

Microsoft Azure Stack HCI – Software Defined Storage auf einem neuen Niveau

Manfred Helber



Twitter: @ManfredHelber

www.manfredhelber.de



Traditional “Three Tier” Infrastructure

Hyper-Converged Infrastructure (HCI)

Everything you need, one familiar product




Hyper-V

Introduced in 2008, ten years ago!
Foundation of our hyperscale Azure cloud



Storage Spaces Direct

Introduced in Windows Server 2016
Foundation of Azure Stack




SDN

Introduced in Windows Server 2016
Foundation of Azure Stack


Included in Windows Server 2016/2019 Datacenter

- ✓ Hypervisor / compute
- ✓ Software-defined storage
- ✓ Software-defined networking
- ✓ Management software
- ✓ Unlimited guest licenses


Start with industry-standard servers



Insert additional drives to each server



Fastest drives provide cache to accelerate the others




Faster



Slower

intel OPTANE™ DC
PERSISTENT MEMORY



Persistent Memory



Non-Volatile Memory Express



Solid-State Drive



Hard Disk Drive

~ 1 μ s

1,000 ns

~ 0.1 ms

100,000 ns

~ 0.1 ms

100,000 ns

~ 10 ms

10,000,000 ns

Deploy with 1, 2, or 3 types of drives

Capacity



With **1**, there is no cache, except by manual configuration.

Cache + Capacity



With **2**, the faster type provides cache, the other type provides capacity.

Cache + Capacity + Capacity



With **3**, the fastest 1 type provides cache; the remaining 2 types provide capacity.

Why cache writes?

Coalescence

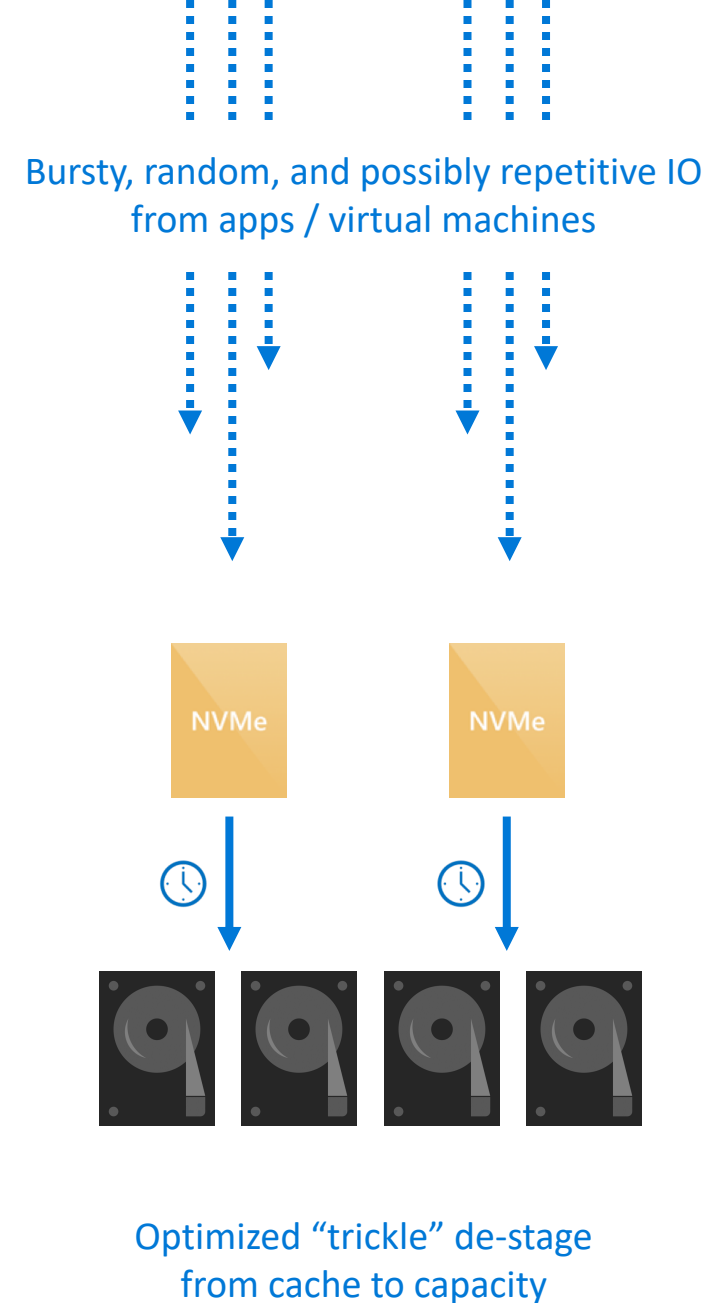
Random IO is hard for an HDD. Squeeze up to 5X more writes by coalescing them in software to minimize arm movement and use every platter rotation effectively.

Absorb bursts

Writes may not happen uniformly over time. Amortize spikes over quieter periods to keep drives working steadily.

Just write once

Apps may overwrite data multiple times in rapid succession. Reduce overall traffic to HDD by only writing the “last” version.



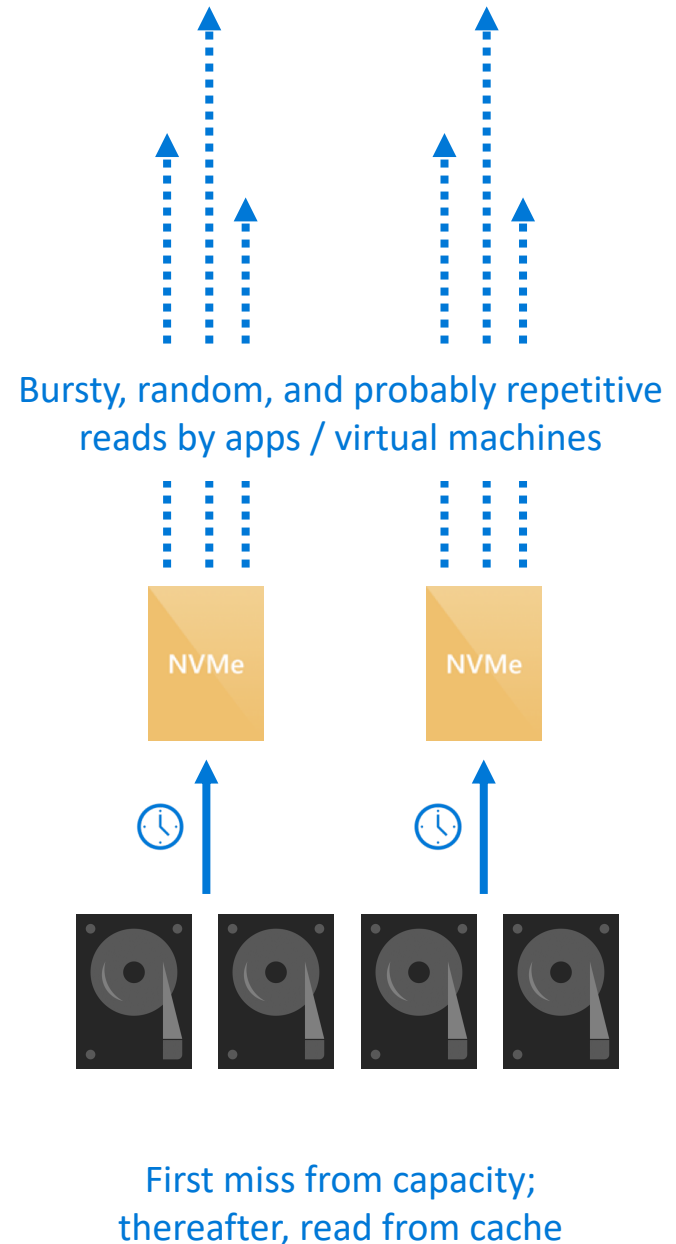
Why cache reads?

Flash is just faster

Anything that's been read multiple times is likely to be read again. The faster you respond, the faster the app proceeds, the more IOPS you get.

Don't interrupt de-stage

Reading from HDD requires seeking and settling, disrupting the optimized de-stage from the previous slide.



Windows Server is on the leading edge



NVMe

Full support for M.2, U.2,
Add-In-Card (AIC)



Intel Xeon® Scalable®

Supported since announcement
on Windows Server 2016+



RDMA

Unique to Windows
software-defined storage



Intel Thunderbolt™ 3

Blazing the trail for
switchless networking

OPTANE™

3D XPoint

Supported since announcement
on Windows Server 2016+



Persistent memory

Use as cache or capacity in
Windows Server 2019

of x86 hardware innovation


The very latest from Intel

12 x Intel® S2600WFT

- ✓ 384 GiB (12 x 32 GiB) memory
- ✓ 2 x 28-core future Intel® Xeon® Scalable processor
- ✓ 1.5 TB Intel® Optane™ DC persistent memory as cache
- ✓ 32 TB NVMe (4 x 8TB Intel® DC P4510) as capacity
- ✓ 2 x Mellanox ConnectX-4 25 Gbps

The very latest from Microsoft

- ✓ Windows Server 2019
- ✓ Insider Preview, build 17713
- ✓ Hyper-V + Storage Spaces Direct







NEW IOPS RECORD

1 3 , 7 9 8 , 6 7 4

Monday, September 24, 2018 | Windows Server 2019 with Intel® Optane™ DC persistent memory



NEW! Larger maximum scale

Limit / Recommendation	Windows Server 2016	Windows Server 2019	Increase
Max <u>servers</u> per cluster	16	16	-
Max <u>drives</u> per cluster	416	416	-
Max raw capacity per <u>cluster</u>	1 PB	4 PB	 4x
Max raw capacity per <u>server</u>	100 TB	400 TB	 4x
Max <u>number</u> of volumes	32	64	 2x
Max <u>size</u> per volume	32 TB	64 TB	 2x

Storage Spaces Direct (S2D) – Hardware requirements

Systems, components, devices, and drivers must be **Windows Server 2016 Certified** per the [Windows Server Catalog](#). In addition, we recommend that servers, drives, host bus adapters, and network adapters have the **Software-Defined Data Center (SDDC) Standard** and/or **Software-Defined Data Center (SDDC) Premium** additional qualifications (AQs), as pictured below. There are over 1,000 components with the SDDC AQs.



The fully configured cluster (servers, networking, and storage) must pass all cluster validation tests per the wizard in Failover Cluster Manager or with the Test-Cluster cmdlet in PowerShell.

Storage Spaces Direct (S2D) – Hardware requirements

Servers

Minimum of 2 servers, maximum of 16 servers

Recommended that all servers be the same manufacturer and model

CPU

Intel Nehalem or later compatible processor; or

AMD EPYC or later compatible processor

Memory

Memory for Windows Server, VMs, and other apps or workloads; plus

4 GB of RAM per terabyte (TB) of cache drive capacity on each server, for Storage Spaces Direct metadata

Boot

Any boot device supported by Windows Server, which now includes SATADOM

RAID 1 mirror is **not** required, but is supported for boot

Recommended: 200 GB minimum size

Storage Spaces Direct (S2D) – Hardware requirements

Networking

Minimum (for small scale 2-3 node)

- 10 Gbps network interface
- Direct-connect (switchless) is supported with 2-nodes

Recommended (for high performance, at scale, or deployments of 4+ nodes)

- NICs that are remote-direct memory access (RDMA) capable, iWARP (recommended) or RoCE
- Two or more NICs for redundancy and performance
- 25 Gbps network interface or higher

Storage Spaces Direct (S2D) – Hardware requirements

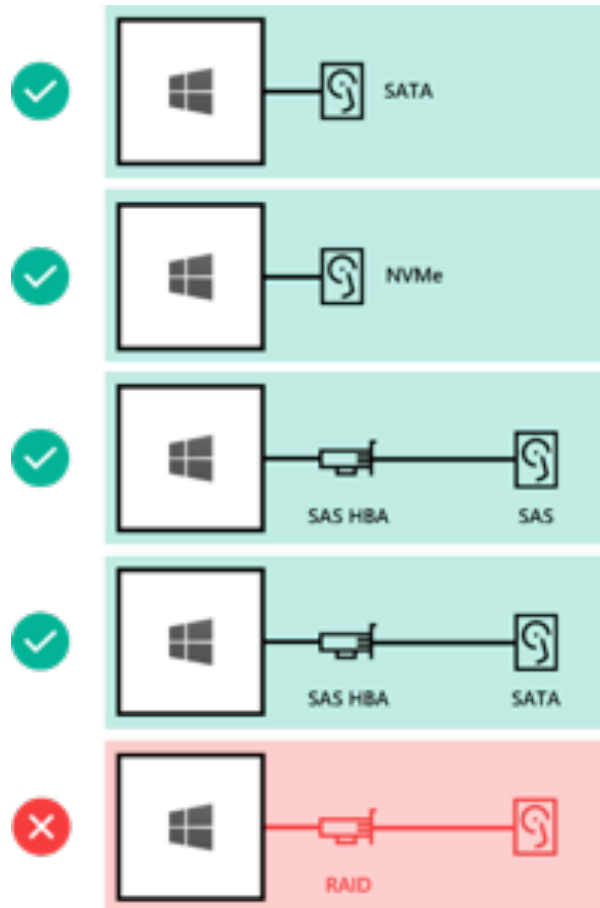
Drives

Storage Spaces Direct works with direct-attached SATA, SAS, or NVMe drives that are physically attached to just one server each. For more help choosing drives, see the Choosing drives topic.

- SATA, SAS, and NVMe (M.2, U.2, and Add-In-Card) drives are all supported
- 512n, 512e, and 4K native drives are all supported
- Solid-state drives must provide power-loss protection
- Same number and types of drives in every server – see Drive symmetry considerations
- NVMe driver is Microsoft's in-box or updated NVMe driver.
- Recommended: Number of capacity drives is a whole multiple of the number of cache drives
- Recommended: Cache drives should have high write endurance: at least 3 drive-writes-per-day (DWPD) or at least 4 terabytes written (TBW) per day – see Understanding drive writes per day (DWPD), terabytes written (TBW), and the minimum recommended for Storage Spaces Direct


Storage Spaces Direct (S2D) – Hardware requirements

How to connect Drives



- Direct-attached SATA drives
- Direct-attached NVMe drives
- SAS host-bus adapter (HBA) with SAS drives
- SAS host-bus adapter (HBA) with SATA drives
- **NOT SUPPORTED:** RAID controller cards or SAN (Fibre Channel, iSCSI, FCoE) storage. Host-bus adapter (HBA) cards must implement simple pass-through mode.

Two options to build a S2D Solution



1

Build Your Own

2

Validated Partner Solutions

Validated Partner Solutions for Storage Spaces Direct

**Validated Partner Solutions for
Windows Server 2016:**

Windows Server Software Defined

<https://Microsoft.com/WSSD>

**Validated Partner Solutions for
Windows Server 2019:**

Azure Stack HCI

<https://Microsoft.com/HCI>

Find Azure Stack HCI solutions from your preferred hardware vendor

Refine results

Hardware partners

Dell EMC

Features

All-NVMe

2-Node Optimized

All-Flash

Persistent Memory

Hybrid Storage

More

All results

Dell EMC

Showing 1 - 9 of 12 results



Dell EMC Storage Spaces
Direct Ready Node -
R640 Hybrid

Scale:
• 2 to 16 nodes



Dell EMC Storage Spaces
Direct Ready Node -
R740xd All-Flash

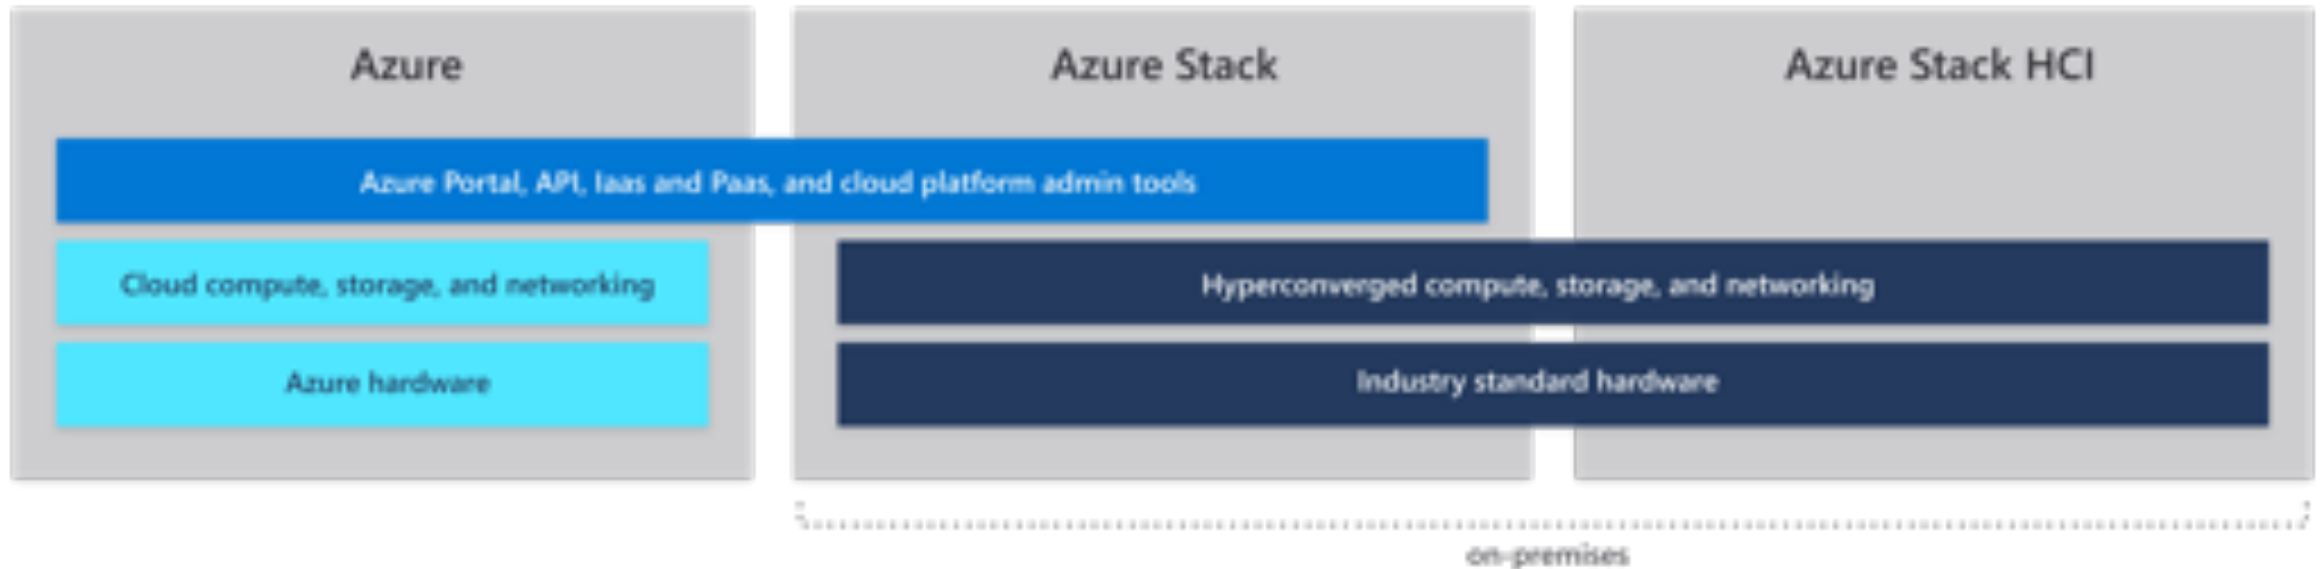
Scale:
• 2 to 16 nodes



Dell EMC Storage Spaces
Direct Ready Node -
R740xd
NVMe(AIC)+HDD

Scale:

Microsoft Azure, Azure Stack and Azure Stack HCI



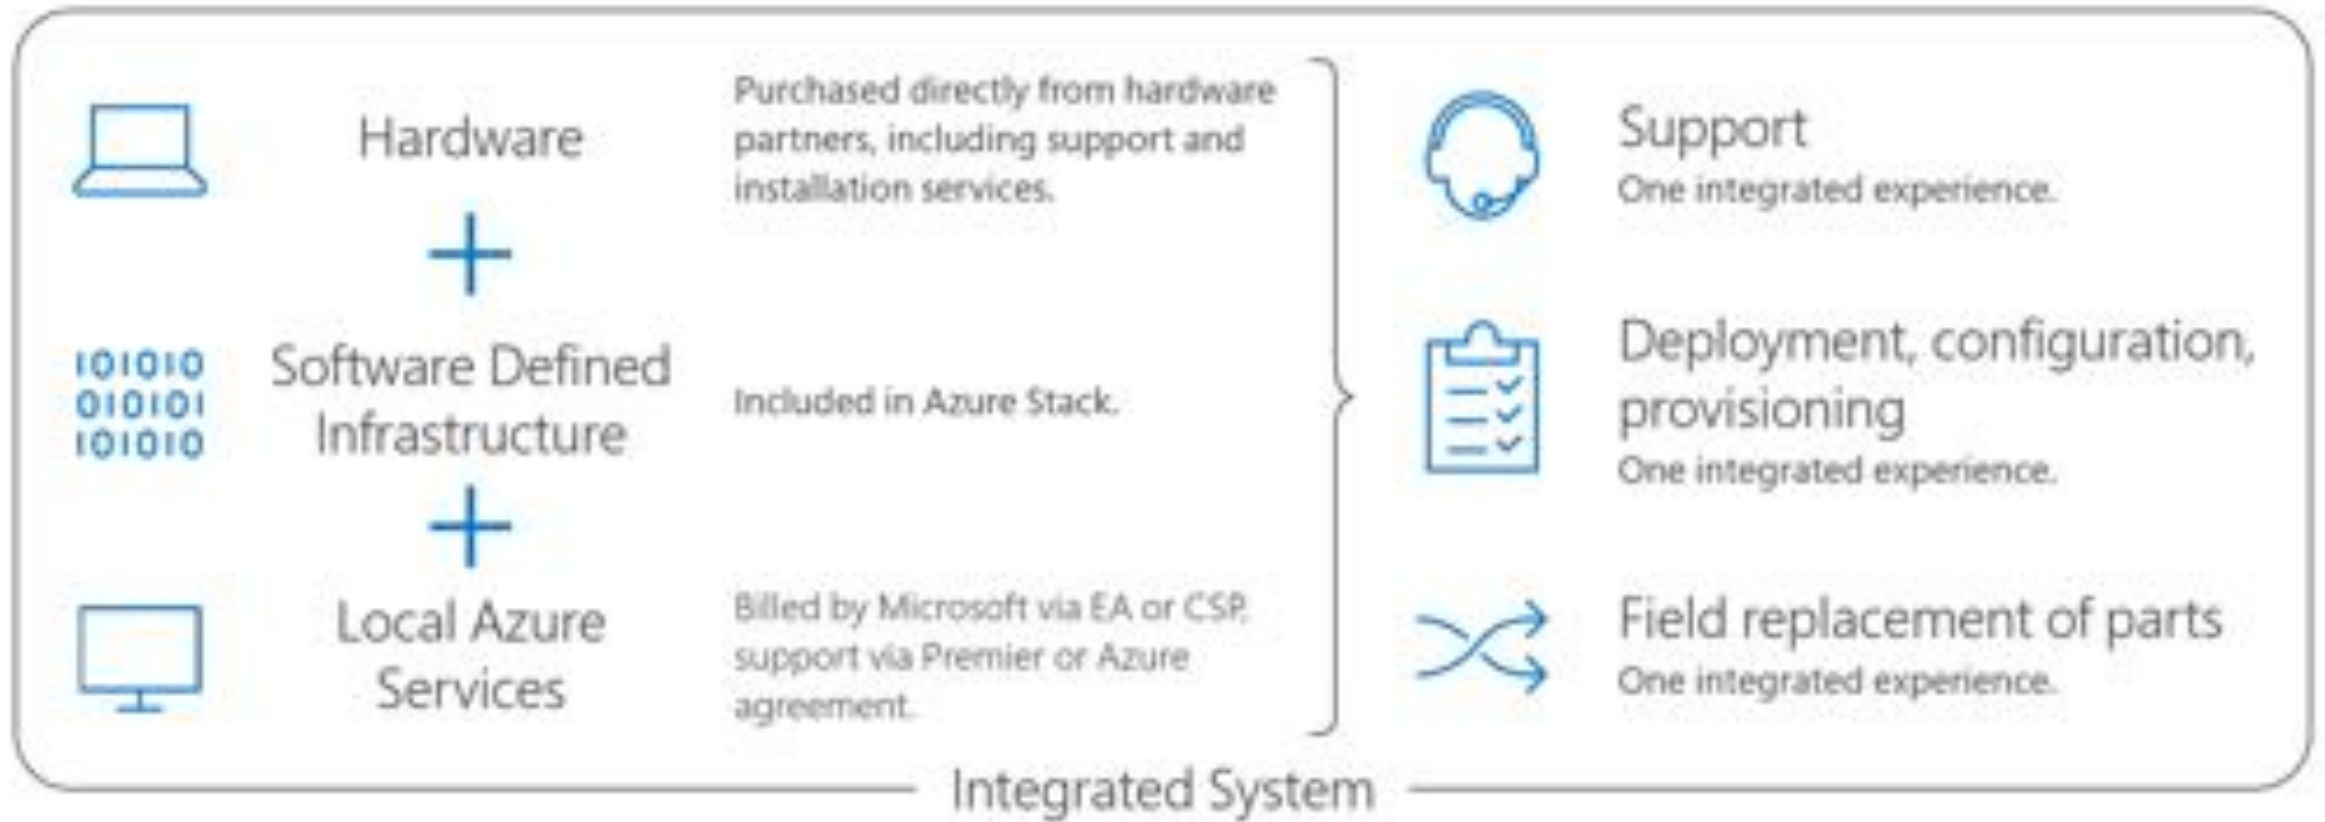
[Azure](#) - Use public cloud services

[Azure Stack](#) - Operate cloud services on-premises

[Azure Stack HCI](#) - Run virtualized apps on-premises, with optional connections to Azure

<https://Microsoft.com/HCI>

How to buy Azure Stack



How to buy Azure Stack HCI



Microsoft Azure Stack HCI - FAQ

What do Azure Stack and Azure Stack HCI solutions have in common?

Azure Stack HCI solutions feature the same Hyper-V based software-defined compute, storage, and networking technologies as Azure Stack. Both offerings meet rigorous testing and validation criteria to ensure reliability and compatibility with the underlying hardware platform.

How are they different?

With Azure Stack, you run cloud services on-premises. You can run Azure IaaS and PaaS services on-premises to consistently build and run cloud applications anywhere, managed with the Azure Portal on-premises.

With Azure Stack HCI, you run virtualized workloads on-premises, managed with Windows Admin Center and familiar Windows Server tools. You can optionally connect to Azure for hybrid scenarios such as cloud-based site recovery, monitoring, and others.

Why is Microsoft bringing its HCI offering to the Azure Stack family?

Microsoft's hyperconverged technology is already the foundation of Azure Stack.

Many Microsoft customers have complex IT environments and our goal is to provide solutions that meet them where they are with the right technology for the right business need. Azure Stack HCI is an evolution of the Windows Server 2016-based Windows Server Software-Defined (WSSD) solutions previously available from our hardware partners. We brought it into the Azure Stack family because we have started to offer new options to connect seamlessly with Azure for infrastructure management services.


Will I be able to upgrade from Azure Stack HCI to Azure Stack?

No, but customers can migrate their workloads from Azure Stack HCI to Azure Stack or Azure.

Manfred Helber



Twitter: @ManfredHelber
LinkedIn: Manfred Helber
www.manfredhelber.de




Windows Business Solutions Club News-Channel

Windows Business Solutions Club Support-Channel

Jetzt zuhause auf

LinkedIn



Schulungsangebote des Windows Business Solutions Club (ehem. Windows Server Kompetenz Club) für Partner zu Windows Server, Microsoft Azure, Microsoft 365, Windows 10 und Modern Devices



Überblick über die aktuellen Schulungsangebote für Partner

Der Windows Business Solutions Club (ehem. Windows Server Kompetenz Club) ist eine Community für Microsoft Partner, die sich für professionelle Infrastruktur- und Anwendungs-Lösungen von Microsoft interessieren. Wir bieten hochkarätige Schulungen auf Level 200 bis 300 mit dem Ziel, dass Sie die Technologien erleben, verstehen und erfolgreich beim Kunden platzieren können. Dabei orientieren wir uns sowohl in Richtung Cloud als auch on-premises.

<https://aka.ms/WBSCEvents>



www.windows-business-solutions-club.de

Vielen Dank!