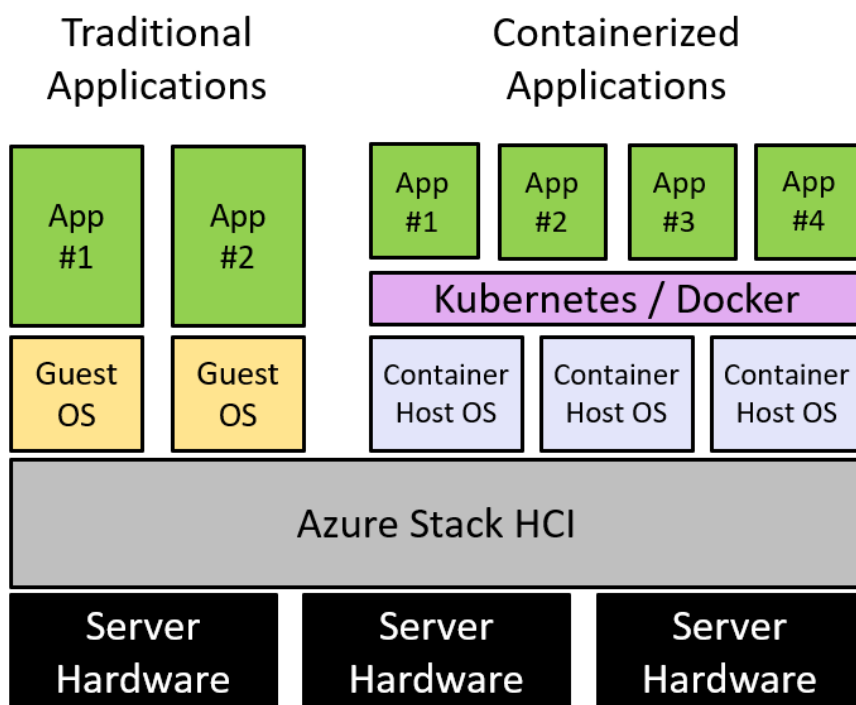


Leverage your Azure Stack HCI investment to run containerized applications alongside virtual workloads in highly-available and resilient patterns. See below for a how-to guide to installing and configuring one of the leading commercially supported Kubernetes Management Platforms available. In this example, [Rancher](#) is utilized for its' ease of use and simple management portal supporting all the latest Kubernetes features and functions.

- Plan and configure QCT hardware and OS/tools to support Azure Stack HCI hosting Rancher
- Step by step documentation to install and configure Rancher including VM setup, Docker installation and Rancher design

High-level architecture of Azure Stack HCI (supports 2 to 16 physical servers) hosting traditional applications in Virtual Machines and a Kubernetes cluster on Linux and Windows Server VMs to support Linux and Windows Containers.



Why Azure Stack HCI?

Flexibility:

- scale VMs (including Container hosts) to optimize workloads across physical hardware resources
- ease of adding physical resources (servers) to increase capacity
- support traditional virtualized applications simultaneously with containerized applications

Resiliency:

- VMs (including Container hosts) can move across the physical server in the event of a physical server or device outage
- Containers and pods can move across the Container hosts (nesting the resiliency on top of VMs)
- all VM and Containers storage operations are highly-available

1. Hardware and OS configuration for Rancher

- ◇ Server: [QuantaGrid D52B-1U](#); 2-16 nodes
 - CPU: Intel® Xeon® Scalable processors with Silver, Gold and Platinum options
 - RAM: 256 GB~768GB
 - NIC: 1x Quanta OCP Mezz CX4, Dual Port 25G or 1x Q.logic 41212, Dual Port 25G
 - Capacity: 4~12x NVMe SSD 3.2TB

- ◇ Server: [QuantaGrid D52BQ-2U](#); 2-4 nodes (Performance)
 - CPU: Intel® Xeon® Scalable processors with Silver, Gold and Platinum options
 - RAM: 256 GB~768GB
 - HBA: SAS 9305-16i
 - NIC: 1x Quanta OCP Mezz CX4, Dual Port 25G or 1x Q.logic 41212, Dual Port 25G
 - Capacity: 4~24x SATA SSD 480GB/960GB/1.92TB
 -

- ◇ Server: [QuantaGrid D52BQ-2U](#); 2-4 nodes (Balanced)
 - CPU: Intel® Xeon® Scalable processors with Silver, Gold and Platinum options
 - RAM: 256 GB~768GB
 - HBA: SAS 9305-16i
 - NIC: 1x Quanta OCP Mezz CX4, Dual Port 25G or 1x Q.logic 41212, Dual Port 25G
 - Cache: 2~4x SATA SSD 480GB/960GB/1.92TB
 - Capacity: 4~8x SATA HDD 8TB/10TB

- ◇ Server: [QuantaGrid D52BM-2U](#); 2-4 nodes
 - CPU: Intel® Xeon® Scalable processors with Silver, Gold and Platinum options
 - RAM: 192 GB~768GB
 - HBA: SAS 9305-16i
 - NIC: 1x Quanta OCP Mezz CX4, Dual Port 25G or 1x Q.logic 41212, Dual Port 25G
 - Cache: 12x Optane DC Persistent Memory
 - Capacity: 4~16x SATA SSD 8TB/10TB

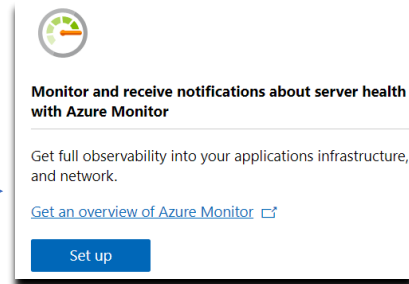
- ◇ Server: [QuantaGrid D52T-1ULH](#); 2-4 nodes
 - CPU: Intel® Xeon® Scalable processors with Silver, Gold and Platinum options
 - RAM: 256 GB~512GB
 - HBA: Avago SAS3 3008 Fury
 - NIC: 2x Mellanox ConnectX-4, Dual Port 25G
 - Cache: 2~4x SATA SSD 480GB/960GB/1.92TB
 - Capacity: 4~12x SATA HDD 8TB/10TB

- ◇ Switch: 2x TOR [QuantaMesh T4048-IX8D](#) and 1x BMC [QuantaMesh T1048-LY4R](#)
- ◇ OS: Windows Server 2019 Datacenter Edition

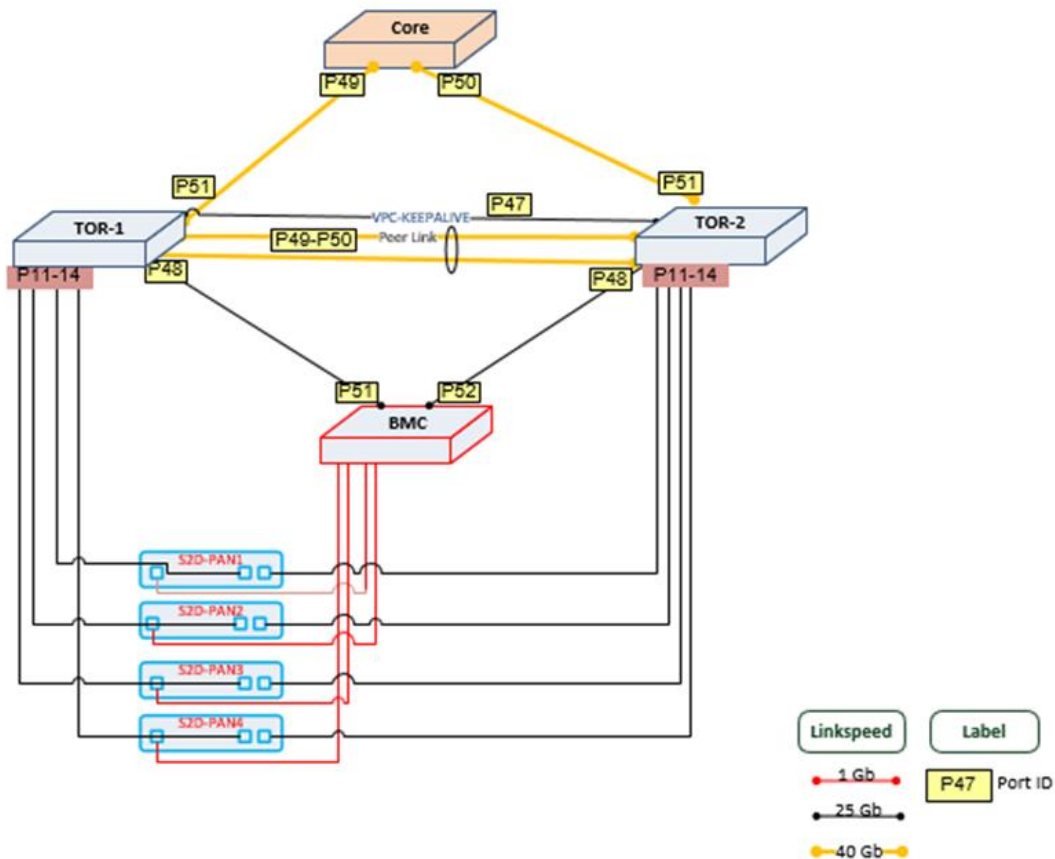
[Step by Step guide to deploy Azure Stack HCI](#)

1. Install Windows Server Datacenter
2. Add Roles and Features
3. Setup Failover Clustering and enable a Cluster Witness
4. Setup Storage Spaces Direct
5. [Install Windows Admin Center \(WAC\)](#)

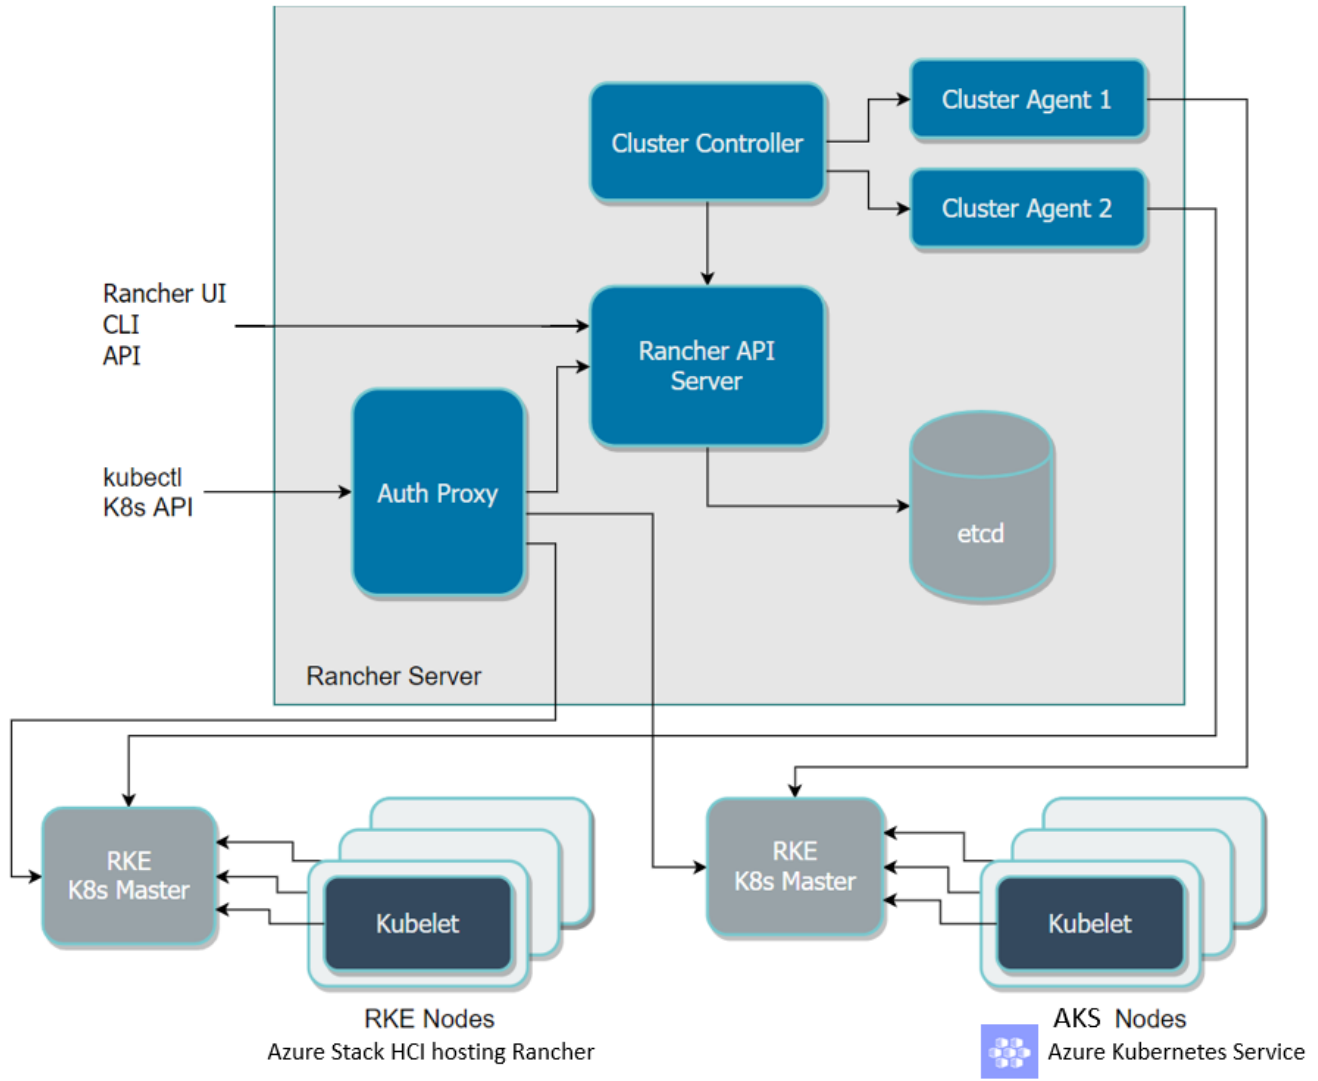
From Windows Admin Center (WAC), Set up **Azure Monitor** to gain insight into your application, VM, container, network and server health of your Azure Stack HCI-based Rancher instantiation.



2. Server Nodes and Network Diagram



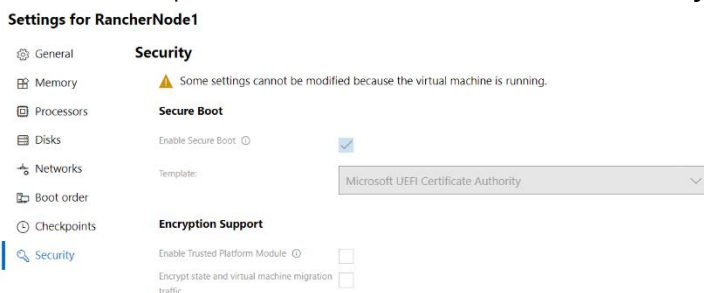
- High-level Architecture of Rancher 2. The figure, below, depicts a Rancher installation that manages two Kubernetes clusters: one created by Rancher Kubernetes Engine (RKE) on Azure Stack HCI and another created by Azure Kubernetes Service (AKS).



4. Prepare Azure Stack HCI for Rancher
 - a. See [Rancher installation requirements](#) to see a list of supported Operating Systems and Rancher roles
 - b. Read the [Best Practices for running Linux on Hyper-V](#) to optimize Hyper-V performance with Linux VMs
 - c. From Windows Admin Center (WAC)
 - i. Navigate to Hyperconverged Cluster Manager
 - ii. Select Virtual Machines and view Inventory
 - d. Create at least one Hyper-V Gen2 VM to install Linux with Docker container support

Note: this Quick Start uses Ubuntu Server 18.04 LTS (the [.ISO file](#) is available for download from [ubuntu.com](#))

 - i. Choose all default Settings for your Gen2 VM except increase the RAM to 8GB and for Operating System, choose "Install an Operating System from an ISO file" and reference the Linux ISO that you downloaded
 - ii. For a Gen2 VM to boot Linux, you must change Settings on the VM for Secure Boot Template to the **Microsoft UEFI Certificate Authority** (see WAC screenshot below).

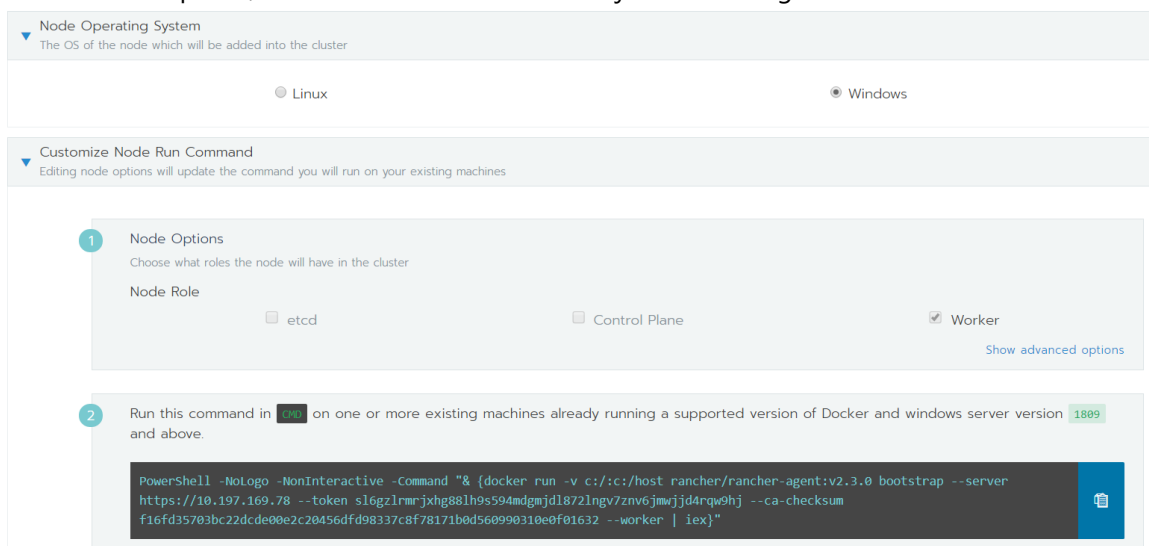


iii. During Linux installation in your VMs, add SSH Server and Docker options (*varies by Linux distribution selected*).

5. [Install Rancher](#) via Docker on the Linux VM(s) by running the following command


```
$ sudo docker run -d --restart=unless-stopped -p 80:80 -p 443:443 rancher/rancher
```

 - e. [Login to Rancher](#) - Access the Rancher portal by browsing to the hostname or IP address of your new Rancher VM
 - f. [Create a Kubernetes cluster](#) in the Rancher portal
 - g. To support Windows Containers, you must enable Windows Server cluster node support as a Worker role. In the Rancher portal, edit the Kubernetes cluster to join and configure the Windows Server VM.



6. Choose a Kubernetes storage option that aligns to your requirements



* QCT, the QCT logo, Rackgo, Quanta, and the Quanta logo are trademarks or registered trademarks of Quanta Computer.
 * Intel, the Intel logo, Optane, and Xeon Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries. All trademarks and logos are the properties of their respective holders.

- h. [Utilize local path provisioner](#). This provides a way for Kubernetes to utilize the local storage in each node.
- i. Utilize one of the CSI (Container Storage Interface) [Drivers](#) to provide persistent storage options to your Kubernetes cluster.

7. Connect to [Azure as a Rancher Cloud Provider](#)
 - j. In the Rancher UI, enable Azure resources to be integrated into your on-prem Rancher instantiation

Cloud Provider

None
 Amazon
 Azure
 Custom

[Read more about the Kubernetes cloud providers](#)

aadClientCertPassword
 The password of the client certificate for an AAD application with RBAC access to talk to Azure RM APIs

aadClientSecret
 The ClientSecret for an AAD application with RBAC access to talk to Azure RM APIs

aadClientCertPath
 The path of a client certificate for an AAD application with RBAC access to talk to Azure RM APIs

aadClientid
 The ClientID for an AAD application with RBAC access to talk to Azure RM APIs

- k. Enable integration with Azure storage services

rancher Cluster Nodes Storage Projects/Namespaces Members Tools

Add Storage Class

Name [Add a Description](#)

AzureStorageClass

Provisioner

Azure Disk

Parameters
Configure the provider-specific parameters for the storage class

Storage Account Type: Standard_LRS

Kind: Managed

Summary

Following this guide, you have installed Azure Stack HCI, deployed Linux VMs (and Windows Server VMs for Windows Container support), installed and configured Docker, installed Rancher and setup a Kubernetes cluster. You also evaluated Kubernetes storage options and are ready to [deploy workloads](#).

For production support of Rancher, Docker and Kubernetes, visit <http://www.rancher.com/support>