

Microsoft Storage Spaces Direct

Hyperconverged Infrastructure

March 2019



Evaluator Group

Making the best decisions for Information Management

Reproduction and distribution permitted to Axellio, Inc.

Microsoft Storage Spaces Direct

Overview

Microsoft Storage Spaces Direct (S2D) is a software-defined storage solution that's included with the Windows Server 2019 Datacenter Edition operating system. When run on industry-standard server hardware with commodity storage devices, S2D can provide the shared capacity for an HCI cluster supporting virtual machines on Hyper-V, which is also included in Windows Server 2019. S2D runs in the Windows Server 2019 kernel, not as a VM like most other SDS solutions, providing potential performance and cost advantages.

HCI solutions using Windows Server 2019 are offered by multiple hardware vendors under the Windows Server Software Defined (WSSD) program, including Dell EMC, HPE, Fujitsu, Lenovo, NEC, Supermicro and others. Although used in that context in this Product Brief, Storage Spaces Direct isn't the name of an HCI product, per se. Technically, it is the SDS layer within the Windows Server 2019 OS that's being used by these vendors to create their specific HCI solutions.

Clusters contain a minimum of two servers, a maximum of sixteen, with a max capacity of 4PB (raw). All servers must have the same number of drives and Microsoft recommends that server nodes be of the same make and model. S2D nodes can be configured with either storage or compute functions on each node, enabling the cluster to scale storage and compute resources independently. Microsoft calls this the "converged" mode. S2D can also be set up with both storage and compute functions residing on the same node, called the "hyperconverged" mode.

S2D uses Cluster Shared Volumes (CSV) as a clustered file system layer. CSV performs metadata synchronization and I/O forwarding between nodes using the SMB (Server Message Block) protocol. S2D is built on the Storage Bus Layer (SBL), a virtual bus architecture that creates a fabric connecting all disks across all nodes with using SMB as the protocol transport. With a feature called SMB Direct, SBL can use RDMA-enabled NICs (iWARP or RoCE) over 25GbE (recommended for 4+ node clusters), although SMB Direct can support smaller and larger bandwidth networks.

In addition to RDMA connectivity, SBL supports persistent memory (Intel Optane DC/3D XPoint) providing extremely low latency storage performance. S2D also derives performance benefits from

Highlights

- Disaggregated HCI solution offers scaling efficiency
- Scales from 2 – 16 nodes, up to 4PB
- Included in Windows Server 2019, Datacenter Edition, with Hyper-V
- S2D offers RDMA performance and support for persistent memory
- Dedupe, compression, synch and asynch replication

running in the Windows Server kernel (“kernel embedded”), compared with HCI software-defined storage layers that run as a VM.

Storage Spaces Direct offers multiple options for creating data resiliency. One uses a distributed RAID model, writing two or three copies (mirrors) of each data block to different devices in different nodes. Another, “nested resiliency”, is similar to a RAID 5+1, providing protection against multiple hardware failures on a two-node cluster. A distributed parity model (erasure coding) offers a more space-efficient way to protect against one or two simultaneous device failures.

S2D has post-process deduplication and compression to optimize capacity, plus asynchronous and synchronous volume replication with Storage Replica to support remote data protection and disaster recovery. S2D does not support stretched clustering. For DR, Storage Replica can provide crash-consistent copies, to another cluster in a remote site or to Azure cloud, but does not automate the failover process. A storage quality of service (QoS) feature allows the creation of policies that specify a minimum and/or maximum performance level for VMs or virtual hard disks.

The Resilient File System (ReFS) is included in Windows Server 2019 and integrated with S2D. It can automatically repair detected data corruption using the appropriate data block copies maintained in the cluster. ReFS has the features real-time tier optimization, block cloning and variable cluster sizes, to improve performance, plus deduplication.

Usage and Deployment

In the converged mode, Storage Spaces Direct can be deployed as a scale-out storage system with compute and storage in separate nodes. With a scale-out file system layer running on top of S2D, it provides SMB file shares to other clusters running Hyper-V or other applications that support SMB. This “disaggregated” architecture enables higher capacity clusters and more efficient expansion as storage nodes can be configured with more storage and smaller CPUs.

With S2D and Hyper-V included in the Windows Server 2019 Datacenter Edition, Microsoft offers an HCI solution with a potential for cost savings, compared with other products that require a license for the HCI software and a hypervisor.

- Characteristics
 - Hyperconverged software solution running in Windows Server 2019 Datacenter Edition
 - Disaggregated architecture enables more efficient scaling
- Performance - SMB 3.0-connected, scale-out storage
 - 10GbE, 25GbE (and higher) RDMA
 - Storage Bus Layer Cache optimizes multiples types of storage devices (SSD, NVMe)

- SBL Cache supports persistent memory
- S2D runs in server OS, not in hypervisor or as VM
- Features and functionality
 - QoS – policies to set min and max performance levels for VMs and vDisks
 - Replication – synchronous and asynchronous between remote locations
 - Storage Optimization - Post-process deduplication and compression
 - Resiliency – supports multiple RAID schemes, incl RAID 5+1 “nested resiliency”
 - Resilient File System (ReFS) – scale-out file system integrated with S2D
- Deployment and Administration
 - Disaggregated architecture can be deployed in two modes
 - Converged Mode – creates scale-out storage systems with storage nodes separate from compute nodes, enabling more efficient and greater capacity scaling
 - Hyperconverged Mode – created HCI with storage and compute on same nodes, using Hyper-V
 - Licensing - S2D included in Windows Server 2019 Datacenter Edition

Axellio Deployment of Microsoft as an HCI

Axellio is a Colorado-based developer of Hyperconverged Infrastructure (HCI) solutions targeting enterprise scalability and simplicity where storage IO-intensive applications and workloads demand more from their infrastructure. Axellio is a certified solution provider in the Windows Server Software-Defined (WSSD) Partner Program with the FabricXpress WSSD, a 2U, dual-node HCI appliance that runs the Windows Storage Spaces Direct software-defined storage platform.

FabricXpress WSSD can be configured in both converged and hyperconverged modes, fully leveraging the disaggregated architecture of Storage Spaces Direct. It supports up to 36 NVMe SSDs per node, as well as persistent memory and RDMA connectivity, enabling its use for high-performance workloads. As a solution provider, Axellio’s focus is on the optimization of hardware and software for greater efficiency and performance, while simplifying deployment and support. The product descriptions, usage and deployment information, EvaluScale table and Evaluator Group opinions about Microsoft Storage Spaces Direct expressed in this Product Brief fully apply to the Axellio FabricXpress WSSD.

Evaluator Group EvaluScale™: Microsoft Storage Spaces Direct

Evaluator Group product review methodology “EvaluScale” assesses each product within a specific technology area. The definitions of the criteria and explanations of how products are reviewed can be found in the [Evaluation Guides](#).

| | Criteria | Description | EG View of MS Storage Spaces Direct | Explanation for MS Storage Spaces Direct |
|----|---------------------------------------|--------------------------------------------------------------------------------------------|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | Performance | Resources and design elements to maximize performance | Exceeds requirements | SMB Direct with RDMA, storage bus cache, disaggregated arch, S2D s/w stack runs in Win kernel, not as VM, persistent memory support |
| 2 | Economics | Effective cost of system and operation (CapEx and OpEx), incl data reduction | Exceeds requirements | Cost savings as a part of Windows OS, Hypervisor included, disaggregated arch allows more efficient capacity growth |
| 3 | Scalability | Total storage and compute capacity (low to high), flexible, efficient expansion | Exceeds requirements | Scales 2 - 16 nodes (4PB) per cluster, disaggregated arch scales storage independently from compute, seamless node upgrade |
| 4 | Management Options | Hypervisor mgmt integration, internal/ISV mgmt package, automation, VM-level mgmt | Meets requirements | Windows Server 2019 management tools, MS Systems Center integration, Windows Admin Center for SMB |
| 5 | Advanced Features | Leverage functionality enabled by SDS, provide unique or adv efficiency or mgmt features | Meets requirements | Resiliency options, HCI and disaggregated configuration, QoS, high performance features, RDMA |
| 6 | Vendor/Product Stability | Company viability, product longevity, impact of OEM partnerships | Meets requirements | New product, relatively new tech, very stable company (Microsoft), evolving deployment options/ecosystem |
| 7 | Data Protection / Business Continuity | Device- and node-level failure protection, BU/DR options, stretched clusters, HA, failover | Area for development | Mirror and multiple parity options, nested resiliency, sync and asynch replication, but no stretched clustering, Win Server BU, DP Manager, Cloud BU |
| 8 | Platform Options | Hypervisor support, container support, related features | Area for development | Runs in Windows Server 2016 kernel or Hyper-V, supports containers, Windows runs Azure Pack cloud platform |
| 9 | Cloud Options | Private cloud features/functionality, public cloud connectivity for hybrid deployments | Exceeds requirements | Connectivity to Azure public cloud, Azure BU, Azure Site Recovery with Hyper-V and Storage Replica, Azure Migrate, Azure Stack, Systems Center self-service portals |
| 10 | Model Options | Variety of available appliance models and configurations | Exceeds requirements | SMB Direct with RDMA, storage bus cache, disaggregated arch, S2D software stack runs in Windows kernel, not as VM |

Evaluator Group Opinion: Differentiating Elements for Microsoft Storage Spaces Direct

This is an innovative product. Its disaggregated architecture addresses one of the fundamental issues with hyperconverged technologies, the inability to scale storage capacity independently from compute. Besides improving capacity efficiency, disaggregation also allows clusters to scale larger, making them better suited for cloud infrastructure environments.

By including the hypervisor and the software-defined storage in the server operating system, S2D could drive down the price points for hyperconverged products and offer better margins for channel partners that provide hardware for S2D solutions. Overall, the concept of running an SDS layer in the server OS, instead of as a virtual machine dependent on a hypervisor, is the direction we see cloud infrastructure designs moving, long term.

The RDMA connectivity of Storage Spaces Direct makes hyperconverged solutions appealing for higher performance use cases, those that IT professionals may not have considered for traditional iSCSI-based HCI products. This increased performance capability also means an S2D-based solution could support more VMs than other HCIs, for a given configuration, improving its economics.

While the performance that S2D brings may open this product to some more extreme use cases, the real opportunity for Microsoft is to show that performance can drive up VM density (the number of VMs that can be supported on each node) for a given workload mix, which will drive down cost proportionally. For users this means fewer nodes to buy to support a given set of workloads and a smaller cluster to manage. The latest version of S2D in Windows 2019 adds more performance, dedupe, improved resiliency and larger max capacity – but still lacks stretched clustering.

Although it was released with Windows Server 2016, S2D hasn't created the disruption we expected. Part of that was probably due to the strength of VMware as the hypervisor standard, but also the lack of a strong support program for getting hardware vendors on board with Windows Server 2016 and S2D as an HCI solution.

The Windows Server Software Defined program may address these issues, as now over a dozen vendors are selling HCI solutions using S2D. Microsoft may still face some challenges competing in the enterprise storage arena, as its history and its ecosystem are more closely aligned with the server market. There are also the obvious limitations of being a Windows-only solution.

Information that is more detailed is available at <http://evaluatorgroup.com>

Copyright 2019 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.