White paper

Microsoft Azure | intel

Azure Stack HCI software architecture

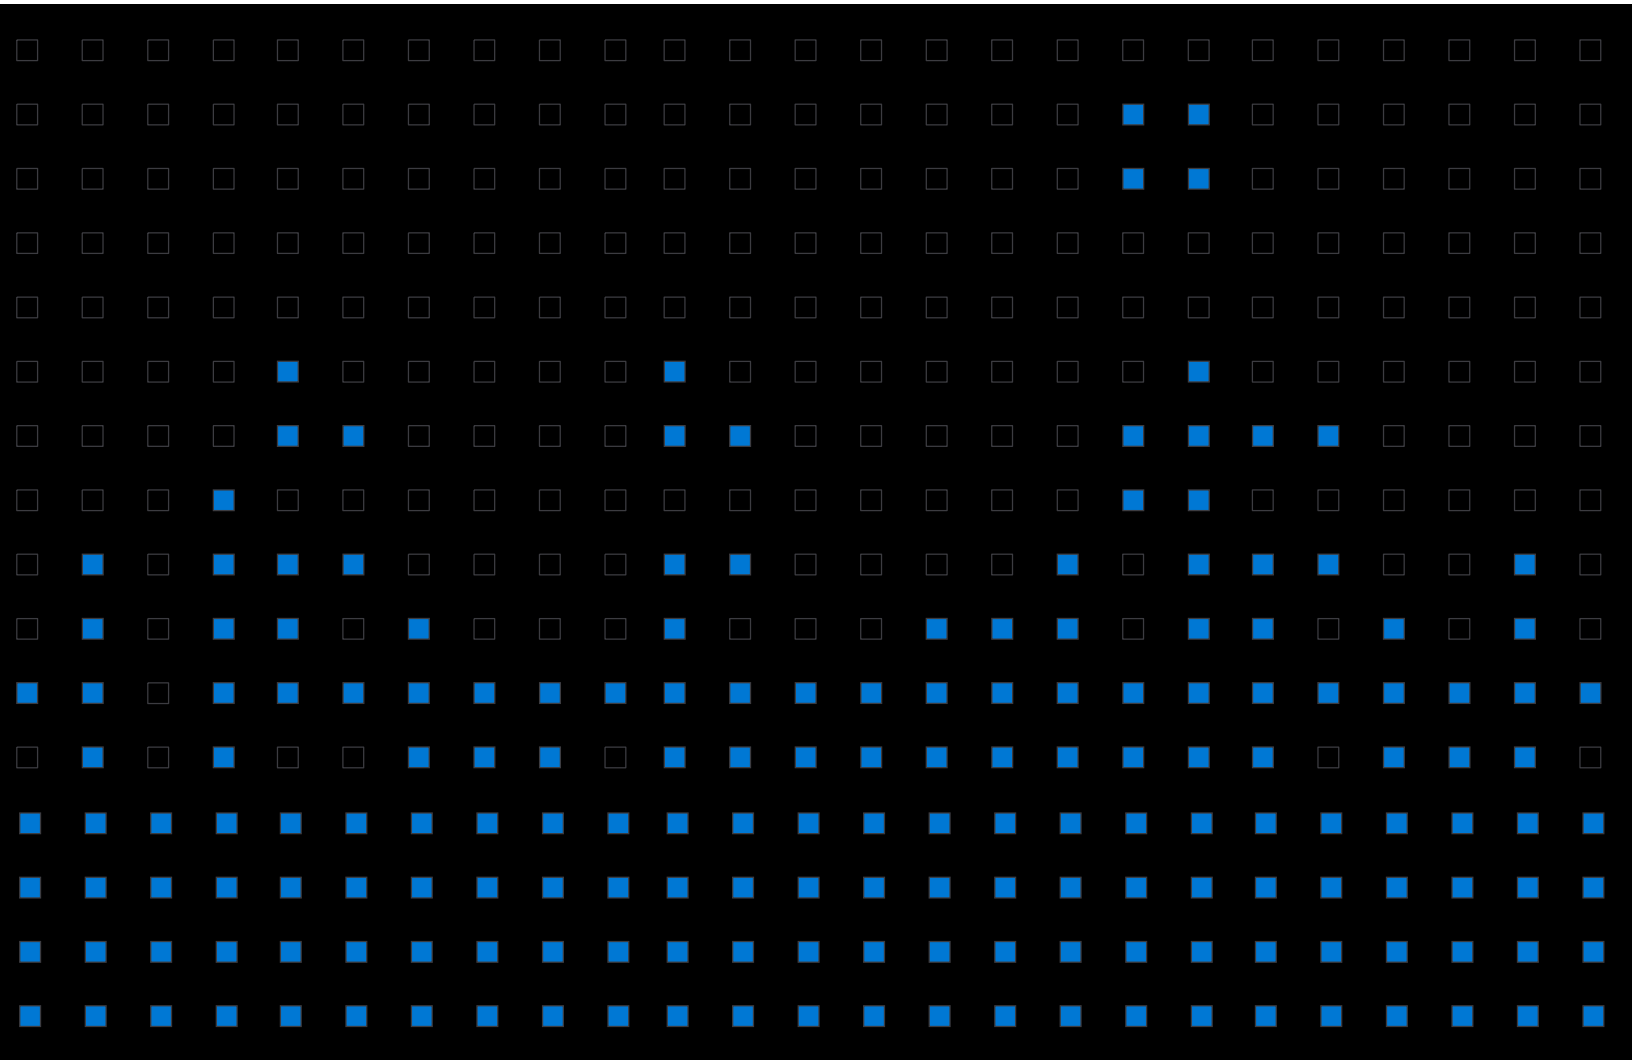

Even as organizations today are moving workloads to the cloud to modernize services and reduce capital expenditures, they are also seeking to gain efficiencies for the virtualized applications they continue to host on premises. Hyperconverged infrastructure (HCI) offers one way for customers to achieve better performance and manage workloads for lower costs in their datacenters. Through HCI, the compute, storage, and network resources of multiple servers are virtualized and consolidated through software and then managed centrally, providing a more efficient and flexible use of resources.

The challenge for these customers is that HCI solutions built to host virtualized workloads on premises typically do not offer integrated cloud services to improve the management of those same workloads. However, Microsoft Azure Stack HCI, the newest solution in the broad set of hybrid capabilities offered by Microsoft, delivers precisely this benefit to customers. With an HCI solution that provides direct access to optional Azure services (such as cloud backup, update management, security monitoring, and disaster recovery), Azure Stack HCI delivers industry-leading performance at an affordable price: the simple cost of a Windows Server 2019 Datacenter license.[1] Given these advantages, Azure Stack HCI offers an excellent choice for an on-premises HCI solution, and it is built completely from components that are included with Windows Server 2019 Datacenter. This white paper will introduce Azure Stack HCI by providing a descriptive overview of the solution, including its hardware specifications, management tools, and software components.

Azure Stack HCI hardware specifications

Azure Stack HCI is offered through Microsoft hardware partners, typically either preconfigured or bundled with simple configuration software, and it can be designed to support a wide array of use cases. Its minimum hardware requirements are low, enabling a small-footprint, two-node deployment if needed. However, it can also scale up to 16 nodes (and much more through a Windows Server 2019 feature called cluster sets). Table 1 shows the solution's minimum hardware requirements and maximum supported specifications.

Minimum hardware requirements	
Number of physical servers per cluster	2
Intel CPU (per node)	• Intel Broadwell (released Q1 2015) or later compatible processor* • 1.4 GHz 64-bit processor • Compatible with x64 instruction set *A 2nd Generation Intel Xeon Scalable processor is required to support Intel Optane DC persistent memory
RAM (per node)	• 32 GB
Maximum supported hardware specifications	
Number of physical servers per cluster	16 (clusters can also be combined using cluster sets to create an HCI platform of hundreds of nodes)
Maximum number of virtual machines (VMs) per host	1,024
Maximum number of disks per VM (SCSI)	256
Maximum storage per cluster	4 PB
Maximum number of logical processors per host	512
Maximum RAM	
Per host	24 TB
Per VM	12 TB for generation-2 VMs 1 TB for generation-1 VMs
Maximum number of virtual processors	
Per host	2,048
Per VM	240 for generation-2 VMs 64 for generation-1 VMs

Table 1. Hardware specifications for Azure Stack HCI

Azure Stack HCI components

The Azure Stack HCI solution is intended to be pre-installed and (partially or completely) preconfigured by Microsoft partners on validated, industry-standard hardware. It is built from the following components, which are also represented graphically in Figure 1:

- **Management tools:**
  - **Windows Admin Center** for central, comprehensive management of local and remote servers through a graphical interface

    - **Azure services (optional)** integrated into Windows Admin Center for offsite backups, site recovery, cloud-based monitoring, and other benefits
  - **PowerShell** for scripting and automation
- **Windows Server 2019 Datacenter roles and features:**
  - **Hyper-V** to run VMs on all physical hosts
  - **Software Defined Networking (SDN) (optional)** for network virtualization
  - **Storage Spaces Direct** for storage virtualization
- **Validated partner hardware**

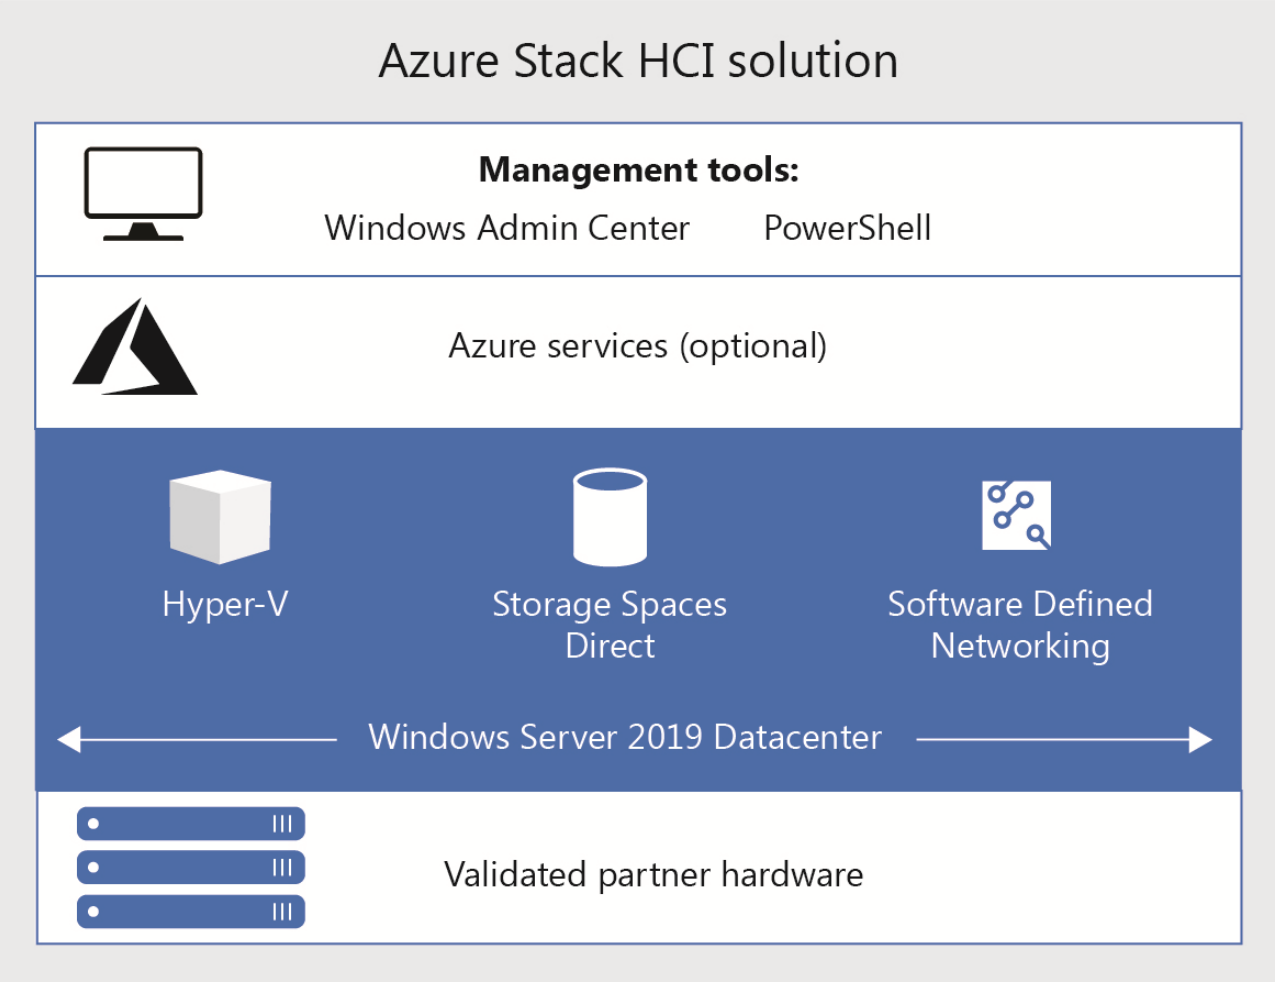

Figure 1. The components of the Azure Stack HCI solution

These same components are described in more detail in the following sections.

Azure Stack HCI management tools

Azure Stack HCI is managed centrally through a new graphical tool, Windows Admin Center, and through PowerShell and other administrative tools.

Windows Admin Center

Windows Admin Center is a new browser-based tool that you can use to centrally monitor and manage your entire Windows Server infrastructure, including your Azure Stack HCI clusters. Some of the features that Windows Admin Center supports for Azure Stack HCI management include:

- Central management of the compute, storage, and network resources through a single tool.
- Central management of virtual resources, including virtual storage and Hyper-V VMs. (For example, you can create, open, resize, and delete volumes, and you can create, start, connect to, and move VMs.)
- Central management of features installed on individual machines in your infrastructure.
- Cluster-wide monitoring and alerts. For example, the dashboard displays memory and CPU usage, storage capacity, input/output (I/O) operations per second (IOPS), throughput, and latency in real time, across every server in the cluster.
- SDN support. For example, you can manage and monitor virtual networks and subnets, connect VMs to virtual networks, and monitor SDN infrastructure.

Figure 2 shows the Hyper-Converged Cluster Manager view in Windows Admin Center, to which an HCI cluster has been added and selected. The left (navigation) pane, highlighted in red, shows the cluster-wide features and components that you can manage with the tool.

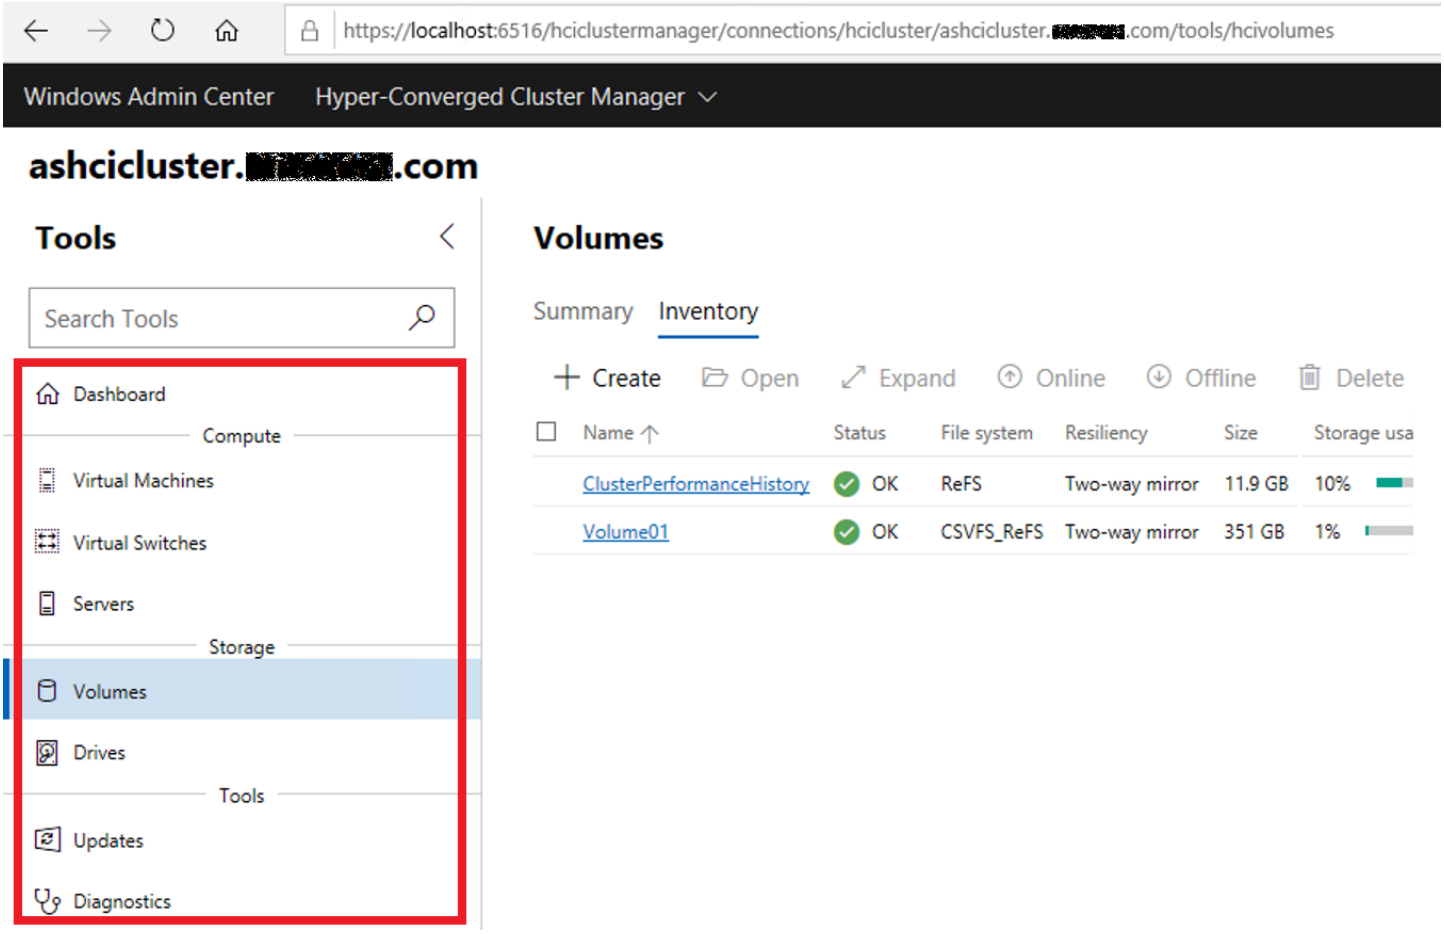

Figure 2. Windows Admin Center enables you to manage compute, storage, network, and other resources in an Azure Stack HCI cluster

You can also use Windows Admin Center to manage individual machines (hosts or VMs) through the Server Manager view. When an individual server is added to the Server Manager view and then selected, you can easily view and manage many of its features, such as its devices, installed apps, registry settings, role and features, and services.

Figure 3 shows a view of Windows Admin Center in which a server has been added and selected. The navigation pane, highlighted in red, shows the many server-specific features you can manage.

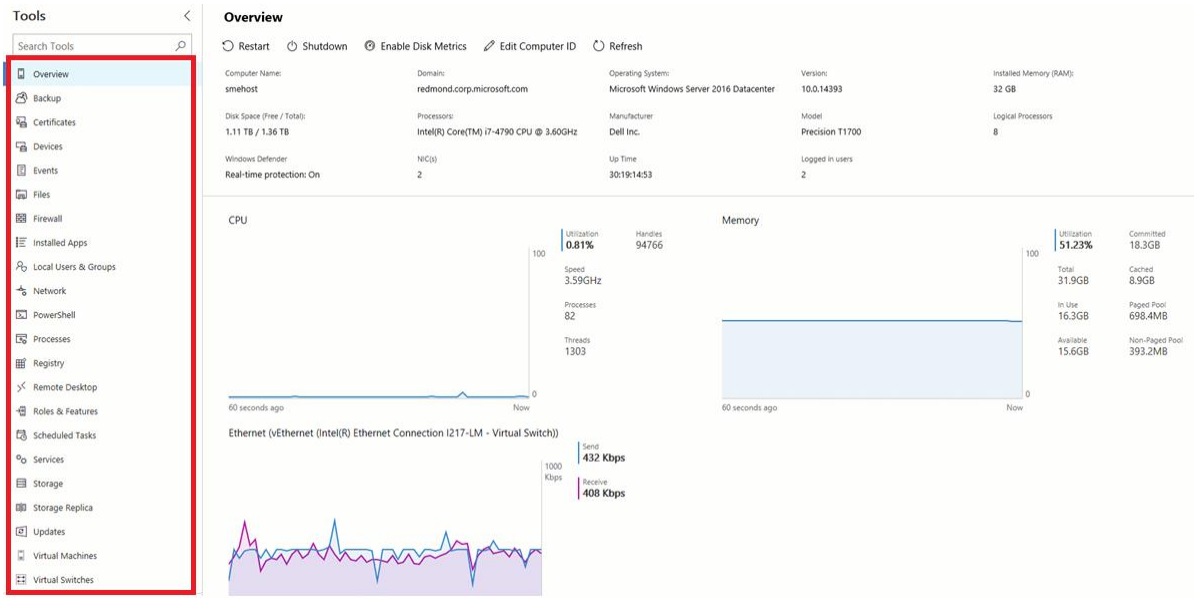

Figure 3. Server management in Windows Admin Center

A particularly powerful feature of Windows Admin Center is that it allows you to easily connect to a remote machine through PowerShell (as shown in Figure 4) or through Microsoft Remote Desktop.

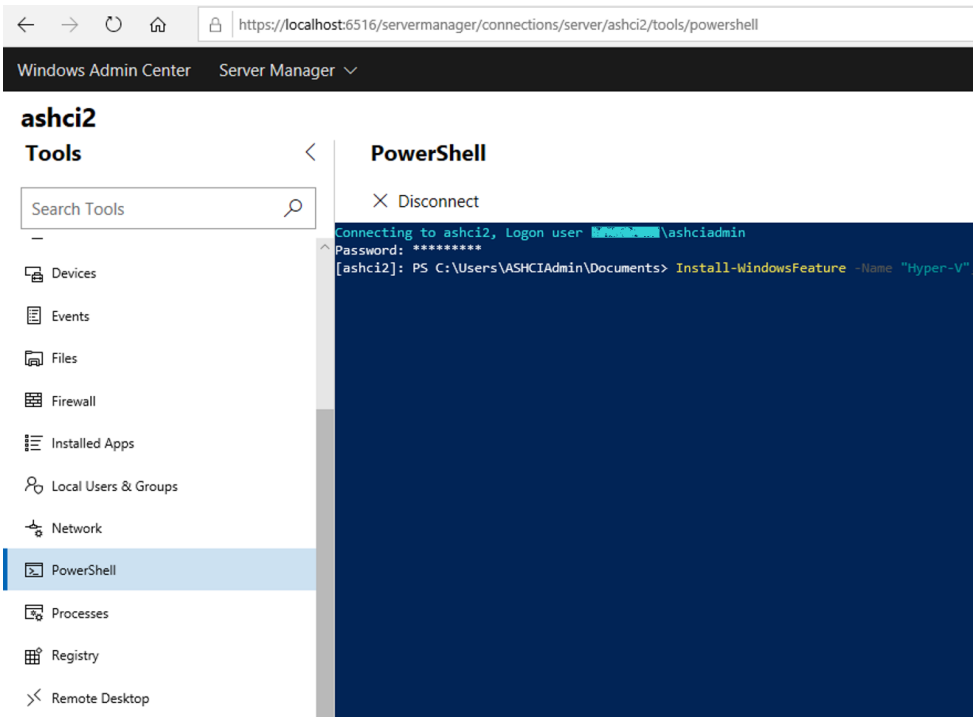

Figure 4. Windows Admin Center allows you to connect to a PowerShell prompt on remote machines

Installing Windows Admin Center

Windows Admin Center is installed independently of Windows Server 2019, and its MSI file is available for download here. For remote administration, the tool can be installed on any machine (except a domain controller) running Windows 10 or Windows Server 2012 or later and running a modern browser such as Microsoft Edge or Google Chrome. Running the MSI file sets up a web server (Internet Information Services [IIS] is *not* required) on the local machine, which communicates over port 6516 by default. It also installs a component called the Windows Admin Center gateway. The gateway enables management of remote servers through remote PowerShell and Windows Management Instrumentation (WMI) over Windows Remote Management (WinRM). If you are not using domain admin credentials to manage the remote servers, you should let Windows Admin Center manage your TrustedHosts lists. Otherwise, you must configure TrustedHosts manually.

When you start Windows Admin Center from the Start menu, it opens in your default browser. It can be configured to connect to your Azure Stack HCI resources either locally or through the internet.

Azure service integrations in Windows Admin Center

The following Azure services are directly available from Windows Admin Center. These features are optional, but they can facilitate administration, improve service availability, help ensure compliance, and provide other benefits to organizations:

- Azure Site Recovery provides high availability and disaster-recovery-as-a-service (DRaaS).
- Azure Monitor is a centralized hub to track what's happening across your applications, network, and infrastructure, with advanced analytics powered by AI. It can send you text messages or email alerts when failures are detected.
- Cloud Witness allows you to use Azure as the lightweight tiebreaker for cluster quorum.
- Azure Backup provides offsite data protection and helps protect against ransomware.
- Azure Update Management performs update assessment and update deployments for Windows VMs running both in Azure and on premises.
- Azure Network Adapter allows you to connect on-premises resources with your VMs in Azure via a point-to-site virtual private network (VPN).
- Azure Security Center offers threat detection and monitoring services for VMs running both in Azure and on premises.

Other administration tools

Windows Admin Center provides a foundational set of administration features to manage Azure Stack HCI, but other tools complement these capabilities.

For example, you can use PowerShell for scripting and automation. And with PowerShell Direct, you can even run PowerShell on a VM from your Hyper-V host, regardless of the network configuration or remote management setting.

Next, you can use Microsoft System Center 2019 to perform the following administrative functions not available through Windows Admin Center:

- Manage heterogeneous networks, including VMware VMs and Linux servers.
- Deploy from bare metal.
- Monitor systems with the help of alerts and notifications.
- Manage systems at datacenter scale.

Azure Stack HCI: Windows Server 2019 Datacenter components

The following section describes the roles and features within Windows Server 2019 Datacenter that make up the Azure Stack HCI solution: Hyper-V, SDN, and Storage Spaces Direct.

Azure Stack HCI virtualization: Hyper-V

Hyper-V, the hypervisor and hardware-virtualization platform for Azure Stack HCI, lets you create and run guest VMs that run either Windows or Linux operating systems. (See here for a list of the Windows and Linux guest operating systems supported on Hyper-V.)

Software architecture of Hyper-V

Hyper-V is based on a thin, micro-kernelized ("type 1") hypervisor that runs directly on the hardware and that uses isolated logical units or "partitions" to separate the host operating system and all guest VMs from each other. Figure 5 provides a high-level overview of the architecture of the Hyper-V environment.

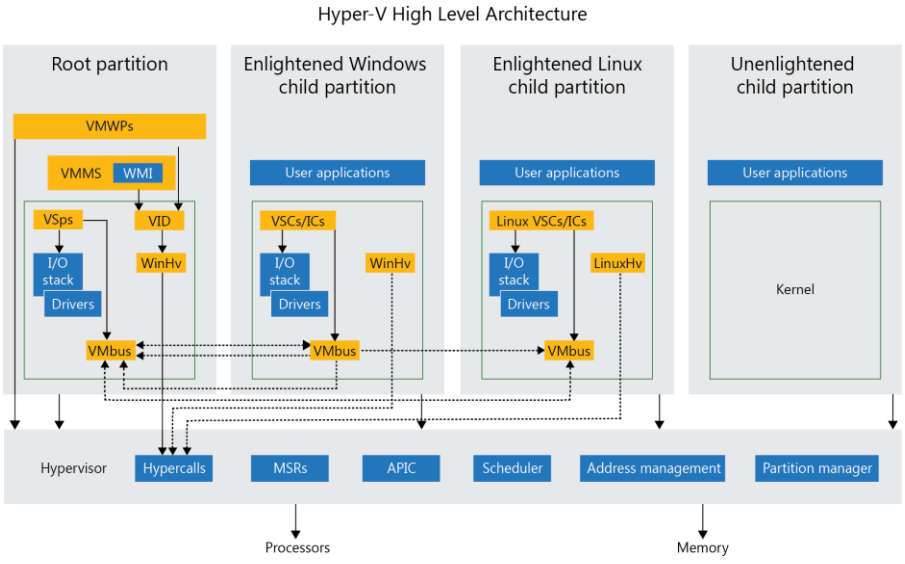

Figure 5. Hyper-V software architecture

The terms used in the diagram are defined in the following list:

- **Root partition.** The partition that runs the host operating system. This partition manages machine-level functions such as device drivers, power management, and the "hot" addition and removal of devices. The root (or parent) partition is the only partition that has direct access to physical memory and devices.
- **Child partitions.** Partitions, or VMs, that host a guest operating system. All access to physical memory and devices by a child partition is provided via Virtual Machine Bus (VMBus) or the hypervisor.
- **Virtual Machine Bus (VMBus).** A channel-based communication mechanism used for inter-partition communication and device enumeration on systems with multiple active virtualized partitions. VMBus is installed with Hyper-V Integration Services and built into Windows, Windows Server, and Linux.
- **Enlightened guest.** A virtualization-aware instance of an operating system that can communicate directly with VMBus and bypass device emulation.
- **Hyper-V Virtual Machine Management service (VMMS).** A service responsible for managing the state of all VMs in child partitions.
- **Virtualization WMI provider.** A set of APIs exposed by VMMS that is responsible for managing and controlling VMs.
- **Virtualization Service Provider (VSP).** A software feature that resides in the root partition and provides synthetic device support to child partitions over VMBus.
- **Virtualization Infrastructure Driver (VID).** A driver that provides partition-management services, virtual processor-management services, and memory-management services for partitions.
- **Virtualization Service Client (VSC).** A synthetic device instance that resides in a VM. VSCs utilize hardware resources that are provided by VSPs in the parent partition. They communicate with the corresponding VSPs in the parent partition over VMBus to satisfy a child partition's device I/O requests.
- **Integration component (IC).** A component that allows child partitions to communication with other partitions and the hypervisor.
- **Hypercall.** An interface for communication with the hypervisor. The hypercall interface accommodates access to the optimizations provided by the hypervisor.
- **Advanced Programmable Interrupt Controller (APIC).** A device that allows priority levels to be assigned to its interrupt outputs.
- **Virtual Machine Worker Process (VMWP).** A process created for each running VM that provides management services from the operating system in the parent partition to the guest operating system.
- **Memory service routine (MSR).** A process used by the CPU to write data to or read data from memory.
- **Windows Hypervisor Interface Library (WinHv).** A bridge that allows guest operating system drivers to call the hypervisor by using standard Windows calling conventions.
- **Linux Hypervisor Interface Library (LinuxHv).** A bridge that allows guest operating system drivers to call the hypervisor by using standard Linux calling conventions.

Hyper-V administration tools

Windows Admin Center enables fine-grained management and monitoring of VMs running in your Azure Stack HCI cluster. Windows Admin Center also allows you to modify Hyper-V settings in a unified way for all hosts in an HCI cluster. Figure 6 shows a view of Windows Admin Center displaying VMs on the network.

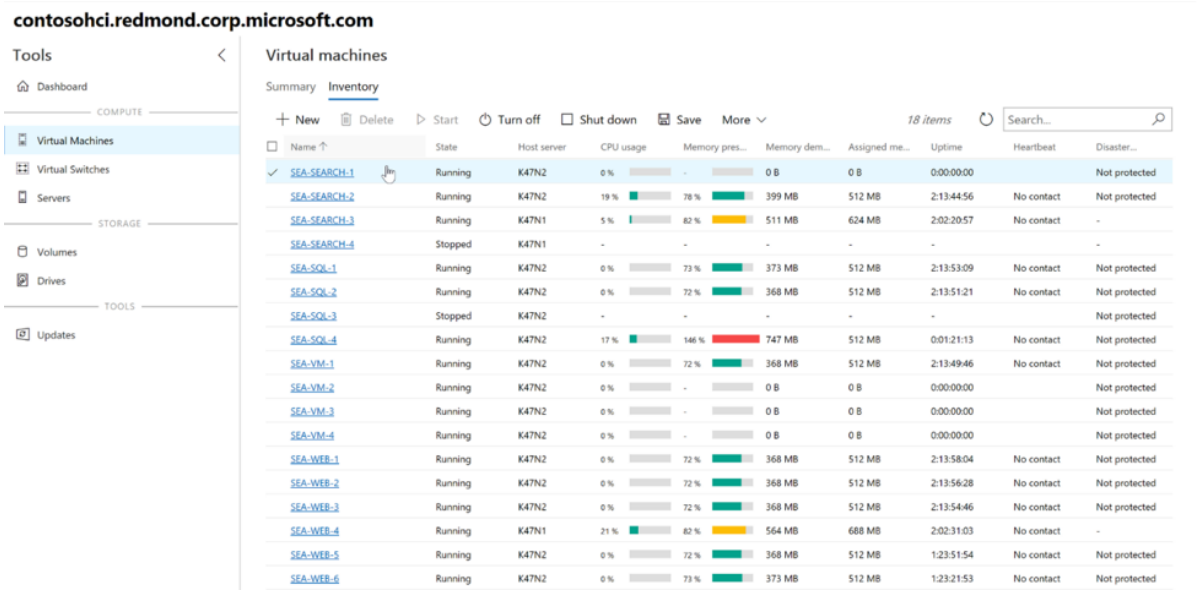

Figure 6. Using Hyper-V to view VMs hosted in an Azure Stack HCI cluster

Windows PowerShell provides another powerful tool for managing Hyper-V guests and hosts on your HCI cluster, and it also allows you to automate management through scripting. For a list of the cmdlets available in the Hyper-V module, see here.

You can also use System Center 2019 – Virtual Machine Manager (VMM) to perform even more powerful and comprehensive administration of your virtualized infrastructure. For more information about VMM, see "What is Virtual Machine Manager?"

New features in Windows Server 2019

The new Hyper-V features that are available for Azure Stack HCI in Windows Server 2019 improve on the shielded VMs feature, which was first introduced in Windows Server 2016. VM shielding lets the owner of a VM add an optional extra layer of security that protects the VM from being read in case the host is compromised. The same feature can also prevent the VM from running on unapproved hosts.

Shielded VMs rely on a service called the Host Guardian Service (HGS). HGS manages the keys used to start up shielded VMs.

Branch office improvements to shielded VMs in Windows Server 2019

You can now run shielded VMs on machines with intermittent connectivity to HGS by using the new fallback HGS and offline-mode features. Fallback HGS allows you to configure a second set of URLs for Hyper-V to try if it can't reach your primary HGS server.

Offline mode allows you to continue to start up your shielded VMs even if HGS can't be reached, as long as the VM has started successfully once and the host's security configuration has not changed.

Linux support for shielded VMs in Windows Server 2019

Windows Server 2019 now supports running Ubuntu, Red Hat Enterprise Linux, and SUSE Linux Enterprise Server inside shielded VMs.

Azure Stack HCI networking (optional): SDN

SDN is an optional component of Azure Stack HCI that allows you to centrally segment, configure, and manage your networks through software. Through SDN, network functions such as switching, routing, firewalling with microsegmentation, and load balancing can all be virtualized and optimized for availability, security, and performance. By decoupling network communication from physical hardware, SDN gives you the freedom to design your networks in a way that improves security and administrative efficiency.

You can use SDN to:

- Dynamically create, secure, and connect your network to meet the evolving needs of your apps.
- Speed up the deployment of your workloads in a nondisruptive manner.
- Contain security vulnerabilities from spreading across your network.
- Define and control policies that govern both physical and virtual networks.
- Consistently implement network policies at scale.

Note that the SDN technology available in Azure Stack HCI is the same SDN technology that is used across all global Azure datacenters. This feature is therefore well-tested at scale and is competitive with other SDN products in terms of performance and scalability.

Note that the SDN technology available in Azure Stack HCI is the same SDN technology that is used across all global Azure datacenters. This feature is therefore well-tested at scale and is competitive with other SDN products in terms of performance and scalability.

The following four Windows Server server roles are used for SDN in Azure Stack HCI:

- **Network Controller**. The Network Controller server role acts as the “brain” of the SDN stack. It’s the programmable point of automation that lets you manage, configure, monitor, and troubleshoot both the virtual and physical network infrastructure in your datacenter. With the help of the network controller, you can use Windows PowerShell, a Representational State Transfer (REST) API, or a management application to manage components such as Hyper-V VMs and virtual switches, physical switches and routers, firewalls, VPN gateways, and load balancers.

  More specifically, network policy is defined through JavaScript Object Notation (JSON) objects and then given to the network controller through a RESTful API. The network controller then sends this policy to the SDN host agent running on each Hyper-V host. The host agent then programs this policy in the Hyper-V virtual switch.
- **Network Virtualization**. The Network Virtualization server role enables you to abstract your workloads away from the physical network. Virtual networks provide multitenant isolation on a shared physical network fabric, thereby improving security while enabling you to use your network resources more fully.

- **Remote Access (RAS)**. The Remote Access server role in Windows Server 2019 includes the RAS gateway component, which routes traffic between the physical network and virtual network resources, regardless of location.
- **Software Load Balancer (SLB) for SDN**. SLB in Azure Stack HCI evenly distributes network traffic among virtual network resources, enabling multiple servers to host the same workload with high availability and scalability. The SLB consists of a *multiplexer (MUX)* that advertises a *virtual IP (VIP)* address to external clients (using Border Gateway Protocol [BGP]) and distributes connections across a set of dynamic IP (DIP) addresses assigned to VMs attached to a network.

SDN architecture

Figure 7 displays a simplified version of SDN architecture.

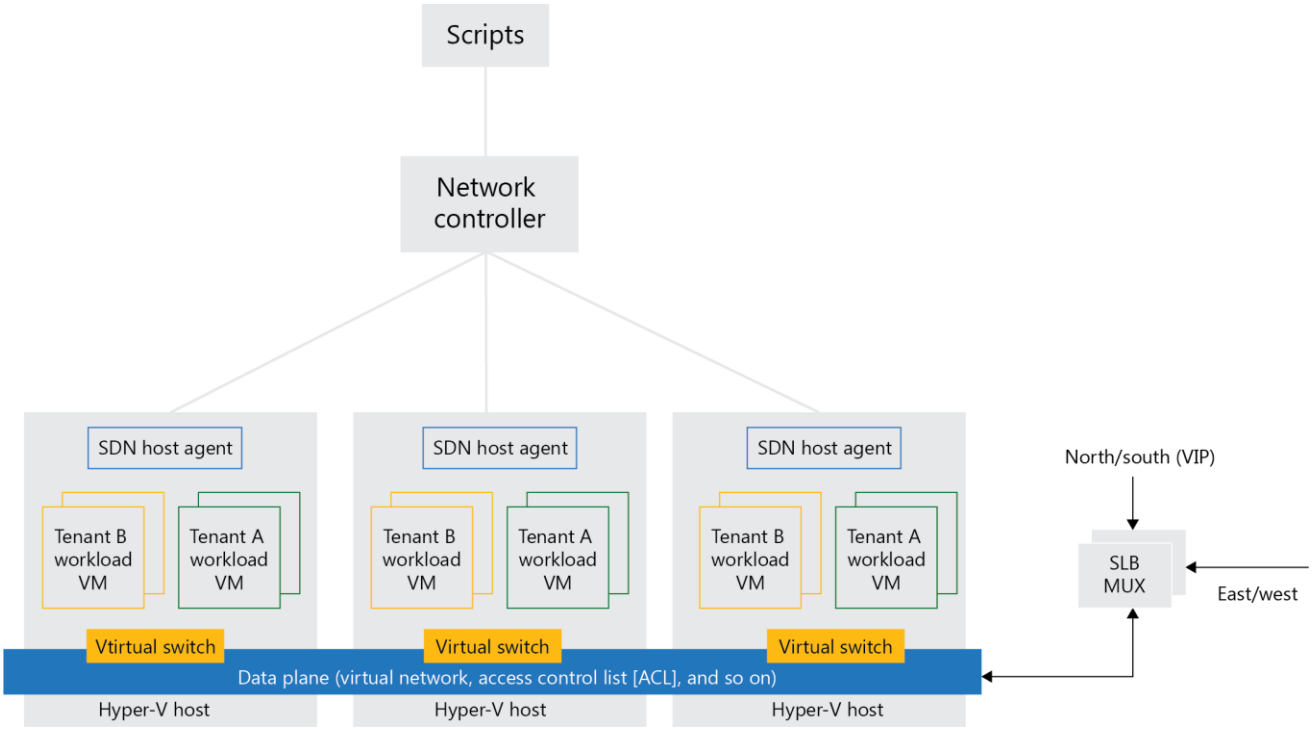

Figure 7. SDN architecture

North/south and east/west traffic

The terms "north/south" and "east/west" in Figure 7 refer to the origin and destination of network traffic. North/south indicates that the network traffic is going outside of the virtual network or datacenter. East/west indicates that the network traffic is coming from inside the virtual network or datacenter.

SDN administration tools

SDN configuration in Azure Stack HCI is typically completed, or at least largely facilitated, by the Microsoft partner supplying the Azure Stack HCI solution. Once Windows Admin Center is set up, you can use it to monitor infrastructure health and manage SDN features such as virtual networking, security, and load balancing.

New SDN features in Windows Server 2019

The following SDN features have been introduced with Windows Server 2019:

- **Encrypted networks**. When the "Encryption Enabled" option is selected on VMs on the same subnet, virtual network traffic is encrypted between them by using the Datagram Transport Layer Security (DTLS) protocol. DTLS protects against eavesdropping, tampering, and forgery by anyone with access to the physical network.
- **Firewall auditing**. Firewall auditing is a new capability for the SDN firewall in Windows Server 2019. When you enable the SDN firewall, any flow processed by SDN firewall rules and access control lists (ACLs) that have logging enabled gets recorded.
- **Virtual network peering**. Virtual network peering lets you connect two virtual networks seamlessly. Once peered, the virtual networks appear as one (for connectivity purposes).
- **Egress metering**. This new feature in Windows Server 2019 enables SDN to offer usage meters for outbound data transfers. With this feature added, the network controller keeps a whitelist per virtual network of all IP ranges used within the SDN, and it considers any packet bound for a destination that is not included in one of these ranges as an outbound data transfer.

Azure Stack HCI storage layer: Storage Spaces Direct

A key component of Azure Stack HCI is software-defined, shared-nothing storage. Storage Spaces Direct in Azure Stack HCI uses industry-standard servers with local-attached drives to provide this feature and create a pool of highly available, highly scalable storage at a cost that is much lower than that of traditional storage area network (SAN) or network-attached storage (NAS) arrays. Its hyperconverged architecture reduces management complexity, while features such as remote direct memory access (RDMA) networking and persistent read/write cache drives are used to maximize performance.

With Azure Stack HCI, Storage Spaces Direct runs on Hyper-V hosts that contribute their own local storage to the cluster storage pool. This eliminates the need to configure file server access and permissions, and it reduces hardware costs for small-to-medium business or remote office/branch office (ROBO) deployments.

Figure 8 shows a three-node Azure Stack HCI cluster in which all servers are contributing local storage. For each server, the two fastest drives (NVM Express [NVMe] drives) are dedicated to read/write caching.

Figure 8. Storage Spaces Direct uses locally attached storage, and the fastest (NVMe) drives are used for read/write caching

Storage Spaces Direct architecture

Storage Spaces Direct is built in part on features in Windows Server such as failover clustering, the Cluster Shared Volume (CSV) file system, Server Message Block (SMB) 3, and Storage Spaces. It also introduces new technology, most notably the Software Storage Bus, which enables servers to see all disks connected to each node in the cluster.

Figure 9 shows the simple Storage Spaces Direct stack used in Azure Stack HCI.

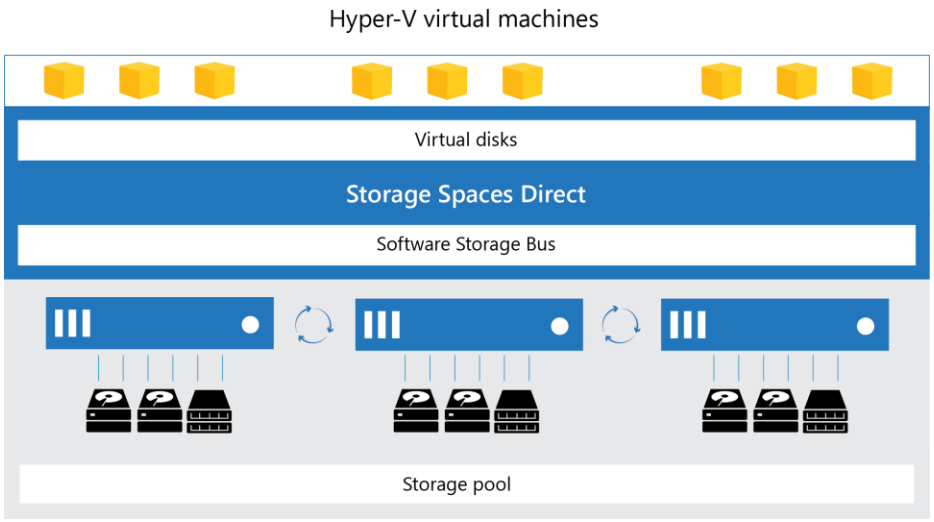

Figure 9. Graphical representation of Storage Spaces Direct

Storage Spaces Direct administration is supported in Windows Admin Center. Figure 10 presents a view of the locally attached drives in an Azure Stack HCI cluster that are used for either capacity or caching in Storage Spaces Direct.

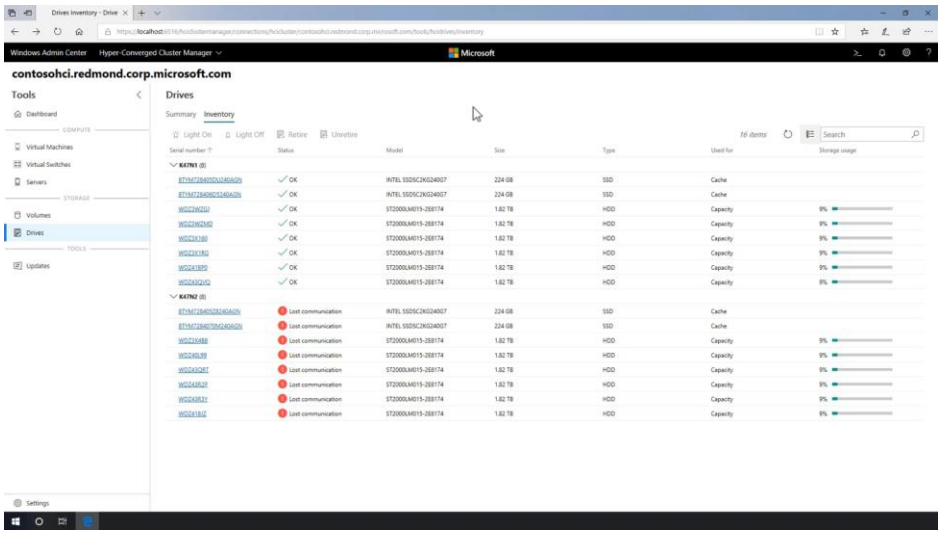

Figure 10. Windows Admin Center lets you fully manage Storage Spaces Direct

Components of Storage Spaces Direct

Storage Spaces Direct is built using the following components:

- **Networking hardware.** Storage Spaces Direct uses SMB3, including SMB Direct and SMB Multichannel, to communicate between servers. A 10-Gbps or faster connection is required, and RDMA is recommended (either through iWARP or RDMA over Converged Ethernet [RoCE]).
- **Storage hardware.** From 2 to 16 servers with local-attached Serial ATA (SATA), serial-attached SCSI (SAS), or NVMe drives. Each server must have at least two solid-state drives (SSDs) and at least four additional drives. The SATA and SAS devices should be behind a host-bus adapter (HBA) and SAS expander.
- **Failover clustering.** The built-in clustering feature of Windows Server is used to connect the servers.
- **Software Storage Bus.** The Software Storage Bus spans the cluster and establishes a software-defined storage fabric whereby all the servers can see all of each other's local drives.
- **Storage Bus Layer Cache.** The Software Storage Bus dynamically binds the fastest drives present (for example, SSDs) to slower drives (for example, hard-disk drives [HDDs]) to provide server-side read/write caching that accelerates I/O and boosts throughput.
- **Storage pool.** A collection of drives that forms the basis of Storage Spaces Direct. The storage pool is automatically created, and all eligible drives are automatically discovered and added to it.
- **Resilient File System (ReFS).** ReFS is the premier file system for virtualization. It includes significant accelerations for .vhdx file operations such as creation, expansion, and checkpoint merging, in addition to built-in checksums to detect and correct bit errors. It also introduces real-time tiers that rotate data between so-called “hot” and “cold” storage tiers in real time, based on usage.
- **Cluster Shared Volumes.** The CSV file system unifies all of the ReFS volumes into a single namespace accessible through any server, so that to each server, every volume looks and acts like it's mounted locally.

What’s new in Storage Spaces Direct in Windows Server 2019

The following improvements to Storage Spaces Direct have been introduced in Windows Server 2019:

- **Deduplication and compression for ReFS volumes**. Data deduplication is now supported on ReFS volumes, thereby allowing you to store up to 10 times more data on the same volume.
- **Nested resiliency**. This feature provides software fault tolerance equivalent to hardware RAID 5+1. With nested resiliency, a two-node Storage Spaces Direct cluster can provide continuously accessible storage for apps and VMs, even if one server node goes down and a drive fails in the other server node.
- **Two-server clusters using a USB flash drive as a witness**. You can use a low-cost USB flash drive plugged into your router to act as a witness in two-server clusters. If a server goes down and then comes back up, the USB drive cluster knows which server has the most up-to-date data.
- **Performance history**. Storage Spaces Direct now provides visibility into historical resource utilization and performance. More than 50 essential counters spanning compute, memory, network, and storage are automatically collected and stored on the cluster for up to one year. You can view this history in Windows Admin Center, and you can query it in PowerShell.

- **Support for up to 4 PB per cluster**. In Windows Server 2019, Storage Spaces Direct supports up to 4 petabytes (PB), or 4,000 terabytes, of raw capacity per storage pool. Related capacity guidelines are increased as well: for example, you can create twice as many volumes (64 instead of 32), each twice as large as before (64 TB instead of 32 TB).
- **Drive-latency outlier detection**. You can now easily identify drives with abnormal latency with proactive monitoring and built-in outlier detection, inspired by the long-standing and successful approach of Azure. Whether it's average latency or something more subtle like 99th percentile latency that stands out, slow drives are automatically labeled in PowerShell and Windows Admin Center with an "abnormal latency" status.
- **Support for Cluster-Aware Updating (CAU)**. CAU is now Storage Spaces Direct-aware, validating and ensuring that data resynchronization completes on each node. CAU inspects updates to restart only if necessary. This enables orchestrating restarts of all servers in the cluster for planned maintenance.

Core Azure Stack HCI technology: failover clustering

The foundation of all Azure Stack HCI clusters is a failover cluster. A failover cluster is a group of independent computers that work together to increase the availability and scalability of hosted VMs or applications, also called clustered roles. The servers (nodes) in the cluster are connected by physical cables and by software. If one or more of the cluster nodes fail, the clustered roles hosted on the node fail over to one of the remaining nodes. In addition, the clustered roles are proactively monitored to verify that they are working properly. If they are not working, they are restarted or moved to another node.

Failover clusters also contribute CSV functionality to Azure Stack HCI. CSVs provide a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes. With the failover clustering feature, users experience few disruptions in service.

What's new in failover clustering in Windows Server 2019

The following enhancements to failover clustering have been introduced in Windows Server 2019:

- **Cluster sets.** Cluster sets enable you to increase the number of servers in a single software-defined datacenter (SDDC) solution beyond the current limits of a cluster. This is accomplished by grouping multiple clusters into a cluster set—a loosely coupled grouping of multiple failover clusters: compute, storage, and hyperconverged. With cluster sets, you can move ("live migrate") online VMs between clusters within the cluster set.
- **Azure-aware clusters**. Failover clusters now automatically detect when they're running in Azure infrastructure-as-a-service (IaaS) VMs and optimize the configuration to provide proactive failover and logging of Azure planned maintenance events to achieve the highest levels of availability. Deployment is also simplified by removing the need to configure the load balancer with a dynamic network name for the cluster name.
- **Cross-domain cluster migration**. Failover clusters can now dynamically move from one Active Directory domain to another, simplifying domain consolidation and allowing clusters to be created by hardware partners and joined to the customer's domain later.
- **USB witness**. You can now use a simple USB drive attached to a network switch as a witness in determining quorum for a cluster. This feature enables a file share to provide witness support for any SMB2-compliant device.
- **Cluster hardening**. Intra-cluster communication over SMB for CSVs and Storage Spaces Direct now makes use of certificates to provide the most secure platform. This feature allows failover clusters to operate with no dependencies on NT LAN Manager (NTLM) and to enable security baselines.

Intel innovations boost Azure Stack HCI

Windows Server is on the leading edge of x86 hardware innovation, and Intel and Microsoft often work together to co-engineer new hardware and software optimizations on the server platform. Consequently, Windows Server is consistently one of the first platforms and hypervisors to support new Intel hardware solutions that can greatly enhance performance.

Intel Optane technology

Azure Stack HCI supports Intel Optane technology in Intel Optane DC persistent memory modules and in Intel Optane DC SSDs. You can use persistent memory in Memory mode to maximize the number of VMs hosted per node, at a lower cost per VM, than is possible with all-DRAM solutions. You can also use Intel Optane DC SSDs as cache drives to deliver heightened throughput for Storage Spaces Direct when SATA-based SSD or HDD drives are used for capacity. The power of Intel Optane technology to accelerate Azure Stack HCI has been demonstrated with a world-record performance of nearly 13.8M IOPS.

Intel Ethernet Network Adapters for RDMA

Internet Wide-Area RDMA Protocol (iWARP) is one of the network protocol options used to implement RDMA, a required technology for Storage Spaces Direct in Azure Stack HCI. The following Intel network adapters support iWARP:

- Intel Ethernet Network Adapter X722-DA2
- Intel Ethernet Network Adapter X722-DA4

In addition, the new Intel Ethernet 800 Series controllers, such as the Intel Ethernet Controller E810, support both iWARP and the second RDMA protocol option, RDMA over Converged Ethernet version 2 (RoCE v2). For more information about Intel Ethernet 800 Series controllers and other Intel innovations for Ethernet, see www.intel.com/ethernet.

Intel Select Solutions for Azure Stack HCI

Intel Select Solutions are predefined combinations of Intel compute, memory, storage, and network products that are designed to support specific workloads in basic ("Base") and advanced ("Plus") configurations. Intel has released Intel Select Solutions for Azure Stack HCI, which are available here.

Hardware partners

More than 150 Azure Stack HCI solutions are available today from 15+ Microsoft hardware partners. These partners offer systems validated by Microsoft to ensure optimal performance and reliability for Azure Stack HCI. Microsoft partners can also help get you up and running by reducing the difficulty of building, installing, configuring, and testing the hardware and software. In addition, many Microsoft partners can offer a single point of contact to support Azure Stack HCI after deployment.

How to find Azure Stack HCI solutions

You can contact your preferred hardware vendor about Azure Stack HCI solutions at www.microsoft.com/en-us/cloud-platform/azure-stack-hci-catalog.

[1] Microsoft. "The new HCI industry record: 13.7 million IOPS with Windows Server 2019 and Intel® Optane™ DC persistent memory." October 2018. https://techcommunity.microsoft.com/t5/Storage-at-Microsoft/The-new-HCI-industry-record-13-7-million-IOPS-with-Windows/ba-p/428314.