

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 Lenovo 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

Lenovo certified Azure Stack HCI solutions can be found at this link – [ThinkAgile MX](#).

Lenovo rack systems feature innovative hardware, software and services that solve customer challenges today and deliver an evolutionary fit-for-purpose, modular design approach to address tomorrow's challenges. These servers capitalize on best-in-class, industry-standard technologies coupled with differentiated Lenovo innovations to provide the greatest possible flexibility in x86 servers. Key advantages of deploying Lenovo rack servers include:

- Highly scalable, modular designs to grow with your business
- Industry-leading resilience to save hours of costly unscheduled downtime
- Expansive storage capacity and flexible storage configurations for optimized workloads

AZURE KUBERNETES SERVICE ON AZURE STACK HCI

Fast flash technologies for lower latencies, quicker response times and smarter data management in real-time for cloud deployments, database, or virtualization workloads, trust Lenovo racks for world-class performance, power-efficient designs and extensive standard features at an affordable price.

The following Lenovo servers have been certified for Microsoft Azure Stack HCI and are equipped to support up to 2x 8-64-core processors, up to 3TB of memory and over 100TB of storage making them ideal candidates for Azure Stack HCI Virtualized Desktop Infrastructure solutions:

- Lenovo ThinkAgile MX3530 Integrated systems and ThinkAgile MX3531 validated nodes (based on SR650 V2)
- Lenovo ThinkAgile MX3330 Integrated systems and ThinkAgile MX3331 validated nodes (based on SR630 V2)
- Lenovo ThinkAgile MX3520 Integrated systems and ThinkAgile MX validated nodes (based on ThinkSystem SR650)
- Lenovo ThinkAgile MX1020 Integrated systems and ThinkAgile MX1021 validated nodes (based on SE350)
- Lenovo ThinkSystem SR630 validated nodes
- Lenovo ThinkSystem SR665 validated nodes
- Lenovo ThinkSystem SR655 validated nodes
- Lenovo ThinkSystem SR645 validated nodes
- Lenovo ThinkSystem SR635 validated nodes

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 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.

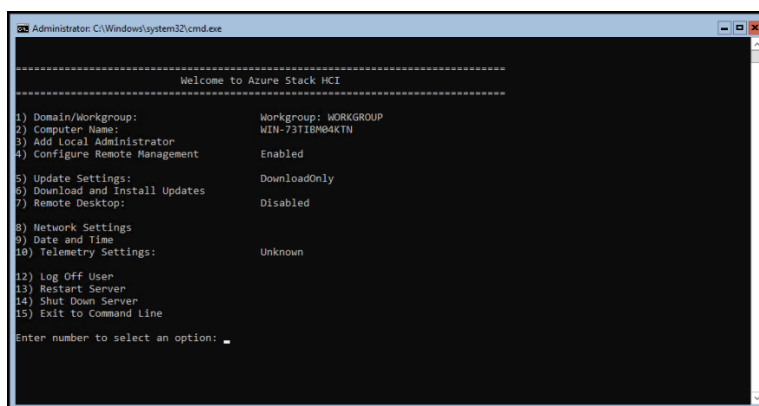


Figure 2 - SCONFIG in Azure Stack HCI

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

AZURE KUBERNETES SERVICE ON AZURE STACK HCI

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

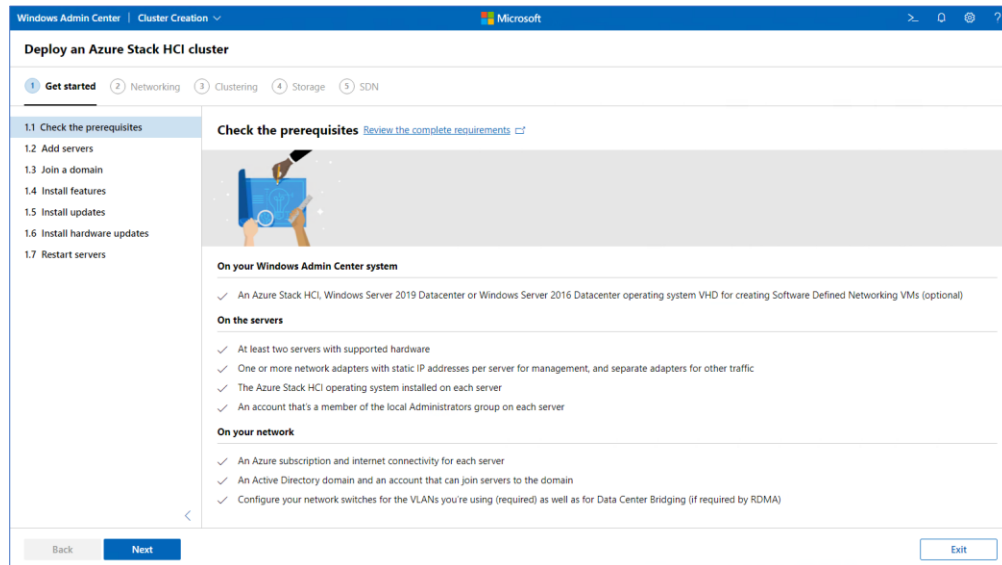


Figure 3 - Deploying an Azure Stack HCI Cluster in Windows Admin Center

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**.

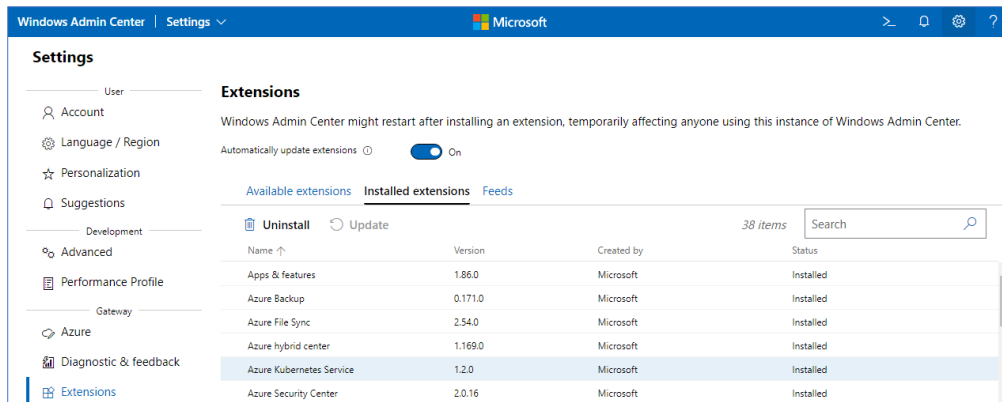


Figure 4 - Validating the AKS-HCI Extension in Windows Admin Center

AZURE KUBERNETES SERVICE ON AZURE STACK HCI

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.

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.

AZURE KUBERNETES SERVICE ON AZURE STACK HCI

Install [Windows Admin Center \(WAC\)](#) for managing Azure Stack HCI.

Microsoft Windows Admin Center (WAC) is a browser-based application that is deployed locally and used to manage Windows Servers, Windows Server Clusters and Azure Stack HCI clusters. Microsoft has made WAC extensible so that hardware partners can build additional features specific to their hardware and firmware. Lenovo XClarity Integrator is an example of one such extension implementation. Lenovo XClarity Integrator is designed to help users to manage and monitor the Lenovo ThinkSystem servers through Lenovo XClarity Administrator in Windows Admin Center. Lenovo XClarity Integrator and Windows Admin Center run in the same environment. Lenovo XClarity Integrator that is integrated with Lenovo XClarity Administrator can be used as an out-of-box management tool and a high-efficiency tool for managing and monitoring the Lenovo servers and components, for example, monitoring the overall status of servers, viewing the inventory of components, checking the firmware consistency of cluster nodes, and launching the management interface.

This [link](#) provides information on features in the Lenovo XClarity Integrator extension and instructions for installing the extension for Windows Admin Center.

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 on Lenovo ThinkAgile systems. 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.