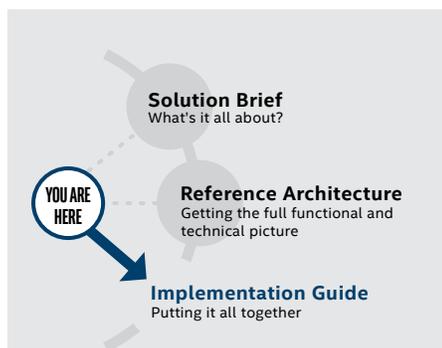


# Intel® Optimized Configurations for Windows Server\* 2016 with Storage Spaces Direct

## Implementing Hyperconverged Software-Defined Storage Solutions



### What You'll Find in This Implementation Guide

This implementation guide provides a starting point for Windows Server\* 2016 with Storage Spaces Direct (S2D) on the latest Intel® technologies for software-defined storage solutions.

#### If you are responsible for:

- **Investment decisions and business strategy ...**  
You'll discover how upgrading to Windows Server 2016 and the latest Intel® Xeon® processor-based servers can help your business become more agile and efficient.
- **Figuring out how to implement and manage data center infrastructure ...**  
You'll learn how an Intel based Windows Server 2016 architecture can transform your data center through software-defined, resilient, flexible, and automated infrastructure.

### Introduction

Today's enterprises are looking to IT to implement technologies that provide greater agility and efficiency, manage the growing volume of data, and quickly roll out new value-added business services. Cloud computing is now a proven delivery model that can address these challenges. As cloud technologies continue to mature, businesses can identify and enable more workloads that fit best in on-premises cloud environments while embracing hybrid models for even greater flexibility and service delivery.

Modern architectures designed to support cloud environments and enable efficient data access make it possible for enterprises to address service delivery on any platform and device, including big data analytics, the Internet of Things, and high-performance computing.

Intel and Microsoft are committed to simplifying delivery of cloud services as part of an overall cloud strategy. Together we're developing more resilient, reliable, and scalable infrastructure solutions that optimize both hardware and software capabilities and drive greater flexibility and automation for the data center through software-defined infrastructure.

### Modern Storage Architectures, More Efficient Data Management

IT decision makers must make smart investments in technologies that can manage the rapid, exponential increases in data that fuel cloud services and the workloads they support. Data storage must become higher performing, more efficient, and cost-effective; provide greater flexibility in the way it manages data; and support dynamic resource allocation and expansion. And IT faces real pressure to evaluate the most cost-effective way to deliver the services they need, through either on-premises environments or third-party clouds.

Windows Server\* 2016 running on the latest Intel® technologies can support the evolution of your cloud strategy and deliver fast, on-demand storage that manages the velocity of today's data-driven business innovation. Higher performance as well as networking and storage enhancements make it possible for data centers to better utilize resources by lowering total cost of ownership (TCO) and ensuring scalability to meet evolving business needs.

This guide provides a starting point for implementing Windows Server 2016 Datacenter Edition on the latest Intel technologies. It focuses on the deployment of Storage Spaces Direct (S2D) for software-defined storage (SDS) solutions.

## Table of Contents

- Software-Defined Storage Solution Overview** ..... 2
- Solution and System Requirements** . . . 5
- Operation and Utilization of the Microsoft\* and Intel Hyperconverged Software-Defined Storage Solutions** 8
- On-Premises Cloud Use Case** . . . . . 15
- Summary** . . . . . 15
- Appendix A: References** . . . . . 16
- Appendix B: Cleaning Your Drives for Deployment** . . . . . 18

## Software-Defined Storage Solution Overview

Intel and Microsoft offer three recommended configurations for software-defined storage (SDS) that take advantage of the new Storage Spaces Direct (S2D) capabilities. They provide options for performance, reliability, and agility from low-latency, input-output operation (IOP)-sensitive business processing to capacity-hungry data warehousing.

All three are hyperconverged configurations for compute and storage that optimize software and hardware. Windows Server 2016 also supports a disaggregated configuration not included in this guide.<sup>1</sup>

### Basic Components

The basic components of the three configurations include:

- The Windows Server 2016 Datacenter Edition operating system
- A standard server based on the Intel® Xeon® processor E5-2600 v4 product family
- Storage media—Intel® Solid State Drive (SSD) Data Center (DC) Family, including SATA drives and PCIe drives using the Non-Volatile Memory Express\* (NVMe\*) interface
- Network interface card (NIC) supporting 10 gigabit Ethernet (10 GbE) or higher (remote direct memory access [RDMA] recommended)

### About Windows Server 2016 Datacenter Edition

Windows Server 2016 provides the foundation for implementing software-defined compute, networking, and storage capabilities in your data center. S2D, an evolution of Storage Spaces first introduced in Windows Server 2012, extends the software-defined capabilities of the earlier version. With S2D, you can build distributed, hyperconverged storage systems using local storage (that is, server-based storage), which simplifies deployment and management of software-defined storage systems and opens up new possibilities for data center innovation. Benefits include:

- Flexibility for pooling storage using SATA and PCIe SSDs and hard drives
- Increased scalability by adding server nodes as needed
- Simplified management of a single storage pool rather than individual devices
- Automatic storage load rebalancing as storage scales
- Greater flexibility for storage in private clouds

### Pools of Shared Storage with Tiering for Performance and Capacity

Storage Spaces Direct creates pools of shared storage accessible to any application running on any of the nodes. Storage deployments can include a mix of storage media (SSDs and hard disk drives [HDDs]) as well as a variety of media protocols such as NVMe, Serial Attached Small Computer System Interface (SAS), and SATA, making it possible to create tiers of physical storage.

## Top New Capabilities in Windows Server\* 2016

### Software-Defined Data Center (Compute, Storage, and Network)

- Storage Spaces Direct
- Simplified Remote Desktop Services deployment in the Microsoft\* Azure\* platform
- Centralized control of network policies
- Hyperscale load balancer from the Azure platform

### Application Platform

- Windows\* Containers
- Nano Server installation option

### Management

- Windows PowerShell\* enhancements

### Security

- Active Directory\* improvements
- Identity enhancements
- Shielded Virtual Machines (VMs)

Find out more. Watch “[The Top 10 Reasons Why You’ll Like Windows Server 2016.](#)”

<sup>1</sup> A disaggregated deployment uses different clusters to separate the Hyper-V\* servers (compute component) from the Storage Spaces Direct servers (storage component). The hyperconverged deployment scenarios run Hyper-V (compute) and S2D (storage) components on the same cluster.

## Enhanced Software-Defined Storage Capabilities

In addition to pools of shared storage, new S2D capabilities include enhanced resilience options, automatic configuration, virtual machine (VM) implementation, Server Message Block (SMB) 3.0 intranode communication, and small deployments.

**Automatic configuration.** Storage pools and tiers are automatically configured into Enable-ClusterS2D (Enable-ClusterStorageSpacesDirect), creating a single storage pool with all eligible storage devices and tiers that reflect the storage configuration of the system. You will receive notification and explanation for any noneligible devices—those without acceptable media types, data-center-class proper bus types, or existing partitions.

**Data resilience from server to chassis and rack.** S2D uses drives that are connected to one node of a Windows Server 2016 cluster and allows Storage Spaces Direct to create pools of capacity using those drives. Virtual disks configured on the pool have redundant data (mirrors or parity) spread across the drives in different nodes of the cluster, allowing access to data even when a node fails or is shut down for maintenance. S2D also extends fault tolerance from server failure to more specific chassis and rack failures to help protect against power loss inside a chassis or rack with shared infrastructure.

**Management.** System Center Virtual Machine Manager (VMM) enables the creation of shared storage pools, carves out volumes, and deploys VMs on the cluster. Simplified workflows within System Center VMM can deploy new clusters or bring existing S2D-enabled clusters under management.

**Server Message Block 3.0.** Storage Spaces Direct uses the SMB 3.0 protocol for intranode (also called east-west) communication. It takes advantage of the powerful features of SMB 3.0, such as SMB Direct (RDMA-enabled NICs) for high-bandwidth and low-latency communication, as well as SMB Multichannel for bandwidth aggregation and network fault tolerance.

**Small deployments.** S2D can be used in smaller deployments (two or three servers), although only mirrored resiliency is supported. Parity resiliency or multiresiliency requires four servers. Deployments with four or more servers deliver greater resiliency and capabilities.

## Virtualization Enhancements

Windows Server 2016 introduces new capabilities for advanced virtual machine management that improve storage performance and efficiency and make it easier to implement software-defined infrastructure.

**Windows Server Containers** isolate applications on a virtual machine by virtualizing the operating system and sharing the kernel for exceptional memory efficiency and high-speed performance. This level of efficiency enables massive density on a single server.

**Shielded Virtual Machines** is a security enhancement that protects against an attack on a single virtualization server running multiple workloads. Shielded VMs safeguard against leaks by only running on a designated fabric. Outside of this fabric, the VMs are encrypted and inaccessible. This capability is a first in the industry.

**Policy for shared storage prioritization** ensures that VMs using the same underlying storage can prioritize applications by specifying the maximum or minimum storage available to any single VM. This keeps higher-priority VMs from being starved for the storage inputs/outputs they need by lower-priority VMs.

Find out more about [Windows Server 2016](#).

## Windows Server\* Software-Defined Program

Microsoft is enabling solution providers to design and certify hardware for enterprise-level software-defined infrastructure solutions running Windows Server\* 2016. The Windows Server Software-Defined Program (WSSD) also provides confirmation that vendors comply with Microsoft best practices for software-defined infrastructure.

## Intel Technologies

The Intel Xeon processor E5-2600 v4 product family offers great flexibility to scale, and enhanced control of your IT environment, with optimizations across compute and storage. Hardware-based Intel technologies help monitor, secure, and orchestrate shared platform resources more effectively.

### Compute Optimizations

The Intel Xeon processor E5-2600 v4 product family accelerates performance across the data center for the largest workloads. With up to 22 cores and 48 threads per processor (a 33 percent increase compared to the prior generation) and up to 60 MB of last-level cache, servers based on this latest processor can provide greater processing power to scale up or out for very large or demanding workloads. Innovative technologies built into the silicon deliver faster transaction speeds, increased bandwidth, lower latency, and reduced power consumption over the previous generation. Combined with Windows Server 2016, servers running the Intel Xeon processor E5-2699 v4 process up to 1.5 times more business transactions per second compared to a 3-year-old server.<sup>2</sup> Windows Server 2016 is also optimized for Intel® Advanced Vector Extensions 2 (Intel® AVX 2), which accelerates the performance of storage efficiency technologies such as tiering.

### Storage Optimizations

S2D pools local storage into highly scalable and available server clusters using an Ethernet network as the storage fabric (RDMA enabled is optional and recommended) and Intel Xeon processor-based servers to power the advanced data services. Intel Xeon processors also support the NVMe interface for faster, lower-latency media connections directly to the CPU, and S2D provides faster input and output operations per second (IOPS), enhanced storage quality of service (QoS), and read-write performance when combined with devices using the NVMe interface. Joining Intel Xeon processors, S2D, and Intel Data Center SSDs using NVMe can significantly cut latency and provide fast IOPS.

### Security Optimizations

Enhanced hardware-assisted capabilities work with Windows Server 2016, software-based key generation, and encryption software to provide increased data protection and security capabilities, including:

- Generation of random, high-quality cryptographic keys and random bits via Intel® Secure Key with RDSEED
- Bulk data encryption and security session initiation protocols, enabling more pervasive encryption without a performance penalty via Intel® Advanced Encryption Set New Instructions (Intel® AES-NI) and cryptographic acceleration
- Secure Boot for guarding hosts with support-host hardware and firmware to deploy Shielded Virtual Machines via Trusted Platform Module (TPM) 2.0 and Unified Extensible Firmware Interface (UEFI) 2.3.1

Find out more information about the [Intel Xeon processor E5-2600 v4 product family](#).

### Intel® Solid State Drive Data Center Family

The Intel SSD Data Center Family for PCIe are NVMe interface storage devices that bring extreme data throughput directly to Intel Xeon processors. Compatible with and optimized for Intel processors, they contribute speed, performance, data integrity, and reliability to keep data highly available and systems up and running.

The Intel SSD Data Center Family for PCIe overcomes SAS/SATA SSD performance limitations by optimizing hardware and software to take full advantage of NVMe SSD technology.

- With NVMe PCIe drives, throughput is increased via a streamlined connection direct to the CPU, delivering up to six-times-faster data transfer speed than 6 Gbps SAS/SATA SSDs.
- The performance of a single drive from the Intel SSD Family for PCIe, specifically the Intel Solid State Drive Data Center P3700 Series (460K IOPS), can replace the performance of seven SATA SSDs aggregated through a host bus adapter (HBA) (approximately 500K IOPS).<sup>3</sup>
- Industry leading data integrity capabilities protect against undetectable errors. Rigorous qualification and compatibility testing ensures a highly reliable SSD with superior annual failure rate levels.

Find out more about the [Intel Solid State Drive Data Center Family for PCIe](#).

<sup>2</sup> Up to 1.5 times more transactions per second comparing an Intel server with two Intel Xeon processor E5-2697 v2, 256 GB memory (source: Intel Technical Report #3049) to an Intel server with two Intel Xeon processor E5-2699 v4, 512 GB memory (source: Intel internal testing).

<sup>3</sup> Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Based on the Intel SSD DC P3500, P3600, and P3700 Series product specifications. Random I/O operations based on thousand (K) IOPS.

## Solution and System Requirements

Windows Server 2016 provides hyperconverged solutions by allowing the same set of servers to provide SDS, through Storage Spaces Direct (S2D), and serve as the hosts for virtual machines using Hyper-V\* technology. To fully take advantage of performance, software-defined storage, and new virtualization features in Windows Server 2016, upgrade to new hardware including the Intel Xeon E5-2600 v4 processor family and high-reliability, high-performance Intel SSDs.

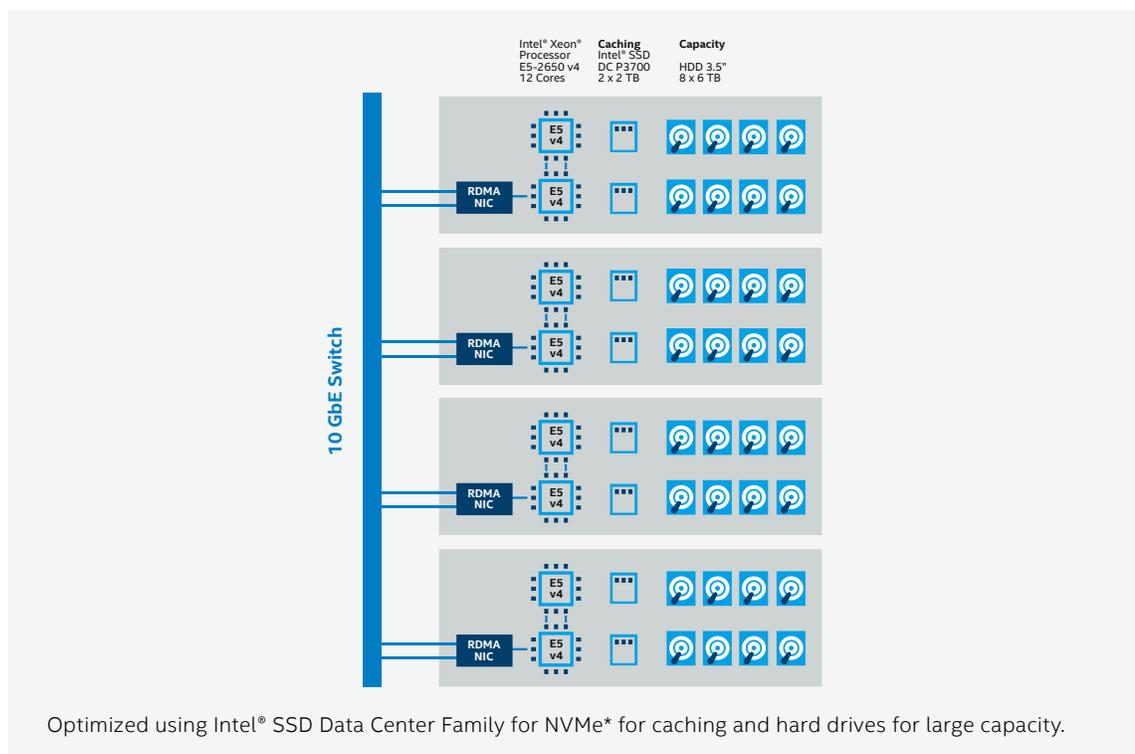
Intel and Microsoft tested three hyperconverged solution configurations. The following lists the solution requirements for each:

- Capacity-optimized hybrid configuration
- Throughput- and capacity-optimized configuration
- IOPS-optimized configuration

For more detail, see the [Intel and Windows Server 2016 reference architecture paper](#).

### 1. Capacity-Optimized Hybrid Configuration (NVMe\* SSD and Hard Drive)

This configuration optimizes dollars per gigabyte (\$/GB) by using PCIe NVMe SSDs for cache and HDDs for high-capacity data storage. Use this approach for workloads such as data warehousing, Microsoft\* Exchange, and Microsoft SharePoint\* technology.

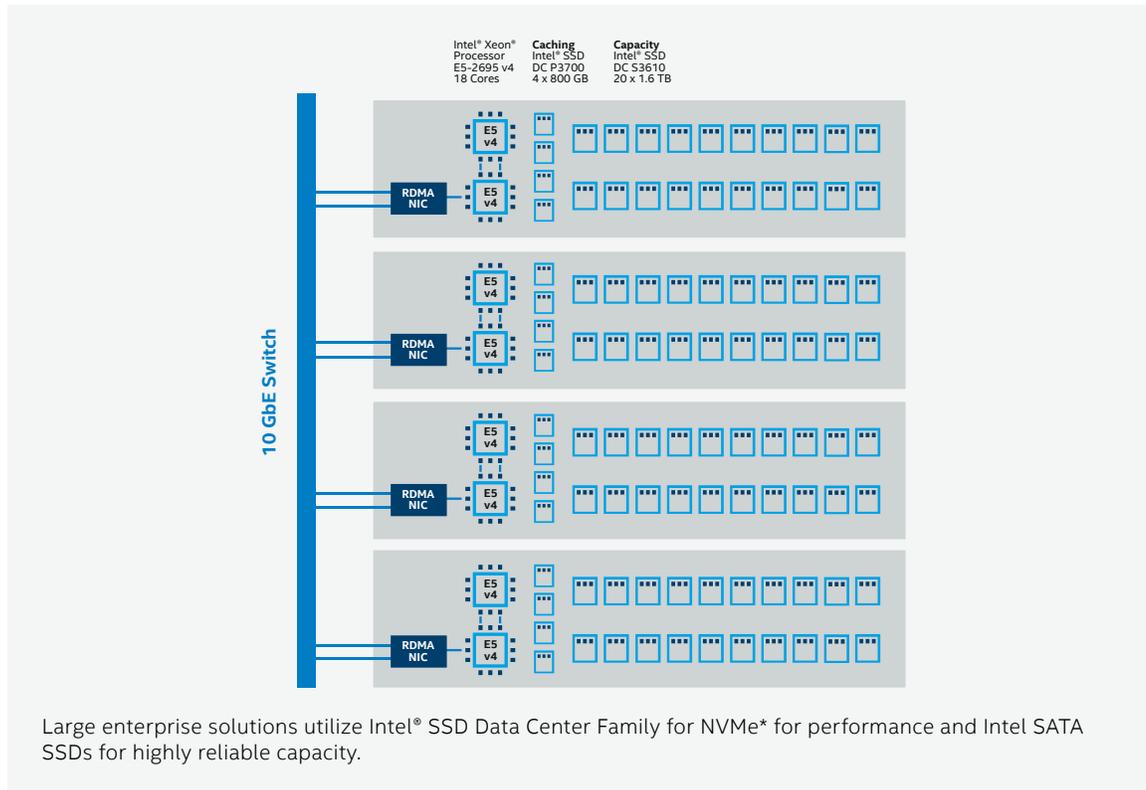


### Compute, Storage, and Network Specifications

- Server: 2U, 1 node
- CPU: Medium-core-count processor, such as Intel Xeon processor E5-2650 v4 with 12 cores
- DRAM: DDR4—16 x 16 GB (256 GB)
- Cache storage: Low-latency, high-endurance SSDs, such as Intel SSD DC P3700 series, 2 x 2.0 TB (4.0 TB)
- Capacity storage: 8 HDD, 3.5 inch with 8 x 6 TB (48 TB)
- NIC: 2 x 10 GbE NIC, RDMA optional
- Switch: 10 GbE switch

## 2. Throughput- and-Capacity-Optimized Configuration (All Flash NVMe and SATA SSD)

This is a balanced configuration that uses both NVMe and SATA SSDs, depending on the performance and latency requirements of the applications and the amount of capacity needed. It works best for supporting a wide range of enterprise applications, including those with larger numbers of users and VMs in Hyper-V virtualized environments.

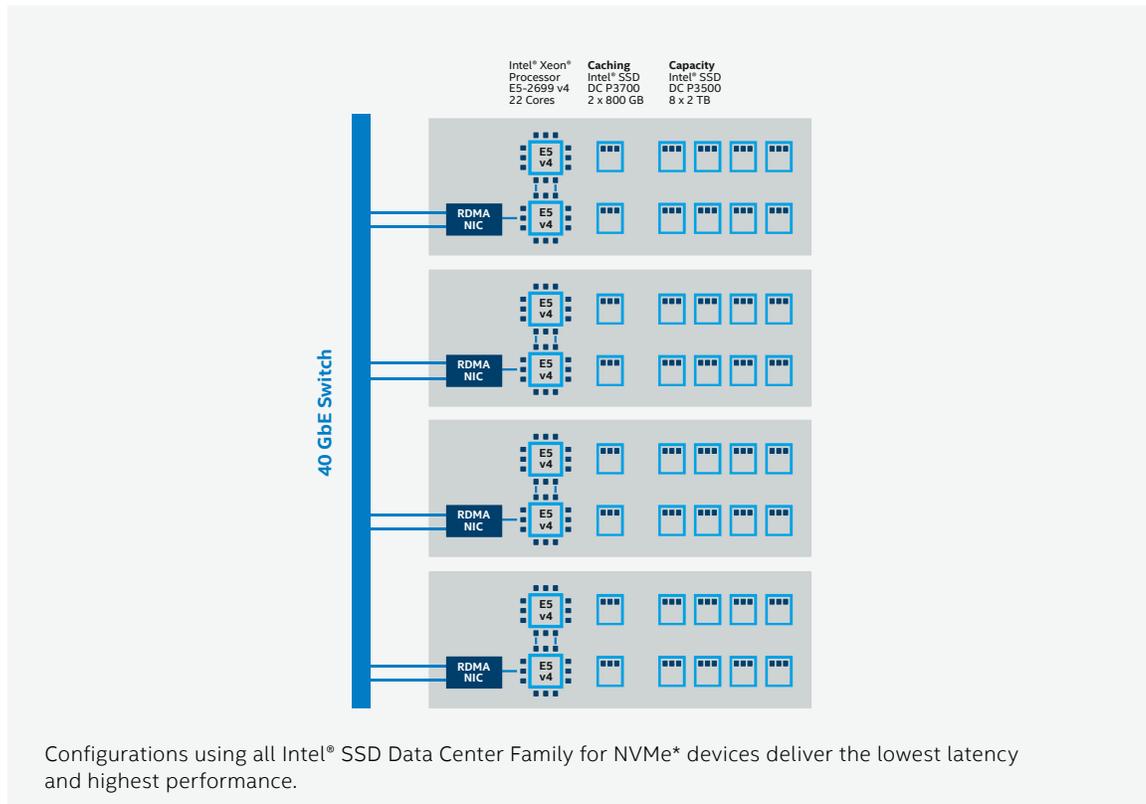


### Compute, Storage, and Network Specifications

- Server: 2U, 1 node or dense 2U, 4 node
- CPU: Optimal higher-core-count processor, such as Intel Xeon processor E5-2695 v4, 18 cores
- DRAM: DDR4—24 x 16 GB (256 GB) minimum; 24 x 32 GB (768 GB) maximum
- Cache storage: Low-latency, high-endurance SSDs, such as Intel SSD DC P3700 series, 4 x 800 GB (3.2 TB)
- Capacity storage: SATA SSDs, such as Intel SSD DC S3610 series, 20 x 1.6 TB (32 TB)
- NIC: 2 x 10 GbE or 40 GbE RDMA NIC
- Switch: 10 GbE or 40 GbE switch

### 3. IOPS-Optimized Configuration (All Flash NVMe SSD)

This configuration is built for workloads that require high performance and low latency, such as online transaction processing (OLTP), virtual desktop infrastructure (VDI), or infrastructure as a service (IaaS). For both the cache tier and capacity tier, we recommend high-endurance, low-latency NVMe SSDs. This contributes to the high reliability of the configuration, with 10-drive-write-per-day endurance on the cache tier, standard endurance on SSDs in the capacity tier, and significantly more CPU efficiency than with SATA SSDs.



#### Compute, Storage, and Network Specifications

- Server: 1U, 1 node or 2U, 1 node
- CPU: High-core-count processor, such as Intel Xeon processor E5-2699 v4 with 22 cores
- DRAM: DDR4—24 x 16 GB (384 GB) minimum; 24 x 32 GB (768 GB) maximum
- Cache storage: Low-latency, high-endurance SSDs, such as Intel SSD DC P3700 series 2 x 800 GB (1.6 TB)
- Capacity storage: Standard-endurance SSDs, such as Intel SSD DC P3520 series or DC P3500 series with 6 x (up to 8 x) 2 TB (12 to 16 TB)
- NIC: 2 x 40 GbE RDMA NIC
- Switch: 40 GbE switch

## Operation and Utilization of the Microsoft\* and Intel Hyperconverged Software-Defined Storage Solutions

All three hyperconverged configurations provide server-based storage and run on industry-standard hardware with local storage media. They deliver high-performance, high-resilience SDS solutions that can replace the need for storage attached networks (SANs). For example, you can use these solutions as the principal storage resource for Windows application environments, as a Structured Query Language (SQL) database, and for data warehousing. The three configurations are appropriate for any use case in which a shared volume makes sense.

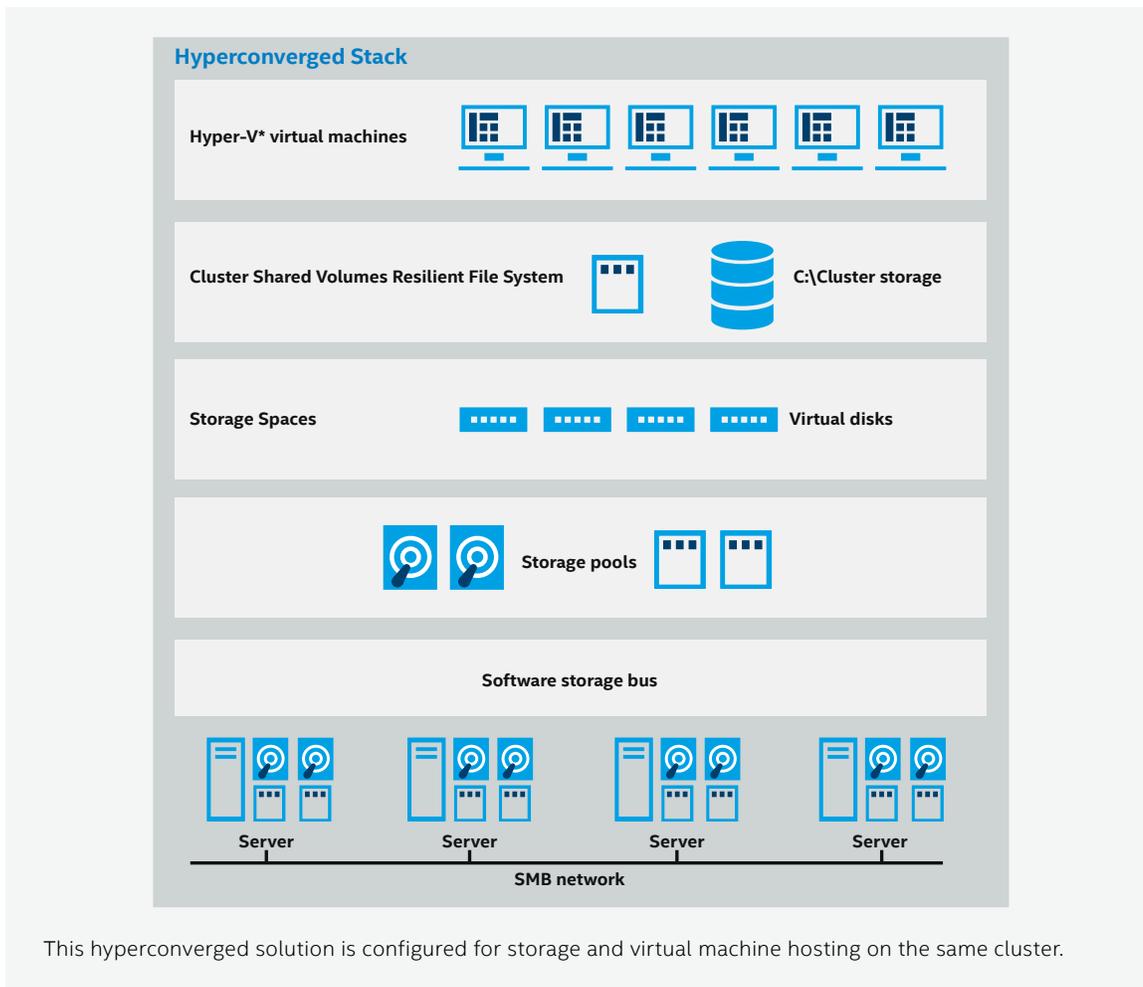
### How the Solution Works

In hyperconverged configurations, Storage Spaces Direct (S2D) integrates with familiar features in the Windows Server SDS stack, including Clustered Shared Volume File System (CSVFS), Storage Spaces, and Failover Clustering.

The S2D hyperconverged deployment uses Hyper-V technology for compute and S2D for storage on the same cluster. VMs are stored on local Cluster Shared Volumes (CSVs), so that both the VMs and the storage they use can be scaled easily. Once S2D is configured and the CSV volumes are available, you can configure Hyper-V technology using the same tools you use for any Hyper-V deployment or failover cluster.

S2D uses an SMB network (SMB 3.0) for cluster communication and data transfer, with each node connected to redundant physical network connections for resilience. S2D does not support drives connected via multiple paths and the Microsoft Multipath MPIO software stack.

You can use in-box tools or other tools to manage the storage and virtual machines, including System Center VMM.



## Options for Installation: Nano Server, Full, or Core Deployment

Windows Server 2016 offers three install types: Full, Core, and Nano Server. Nano Server is a remotely administered server operating system optimized for private clouds and data centers. Similar to Windows Server in Server Core mode, but significantly smaller, Nano Server has no local logon capability and only supports 64-bit applications, tools, and agents. Nano Server takes up far less disk space, sets up significantly faster, and requires far fewer updates and restarts than the other options.

Core installations are typically used for lower-maintenance server environments—for example, traditional VM workloads. Full installations are best used for specialized workloads such as third-party applications and relational database services.

This guide highlights the Full installation option. Find out about the [advantages of using and deploying Nano Server](#).

## Additional Considerations

Your specific implementation will dictate how you address the following additional considerations for capacity growth, performance, and resilience:

**Growing storage capacity and performance.** To increase capacity and performance, add more servers with the required amount of storage.

**Performance.** S2D enables storage tiering and automatically creates a single large shared pool with two storage tiers based on the storage devices in the system. With auto tiering, S2D assigns the highest performing media to the cache and configures the capacity tier with the available mix of device types (for example, read-write, write, write-through, or write-back).

**High availability and resilience.** High availability in S2D is achieved with non-RAID adapters combined with redundancy. In Windows Server 2016, redundancy can be configured in three ways, depending on your workload requirements:

- **Simple spaces:** Simple spaces distributes data across a set of pool drives but is not resilient to drive failures and can be used for high-performance workloads where resiliency is provided by the application or is not required.
- **Mirror spaces:** This type of redundancy distributes and mirrors data across a set of pool drives to support a two-way mirror (resilient to single drive failures) or three-way mirror (resilient to double drive failures). Mirror spaces are appropriate for most workloads in clustered and nonclustered deployments.
- **Parity spaces:** Parity spaces distributes data across a set of pool drives with a single drive write block used to store parity information. Parity is the most capacity-efficient resiliency option. Configure using parity spaces for large-block append-style workloads such as archiving in nonclustered deployments.

## Step-by-Step Deployment of an Intel Optimized Hyperconverged Solution with S2D

Deployment of a hyperconverged solution involves four major steps that must be completed in order, although not all at one sitting. Each step contains smaller tasks. The following step-by-step guide is based on the hardware and software specified in the “Basic Components” section of this guide:

## Intel® Optimized Hyperconverged Solution with S2D: Step by Step

### STEP 1:

#### Prepare the servers and storage media.

- Confirm the latest platform firmware, drive firmware, and drivers.
- Set up and configure your system.

### STEP 2:

#### Deploy the operating system.

- Install Windows Server\* 2016.
- Install Windows Server 2016 roles and features.

### STEP 3:

#### Configure the network.

- Configure top-of-rack (ToR) switch.
- Enable network quality of service.
- Create a Hyper-V\* virtual switch with Switch Embedded Teaming (SET) and remote direct memory access (RDMA) network interface card (NIC).

### STEP 4:

#### Configure Storage Spaces Direct (SD2).

- Run cluster validation tool.
- Create a cluster.
- Configure a cluster witness.
- Use clean drives.
- Enable Storage Spaces Direct.
- Create virtual disks.
- Create or deploy virtual machines.

## Gather the Information You Need

Make sure you have the following information available during your installation:

- **Server names.** You will need to know your organization's naming policies for computers, files, paths, and other resources so you can provide unique names as you provision your servers.
- **Domain name.** You will be joining your newly provisioned servers to your domain. Find out how to set up a new [Windows Active Directory\\* domain](#).
- **Administrator password for the new servers.** You will be prompted to input the password for the local administrator account.
- **For RDMA configurations.** You will also need to know the top-of-rack (ToR) switch make and model; the network adapter make, model, and protocol; and the virtual LAN (VLAN) ID for the two network interfaces used by the management operating system on the hyperconverged hosts.

## Step 1: Prepare the Servers and Storage Media

### *Confirm the latest platform firmware, drive firmware, and drivers.*

For each node, make sure you have the latest versions of the firmware and drivers needed for Windows Server 2016.

You will deploy and configure individual nodes with the local administrator account <ComputerName>\Administrator. Follow these steps.

1. Update Intel SSD firmware. Download the latest copy of the Intel SSD Data Center Tool from [the Intel download center](#).

Note: Make sure you have the latest .zip file. This is updated regularly.

2. Download the latest platform BIOS (check with your server OEM to be sure of the appropriate settings).
3. Download the latest network drivers for your NIC hardware.

### *Set up and configure your system.*

1. Ensure that no external hardware is connected to your system under test (network cards, storage, Ethernet cables, and so forth) other than the operating system disk.
2. Make sure you understand and narrow down all necessary BIOS settings to:
  - Enable virtualization.
  - Manage power (on or off; you may need to experiment).
  - Intel® Hyper-Threading Technology (Intel® HT Technology) (on or off depending on use case).
3. Update the platform BIOS, field replaceable unit/sensor data record (FRUSDR), baseboard management controller (BMC) firmware, and Intel® System Support Utility (Intel® SSU) firmware for all nodes using the recommended BIOS package.
  - Use the Extensible Firmware Interface (EFI) shell to upgrade the BIOS:
    - Copy the BIOS package onto a USB drive.
    - Boot to EFI shell.
    - Map -r to find the mapped file system.
    - Your USB device found should map as file system 0 (Fs0). Run the start-up script.

*Cd to directory*

*Run Startup.nsh*

- Reboot the nodes once with the updated BIOS.
- Reboot again, go into system BIOS, and apply all the necessary BIOS settings.

Note: Make sure the BIOS is in performance mode and that you are applying consistent settings across every node.

## Step 2: Deploy the Operating System (Full Installation)

### Install Windows Server 2016.

The Full option installs the standard user interface (Server Graphical Shell) and all tools, including client experience features. You can use either remote or local management for your installation.

You can install Windows Server 2016 like any other Failover Cluster deployment, including joining the cluster to an Active Directory domain and installing the Hyper-V role and the Failover Cluster feature. Use a remote ISO media mount via the Integrated Management Module (IMM) or a bootable USB or DVD containing the installation content.

Use the Setup wizard to choose between Windows Server 2016 and Windows Server (Server with Desktop Experience). Select Server with Desktop Experience for the Full installation option. (Note: The default is for Windows Server 2016, which installs the Core option.)

### Install Windows Server roles and features.

You can install server roles and features using Server Manager or the Windows\* PowerShell\* framework. With PowerShell, use the Get-WindowsFeature command to determine what is available or may already be installed. For the purposes of an S2D implementation, be sure to install (or confirm) the following three required features in addition to the Windows Server 2016 defaults:

- Failover Clustering
- File Services
- Hyper-V technology

### Configure the operating system.

The next tasks configure the operating system, including IP address assignments, Windows Update, Active Directory Domain Join, and internal drive verification.

1. Disable all network interfaces other than the SMB1 and SMB2 virtual NICs (vNICs).
2. Configure a static IP address for the operating system or public-facing interface on the SMB1 vNIC (for example, 10.10.10.x).
3. Configure the storage heartbeat interface with a static IP address on the SMB2 vNIC using a different subnet (for example, 10.10.11.x).
4. Confirm that connections are functioning properly, pinging from each interface (public and heartbeat) to the corresponding node.
5. Make sure you have the latest fixes and patches. Update the Windows Server components via Windows Update.

Note: You may have to reboot each node to fully install updates.

### Add domain accounts.

Up to this point, you have been deploying and configuring nodes with the local administrator account <ComputerName>\Administrator. Managing a hyperconverged system will generally require a domain account in the Administrators group on each node. After this step, be sure to log onto the systems with a domain account whenever you work with cluster services.

1. Use the following PowerShell command to join each server node to the Windows Active Directory domain:

```
Add-Computer -DomainName <DomainName> -Reboot
```

The login to your cluster services will now be with a domain account rather than as the local administrator.

2. Add the domain account to the local Administrators Security Group on each node using a PowerShell console opened with Administrator privileges. Issue to the following:

```
Net localgroup Administrators <DomainAccount> / add
```

3. Verify that the internal drives are online by going to Server Manager > Tools > Computer Management > Disk Management. If any drives are offline, select the drive, right-click it, and click Online (or you can use a single PowerShell command to bring all drives in each host online).

### Step 3: Configure the Network

Redundant physical network connections are a best practice for resiliency. The following assumes two RDMA NIC ports (one dual port, or two single port). Microsoft also recommends at least two 10 Gb connections between nodes. For Storage Spaces Direct, deploy the Hyper-V switch with RDMA-enabled (recommended, not required) host vNICs.

#### Configure the ToR switch.

Microsoft recommends using a network adapter that implements RDMA. Network QoS for this type of RDMA requires that the ToR switch has specific capabilities set for the network ports to which the NICs are connected.

#### Enable network QoS.

Network QoS ensures that the SDS system has enough bandwidth to communicate between nodes for resiliency and performance. Use [Enter-PSSession](#) to connect from a management system, and do the following to each server:

1. Set a network QoS policy for SMB-Direct, the protocol used by the SDS system.

```
New-NetQosPolicy "SMB" -NetDirectPortMatchCondition 445 -PriorityValue8021Action 3
```

2. Turn on Flow Control for SMB.

```
Enable-NetQosFlowControl -Priority 3
```

3. Disable flow control for other traffic.

```
Disable-NetQosFlowControl -Priority 0,1,2,4,5,6,7
```

4. From a list of network adapters, identify the target adapters.

```
Get-NetAdapter | FT Name,InterfaceDescription,Status,LinkSpeed
```

5. Apply network QoS to the target adapters using the target adapter name.

```
Enable-NetAdapterQos -InterfaceAlias "<adapter1>","<adapter2>"
```

6. Create a Traffic class and give SMB Direct 30 percent of the bandwidth at a minimum. The name of the class will be "SMB."

```
New-NetQosTrafficClass "SMB" -Priority 3 -BandwidthPercentage 30 -Algorithm ETS
```

#### Create a Hyper-V virtual switch with SET and RDMA (recommended) vNIC.

Physical NIC ports can be used for both the host and virtual machines via the Hyper-V virtual switch. The Hyper-V virtual switch also enables RDMA (recommended, not required) from the host, which gives you higher throughput, lower latency, and less system (CPU) impact. The physical network interfaces are teamed using the Switch Embedded Teaming (SET) feature, new in Windows Server 2016.

Use an [Enter-PSSession](#) to connect to each of the servers.

1. Identify the network adapters (you will use this info in step 2).

```
Get-NetAdapter | FT Name,InterfaceDescription,Status,LinkSpeed
```

2. Create the virtual switch connected to both of the physical network adapters, and enable the SET. You may notice a message that your PSSession lost connection. This is expected, and your session will reconnect.

```
New-VMSwitch -Name SETswitch -NetAdapterName "<adapter1>","<adapter2>" -EnableEmbedded  
Teaming $true
```

3. Add host vNICs to the virtual switch. This configures a vNIC from the virtual switch that you just configured for the management operating system to use.

```
Add-VMNetworkAdapter -SwitchName SETswitch -Name SMB_1 -managementOS
```

```
Add-VMNetworkAdapter -SwitchName SETswitch -Name SMB_2 -managementOS
```

- Configure the host vNIC to use a VLAN. They can be on the same or different VLANs.

```
Set-VMNetworkAdapterVlan -VMNetworkAdapterName "SMB_1" -VlanId <vlan number>
-Access-ManagementOS
```

```
Set-VMNetworkAdapterVlan -VMNetworkAdapterName "SMB_2" -VlanId <vlan number>
-Access-ManagementOS
```

- Verify that the VLAN ID is set.

```
Get-VMNetworkAdapterVlan -ManagementOS
```

- Restart each host vNIC adapter so that the VLAN is active.

```
Restart-NetAdapter "vEthernet (SMB_1)"
```

```
Restart-NetAdapter "vEthernet (SMB_2)"
```

- Associate each vNIC configured for RDMA to a physical adapter connected to the virtual switch.

```
Set-VMNetworkAdapterTeamMapping -VMNetworkAdapterName 'SMB_1' -ManagementOS -
PhysicalNetAdapterName 'SLOT 2'
```

```
Set-VMNetworkAdapterTeamMapping -VMNetworkAdapterName 'SMB_2' -ManagementOS -
PhysicalNetAdapterName 'SLOT 2 2'
```

- Verify RDMA capabilities. Values should show True for RDMA Capable for the RDMA-enabled interfaces.

```
Get-SmbClientNetworkInterface
```

#### Step 4: Configure Storage Spaces Direct

Be sure to configure Storage Spaces Direct using a management system of the same version as the target servers. Do not use PSSession. Instead run in a Windows PowerShell session, opened as Administrator.

##### Run cluster validation tool.

Run the cluster validation tool to verify that the server nodes are configured correctly for a failover cluster.

```
Test-Cluster -Node <MachineName1,MachineName2,MachineName3,MachineName4>
-Include "Storage Spaces Direct",Inventory,Network,"System Configuration"
```

##### Create a cluster.

Use the PowerShell framework to create a cluster with four to 16 validated nodes. Be sure to add the -NoStorage parameter to the cmdlet. Otherwise, drives may be automatically added to the cluster that will need to be removed before enabling Storage Spaces Direct. Once the cluster is created, you may have to wait for Domain Name Service (DNS) entry for the cluster name to be replicated.

```
New-Cluster -Name <ClusterName> -Node <MachineName1,MachineName2,
MachineName3,MachineName4> -NoStorage
```

Note: When creating the cluster, you will get a warning that states, "There were issues while creating the clustered role that may prevent it from starting. For more information, view the report file below." You can safely ignore this warning.

Note: If the servers are using static IP addresses, modify the command to reflect the static IP address by adding the following parameter and specifying the IP address: -StaticAddress.

Note: DNS entry for the cluster name may take some time to be replicated, depending on the environment and DNS replication configuration. If you are unable to resolve the cluster, use the machine name of a node that is an active member of the cluster instead of the cluster name.

**Configure a cluster witness.**

Configure a witness for the cluster. This will keep your system running if nodes fail or are offline. For example, a four-node system would continue to operate even if two nodes failed. You can use a file share as a witness or use a cloud witness.

For more information about configuring a cloud witness, see [Introducing Cloud Witness in Windows Server 2016](#).

For more information about configuring a file-share witness, see [Configuring a File Share Witness on a Scale-Out File Server](#).

**Use clean drives.**

The drives intended to be used for S2D must be empty and without partitions or other data, or they will not be included in the S2D system. If your hardware is new, your drives should be clean.

To clean your drives, see “Appendix B: Cleaning Your Drives for Deployment.”

**Enable Storage Spaces Direct.**

After creating the cluster, use the Enable-ClusterStorageSpacesDirect PowerShell cmdlet to put your storage system into S2D mode. This automatically creates a single large shared pool. S2D configures the cache using the highest performing media available and creates default tiers for capacity and performance configured with the available mix of device types and resiliency you have specified.

Use the PowerShell framework opened with Administrator privileges for the following command:

```
Enable-ClusterStorageSpacesDirect -CimSession <ClusterName>
```

Note: You can also use the node name instead of the cluster name to enable Storage Spaces Direct. Using the node name may be more reliable since DNS replication delays may occur with the newly created cluster name.

**Create virtual disks.**

Create a virtual disk with both mirror and parity resiliency on the storage pool created in the previous step using PowerShell software. You can use the management system with or without using a PSSession. If using a PSSession, remove the -CimSession parameter below. The following examples are meant to help guide you in creating volumes with the resiliency and size you need for your installation:

For mirror and parity resiliency:

```
New-Volume -StoragePoolFriendlyName "S2D*" -FriendlyName <VirtualDiskName>
-FileSystem CSVFS_ReFS -StorageTierfriendlyNames Capacity,Performance
-StorageTierSizes <Size of Capacity tier in size units, example: 800GB>, <Size
of Performance tier in size units, example: 80GB> -CimSession <ClusterName>
```

For mirror resiliency only:

```
New-Volume -StoragePoolFriendlyName "S2D*" -FriendlyName <VirtualDiskName>
-FileSystem CSVFS_ReFS -StorageTierfriendlyNames Performance -StorageTierSizes
<Size of Performance tier in size units, example: 800GB> -CimSession <ClusterName>
```

**Create or deploy virtual machines.**

Provision virtual machines onto the nodes of the hyperconverged S2D cluster. As with clustered VMs on failover clusters, store the virtual machine's files on the systems' CSV namespace (for example, C:\ClusterStorage\Volume1).

## On-Premises Cloud Use Case

Windows Server 2016 is a cloud-ready operating system that delivers new layers of security and innovation inspired by Microsoft Azure\* technology. The software-defined capabilities in Windows Server 2016 provide a strong foundation for flexible, highly efficient cloud services, reducing risk with multiple layers of protection and providing a platform optimized for your current, as well as native, cloud applications.

With Windows Server 2016 and Storage Spaces Direct (S2D), enterprises can deploy SDS in their private clouds using industry-standard servers and local storage using a range of media types. The SDS solution is deployed, managed, and monitored using System Center VMM, System Center Operations Manager (SCOM), and PowerShell software. Use cases include hyperconverged or converged solutions as we've described in this guide, storage for Hyper-V IaaS, and backup and replication targets.

Windows Server 2016 also enables software-defined networking (SDN). A new network controller role, the same components used in Azure, provides capabilities such as a network load balancer, distributed firewall, and VPN gateway for clusters.

## Summary

IT is under pressure to implement technologies that can provide greater agility and efficiency, manage the rapidly growing volume of data, and quickly roll out new value-added business services. Modern architectures designed to support cloud environments and enable efficient resource use help enterprises address these challenges.

Together, Intel and Microsoft are simplifying the delivery of cloud services. We are developing resilient, reliable, and scalable infrastructure solutions that optimize both hardware and software capabilities and enable software-defined technologies that bring even greater flexibility and efficiency to the data center.

Siloed storage systems have become a limitation for data center transformation. An Intel optimized solution for Windows Server 2016 with Storage Spaces Direct is a smart investment in technologies that can efficiently manage the rapid, exponential increase in data volume. This guide presented three tested hyperconverged configurations for SDS using Windows Server 2016 running on the latest Intel Xeon processors and SSDs. By deploying one or more of these configurations, you can build highly available, distributed, server-based storage systems that use multiple drive types. The result is data storage that is high performing, efficient, and cost-effective, providing greater flexibility in data management and supporting dynamic resource allocation for a wide range of data center workloads.

See [performance comparisons on the three recommended configurations](#).

Learn more about [Windows Server 2016](#).

Learn more about the [Intel Xeon processor E5 v4 family](#).

Learn more about the [Intel Solid State Drive Data Center Family for PCIe](#).

## Appendix A: References

### Intel and Microsoft Solutions

#### ***Intel and Windows\* Server 2016: Better Together Solution Brief***

A solution brief describing advanced data center functionalities for Windows Server 2016, optimized for Intel's best-in-class storage, compute, and networking technologies.

[intel.com/content/www/us/en/big-data/intel-and-windows-server-2016-paper.html](http://intel.com/content/www/us/en/big-data/intel-and-windows-server-2016-paper.html)

#### ***Microsoft and Intel Team to Advance Software Defined Storage for Private Clouds***

In this podcast, Spencer Shepler, private cloud solutions architect at Microsoft, joins the Intel team to discuss advances in software-defined storage using Storage Spaces Direct in Windows Server 2016. Part of Intel's Conversations in the Cloud series. (7:53 min.)

<http://connectedsocialmedia.com/14923/microsoft-and-intel-team-to-advance-software-defined-storage-for-private-clouds-conversations-in-the-cloud-episode-61/>

### Intel Blogs

#### ***3 Ready to Go Configurations for Windows Server 2016 with Storage Spaces Direct***

Learn about the three recommended configurations with Intel Xeon processors and Intel storage devices.

<http://itpeernetwork.intel.com/3-ready-to-go-configurations-for-windows-server-2016-with-storage-spaces-direct/>

#### ***IOPS Performance on NVMe + HDD Configuration with Windows Server 2016 and Storage Spaces Direct***

Intel presents internal testing results of a hybrid storage configuration using Intel NVMe and HDD, for a capacity-optimized solution.

<http://itpeernetwork.intel.com/iops-performance-nvme-hdd-configuration-windows-server-2016-storage-spaces-direct/>

#### ***IOPS Performance of All Flash NVMe + SATA Configuration—Windows Server 2016 and Storage Spaces Direct***

For customers interested in higher throughput, Intel demonstrates consistent, balanced performance across multiple servers using an all Intel SSD configuration with a mix of NVMe and SATA SSDs.

<http://itpeernetwork.intel.com/iops-performance-flash-nvme-sata-configuration-windows-server-2016-storage-spaces-direct/>

#### ***Record Performance of All Flash NVMe Configuration—Windows Server 2016 and Storage Spaces Direct***

Instructions for creating a hyperconverged solution using S2D in Windows Server 2016.

<http://itpeernetwork.intel.com/record-performance-flash-nvme-configuration-windows-server-2016-storage-spaces-direct-2/>

#### ***OLTP Performance on an All Flash NVMe + SATA Configuration—Storage Spaces Direct Cluster with 144 SQL Server VMs***

Intel testing results demonstrate leading-edge OLTP performance using Microsoft SQL Server\* on Windows Server 2016 with Storage Spaces Direct.

<https://itpeernetwork.intel.com/oltp-performance-flash-nvme-sata-configuration-storage-spaces-direct-cluster-144-sql-server-vm/>

## **More about Storage Spaces Direct**

### ***Hyper-Converged Solution Using Storage Spaces Direct in Windows Server 2016***

Instructions for creating a hyperconverged solution using S2D in Windows Server 2016.

<https://technet.microsoft.com/windows-server-docs/storage/storage-spaces/hyper-converged-solution-using-storage-spaces-direct>

### ***Microsoft Storage Spaces Direct Deployment Guide***

Implementation guide for a hyperconverged server-based storage solution from Lenovo.

<https://lenovopress.com/lp0064.pdf>

### ***Storage Spaces Direct in Windows Server 2016***

A great overview of S2D, covering benefits, deployment options, and how it works. Includes a video slide presentation (4:57 min.) that demonstrates how easy it is to enable S2D, form a server cluster, pool resources, create a volume, and scale out by adding an additional server.

<https://technet.microsoft.com/en-us/windows-server-docs/storage/storage-spaces/storage-spaces-direct-overview>

### ***Ten Reasons You'll Love Windows Server 2016: Software-Defined Storage***

The first in a series of five videos on the advanced storage capabilities in Windows Server 2016. (14:40 min.)

<https://channel9.msdn.com/Blogs/windowsserver/Software-Defined-Storage>

### ***Volume Resiliency and Efficiency in Storage Spaces Direct***

A detailed description of mirror and parity resiliency in Windows Server 2016, including how they work in S2D, how data is laid out across servers, and how data is reconstructed if a drive fails.

<https://blogs.technet.microsoft.com/filecab/2016/09/06/volume-resiliency-and-efficiency-in-storage-spaces-direct/>

## Appendix B: Cleaning Your Drives for Deployment

The drives intended to be used for S2D must be empty and without partitions or other data, or they will not be included in the S2D system. If you are repurposing existing drives, follow this procedure to make sure they are clean and ready for deployment.

**Warning:** Make sure you have no data on the cluster drives before running these commands.

```
icm (Get-Cluster -Name <ClusterName> | Get-ClusterNode) {
    Update-StorageProviderCache
    Get-StoragePool |? IsPrimordial -eq $false | Set-StoragePool -IsReadOnly:$false
    -ErrorAction SilentlyContinue
    Get-StoragePool |? IsPrimordial -eq $false | Get-VirtualDisk | Remove-
    VirtualDisk -Confirm:$false- ErrorAction SilentlyContinue
    Get-StoragePool |? IsPrimordial -eq $false | Remove-StoragePool -Confirm:$false
    -ErrorAction SilentlyContinue
    Get-PhysicalDisk | Reset-PhysicalDisk -ErrorAction SilentlyContinue
    Get-Disk |? Number -ne $null |? IsBoot -ne $true |? IsSystem -ne $true |?
    PartitionStyle -ne RAW |% {
        $_ | Set-Disk -isoffline:$false
        $_ | Set-Disk -isreadonly:$false
        $_ | Clear-Disk -RemoveData -RemoveOEM -Confirm:$false
        $_ | Set-Disk -isreadonly:$true
        $_ | Set-Disk -isoffline:$true
    }
    Get-Disk |? Number -ne $null |? IsBoot -ne $true |? IsSystem -ne $true |?
    PartitionStyle -eq RAW | Group -NoElement -Property FriendlyName
} | Sort -Property PsCompute
```



All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps. Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer, or learn more at [intel.com](http://intel.com).

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

© 2016 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

\*Other names and brands may be claimed as the property of others.

1216/RF/MRM/PDF-USA

♻ Please Recycle

335269-001