

 **OUTDATED**

Microsoft Clustered Storage Spaces and **StarWind Virtual SAN[®]** for a Complete Software-Defined Storage Solution

JANUARY 2015

WHITE PAPER BY JON TOIGO

Trademarks

“StarWind”, “StarWind Software” and the StarWind and the StarWind Software logos are registered trademarks of StarWind Software. “StarWind LSFS” is a trademark of StarWind Software which may be registered in some jurisdictions. All other trademarks are owned by their respective owners.

Changes

The material in this document is for information only and is subject to change without notice. While reasonable efforts have been made in the preparation of this document to assure its accuracy, StarWind Software assumes no liability resulting from errors or omissions in this document, or from the use of the information contained herein. StarWind Software reserves the right to make changes in the product design without reservation and without notification to its users.

Technical Support and Services

If you have questions about installing or using this software, check this and other documents first - you will find answers to most of your questions on the [Technical Papers](#) webpage or in [StarWind Forum](#). If you need further assistance, please [contact us](#).

Copyright ©2009-2015 StarWind Software Inc.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written consent of StarWind Software.

In 2016, Gartner named StarWind “Cool Vendor for Compute Platforms”.

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

About StarWind

StarWind is a pioneer in virtualization and a company that participated in the development of this technology from its earliest days. Now the company is among the leading vendors of software and hardware hyper-converged solutions. The company's core product is the years-proven StarWind Virtual SAN, which allows SMB and ROBO to benefit from cost-efficient hyperconverged IT infrastructure. Having earned a reputation of reliability, StarWind created a hardware product line and is actively tapping into hyperconverged and storage appliances market. In 2016, Gartner named StarWind “Cool Vendor for Compute Platforms” following the success and popularity of StarWind HyperConverged Appliance. StarWind partners with world-known companies: Microsoft, VMware, Veeam, Intel, Dell, Mellanox, Citrix, Western Digital, etc.

Summary

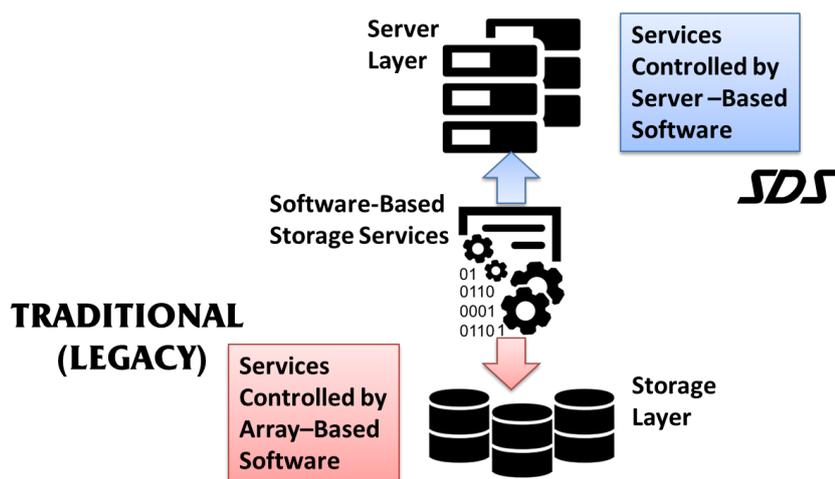
Software-defined storage (SDS) architecture continues to gain mind share in organizations that are virtualizing their storage infrastructure. Within the past 18 months, leading server hypervisor software vendors, including Microsoft, have introduced their own integral SDS offerings that are intended to deliver a “one stop shop” for consumers of their server virtualization wares. However, the SDS technologies from hypervisor vendors have tended to lag behind the wares of third party SDS innovators, such as StarWind Software. The smart innovator seeks to work with, rather than to compete with, the much larger operating system and server virtualization vendor. So, StarWind Software is positioning its Virtual SAN as a technology that augments and complements Microsoft’s integral SDS offering. In the process, StarWind resolves some of the technical issues that continue to create problems for Redmond’s software-defined storage approach.

Introduction

Software-defined storage (SDS) has captured the attention of the analyst community, the trade press and the server virtualization market, where consumers are looking for less expensive and more responsive storage solutions to support virtualized application workload. That said, there is still considerable disagreement regarding the meaning of software-defined storage itself.

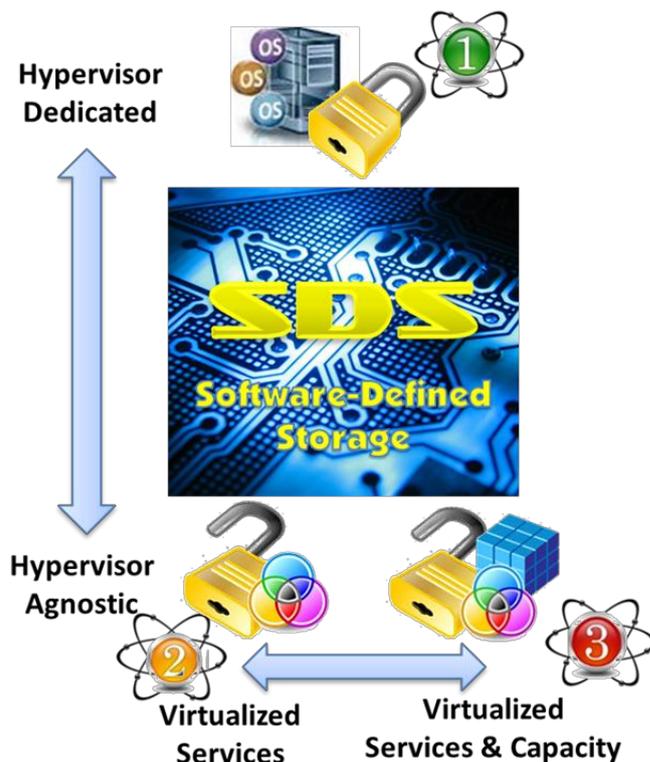
At a high level, SDS refers to the creation of a universal storage controller in a server-based software layer – a strategy that moves the value-add storage services, functionality that traditional array vendors prefer to host on their storage array controllers, out of the array kit and into the server software stack. By doing so, SDS advocates say, the storage array becomes a box of commodity storage devices – just a bunch of disks (JBOD).

This, they claim, reduces the acquisition cost of storage hardware and, in theory at least, enables the easier allocation of capacity and value-add functionality – or storage services – to specific applications and workloads.



Software-defined storage is also supposed to make it easier for a virtualization administrator, who may have little or no knowledge of the technical details of storage, to manage storage resources and to ensure that application performance is not constrained by storage-related issues.

Unfortunately, there is no agreed upon strategy, architecture, or topology for delivering software-defined storage to market. Every vendor in the SDS space has its own interpretation of how to build SDS infrastructure. That has led to the introduction of a wide variety of software-defined solutions in the market today. For purposes of analysis, the most important differentiators are two-fold: hypervisor dependency and hardware dependency.



Some server hypervisor software vendors, like VMware for example, have built SDS solutions that are designed to support only workloads that are hosted using that hypervisor. Their SDS architecture is “hypervisor dependent” and creates what can be called an “isolated island” of storage – one that can only be accessed, used or shared among workloads running under the same hypervisor software stack.

Other approaches to delivering software defined storage, usually products from third party SDS vendors, are “hypervisor-agnostic.” Independent software developers like StarWind Software have gone to great effort to create virtual SAN software defined storage solutions that can host the data from applications and workload operating under multiple hypervisors. That is what we call hypervisor-agnostic SDS and it is a good general differentiator for planners seeking to make sense of the various SDS offerings in the market.

The other big differentiator in SDS is linked to hardware dependency. The ideal SDS solution is supposed to be completely “hardware agnostic.” Having removed all of the value-add software from proprietary array controllers and placed this functionality into a software layer running on a server, SDS should be capable of working with any commodity array components that remain with no concern about the vendor brand name on the outside of the array cabinet. The core idea of SDS is that stripping away the value add software from the array creates commodity storage: as a result, every array is simply a box of Seagate or Western Digital disk drives and perhaps whatever flash storage components that may be popular at any given time. Vendors of these commodity components all conform to the same protocols for operation and access – SATA or SAS, or maybe Fibre Channel or iSCSI, so the SDS solution shouldn't prefer one vendor's box of drives over any other vendor's box of drives.

However, some SDS solutions remain “hardware dependent.” This is often the case with respect to hardware-based “hyper-converged SDS” storage hardware solutions, most of which are simply proprietary storage appliances, still implementing proprietary on-controller value-add software, and designed to work with a specific hypervisor. Hardware dependency can also be found in some of the so-called “open” SDS offerings from server hypervisor vendors. For example, a leading hypervisor vendor’s Virtual SAN requires the exclusive use of SAS drives in its infrastructure, or, optionally, the use of storage node hardware that is “pre-certified” and branded by the vendor.

By contrast, third party SDS technologies such as StarWind Software’s Virtual SAN, in addition to being hypervisor agnostic, are hardware agnostic as well. In the case of StarWind, SATA as well as SAS disk may be used as primary storage, and even hyper-converged appliances are supported for use in the StarWind SDS infrastructure.

Bottom line: while virtualization evangelists may stress hypervisor and hardware agnosticism as key characteristics of software-defined storage nirvana, such characteristics are by no means assured in current SDS products. Such agnosticism isn’t necessarily in line with a particular hypervisor software vendor’s or storage hardware vendor’s revenue targets. The result is a proliferation of SDS architectures that work with only one hypervisor, or with only a select few, or that use the software-defined storage software layer to parse and manage only value-add functions, while leaving the hardware vendor to control capacity, vs controlling services and capacity from the software layer. As a result, for virtualization planners, a key challenge is often just figuring out the vendor’s peculiar definition of software defined storage, what it actually delivers in terms of functionality, and what it relies on the storage array vendor to provide. Like VMware, Microsoft has its own SDS architecture, which this paper will summarize and assess.

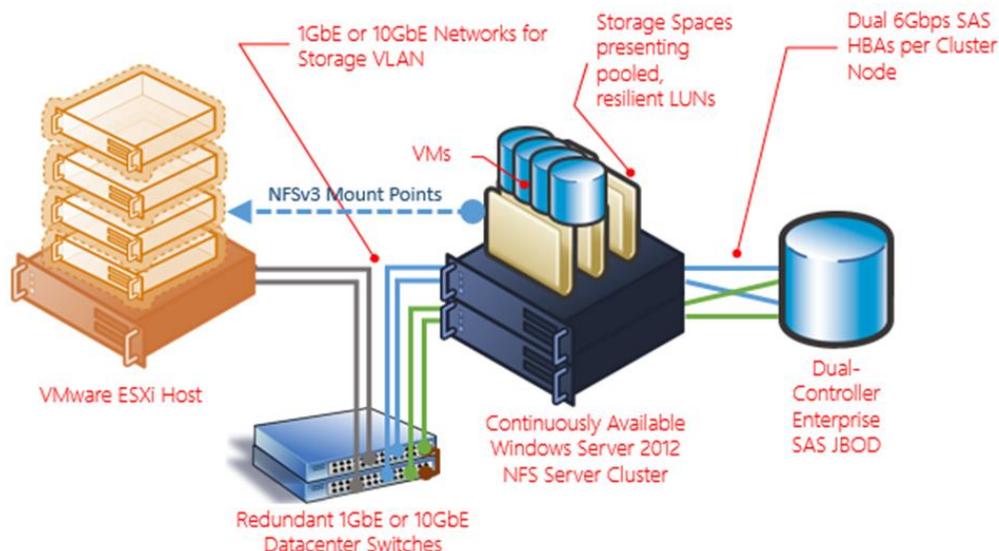
In a Hyper-V World

Microsoft's current software-defined storage solution is aimed at server infrastructure that has been virtualized using its Hyper-V technology. In 2012, Redmond began discussing its architecture for storage behind virtual servers mainly in the context of VMware's many approaches for dealing with its storage problems. Interestingly, Redmond's marketing materials villainized rival VMware's server hypervisor as the source of its application performance and storage I/O processing problems, claiming that software problems required customers to purchase more expensive storage solutions.

For its own part, Microsoft introduced a new storage technology, based on a storage concept introduced for non-virtualized servers in Windows Server 2012 called "Storage Spaces." They called their SDS variant of the technology "Clustered Storage Spaces" and claimed that it would dramatically reduce storage costs by eliminating the need for expensive brand name hardware.

Clustered Storage Spaces, Microsoft claimed, could be used to build massive storage behind Hyper-V servers that could scale without incurring the high price tag of brand name storage. The resulting infrastructure would use SAS disk and would support server side Flash and DRAM memory storage. Moreover, it would avail itself of Microsoft's SMB 3.0 protocol for ease of application access while providing nearly the same performance as Fibre Channel. As a plus, it would support certain functionality that was usually found on enterprise array controllers, like thin provisioning and de-duplication.

Basically, the resulting storage could be presented as a set of pooled resilient LUNs (or storage spaces), and you could leverage Windows clustering services and Windows file systems, such as NTFS, to make the infrastructure resilient and highly available and to present its storage as a general purpose file server repository.

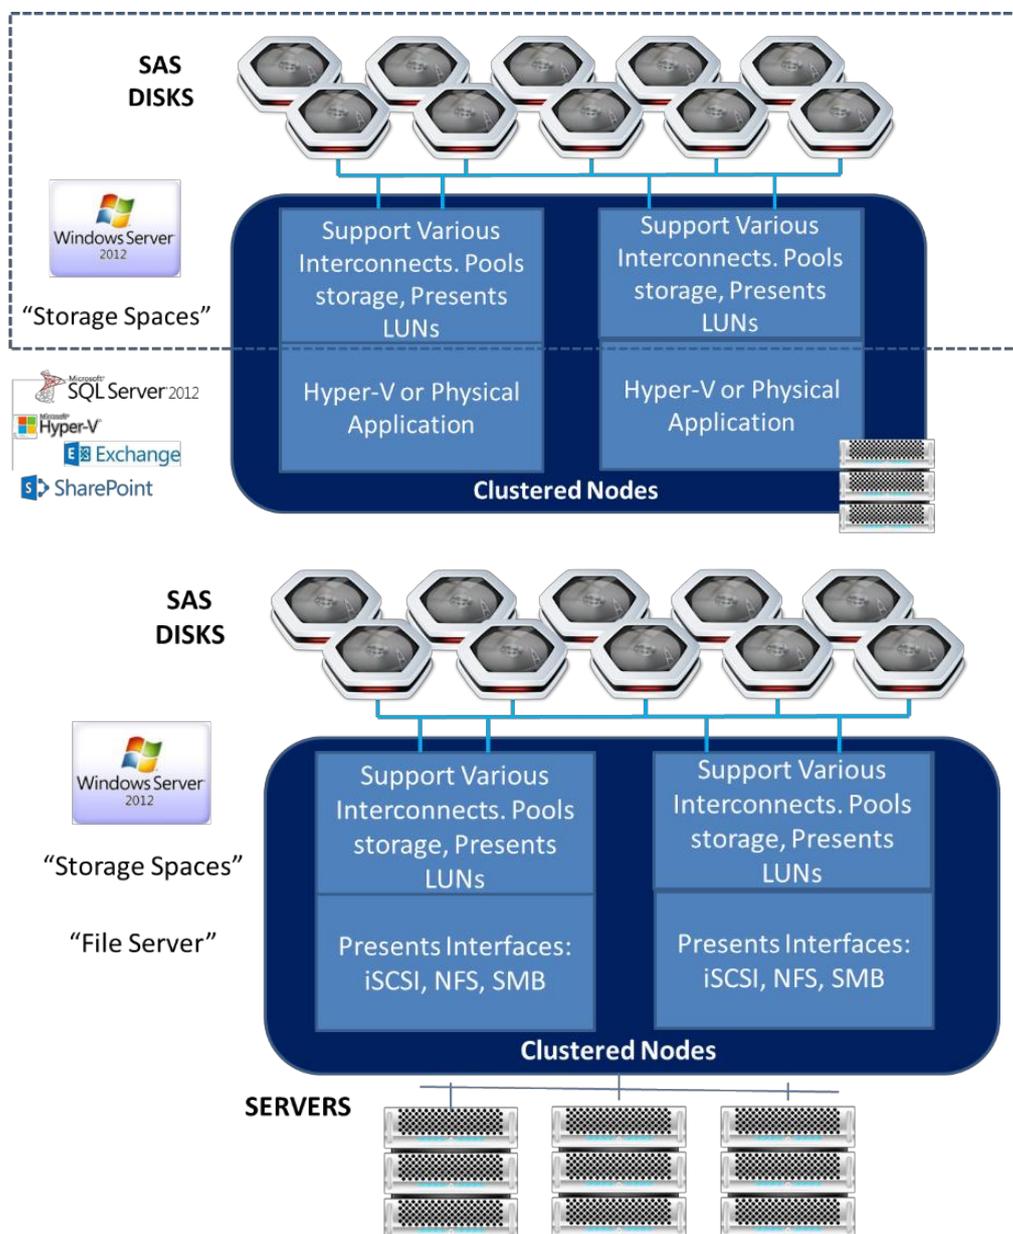


An early illustration. Source: blogs.technet.net

The resulting infrastructure, in one writer's words, was a "virtual storage appliance" that could be built using existing server hardware, Windows server software, and "any storage that can be mounted to the server." As a bonus, Microsoft provided a migration tool for sharing its storage with VMware servers – actually, for converting and migrating Virtual Machines formatted to run with VMware hypervisors to Microsoft Hyper-V (and with no additional licensing fees!).

Said Redmond's evangelists, the resulting storage avails itself in ongoing storage improvements that Redmond is making to its SMB 3.0 protocol, including SMB Transparent Failover, SMB Direct and SMB Multichannel. The chorus of analysts engaged by the vendor added their estimates of Clustered Storage Spaces business value, stating that Storage Spaces over SMB could deliver nearly the same performance (within 1-4%) as that of traditional FC storage solutions, but at approximately 50% of the cost.

There are different views of this total cost of ownership claim, of course, but the important points about Clustered Storage Spaces was that the architecture supported both block and file data storage. Moreover, in its role as a scalable file server, its storage could store virtually any file data from any application, whether virtualized by Microsoft or not.



Providing both block storage (to accommodate applications such as Microsoft SQL Server) and file storage, Microsoft claimed that Clustered Storage Spaces qualified as software-defined storage.

Their strategy, after all, created a software-based controller that featured a robust set of storage services once found only on array controllers, including

- Different resiliency levels (simple or stripe, mirror or parity) per Storage Space
- Thin provisioning
- Sharable storage
- Memory-based read caching
- Active-Active clustering with failover
- SMB copy offload
- Snapshots and clones of Microsoft Virtual Hard Disks (VHDs) on storage pool
- Two and three way mirroring
- Fast rebuild of failed disk
- Tiering of “hot data” to solid state disk (SSD)
- Write back caching - write to solid state disk (SSD) first, then to SAS
- Data de-duplication

Pushback Reveals Issues and Limitations

Almost immediately following the release of Clustered Storage Spaces, pushback from engineers and operators commenced. Several issues have garnered significant attention.

For one, according to software-defined storage “purists,” Microsoft’s Clustered Storage Spaces didn’t really qualify as SDS, since it did not fully divorce the storage control layer from the hardware. Detractors observed that the technology placed a heavy hardware requirement on infrastructure design. Only SAS (Serial Attached Storage protocol) disks and disk arrays could be used. Technically speaking, this doesn’t completely segregate storage services from storage hardware, which makes this cobble only partly hardware agnostic.

Another criticism held that the active-active high availability clustering touted by Redmond was only possible if the clustered nodes were using only external SAS JBOD arrays as their storage. If a node used internal storage – storage on a server – as part of its storage pool, active-active clustering was not possible, only the less available active-passive clustering could be used. While active-passive clustering can provide reasonably effective availability guarantees, it does not enable the “seamless” or “transparent” re-hosting of virtual workload that the active-active cluster advocates want. This situation also reflected a hypervisor dependency of the SDS solution.

Other issues had more to do with the efficiency of the solution than to its SDS purity. For example, many criticized how Clustered Storage Spaces used flash storage, both in the context of write caching and especially in performing data de-duplication processing. In both cases, flash memory was not being used in a manner that minimized wear. It is well understood that flash technology is vulnerable to memory wear: when a memory cell in a cell group on a flash device reaches its maximum writes – somewhere around 250,000 in multilayer cell flash memory – the cell fails and the entire group of cells to which it belongs must be retired. Over time, enough dead or unusable cell groups translates to a flash device of extremely limited capacity that needs to be replaced. The only workaround to this vulnerability is to limit the number and frequency of small block writes to the flash device, usually through write coalescence technology, which Microsoft does not offer.

Given that Microsoft Clustered Storage Spaces technology offers no write management, it can be argued that it does not use flash technology efficiently. Moreover, it writes small block data to silicon in preparation for de-duplication, a further burden on flash durability that caused many engineers to question the economics of Clustered Storage Spaces in server environments using substantial flash technology.

In addition to its implementation of de-duplication functionality, there are also concerns over Microsoft's thin provisioning methodology, which requires storage to be overprovisioned – for planners to buy and deploy more than they need – so that they can respond to a burgeoning space problem when it is reported. Similarly, Microsoft's much touted parity striping functionality is reported to require operators to embrace two-way mirroring for performance reasons. Two-way mirroring, critics argue, will most certainly increase storage capacity demand requiring the purchase of additional SAS disk and increasing the cost of storage infrastructure by as much as 300 to 650 percent, depending on which industry analyst one consults.

The Fix

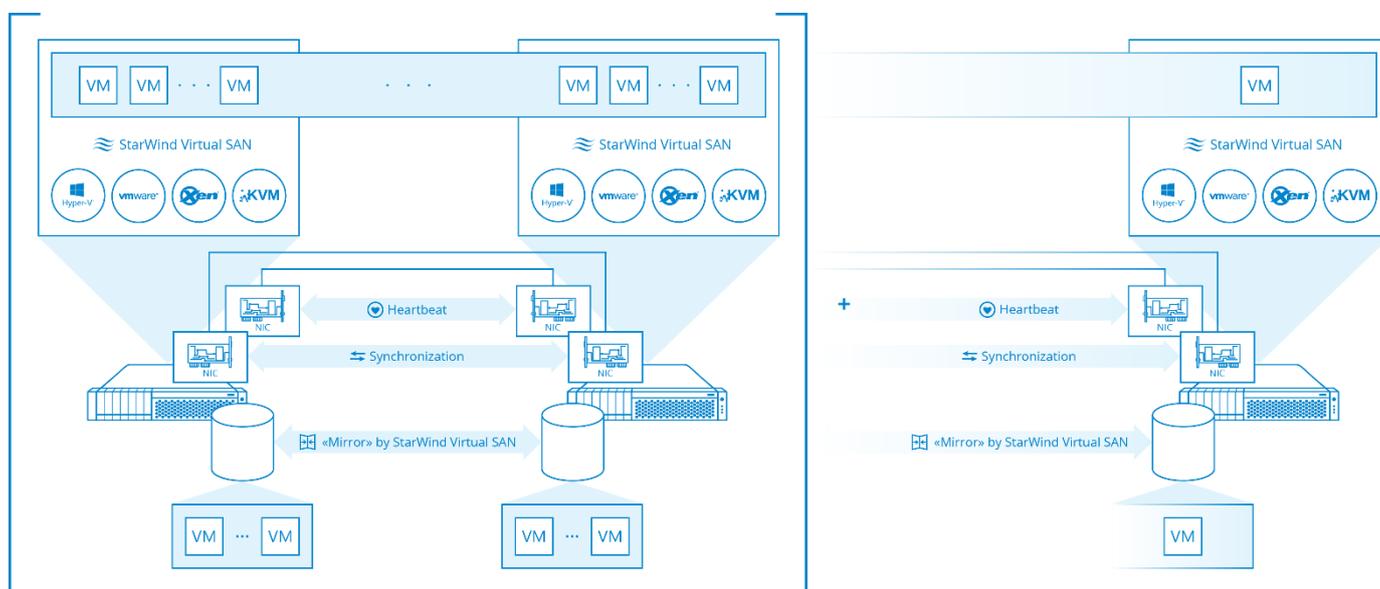
Microsoft recently announced that it is making changes to Clustered Storage Spaces as part of Windows 10 Server. Spokespersons are making cryptic references in the trade press to “Storage Spaces-Shared Nothing” architecture. Details are sparse, but early statements suggest that the new approach will support not only SAS, but also SATA disk, plus PCIe and NVMe flash devices. However, according to one account, deployment will be in a “compute and storage separated model” with a fairly rigid topology requirement: a minimum of four storage nodes, plus two Hyper-V server nodes, will be required to build a shared-nothing SDS infrastructure, which by itself will increase the cost of the resulting infrastructure. Microsoft has not offered any details regarding new storage services or improvements to storage service issues cited above in the new SDS offering, which is now slated for delivery in 2016.

By contrast, an SDS solution that exists today and that complements the Microsoft SDS approach while obviating many of the limitations of Clustered Storage Spaces is the Virtual SAN from StarWind Software.

StarWind's offering deploys effortlessly into a Microsoft Server/Hyper-V environment where it delivers enhancements to Clustered Storage Spaces that include

- Support for SATA, PCIe Flash, and 40+GbE connections
- Real active-active clustering of nodes configured with mixed internal storage and external JBODs
- A log structured file system to enable efficient write caching ahead of flash storage, reducing write frequency and wear rates and preserving flash investments
- Support for legacy storage gear

Essentially, with StarWind Software's Virtual SAN, the storage used by Clustered Storage Spaces can be configured and presented by StarWind, alleviating many of the known constraints in Microsoft's current storage model.



Going forward, StarWind Software's Virtual SAN already provides the capabilities that Microsoft is promising with its next generation SDS solution. StarWind already provides shared nothing clustering, and in more economical two and three node footprints.

While Microsoft (like VMware) appears to be abandoning the small to medium business and remote office/branch office markets, which require simple yet scalable storage solutions, in favor of more complex offerings better suited to "large enterprise" and "cloud" hosting environments, StarWind supports the smaller shops with tighter budgets by delivering a Virtual SAN today that is affordable, easily deployed, and open to both white box and branded storage products. The StarWind offering also supports hyper-converged storage appliances that may be appropriate to departmental settings or other use cases where storage requirements are growing at a slower pace.

If you are considering software-defined storage for your Microsoft computing environment, StarWind Software's Virtual SAN is worth a look. For more information, or to try out a complementary technology for software-defined storage in a Microsoft Hyper-V environment, contact StarWind Software.

Contacts

US Headquarters	EMEA and APAC
 1-617-449-7717	 +44 20 3769 1857 (UK)
 1-617-507-5845	 +49 302 1788 849 (Germany)
	 +33 097 7197 857 (France)
	 +7 495 975 94 39 (Russian Federation and CIS)
	 1-866-790-2646

Customer Support Portal:

<https://www.starwind.com/support>

Support Forum:

<https://www.starwind.com/forums>

Sales:

sales@starwind.com

info@starwind.com

General Information:



www.starwind.com **StarWind Software, Inc.** 35 Village Rd., Suite 100, Middleton, MA 01949 USA
©2015, StarWind Software Inc. All rights reserved.