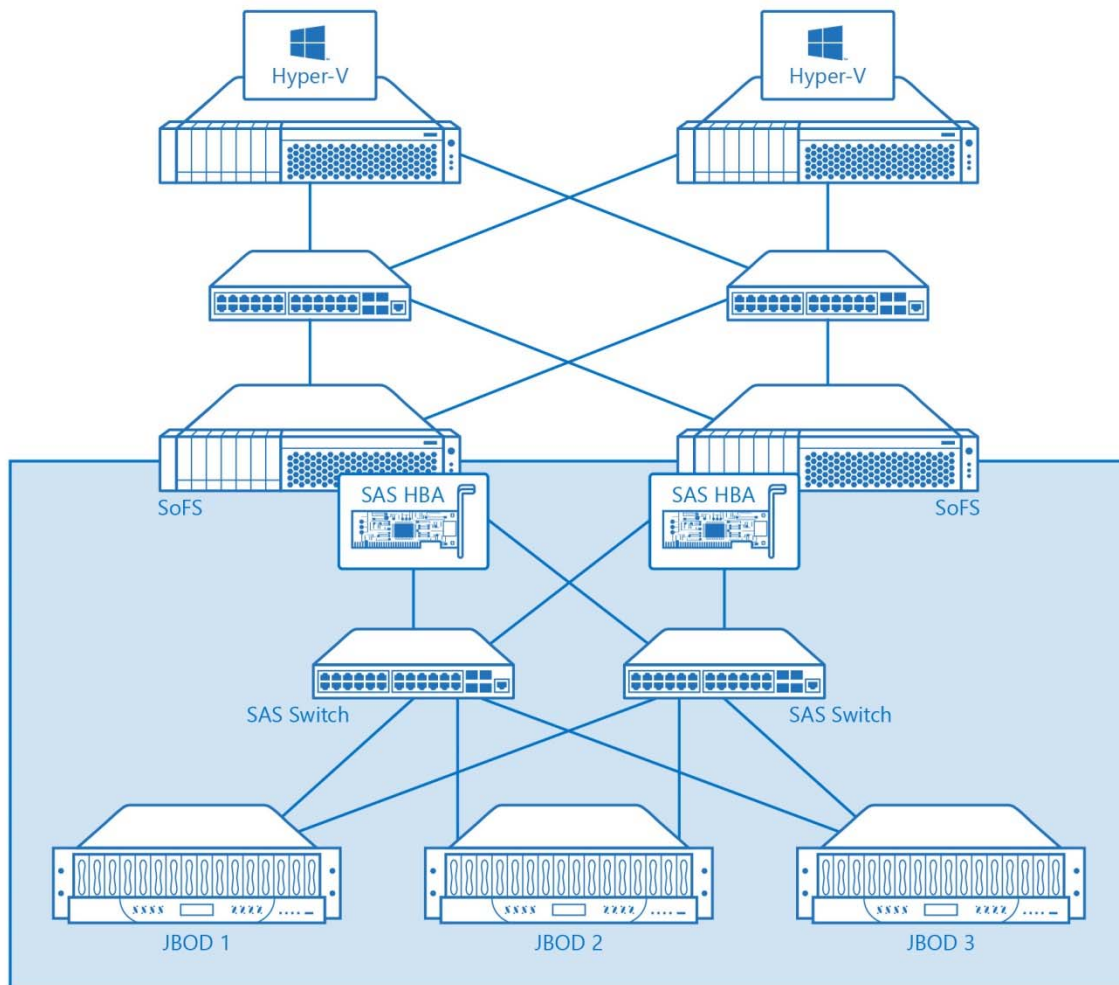


Figure-1 Large Microsoft SOFS architecture (Image via StarWind)

Not shown in figure-1 is lack of support for PCIe SSD's or other local SAS, SATA SSD and HDD's in SOFS storage environment. While the SOFS CiB reference architecture shown in figure-1 can be applicable for larger enterprise environments that need to scale-out and scale-up, it does not support scaling-down to meet the needs of smaller SMB and ROBO environments.



StarWind Virtual SAN and Microsoft SOFS

Cutting down SMB and ROBO virtualization cost by using less hardware with Microsoft Scale-Out File Server

StarWind Virtual SAN and Microsoft SOFS

As mentioned above, the combination of Microsoft Server 2012 [R2], CSV, SMB 3.0 and SOFS provides a good option for scaling larger environments, some additional functionality is needed to scale-down without compromise.

The Microsoft SOFS CiB approach supports scale-up and scale-out using external shared storage, however it has a challenge with using internal storage while providing failover resiliency. This is where StarWind comes into the solution enabling resilient sharing of local PCIe, SAS and SATA SSD's as well as HDD's across the SOFS cluster.

For smaller SMBs and ROBOs this means being able to:

- Use local JBOD and other storage shared in SOFS cluster
- Use existing server, storage and I/O network hardware
- Avoid overhead of a larger scale-out architecture
- Freedom to choose your hardware and from whom
- Enable read and write cache using local SSD's
- Use SATA/SAS/PCIe SSD as cache or cache and storage
- Leverage DRAM as high-performance write cache
- Use wide-area asynchronous replication for BC and DR

What this all means is that you can scale-down a Microsoft Windows Server 2012 [R2] + SMB 3.0 + SOFS combination and then add StarWind. This combined solution scales-down the amount additional hardware, complexity and cost while providing reliable (resilient) high-performance storage for smaller SMBs and ROBO environments.

Figure-2 on the next page shows Microsoft Server 2012[R2] + SMB 3.0 + SOFS + StarWind enabled configuring supporting SQL Server and Hyper-V that has been scaled-down to meet needs of smaller SMBs and ROBO environments. The "Scaling-down" means using less hardware that results in lower cost and complexity while enabling all the benefits of found in larger SOFS and CiB deployments.

SOFS and Internal DAS Storage

By default, SOFS works well for scaling-up and scaling-out to meet the needs of larger environments.

Using shared storage such as switched SAS JBOD enables larger environments to scale, however what about scaling-down? For smaller SMBs and ROBOs SOFS CiBs designed for larger environments, bring too much cost, complexity and hardware overhead.

What is needed is a simpler approach that scales-own the hardware complexity overhead (and cost) for smaller SMBs and ROBOs.

This means being able to use internal and direct attached storage (DAS) such as PCIe SSD along with SAS and SATA SSD as well as HDD's something that SOFS needs help from third parties such as StarWind to do.

StarWind Virtual SAN and Microsoft SOFS

Cutting down SMB and ROBO virtualization cost by using less hardware with Microsoft Scale-Out File Server




Figure-2 SOFS with StarWind (via StarWind)

As part of the example in figure-2, local internal PCIe flash SSD are used for read and write-back cache. In addition, local standard SAS and SATA HDD's (or SSD's) are mirrored across the nodes to provide resilient space capacity for storing Hyper-V, SQL Server and other data. For server and storage I/O connectivity, multiple standard 10GbE ports are used including switch-less (as shown) or with switched. Various industry standard networking interfaces are supported for enabling connectivity between the servers including multiple adapters to meet specific performance needs. Not shown in figure-2 are using DRAM memory as distributed read or write cache along with StarWind enabled dedupe to reduce data footprint overhead impact (space capacity).

Unlike the Microsoft SOFS reference architecture shown in figure-1 that may be more applicable to a larger environment looking to scale-out and scale-up, the StarWind enabled approach showing in figure-2 enables scaling-down to meet needs of smaller SMB and ROBO environments. The StarWind enabled configuration also enables use of PCIe, USB and other local devices that are not normally supported for shared-storage in a SOFS environment. In addition to enabling scaling-down to two-nodes for SMB and ROBO resiliency, StarWind also supports scale-out for larger environments without requiring additional proprietary external shared hardware to implement a shared anything cluster. The benefit of the above is a simpler Microsoft SOFS environment enabled by log-structured StarWind software is to remove complexity, reduce hardware overhead, while maintaining or enhancing resiliency without compromise, not to mention also cutting costs and simplifying management.



StarWind Virtual SAN and Microsoft SOFS

Cutting down SMB and ROBO virtualization cost by using less hardware with Microsoft Scale-Out File Server

Summary

There are various challenges, issues and opportunities for different types of environment from small SMB, enterprise ROBO to service provider among others. There are various storage options for supporting Microsoft SOFS and Windows Server storage environments and one size or architecture approach does not address all environments. However, some solutions may be stretched or squeezed resulting in compromise in other areas.

For smaller SMBs and ROBOs, resilient high performance solutions are possible to support SQL Server and Hyper-V using internal direct attached storage using StarWind + SMB 3.0 + Server 2012 [R2] + SOFS. StarWind Virtual SAN is an example of a solution that compliments Microsoft SOFS and Windows Server 2012 R2 that is hardware agnostic providing software defined storage management. In addition to complimenting Microsoft Windows as well as other environments, StarWind Virtual SAN also cuts complexity and reduces cost. Cost savings are the result of having a lower hardware footprint overhead vs. other solutions and implementation approaches.

StarWind supports scale-up, scale-out, as well as scale-down to meet the needs of smaller SMB and ROBO environments including performance for SQL Server, Hyper-V and file sharing along with resiliency. For example, StarWind can enable resiliency including replication with as few as two physical machine (PM) servers without introducing extra hardware costs or increasing complexity.

Learn more at StarWind landing page: www.starwindsoftware.com

About the author

Greg Schulz is Founder and Sr. Analyst of independent IT advisory consultancy firm Server and StorageIO (StorageIO). He has worked in IT at an electrical utility, financial services and transportation firms in roles ranging from business applications development to systems management, architecture, strategy and capacity planning with over three decades of applied experience. Mr. Schulz is author of the Intel Recommended Reading List books “Cloud and Virtual Data Storage Networking” and “The Green and Virtual Data Center” via CRC Press and “Resilient Storage Networks” (Elsevier). He is a five-time VMware vExpert. Learn more at www.storageio.com.



All trademarks are the property of their respective companies and owners. The Server and StorageIO (StorageIO) Group makes no expressed or implied warranties in this document relating to the use or operation of the products and techniques described herein. StorageIO in no event shall be liable for any indirect, inconsequential, special, incidental or other damages arising out of or associated with any aspect of this document, its use, reliance upon the information, recommendations, or inadvertent errors contained herein. Information, opinions and recommendations made by StorageIO are based upon public information believed to be accurate, reliable, and subject to change. Refer to StorageIO privacy and [Disclosure policy here](#). This industry trends and perspective white paper is compliments of StarWind <http://starwindsoftware.com>