

AZURE KUBERNETES SERVICE ON AZURE STACK HCI

Leverage your Azure Stack HCI investment to run containerized applications alongside traditional virtual workloads in highly available and resilient patterns. In this quick start guide, you will learn how to install and configure the Azure Kubernetes Service for Azure Stack HCI (AKS-HCI), using Windows Admin Center. From there, you will be able to explore additional scenarios around deployment, monitoring and management of containerized applications and workloads. You will perform the following tasks as part of this guide:

- Plan and configure Supermicro hardware and OS/tools to support the AKS-HCI deployment
- Step-by-step documentation to install and configure AKS-HCI including integration with Azure Arc

Introducing Azure Kubernetes Service on Azure Stack HCI

Azure Kubernetes Service on Azure Stack HCI (AKS-HCI) is a Kubernetes-based orchestrator that automates running containerized applications on clusters that use Azure Stack HCI. Orchestrators such as the open-source Kubernetes, automate much of the work involved with deploying and managing multiple containers. However, Kubernetes can be complex to set up and maintain. AKS-HCI helps simplify setting up Kubernetes on-premises, making it quicker to get started hosting Linux and Windows containers. You can evaluate both Azure Stack HCI and AKS-HCI free for the first 60 days.




Figure 1 - High-level AKS-HCI Architecture

From an architectural perspective, as shown above, AKS-HCI consists of several key components. At the base, is your Azure Stack HCI cluster, which not only serves as the platform for the Azure Kubernetes Service and containerized workloads, but also continues to be a powerful enterprise platform for traditional virtualized workloads.

On the left, you have a management cluster, running on a single virtual machine (VM). When you install AKS-HCI, a management cluster is automatically created and configured. This management cluster is responsible for provisioning and managing target clusters where workloads run. It includes an API server, and a Load Balancer. On the right-hand side, you have your target (workload) cluster. The target cluster is a highly available deployment of Kubernetes using Linux VMs for running Kubernetes control plane components as well as Linux worker nodes. Windows Server based VMs are used for the Windows worker nodes. All of this can be quickly and easily deployed using Windows Admin Center, or PowerShell.



AZURE KUBERNETES SERVICE ON AZURE STACK HCI

Why Azure Kubernetes Service on Azure Stack HCI?

Azure Stack HCI provides enterprise customers with a highly available, cost-efficient, flexible platform to run high-performance workloads. These workloads could run within traditional virtual machines, but also, within containers, ensuring you get the best utilization from your hyperconverged infrastructure. Azure Stack HCI provides a compelling platform for Kubernetes for a number of reasons; firstly, it allows you to quickly scale virtual machines, including container hosts, to optimize workloads across your physical hardware resources, and it's quick and easy to add additional physical resources to increase capacity.

In addition, Azure Stack HCI allows VMs, including container hosts, to migrate across physical server hosts, to optimize utilization and account for infrastructure maintenance. It also provides automated VM failover, within a local cluster, or a cluster stretched across sites. Within the VMs themselves, containers and pods can move across the container hosts, offering an extra level of redundancy across physical and virtual. And all of this is running on a stack, completely supported by Microsoft.

Before you get started

The following summarizes, at a high-level what you will need to run Azure Kubernetes Service on Azure Stack HCI.

- An Azure Stack HCI version 20H2 cluster with a minimum of 2 nodes. From a CPU perspective, it's recommended to have at least 16 cores, with 256GB memory, but this is largely determined by your workload requirements. A minimal install of AKS-HCI with a single target cluster will likely consume approximately 24-30GB RAM.
- We recommend 1TB storage capacity in your Azure Stack HCI storage pool used to store the AKS-HCI virtual machines, however again, this will be determined by your workload requirements.
- A system running Windows Admin Center. This should be registered with Azure, and in the same domain as the Azure Stack HCI cluster.
- External internet connectivity for both the Azure Stack HCI nodes, and the AKS-HCI components.
- An Azure subscription with the appropriate permissions.

You can also review [more detailed requirements information within our documentation](#).

Deploying Azure Kubernetes Service on Azure Stack HCI

The first step in the process is to obtain and setup your hardware that will support Azure Stack HCI, and subsequently, AKS-HCI.

Step 1: Hardware and OS configuration for Azure Kubernetes Service on Azure Stack HCI

Supermicro offers next generation X12 BigTwin 2U 4-node, Ultra, UP single processor server, X11 Ultra, as well as AMD based A+ H12 WIO 1U systems, with highest performance and flexibility, in 1U and 2U form factors, for supporting the AKS-HCI scenario.

Below are the supported configurations:



AZURE KUBERNETES SERVICE ON AZURE STACK HCI

A+ H12 WIO 1U All-Flash:








Supermicro AS -1114S-WN10RT-HCI – All Flash NVMe

Scale

2 to 16 nodes

Single Node Data

-  8 to 64 cores (AMD 3rd Gen)
-  128GB to 4TB memory
-  15.36TB to 76.8TB raw storage
-  NVMe (Gen4)
-  Up to 100GbE








Supermicro AS -1114S-WN10RT-HCI – All Flash SSD

Scale

2 to 16 nodes

Single Node Data

-  8 to 64 cores (AMD 3rd Gen)
-  128GB to 4TB memory
-  960GB to 76.8TB raw storage
-  SATA SSD
-  Up to 100GbE

X12 BigTwin 2U 4-node All-Flash:








Supermicro SYS-220BT-HNC8R-HCI – All Flash NVMe

Scale

4 to 16 nodes (1 to 4 systems)

Single Node Data

-  16 to 72 cores (intel 3rd Gen)
-  128GB to 4TB memory
-  15.36TB to 46.08TB raw storage
-  NVMe (Gen4)
-  Up to 100GbE

AZURE KUBERNETES SERVICE ON AZURE STACK HCI

X12 Ultra 1U All-Flash:








Supermicro SYS-120U-TNR- HCI – All Flash NVMe + SSD

Scale

2 to 16 nodes

Single Node Data

-  16 to 80 cores (intel 3rd Gen)
-  128GB to 4TB memory
-  3.75TB to 75TB raw storage
-  NVMe + SSD
-  Up to 100GbE

X12 UP 1U All-Flash:








Supermicro SYS-510P-MR – HCI – All Flash NVMe

Scale

2 to 4 nodes

Single Node Data

-  8 to 32 cores (intel 3rd Gen)
-  64GB to 256GB memory
-  4TB to 32TB raw storage
-  NVMe
-  25GbE








Supermicro SYS-510P-MR – HCI – All Flash SSD

Scale

2 to 4 nodes

Single Node Data

-  8 to 32 cores (intel 3rd Gen)
-  64GB to 256GB memory
-  3.84TB to 30.4TB raw storage
-  SSD
-  25GbE

AZURE KUBERNETES SERVICE ON AZURE STACK HCI

X11 Ultra 1U/2U All-flash and Hybrid:



Supermicro SYS-2029U-TN24R4T- HCI

Scale

2 to 4 nodes

Single Node Data

- 8 to 56 cores (intel)
- 128GB to 6TB memory
- 4TB to 367TB raw storage
- NVMe
- Up to 100GbE



Supermicro SYS-6029U-E1CR4-HCI – ALL Flash

Scale

2 to 16 nodes

Single Node Data

- 8 to 56 cores (intel)
- 128GB to 6TB memory
- 12TB to 96TB raw storage
- NVMe + SSD
- Up to 100GbE



Supermicro SYS-6029U-E1CR4-HCI - Hybrid

Scale

2 to 16 nodes

Single Node Data

- 8 to 56 cores (intel)
- 128GB to 6TB memory
- 12TB to 96TB raw storage
- NVMe + SSD + HDD
- Up to 100GbE



Supermicro SYS-1029U-TN10RT-HCI

Scale

2 to 16 nodes

Single Node Data

- 8 to 56 cores (intel)
- 128GB to 6TB memory
- 4TB to 153TB raw storage
- NVMe
- Up to 100GbE



Supermicro SYS-1029U-TRT – HCI - Hybrid

Scale

2 to 4 nodes

Single Node Data

- 8 to 56 cores (intel)
- 128GB to 6TB memory
- 4TB to 60.8TB raw storage
- NVMe + SSD + HDD
- Up to 40GbE



Supermicro SYS-1029U-TRT – HCI – All Flash

Scale

2 to 4 nodes

Single Node Data

- 8 to 56 cores (intel)
- 128GB to 6TB memory
- 4TB to 60.8TB raw storage
- NVMe + SSD
- Up to 40GbE

With your hardware racked, configured, and connected, you are ready to deploy the Azure Stack HCI OS. The first step in deploying Azure Stack HCI is to [download Azure Stack HCI](#) and install the operating system on each server that you want to cluster. You can [deploy Azure Stack HCI using your preferred method](#) – this could be via USB, network deployment, ISO boot over a dedicated OOB management port etc, Step through the simple Azure Stack HCI OS installation wizard, and



AZURE KUBERNETES SERVICE ON AZURE STACK HCI

once complete, you should be at the Server Configuration Tool (SCONFIG) interface. If you need to, make any simple changes here, but all that should be required is a single NIC with an IP address on your management network.




Figure 2 - SCONFIG in Azure Stack HCI

Step 2: Deploy and Configure Windows Admin Center

With your Azure Stack HCI nodes deployed, and accessible over the network, the next step is to deploy the Windows Admin Center. If you haven't already, [download the Windows Admin Center software](#). This should be installed on a Windows 10 or Windows Server 2016/2019 machine. This machine should also be joined to your management domain. This should be the same domain that your Azure Stack HCI nodes will be joined to.

Step 3: Create Azure Stack HCI Cluster

With the Windows Admin Center installed, open the Windows Admin Center, and [step through the process of creating an Azure Stack HCI cluster](#).

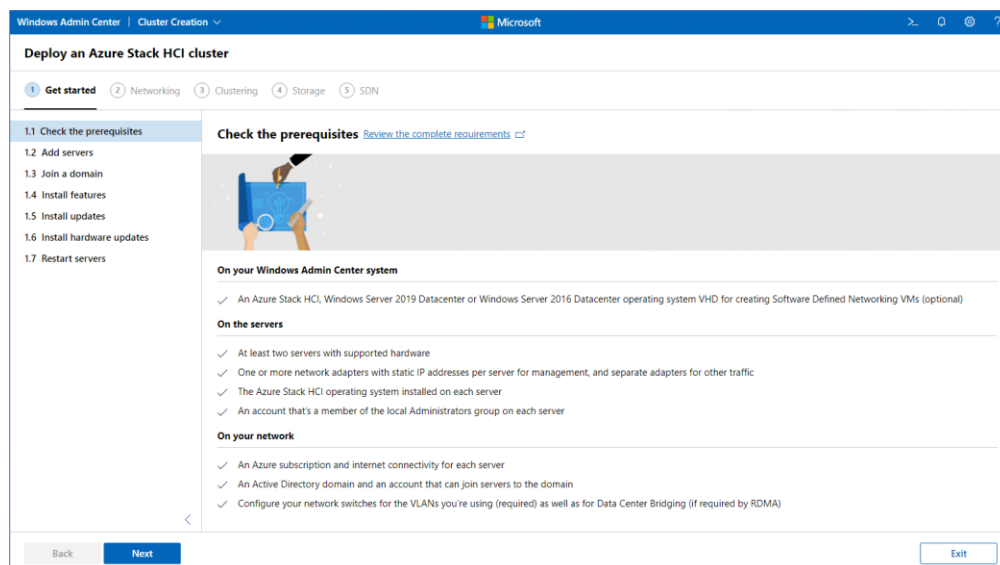


Figure 3 - Deploying an Azure Stack HCI Cluster in Windows Admin Center



AZURE KUBERNETES SERVICE ON AZURE STACK HCI

The wizard will walk you through selecting your nodes, joining the nodes to the domain, installing required roles and features, and updates, before moving on to configuring the physical and virtual networks, clustering and software defined storage. When the wizard is complete, you should see your new cluster in your **All connections** view within Windows Admin Center.

Step 4: Check for AKS-HCI Extension Updates

With your Azure Stack HCI cluster under management of Windows Admin Center, the next step is to check that the latest version of the Azure Kubernetes Service on Azure Stack HCI extension is installed. In **Settings**, under **Extensions**, then **Installed Extensions**, you should see the Azure Kubernetes Service extension installed, and not requiring an update. If an update is required, select the extension, and click **Update**.




Figure 4 – Validating the AKS-HCI Extension in Windows Admin Center

Step 5: Set up the Azure Kubernetes Service Host

There is one final step that should be completed before you create a Kubernetes cluster for running workloads, specifically, you need to set up an Azure Kubernetes Service management cluster. Firstly, within Windows Admin Center, make a connection to your Azure Stack HCI cluster, and once connected, under **Compute**, click on **Azure Kubernetes Service**, and in the central pane, click **Set up**, to launch the wizard.




Figure 5 - Azure Kubernetes Service in Windows Admin Center

[Follow the documented steps](#) to complete the simple wizard. The wizard will check your Azure Stack HCI cluster and your Windows Admin Center system, then allow you to configure the Azure Kubernetes Service management cluster, defining storage and networking, along with Azure Arc integration. **Note, the first 60 days usage of both Azure Stack HCI and AKS on Azure Stack HCI is free.** However, after this point, usage will be charged. Click for pricing information for [Azure Stack HCI](#) and [AKS on Azure Stack HCI](#). Upon completing the wizard, the deployment will begin, and once complete, you will be ready to deploy a Kubernetes cluster.



AZURE KUBERNETES SERVICE ON AZURE STACK HCI

Step 6: Creating a Kubernetes cluster

Now that the Azure Kubernetes Service is set up on your Azure Stack HCI cluster, it works similarly to the hosted Azure Kubernetes Service: you use the service to create Kubernetes clusters that run your containerized applications. These Kubernetes clusters are groups of VMs that act as worker nodes, running your application containers. These workers could be Linux or Windows-based. The Kubernetes cluster also contains a control plane, which consists of Kubernetes system services used to orchestrate the application containers, and a load balancer.

To start the process, on the homepage for Windows Admin Center, click **Add**, navigate to the Kubernetes clusters tile, and click **Create new** to launch the Kubernetes cluster wizard.




Figure 6 - Create a new Kubernetes cluster in Windows Admin Center

Walk [through the wizard](#), providing information about your desired Kubernetes cluster. You'll firstly optionally choose to [integrate the new Kubernetes cluster with Azure Arc](#), provide the details of your Azure Kubernetes Service management cluster, created earlier, and define your primary node pool size, which will run key system services. From there, you will walk through creating additional workload pools, running Windows, or Linux. You will define some networking settings, optionally enable Active Directory integration, then review your settings and create your cluster. Once completed, you can download your Kubernetes kubeconfig file, and if you chose to integrate with Azure Arc, click the link to be taken directly to the Azure portal for more configuration options.

Next Steps

With your Kubernetes cluster deployed, you are at a point where you can explore some more advanced scenarios, including:

- [Running Linux applications in AKS-HCI](#) - Run a multi-container application that includes a web front end and a Redis instance. You will also scale and test your application.
- [Running Windows applications in AKS-HCI](#) - run an ASP.NET sample application in a Windows Server container to the cluster. You then see how to test and scale your application.
- [Integrate your Kubernetes cluster with Azure Arc](#) - assuming you chose not to integrate during the Windows Admin Center wizard, this guide will walk you through integrating your Kubernetes cluster with Azure Arc.
- Once integrated with Azure Arc, you can proceed to [use Azure Policy to apply cluster configurations](#), [deploy configurations using GitOps](#), and [enable monitoring of your Kubernetes cluster](#).

Summary

Following this guide, you have installed Azure Stack HCI, deployed Windows Admin Center and integrated the Azure Kubernetes Service extension. You have then deployed the Azure Kubernetes Service management cluster onto your Azure Stack HCI cluster, and setup a Kubernetes cluster for your workloads. At this point, you are ready to continue your Kubernetes learnings by exploring the resources covered in the next steps above.

