

Managed Clustered Containers Platform (AKS Hybrid) Planning & Deployment

Description

A managed, clustered, containerised platform based on the Kubernetes clustering technology stack, with support for Linux and Windows containers, enabling a wide range of containerised workloads that operate on Azure Stack HCI and Windows Server 2022-based platforms.

The platform provides transparent patching, security and feature updates, with support for management via local and cloud-based interfaces.

Capabilities

- Fully automated deployment of Kubernetes clustered nodes onto base Windows Server 2022 or AS HCI 22H2 infrastructure using The Microsoft AKS platform, with remote management of workloads enabled via ARC for Kubernetes.
- Performance and scale of nodes (and therefore the number of containers that can be run) may be controlled via the number of resources (CPU's cores, RAM, Network capacity and storage) allocated to the virtual machines used to host each Kubernetes's node.
- Updates to Linux and Windows base operating systems, including feature and security updates, are applied automatically.
- Deploy and manage both Linux and Windows-based containerised apps.
- Scale up or down by adding or removing nodes to the Kubernetes cluster.
- Manage storage and networking on the Kubernetes cluster.
- Persistent storage is accessible to container-based workloads.

Remote management of the Kubernetes platform via the Azure portal using the following services:

- Azure Monitor
- Azure Policy
- Role-Based Access Control
- GitOps
- Windows Admin Center (located in Azure)
- Azure ARC for Kubernetes (in cluster management)

Local (on-prem) management of the Kubernetes platform is supported via the most common Kubernetes management solutions, including kubectl or other open-source tools.

Full support for Linux or Windows-based containers.

- Node pools (groups of identical Kubernetes cluster nodes) can support Linux containers, Windows containers, or both.

Linux and Windows nodes (virtual machines) are created as part of the solution, removing the requirement to directly manage the base Linux or Windows operating systems.



Support for network segmentation using Microsoft Software Defined Networking (SDN) on a per pod (application) level.

Dynamic scaling is supported if appropriate resources are available within the base clustered platform.

Limitations

Deployments that exceed the following specifications are not supported:

Resource	Maximum
Physical servers per cluster	8
Kubernetes Clusters	8
Total number of VMs	200

Requirements

- A maximum of eight servers in the cluster
- 1 TB of available capacity in the storage pool for Azure Kubernetes Service
- At least 30 GB of available memory (RAM) for running Azure Kubernetes Service VMs
- All servers in the cluster must use the EN-US region and language selection.