



Use Rancher on NetApp HCI

HCI

NetApp
September 16, 2021

This PDF was generated from https://docs.netapp.com/us-en/hci/docs/concept_rancher_product_overview.html on September 16, 2021. Always check docs.netapp.com for the latest.

Table of Contents

- Use Rancher on NetApp HCI 1
 - Rancher on NetApp HCI overview 1
 - Rancher on NetApp HCI concepts 3
 - Requirements for Rancher on NetApp HCI 4
 - Deploy Rancher on NetApp HCI 6
 - Post-deployment tasks 11
 - Deploy user clusters and applications 16
 - Manage Rancher on NetApp HCI 17
 - Monitor a Rancher on NetApp HCI implementation 17
 - Upgrade Rancher on NetApp HCI 19
 - Remove a Rancher installation on NetApp HCI 25

Use Rancher on NetApp HCI

Rancher on NetApp HCI overview

Rancher is a complete software stack for teams adopting containers. Rancher addresses the operational and security challenges of managing multiple Kubernetes clusters across different infrastructures, while providing DevOps teams with integrated tools for running containerized workloads.

Deploying Rancher on NetApp HCI deploys the Rancher control plane, also referred to as the *Rancher server*, and enables you to create on-premises Kubernetes clusters. You deploy the Rancher control plane by using the NetApp Hybrid Cloud Control.

After deployment, using the Rancher control Plane, you provision, manage, and monitor Kubernetes clusters used by Dev and Ops teams. Dev and Ops teams can use Rancher to perform activities on user clusters that reside on NetApp HCI itself, a public cloud provider, or any other infrastructure that Rancher enables.

Benefits of Rancher on NetApp HCI

- **Ease of installation:** You do not need to learn how to install and configure Rancher. You can deploy a template-based implementation, which was jointly developed by NetApp HCI and Rancher.
- **Lifecycle management:** In a manual Rancher implementation, updates for the Rancher server application or the Rancher Kubernetes Engine (RKE) cluster are not automated. Rancher on NetApp HCI provides the ability for updates to the management cluster, that includes the Rancher server and the RKE.

What you can do with Rancher on NetApp HCI

With Rancher on NetApp HCI, you can:

- Deploy services across cloud providers and your private cloud.
- Port the apps and data across a hybrid cloud architecture regardless of cloud location without compromising service-level agreements.
- Spin up cloud-native applications yourself.
- Centralize management of multiple clusters (new and existing).
- Perform orchestration of hybrid cloud Kubernetes-based applications.

Technical Support option

Using Rancher on NetApp HCI and Kubernetes open-source software includes free deployment and usage. License keys are not required.

You can choose a NetApp Rancher Support option to obtain core-based, Rancher enterprise support.




Rancher Support is not included in your NetApp Support Edge agreement. Contact NetApp Sales or your reseller for options. If you purchase Rancher Support from NetApp, you will receive an email with instructions.

Rancher on NetApp HCI architecture and components

Here is an overview of the various components of Rancher on NetApp HCI:

Rancher on NetApp HCI



- **NetApp Hybrid Cloud Control:** This interface enables you to deploy Rancher on NetApp HCI and NetApp Element software, required for Rancher on NetApp HCI.



You can use NetApp Hybrid Cloud Control also to upgrade management services, expand your system, collect logs, and monitor your installation.

- **Management services:** Management services run on the management node and enable you to deploy Rancher on NetApp HCI using NetApp Hybrid Cloud Control.
- **Management cluster:** Rancher on NetApp HCI deploys three virtual machines on the Rancher management cluster, which you can see using NetApp Hybrid Cloud Control, vCenter Server, or the Rancher user interface. The management cluster virtual machines host the Rancher server, the Rancher Kubernetes Engine (RKE), and the Linux OS.



For the best performance and greater security, consider using a dedicated Kubernetes cluster for the Rancher management server. You should not run your user workloads on the management cluster.

- **User clusters:** The downstream Kubernetes user clusters run your apps and services. Any cluster that you deploy from Rancher or import into Rancher is a user cluster.
- **Trident:** A Trident catalog is available to Rancher on NetApp HCI and runs in the user clusters. Inclusion of this catalog simplifies the Trident deployment to user clusters.

Find more information

- [Rancher documentation about architecture](#)

- [NetApp HCI Resources page](#)

Rancher on NetApp HCI concepts

Learn basic concepts related to Rancher on NetApp HCI.

- **Rancher server** or **Control plane**: The Rancher control plane, sometimes called the *Rancher Server*, provisions, manages, and monitors Kubernetes clusters used by Development and Operations teams.
- **Catalogs**: Catalogs are GitHub repositories or Helm Chart repositories filled with applications that are ready-made for deployment. Rancher provides the ability to use a catalog of Helm charts that make it easy to deploy applications repeatedly. Rancher includes two types of catalogs: built-in global catalogs and custom catalogs. Trident is deployed as a catalog. See [Rancher documentation about catalogs](#).
- **Management cluster**: Rancher on NetApp HCI deploys three virtual machines on the Rancher management cluster, which you can see using Rancher, Hybrid Cloud Control, and the vCenter Plug-in. The management cluster virtual machines host the Rancher server, the Rancher Kubernetes Engine (RKE), and the Linux OS.
- **User clusters**: These downstream Kubernetes clusters run your apps and services. In Kubernetes installations of Rancher, the management cluster should be separate from the user clusters. Any cluster that a Rancher user deploys from Rancher, or imports into Rancher, is considered a user cluster.
- **Rancher node template**: Hybrid Cloud Control uses a Rancher node template to make deployment simpler.

See [Rancher documentation about node templates](#).

Trident software and persistent storage concepts

Trident, itself a Kubernetes-native application, runs directly within a Kubernetes cluster. With Trident, Kubernetes users (such as developers, data scientists, and Kubernetes administrators) can create, manage, and interact with persistent storage volumes in the standard Kubernetes format that they are already familiar with. With Trident, NetApp solutions can meet persistent volume claims that are made by Kubernetes clusters.

With Rancher, you can use a persistent volume, one that exists independently of any specific pod and with its own lifetime. Using Trident to manage persistent volume claims (PVCs) insulates the developers creating pods from the lower-level implementation details of the storage that they are accessing.

When a containerized application issues a persistent volume claim (PVC) request, Trident dynamically provisions storage per the parameters requested against the NetApp Element software storage layer in NetApp HCI.

A Trident catalog is available to Rancher on NetApp HCI and runs in the user clusters. As part of the Rancher on NetApp HCI implementation, a Trident installer is available in the Rancher catalog by default. Inclusion of this catalog simplifies the Trident deployment to user clusters.

See [Install Trident with Rancher on NetApp HCI](#).

For details, visit the [Trident documentation](#).

Find more information

- [Rancher documentation about architecture](#)

- [Kubernetes terminology for Rancher](#)
- [NetApp HCI Resources page](#)

Requirements for Rancher on NetApp HCI

Before you install Rancher on NetApp HCI, ensure your environment and your NetApp HCI system meet these requirements.



If you accidentally deploy Rancher on NetApp HCI with incorrect information (such as an incorrect Rancher server FQDN), there is no way to correct the deployment without removing it and redeploying. You will need to remove the Rancher on NetApp HCI instance and then redeploy Rancher on NetApp HCI from NetApp Hybrid Cloud Control UI. See [Remove a Rancher installation on NetApp HCI](#) for more information.

Node requirements

- Ensure that your NetApp HCI system has at least three compute nodes; this is required for full resiliency. Rancher on NetApp HCI is not supported on storage-only configurations.
- Ensure that the datastore you intend to use for the Rancher on NetApp HCI deployment has at least 60GB of free space.
- Ensure that your NetApp HCI cluster is running management services version 2.17 or later.

Node details

Rancher on NetApp HCI deploys a three-node management cluster.

All nodes have the following characteristics:

vCPU	RAM (GB)	Disk (GB)
2	8	20

Network requirements

- Ensure that the network that you intend to deploy the Rancher on NetApp HCI management cluster has a route to the management node management network.
- Rancher on NetApp HCI supports DHCP addresses for the control plane (Rancher server) and user clusters, but we recommend static IP addresses for production environments. Ensure that you have allocated the necessary static IP addresses if you are deploying in a production environment.
 - Rancher server requires three static IP addresses.
 - Each user cluster requires as many static IP addresses as nodes in the cluster. For example, a user cluster with four nodes requires four static IP addresses.
 - If you plan on using DHCP addressing for the Rancher control plane or user clusters, ensure that the DHCP lease duration is at least 24 hours.
- If you need to use an HTTP proxy to enable internet access for Rancher on NetApp HCI, you need to make a pre-deployment change to the management node. Log in to your management node using SSH and follow the [instructions](#) in the Docker documentation to manually update the proxy settings for Docker.
- If you enable and configure a proxy server during deployment, the following IP address ranges and

domains are automatically added to the Rancher server noProxy settings:

```
127.0.0.0/8, 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16, .svc,  
.cluster.local
```

- Ensure that your management node can use DNS to resolve the host name `<any IP address>.nip.io` to an IP address. This is the DNS provider used during deployment; if the management node cannot resolve this URL, deployment will fail.
- Ensure that you have set up DNS records for each static IP address you need.

VMware vSphere requirements

- Ensure that the VMware vSphere instance you are using is version 6.5, 6.7, or 7.0.
- You can use a vSphere Standard Switch (VSS) networking configuration, but if you do, ensure that the virtual switches and physical hosts used for Rancher VMs can access all the same port groups, in the same way that you would ensure for regular VMs.

Deployment considerations

You might want to review the following considerations:

- Types of deployments
 - Demo deployments
 - Production deployments
- Rancher FQDN



Rancher on NetApp HCI is not resilient to node failures unless you configure some type of network load balancing. As a simple solution, create a round robin DNS entry for the three static IP addresses reserved for Rancher server. These DNS entries should resolve to the Rancher server FQDN that you will use to access the Rancher server host, which serves the Rancher web UI once deployment is complete.

Types of deployments

You can deploy Rancher on NetApp HCI in the following ways:

- **Demo deployments:** If DHCP is available in the targeted deployment environment and you want to demo the Rancher on NetApp HCI capability, then a DHCP deployment makes the most sense.

In this deployment model, the Rancher UI is accessible from each of the three nodes in the management cluster.

If your organization does not use DHCP, you can still try it out by using four static IP addresses allocated prior the deployment, similar to what you would do for a production deployment.

- **Production deployments:** For production deployments or when DHCP is not available in the targeted deployment environment, a little more pre-deployment work is required. The first step is to obtain three consecutive IP addresses. You enter the first during the deployment.

We recommend using L4 load balancing or round-robin DNS configuration for production environments.

This requires a fourth IP address and separate entry in your DNS configuration.

- **L4 load balancing:** This is a technique where a virtual machine or container hosting an application like nginx is configured to distribute requests among the three nodes of the management cluster.
- **Round-robin DNS:** This is a technique where a single host name is configured in the DNS system that rotates requests among the three hosts that form the management cluster.

Rancher FQDN

The installation requires assignment of a Rancher URL, which includes the fully qualified domain name (FQDN) of the host where the Rancher UI will be served after the installation is complete.

In all cases the Rancher UI is accessible in your browser over https protocol (port 443).

Production deployments require an FQDN configured that load balances across the management cluster nodes. Without using FQDN and load balancing, the environment is not resilient and is suitable only for demo environments.

Required ports

Ensure that the list of ports in the "Ports for Rancher Server Nodes on RKE" section of the **Rancher Nodes** section of the official [Rancher documentation](#) are open in your firewall configuration to and from the nodes running Rancher server.

Required URLs

The following URLs should be accessible from the hosts where the Rancher control plane resides:

URL	Description
https://charts.jetstack.io/	Kubernetes integration
https://releases.rancher.com/server-charts/stable	Rancher software downloads
https://entropy.ubuntu.com/	Ubuntu entropy service for random number generation
https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo/v1.3.1/install.sh	VMware guest additions
https://download.docker.com/linux/ubuntu/gpg	Docker Ubuntu GPG public key
https://download.docker.com/linux/ubuntu	Docker download link
https://hub.docker.com/	Docker Hub for NetApp Hybrid Cloud Control

Deploy Rancher on NetApp HCI

To use Rancher on your NetApp HCI environment, you first deploy Rancher on NetApp HCI.



Before starting the deployment, be sure to check the datastore free space and other [requirements for Rancher on NetApp HCI](#).



Rancher Support is not included in your NetApp Support Edge agreement. Contact NetApp Sales or your reseller for options. If you purchase Rancher Support from NetApp, you will receive an email with instructions.

What happens when you deploy Rancher on NetApp HCI?

The deployment involves the following steps, each described further:

- Use the NetApp Hybrid Cloud Control to initiate the deployment.
- The Rancher deployment creates a management cluster, which includes three virtual machines.

Each virtual machine is assigned all Kubernetes roles for both the Control Plane and Worker. This means that the Rancher UI is available on each node.

- The Rancher Control Plane (or *Rancher Server*) is also installed, using the NetApp HCI node template in Rancher for easier deployment. The Rancher Control Plane automatically works with the configuration used in the NetApp Deployment Engine, which was used to build the NetApp HCI infrastructure.
- After deployment, you will receive an email from NetApp providing you with the option to register for NetApp Support on Rancher deployments on NetApp HCI.
- After deployment, Dev and Ops teams can then deploy their user clusters, similar to any Rancher deployment.

Steps to deploy Rancher on NetApp HCI

- [Access the NetApp Hybrid Cloud Control](#)
- [Deploy Rancher on NetApp HCI](#)
- [Verify your deployment by using vCenter Server](#)

Access the NetApp Hybrid Cloud Control

To begin the deployment, access the NetApp Hybrid Cloud Control.

1. Open a web browser and browse to the IP address of the management node. For example:

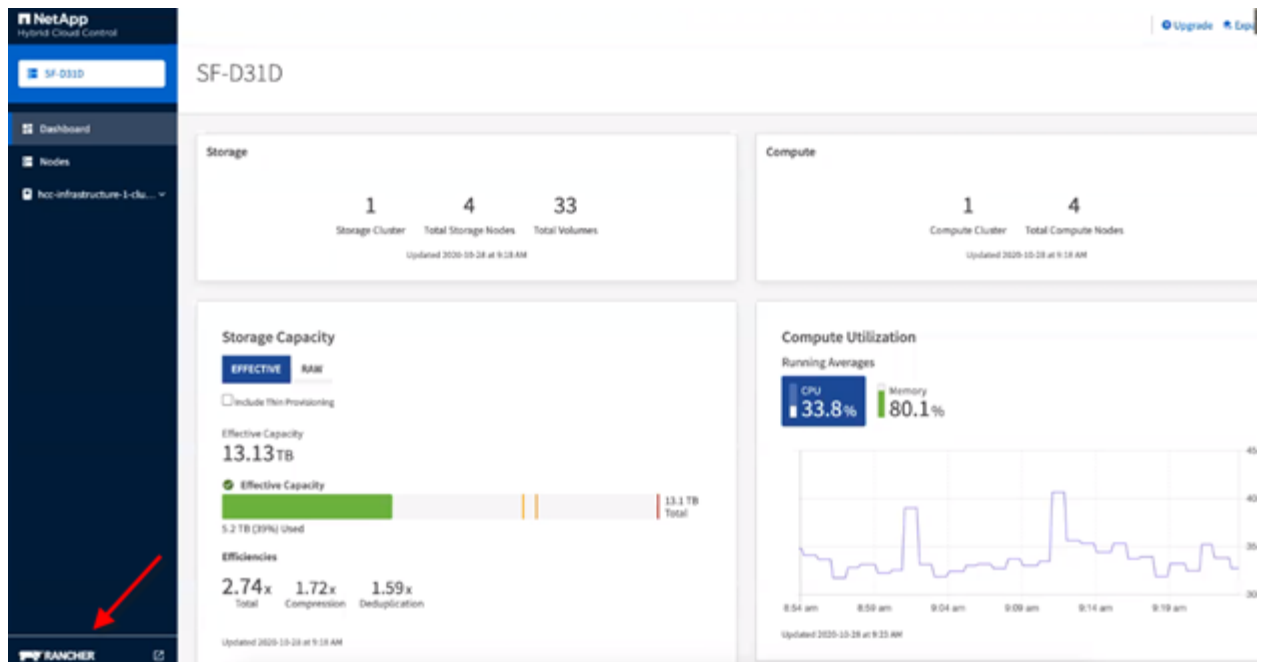
`https://<ManagementNodeIP>`

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.

The NetApp Hybrid Cloud Control interface appears.

Deploy Rancher on NetApp HCI

1. From the Hybrid Cloud Control, click the **Rancher** icon in the lower left of the navigation bar.



A popup window shows a message about getting started with Rancher.

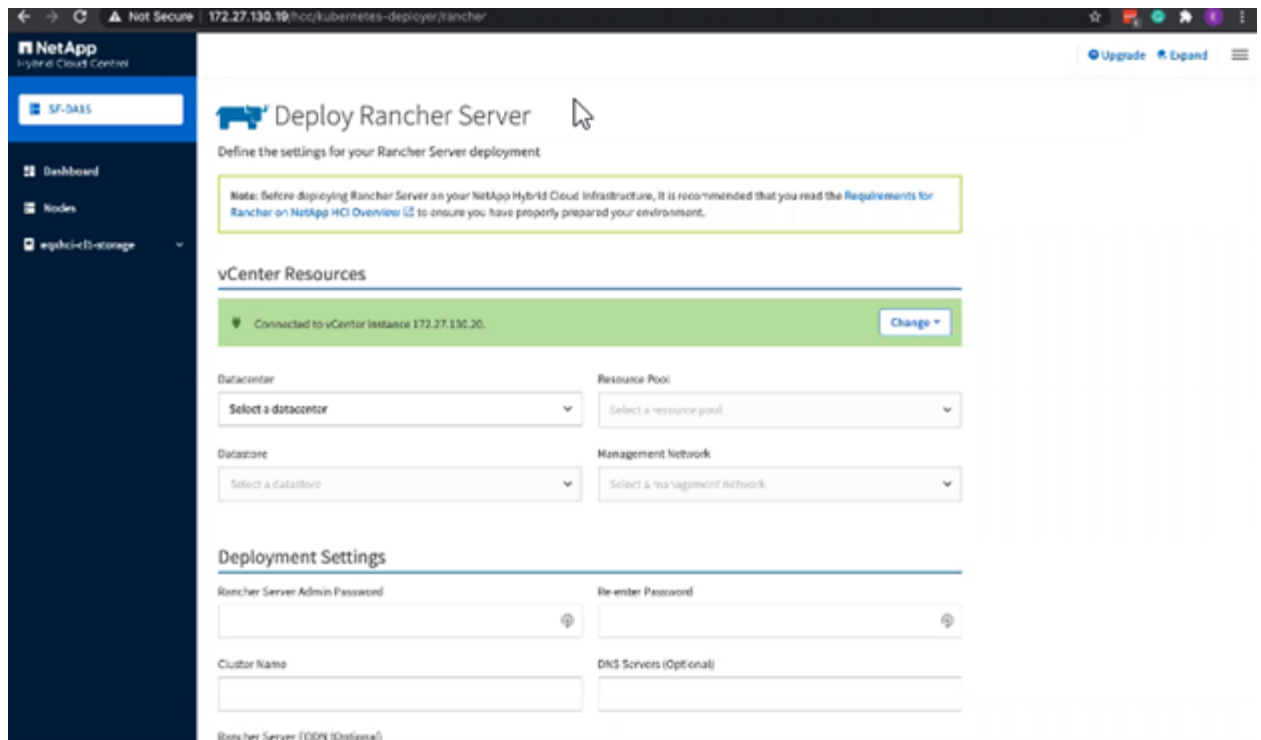
Getting Started with Rancher

Rancher is a complete software stack for teams adopting containers. It addresses the operational and security challenges of managing multiple Kubernetes clusters across any infrastructure, while providing DevOps teams with integrated tools for running containerized workloads. Deploy Rancher Server on your NetApp HCI to start creating on-premises Kubernetes clusters using Rancher and NetApp's proven enterprise technology.

[Deploy Rancher](#) [Cancel](#)

2. Click **Deploy Rancher**.

The Rancher UI appears.



Your vCenter credentials are collected based on your NetApp Deployment Engine installation.

3. Enter **vCenter Resources** information. Some fields are described next.
 - **Datacenter:** Select a datacenter. After you select the datacenter, all other fields are prepopulated, although you can change them.
 - **Datastore:** Select a datastore on the NetApp HCI storage nodes. This datastore should be resilient and accessible to all of the VMware hosts. Do not select a local datastore that is accessible to only one of the hosts.
 - **Management network:** This should be accessible from the management stations and from the virtual machine network where the user clusters will be hosted.
4. Enter **Deployment Settings** information:
 - **DNS Servers:** Optional. If you use load balancing, enter the internal DNS server information.
 - **Rancher Server FQDN:** To ensure that the Rancher Server remains available during node failures, provide a fully-qualified domain name (FQDN) that your DNS server can resolve to any of the IP addresses assigned to the Rancher Server cluster's nodes. This FQDN with the "https" prefix becomes the Rancher URL that you will use to access your Rancher implementation.

If no domain name is provided, wildcard DNS will be used instead and you will be able to access the Rancher Server using one of the URLs presented after the deployment completes.
5. Enter **Advanced Settings** information:
 - **Assign Static IP Addresses:** If you enable static IP addressing, provide starting IP addresses for three IPv4 addresses in sequence, one for each management cluster virtual machine. Rancher on NetApp HCI deploys three management cluster virtual machines.
 - **Configure Proxy Server:**
6. Review and select the checkbox for the Rancher End User License Agreement.
7. Review and select the checkbox to acknowledge information about Rancher software.


8. Click **Deploy**.

A bar indicates the deployment progress.



The Rancher deployment could take about 15 minutes.

When the deployment is complete, Rancher displays a message about the completion and provides a Rancher URL.



9. Record that Rancher URL that Sdisplays at the end of the deployment. You will use this URL to access the Rancher UI.

Verify your deployment by using vCenter Server

In your vSphere client, you can see the Rancher management cluster, which includes the three virtual machines.



Once you have finished deployment, do not modify the configuration of the Rancher server virtual machine cluster or remove the virtual machines. Rancher on NetApp HCI relies on the deployed RKE management cluster configuration to function normally.

What's next?

After deployment, you can do the following:

- [Complete post-deployment tasks](#)
- [Install Trident with Rancher on NetApp HCI](#)
- [Deploy user clusters and applications](#)
- [Manage Rancher on NetApp HCI](#)
- [Monitor Rancher on NetApp HCI](#)

Find more information

- [Rancher deployment troubleshooting](#)
- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp HCI Resources page](#)

Post-deployment tasks

Post-deployment tasks overview

After you deploy Rancher on NetApp HCI, you should continue with post-deployment activities.

- [Ensure Rancher Support parity](#)
- [Improve Rancher VM resiliency](#)
- [Configure monitoring](#)
- [Install Trident](#)
- [Enable Trident support for user clusters](#)

Find more information

- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources page](#)

Ensure Rancher Support parity

After you deploy Rancher on NetApp HCI, you need to ensure that the number of Rancher Support cores you purchased matches the number of CPU cores you are using for Rancher management VMs and user clusters.

If you purchased Rancher Support for only part of your NetApp HCI compute resources, you need to take action in VMware vSphere to ensure that Rancher on NetApp HCI and its managed user clusters are only running on hosts for which you have purchased Rancher Support. See the VMware vSphere documentation for information about how to help ensure this by confining compute workloads to specific hosts.

Find more information

- [vSphere HA and DRS Affinity Rules](#)
- [Create VM Anti-Affinity Rules](#)
- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources page](#)

Improve Rancher VM resiliency

After you deploy Rancher on NetApp HCI, your vSphere environment will contain three new nodes as virtual machines to host the Rancher environment. The Rancher web UI is available from each of these nodes. For full resiliency, each of the three virtual machines along with the corresponding virtual disks should reside on a different physical host after

events like power cycles and failovers.

To ensure that each VM and its resources remain on a different physical host, you can create VMware vSphere Distributed Resource Scheduler (DRS) anti-affinity rules. This is not automated as part of Rancher on NetApp HCI deployment.

For instructions on how to configure DRS anti-affinity rules, see the following VMware documentation resources:

[Create VM Anti-Affinity Rules](#)

[vSphere HA and DRS Affinity Rules](#)

Find more information

- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources page](#)

Enable monitoring

After you deploy Rancher on NetApp HCI, You can enable Active IQ storage monitoring (for SolidFire all-flash storage and NetApp HCI) and NetApp HCI compute monitoring (for NetApp HCI only) if you did not already do so during installation or upgrade.

For instructions on how to enable monitoring, see [Enable Active IQ and NetApp HCI monitoring](#).

Find more information

- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources page](#)

Install Trident

Learn about how to install Trident after you install Rancher on NetApp HCI. Trident is a storage orchestrator, which integrates with Docker and Kubernetes, as well as platforms built on these technologies, such as Red Hat OpenShift, Rancher, and IBM Cloud Private. The goal of Trident is to make the provisioning, connection, and consumption of storage transparent and frictionless for the applications. Trident is a fully supported open source project maintained by NetApp. Trident enables you to create, manage, and interact with persistent storage volumes in the standard Kubernetes format that you are familiar with.



For more information about Trident, see the [Trident documentation](#).

What you'll need

- You have installed Rancher on NetApp HCI.
- You have deployed your user clusters.
- You have configured your user cluster networks for Trident. See [Enable Trident support for user clusters](#) for instructions.
- You have completed the necessary prerequisite steps for work node preparation for Trident. See the [Trident documentation](#).

About this task

The Trident installer catalog is installed as part of the Rancher installation using NetApp Hybrid Cloud Control. In this task, you use the installer catalog to install and configure Trident.

As part of the Rancher installation, NetApp provides a node template. If you are not planning to use the node template that NetApp provides, and you want to provision on RHEL or CentOS, there might be additional requirements. If you change your worker node to RHEL or CentOS, there are several prerequisites that should be met. See the [Trident documentation](#).

Steps

1. From the Rancher UI, select a project for your user cluster.





For information about projects and namespaces, see the [Rancher documentation](#).

2. Select **Apps**, and select **Launch**.



3. On the **Catalog** page, select the Trident installer.



On the page that opens, you can select the **Detailed Descriptions** arrow to learn more about the Trident app and also to find the link to the [Trident documentation](#).

4. Select the **Configurations Options** arrow, and enter the credentials and storage configuration information.

<p>Storage Tenant *</p> <input type="text" value="NetApp-HCI"/> <p>The name of the tenant that is already present on the SolidFire AFA.</p>	<p>SVIP *</p> <input type="text" value=""/> <p>The virtual/cluster IP address for data (I/O).</p>
<p>MVIP *</p> <input type="text" value=""/> <p>The virtual/cluster IP address for management.</p>	<p>Trident Backend Name *</p> <input type="text" value="solidfire"/> <p>The name of this Trident backend configuration.</p>
<p>Trident Storage Driver *</p> <input type="text" value="solidfire-san"/> <p>The name of the Trident storage driver.</p>	



The default storage tenant is NetApp HCI. You can change this value. You can also change the backend name. However, do not change the default storage driver value, which is **solidfire-san**.

5. Select **Launch**.

This installs the Trident workload on the **trident** namespace.

6. Select **Resources > Workloads**, and verify that the **trident** namespace includes the following components:

Namespace: trident		
<input type="checkbox"/>	▶ Active	trident-csi
<input type="checkbox"/>	▶ Active	trident-csi
<input type="checkbox"/>	▶ Active	trident-installer
<input type="checkbox"/>	▶ Active	trident-operator

7. (Optional) Select **Storage** for the user cluster to see the storage classes that you can use for your persistent volumes.



The three storage classes are **solidfire-gold**, **solidfire-silver**, and **solidfire-bronze**. You can make one of these storage classes the default by selecting the icon under the **Default** column.

Find more information

- [Enable Trident support for user clusters](#)
- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp Element Plug-in for vCenter Server](#)

- [NetApp HCI Resources page](#)

Enable Trident support for user clusters

If your NetApp HCI environment does not have a route between the management and storage networks, and you deploy user clusters that need Trident support, you need to further configure your user cluster networks after installing Trident. For each user cluster, you need to enable communication between the management and storage networks. You can do this by modifying the networking configuration for each node in the user cluster.

About this task

Follow these general steps to modify the networking configuration for each node in the user cluster. These steps assume that you created the user cluster with the default node template that is installed with Rancher on NetApp HCI.



You can make these changes as part of a custom node template to use for future user clusters.

Steps

1. Deploy a user cluster with existing default template.
2. Connect the storage network to the user cluster.
 - a. Open the VMware vSphere web client for the connected vCenter instance.
 - b. In the Hosts and Clusters inventory tree, select a node in the newly deployed user cluster.
 - c. Edit the node's settings.
 - d. In the settings dialog, add a new network adapter.
 - e. In the **New Network** drop down list, browse for a network and select **HCI_Internal_Storage_Data_Network**.
 - f. Expand the network adapter section and record the MAC address for the new network adapter.
 - g. Click **OK**.
3. In Rancher, download the SSH private key file for each node in the user cluster.
4. Connect using SSH to a node in the user cluster, using the private key file that you have downloaded for that node:

```
ssh -i <private key filename> <ip address>
```

5. As the superuser, edit and save the `/etc/netplan/50-cloud-init.yaml` file so that it includes the `ens224` section, similar to the following example. Replace `<MAC address>` with the MAC address you recorded earlier:

```
network:
  ethernets:
    ens192:
      dhcp4: true
      match:
        macaddress: 00:50:56:91:1d:41
        set-name: ens192
    ens224:
      dhcp4: true
      match:
        macaddress: <MAC address>
        set-name: ens224
  version: 2
```

6. Use the following command to reconfigure the network:

```
`netplan try`
```

7. Repeat steps 4 through 6 for each remaining node in the user cluster.
8. When you have reconfigured the network for each node in the user cluster, you can deploy applications in the user cluster that utilize Trident.

Deploy user clusters and applications

After deploying Rancher on NetApp HCI, you can set up user clusters and add applications to those clusters.

Deploy user clusters

After deployment, Dev and Ops teams can then deploy their Kubernetes user clusters, similar to any Rancher deployment, on which they can deploy apps.

1. Access the Rancher UI using that URL provided to you at the end of the Rancher deployment.
2. Create user clusters. See Rancher documentation about [deploying workloads](#).
3. Provision user clusters in Rancher on NetApp HCI. See Rancher documentation about [setting up Kubernetes clusters in Rancher](#).

Deploy applications on user clusters

Similar to any Rancher deployment, you add applications on Kubernetes clusters.

See Rancher documentation about [deploying applications across clusters](#).

Find more information

- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp HCI Resources page](#)

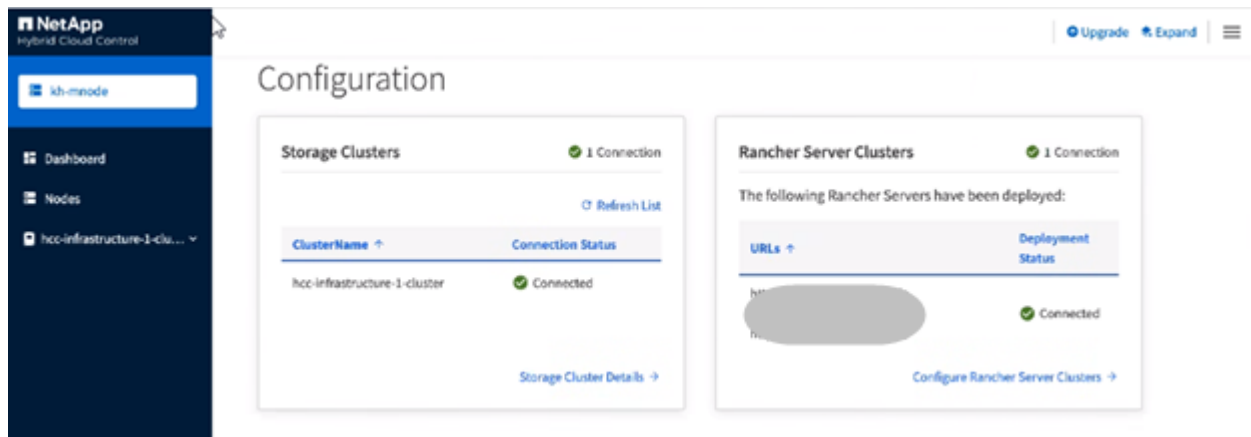
Manage Rancher on NetApp HCI

After deploying Rancher on NetApp HCI, you can view the Rancher server cluster URLs and status. You can also delete the Rancher server.

Identify Rancher server cluster URLs and status

You can identify Rancher server cluster URLs and determine server status.

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, select the top right Options icon and select **Configure**.



The Rancher Server Clusters page displays a list of Rancher server clusters that have been deployed, the associated URL, and status.

Find more information

- [Remove Rancher](#)
- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp HCI Resources page](#)

Monitor a Rancher on NetApp HCI implementation

There are multiple ways to monitor Rancher server, management clusters, and other details.

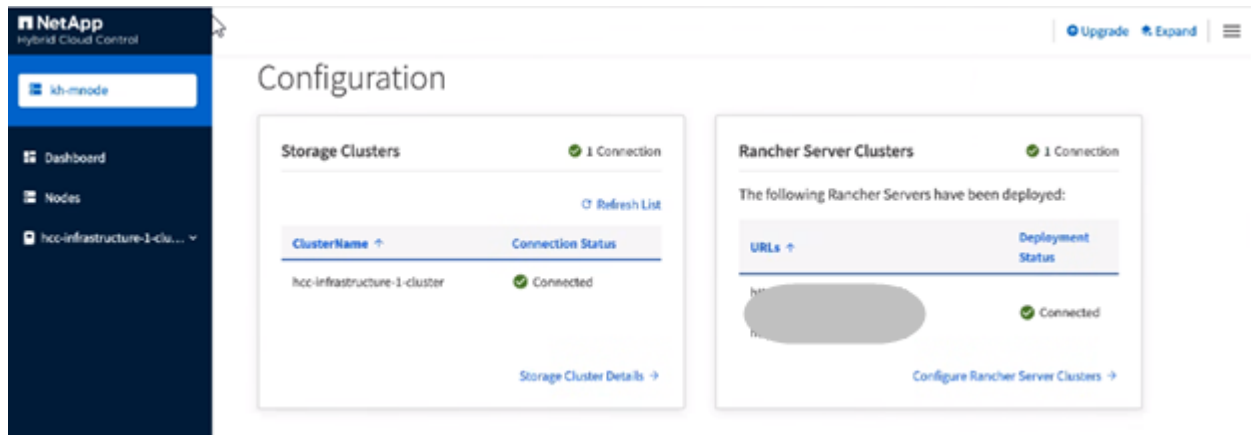
- NetApp Hybrid Cloud Control
- Rancher UI

- NetApp Active IQ
- vCenter Server

Monitor Rancher using the NetApp Hybrid Cloud Control

Using the NetApp Hybrid Cloud Control, you can view the Rancher URL and Rancher server cluster status. You can also monitor the nodes in which Rancher is running.

1. Log in to NetApp Hybrid Cloud Control by providing Element storage cluster administrator credentials.
2. From the Dashboard, click on the top right Options icon and select **Configure**.



3. To view nodes information, from the Hybrid Cloud Control Dashboard, expand the name of your storage cluster and click **Nodes**.

Monitor Rancher using the Rancher UI

Using the Rancher UI, you can see information about Rancher on NetApp HCI management clusters and user clusters.



In the Rancher UI, management clusters are referred to as "local clusters."

1. Access the Rancher UI using that URL provided to you at the end of the Rancher deployment.
2. See [Monitoring in Rancher v2.5](#).

Monitor Rancher using NetApp Active IQ

Using NetApp Active IQ, you can view Rancher telemetry, such as installation information, nodes, clusters, status, namespace information, and more.

1. Log in to NetApp Hybrid Cloud Control by providing Element storage cluster administrator credentials.
2. From the top right menu, select **NetApp Active IQ**.

Monitor Rancher using vCenter Server

Using vCenter Server, you can monitor the Rancher virtual machines.

Find more information

- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources page](#)

Upgrade Rancher on NetApp HCI

To upgrade Rancher software, you can use the NetApp Hybrid Cloud Control (HCC) UI or REST API. HCC provides an easy button process to upgrade the components of your Rancher deployment, including Rancher server, Rancher Kubernetes Engine (RKE), and the management cluster's node OS (for security updates). You can alternatively use the API to help automate upgrades.

Upgrades are available by component instead of a cumulative package. As such, some component upgrades such as the Ubuntu OS come available on a more rapid cadence. Upgrades affect only your Rancher server instance and the management cluster that Rancher Server is deployed on. Upgrades to the management cluster node's Ubuntu OS are for critical security patches only and do not upgrade the operating system. User clusters cannot be upgraded from NetApp Hybrid Cloud Control.

What you'll need

- **Admin privileges:** You have storage cluster administrator permissions to perform the upgrade.
- **Management services:** You have updated your management services bundle to the latest version.



You must upgrade to the latest management services bundle 2.17 or later for Rancher functionality.

- **System ports:** If you are using NetApp Hybrid Cloud Control for upgrades, you have ensured that the necessary ports are open. See [Network ports](#) for more information.

Upgrade options

Choose one of the following upgrade processes:

- [Use NetApp Hybrid Cloud Control UI to upgrade a Rancher deployment](#)
- [Use NetApp Hybrid Cloud Control API to upgrade a Rancher deployment](#)

Use NetApp Hybrid Cloud Control UI to upgrade a Rancher deployment

Using the NetApp Hybrid Cloud Control UI, you can upgrade any of these components in your Rancher deployment:

- Rancher server
- Rancher Kubernetes Engine (RKE)
- Node OS security updates

What you'll need

- A good internet connection. Dark site upgrades (upgrades at a site without external connectivity) are not available.

Steps

1. Open a web browser and browse to the IP address of the management node:

```
https://<ManagementNodeIP>
```

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. On the **Upgrades** page, select **Rancher**.
5. Select the **Actions** menu for the software you want to upgrade.
 - Rancher server
 - Rancher Kubernetes Engine (RKE)
 - Node OS security updates
6. Select **Upgrade** for Rancher server or RKE upgrades or **Apply Upgrade** for Node OS security updates.



For node OS, unattended upgrades for security patches are run on a daily basis but the node is not rebooted automatically. By applying upgrades, you are rebooting each node for the security updates to take effect.

A banner appears indicating the component upgrade is successful. There could be up to a 2 minute delay before NetApp Hybrid Cloud Control UI shows the updated version number.

Use NetApp Hybrid Cloud Control API to upgrade a Rancher deployment

You can use APIs to upgrade any of these components in your Rancher deployment:

- Rancher server
- Rancher Kubernetes Engine (RKE)
- Node OS (for security updates)

You can use an automation tool of your choice to run the APIs or the REST API UI available on the management node.

Options

- [Upgrade Rancher Server](#)
- [Upgrade RKE](#)
- [Apply node OS security updates](#)



For node OS, unattended upgrades for security patches are run on a daily basis but the node is not rebooted automatically. By applying upgrades, you are rebooting each node for the security updates to take effect.

Upgrade Rancher Server

API commands

1. Initiate the list upgrade versions request:

```
curl -X POST "https://<managementNodeIP>/k8sdeployer/1/upgrade/rancher-versions" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```



You can find the bearer `${TOKEN}` used by the API command when you [authorize](#). The bearer `${TOKEN}` is in the curl response.

2. Get task status using task ID from previous command and copy the latest version number from the response:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/task/<taskID>" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```

3. Initiate Rancher server upgrade request:

```
curl -X PUT "https://<mNodeIP>/k8sdeployer/1/upgrade/rancher/<version number>" -H "accept: application/json" -H "Authorization: Bearer"
```

4. Get task status using task ID from upgrade command response:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/task/<taskID>" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```

REST API UI steps

1. Open the management node REST API UI on the management node:

```
https://<managementNodeIP>/k8sdeployer/api/
```

2. Select **Authorize** and complete the following:
 - a. Enter the cluster user name and password.
 - b. Enter the client ID as `mnode-client`.
 - c. Select **Authorize** to begin a session.
 - d. Close the authorization window.
3. Check for the latest upgrade package:
 - a. From the REST API UI, run **POST /upgrade/rancher-versions**.
 - b. From the response, copy the task ID.
 - c. Run **GET /task/{taskID}** with the task ID from the previous step.
4. From the **/task/{taskID}** response, copy the latest version number you want to use for the upgrade.

5. Run the Rancher Server upgrade:

- a. From the REST API UI, run **PUT /upgrade/rancher/{version}** with the latest version number from the previous step.
- b. From the response, copy the task ID.
- c. Run **GET /task/{taskID}** with the task ID from the previous step.

The upgrade has finished successfully when the `PercentComplete` indicates 100 and `results` indicates the upgraded version number.

Upgrade RKE

API commands

1. Initiate the list upgrade versions request:

```
curl -X POST "https://<mNodeIP>/k8sdeployer/1/upgrade/rke-versions" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```



You can find the bearer `${TOKEN}` used by the API command when you [authorize](#). The bearer `${TOKEN}` is in the curl response.

2. Get task status using task ID from previous command and copy the latest version number from the response:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/task/<taskID>" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```

3. Initiate the RKE upgrade request

```
curl -X PUT "https://<mNodeIP>/k8sdeployer/1/upgrade/rke/<version number>" -H "accept: application/json" -H "Authorization: Bearer"
```

4. Get task status using task ID from upgrade command response:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/task/<taskID>" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```

REST API UI steps

1. Open the management node REST API UI on the management node:

```
https://<managementNodeIP>/k8sdeployer/api/
```

2. Select **Authorize** and complete the following:

- a. Enter the cluster user name and password.
 - b. Enter the client ID as `mnode-client`.
 - c. Select **Authorize** to begin a session.
 - d. Close the authorization window.
3. Check for the latest upgrade package:
 - a. From the REST API UI, run **POST /upgrade/rke-versions**.
 - b. From the response, copy the task ID.
 - c. Run **GET /task/{taskID}** with the task ID from the previous step.
 4. From the **/task/{taskID}** response, copy the latest version number you want to use for the upgrade.
 5. Run the RKE upgrade:
 - a. From the REST API UI, run **PUT /upgrade/rke/{version}** with the latest version number from the previous step.
 - b. Copy the task ID from the response.
 - c. Run **GET /task/{taskID}** with the task ID from the previous step.

The upgrade has finished successfully when the `PercentComplete` indicates 100 and `results` indicates the upgraded version number.

Apply node OS security updates

API commands

1. Initiate the check upgrades request:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/upgrade/checkNodeUpdates"
-H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```



You can find the bearer `${TOKEN}` used by the API command when you [authorize](#). The bearer `${TOKEN}` is in the curl response.

2. Get task status using task ID from previous command and verify a more recent version number is available from the response:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/task/<taskID>" -H "accept:
application/json" -H "Authorization: Bearer ${TOKEN}"
```

3. Apply the node updates:

```
curl -X POST "https://<mNodeIP>/k8sdeployer/1/upgrade/applyNodeUpdates"
-H "accept: application/json" -H "Authorization: Bearer"
```



For node OS, unattended upgrades for security patches are run on a daily basis but the node is not rebooted automatically. By applying upgrades, you are rebooting each node sequentially for the security updates to take effect.

4. Get task status using task ID from the upgrade `applyNodeUpdates` response:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/task/<taskID>" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```

REST API UI steps

1. Open the management node REST API UI on the management node:

```
https://<managementNodeIP>/k8sdeployer/api/
```

2. Select **Authorize** and complete the following:
 - a. Enter the cluster user name and password.
 - b. Enter the client ID as `mnode-client`.
 - c. Select **Authorize** to begin a session.
 - d. Close the authorization window.
3. Verify if an upgrade package is available:
 - a. From the REST API UI, run **GET /upgrade/checkNodeUpdates**.
 - b. From the response, copy the task ID.
 - c. Run **GET /task/{taskID}** with the task ID from the previous step.
 - d. From the **/task/{taskID}** response, verify that there is a more recent version number than the one currently applied to your nodes.
4. Apply the node OS upgrades:



For node OS, unattended upgrades for security patches are run on a daily basis but the node is not rebooted automatically. By applying upgrades, you are rebooting each node sequentially for the security updates to take effect.

- a. From the REST API UI, run **POST /upgrade/applyNodeUpdates**.
- b. From the response, copy the task ID.
- c. Run **GET /task/{taskID}** with the task ID from the previous step.
- d. From the **/task/{taskID}** response, verify that the upgrade has been applied.

The upgrade has finished successfully when the `PercentComplete` indicates 100 and `results` indicates the upgraded version number.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Remove a Rancher installation on NetApp HCI

If you accidentally deploy Rancher on NetApp HCI with incorrect information (such as an incorrect Rancher server FQDN), you need to remove the installation and then redeploy. Follow these steps to remove the Rancher installation on NetApp HCI instance.

This action does not delete the user clusters.



You might want to retain the user clusters. If you do retain them, you can later migrate them to another Rancher implementation. If you want to delete the user clusters, you should do that first before deleting the Rancher server; otherwise, deleting the user clusters after the Rancher server is deleted is more difficult.

Options

- [Remove Rancher on NetApp HCI using NetApp Hybrid Cloud Control](#) (Recommended)
- [Remove Rancher on NetApp HCI using the REST API](#)

Remove Rancher on NetApp HCI using NetApp Hybrid Cloud Control


You can use NetApp Hybrid Cloud Control web UI to remove the three virtual machines that were set up during deployment to host the Rancher server.

Steps

1. Open a web browser and browse to the IP address of the management node:

```
https://<ManagementNodeIP>
```

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. From the Dashboard, click the menu on the upper right.
4. Select **Configure**.



5. In the **Rancher Server Clusters** pane, click **Configure Rancher Server Clusters**.
6. Select the **Actions** menu for the Rancher installation you need to remove.



Clicking **Delete** immediately removes the Rancher on NetApp HCI management cluster.

7. Select **Delete**.

Remove Rancher on NetApp HCI using the REST API

You can use the NetApp Hybrid Cloud Control REST API to remove the three virtual machines that were set up during deployment to host the Rancher server.

Steps

1. Enter the management node IP address followed by `/k8sdeployer/api/`:

```
https://[IP address]/k8sdeployer/api/
```

2. Click **Authorize** or any lock icon and enter cluster admin credentials for permissions to use APIs.
 - a. Enter the cluster user name and password.
 - b. Select **Request body** from the type drop-down list if the value is not already selected.
 - c. Enter the client ID as `mnode-client` if the value is not already populated.
 - d. Do not enter a value for the client secret.
 - e. Click **Authorize** to begin a session.
 - f. Close the window.
3. Close the **Available authorizations** dialog box.
4. Click **POST/destroy**.
5. Click **Try it out**.
6. In the request body text box, enter the Rancher server FQDN as the `serverURL` value.
7. Click **Execute**.

After several minutes, the Rancher server virtual machines should no longer be visible in the Hosts and Clusters list in vSphere Client. After removal, you can use NetApp Hybrid Cloud Control to redeploy Rancher on NetApp HCI.

Find more Information

- [Rancher deployment troubleshooting](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources Page](#)

Copyright Information

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system-without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.