THALES

CTE for Kubernetes Integrations and Solutions

FOR THIRD-PARTY PRODUCTS

Integrating CTE for Kubernetes with Rancher Kubernetes Management Platform

Rancher is an open source software platform that enables Enterprise Kubernetes Container Management. It provides a CNCF-certified (Cloud Native Computing Foundation) Kubernetes distribution that can run on any Kubernetes cluster or distribution. It solves common Kubernetes installation complexities by removing most host dependencies, thereby presenting a stable path for deployment, upgrades & rollbacks.

This document describes installing Rancher and deploying a few workloads to a Kubernetes cluster provisioned by Rancher.

Test Environment

- **CTE for Kubernetes Agent**: 1.2.0
- **CipherTrust Manager**: 2.12.0
- **OS**: Ubuntu 20.04
- **Rancher**: 1.5.1
- **Kubernetes cert-manager**: 1.5.1

Prerequisite

Create a 3-node Kubernetes cluster in which all of the worker nodes and master node are in the **Ready** state.

Policy and Policy Elements

CTE for Kubernetes Policy

Name	Access
	all_ops with Permit, Apply Key, Audit
	all_ops with Deny, Audit
Default	No Access

Steps to integrate Rancher with CTE for Kubernetes Cluster

You will need to use K3s for installation. K3s is a lightweight Kubernetes distribution created by Rancher Labs.

Create a Kubernetes cluster for Rancher

- Use the default K3s installation script, type:

```
curl -sfL https://get.k3s.io INSTALL_K3S_CHANNEL="v1.21"
K3S_KUBECONFIG_MODE="644" sh -
```

Testing the cluster

In the previous step, K3s created a new Kubernetes cluster and installed the `kubectl CLI` which you can use to directly interact with the Kubernetes API.

- To list all the nodes in the cluster and check their status, type:

```
kubectl get nodes
```

System Response

```
All Pods are Running
```

Install Helm

Helm is a package manager for Kubernetes. It is used as the installation tool for Rancher when deploying Rancher onto a Kubernetes cluster.

1. Download Helm CLI, type:

```
curl https://raw.githubusercontent.com/helm/helm/master/scripts/
get-helm-3 \ | bash
```

2. Check for the version, type:

```
helm version --client
```

3. Create a soft symlink between rancher, `/etc/rancher/k3s/k3s.yaml` and the Kubernetes config file, `~/.kube/config` so that helm can interact with your cluster, type:

```
mkdir -p ~/.kube ln -s /etc/rancher/k3s/k3s.yaml ~/.kube/config
```

4. Validate the connection, type:

```
helm ls --all-namespaces
```

Install Kubernetes Cert-Manager

Kubernetes cert-manager is an add-on that automates the management and issuance of TLS certificates from various issuing sources.

1. Obtain the latest version of the `cert-manager`, type:

```
kubectl apply --validate=false -f https://github.com/jetstack/certmanager/releases/download/v1.5.1/cert-manager.crds.yaml
```

2. Add the Helm repository for Jetstack:

```
helm repo add jetstack https://charts.jetstack.io
```

3. Install `cert-manager`.

```
helm install cert-manager jetstack/cert-manager --namespace cert-manager --version v1.5.1 --create-namespace
```

4. Once the Helm chart is installed, you can monitor the rollout status of `cert-manager`.

```
kubectl -n cert-manager rollout status deploy/cert-manager
```

System Response

```
Waiting for deployment cert-manager rollout to finish: 0 of 1 updated replicas are available... deployment "cert-manager" successfully rolled out
```

5. You can also monitor the rollout status of `cert-manager-webhook`.

```
kubectl -n cert-manager rollout status deploy/cert-manager-webhook
```

Install Rancher

Install Rancher in HA mode on your Rancher01 Kubernetes cluster.

1. Add rancher-stable as a Helm repository, type:

```
helm repo add rancher-stable https://releases.rancher.com/servercharts/stable
```

2. Install Rancher using the Helm install command, type:

```
helm install rancher rancher-stable/rancher \
--namespace cattle-system  \
--set hostname=rancher.${vminfo:rancher01:public_ip}.sslip.io  \
--set replicas=1  \
--set bootstrapPassword=RancherOnK3s  \
--create-namespace \
```

3. Verify that Rancher is successfully deployed, type:

```
kubectl -n cattle-system rollout status deploy/rancher
```

4. Get container details for the cattle-system namespace, type:

```
kubectl get pod -n cattle-system
```

5. List all pods in all namespaces, type:

```
kubectl get pods --all-namespaces
```

6. Verify Rancher is ready to access, type:

```
curl -kv https://rancher.${vminfo:Rancher01:public_ip}.sslip.io
2>&1 | grep -q "dynamiclistener-ca"; if [ $? != 0 ]; then echo "
Rancher isn't ready yet"; sleep 5; continue; fi; break; done;
echo "Rancher is Ready";
```

Integrating CTE for Kubernetes with Tanzu Kubernetes Grid

Tanzu Kubernetes Grid is an enterprise-ready Kubernetes runtime that unifies management of Kubernetes clusters, vSphere-optimized pods, containerized workloads, and VMware virtual machines.

This document describes installing Tanzu and deploying a few workloads to a Kubernetes cluster.

Test Environment

- **CTE for Kubernetes Agent**: 1.4.0
- **CipherTrust Manager**: 2.15.0
- **OS**: ubuntu-2204-kube-v1.28.4+vmware.1-tkg.1.ova
- **VMware Tanzu Kubernetes Grid**: 2.5.0
- **VMware Tanzu CLI**: 1.1.0

Prerequisites

Prior to Setting up CipherTrust Transparent Encryption with Tanzu, you must:

- Download the OS `.ova` from the VMware Tanzu product page
- Install Tanzu CLI for use with Tanzu Kubernetes Grid v2.5
- Install Tanzu CLI Plugins for TKG v2.5
- Install Kubernetes CLI (kubectl)
- Prepare VMware vSphere Server
- Create an SSH Key Pair
- DeployStandalone Management Clusters

References

For more information about deploying Tanzu, reference the following documents:

- Prepare to Deploy Management Clusters to vSphere
- Interoperability Result

Deploying Standalone Management Clusters

1. On the system on which you downloaded and installed the Tanzu CLI, type:

   ```
   tanzu management-cluster create –ui
   ```

2. Click **Deploy** for VMware vSphere.

3. Enter the vCenter Single Sign-On username and password for a user account that has the required privileges for the Tanzu Kubernetes Grid operation.

   > **Note**
   >
   > The account name must include the domain, for example administrator@vsphere.local.

4. Click **Connect**.

5. If the **Disable Verification** option is deselected in the previous screen, then you must manually verify the SSL thumbprint of the vCenter Server certificate and click **Continue**.

6. If you are deploying a management cluster to a vSphere 7 or vSphere 8 instance, confirm whether you want to proceed with the deployment.

7. Select the Datacenter in which to deploy the management cluster from the Datacenter drop-down menu.

8. Manually paste the contents of the key into the text box to add your SSH public key and click **Next**.

9. In the Management Cluster Settings:

   a. Select a size from the predefined CPU, memory, and storage configurations.

   > **Note**
   >
   > **Minimum configuration**: 2 CPUs and 4 GB memory

   b. Enter a name for your management cluster.

   c. Select the **Machine Health Checks** option if you want to activate it.

   > **Note**
   >
   > You can activate, or deactivate, Machine Health Checks on clusters after deployment using the CLI.

   d. Select the **Enable Audit Logging** option if you want to record requests made to the Kubernetes API server.

   e. For **Worker Node Instance Type**, select the configuration for the worker node virtual machine (VM).

   f. For **Control Plane Endpoint Provider**, select **Kube-vip**.

   e. For **Control Plane Endpoint**, enter a static virtual IP address, or FQDN, for API requests to the management cluster.

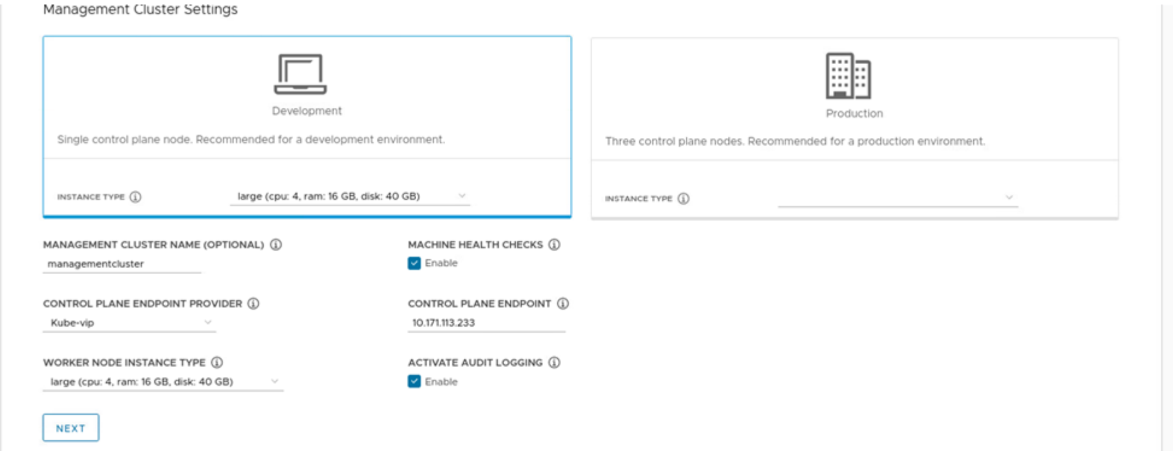

Management Cluster Settings

Single control plane node. Recommended for a development environment.

INSTANCE TYPE: large (cpu: 4, ram: 16 GB, disk: 40 GB)

Three control plane nodes. Recommended for a production environment.

INSTANCE TYPE:

MANAGEMENT CLUSTER NAME (OPTIONAL): managementcluster

MACHINE HEALTH CHECKS: ☑ Enable

CONTROL PLANE ENDPOINT PROVIDER: Kube-vip

CONTROL PLANE ENDPOINT: 10.171.113.233

WORKER NODE INSTANCE TYPE: large (cpu: 4, ram: 16 GB, disk: 40 GB)

ACTIVATE AUDIT LOGGING: ☑ Enable

10. Click **Next** two times until you get to the **vSphere Resources** screen.

a. Select the VM folder in which to place the management cluster VMs.

b. Select the vSphere datastores for the management cluster to use.

Note

The storage policy for the VMs can **only** be specified when you deploy the management cluster from a configuration file.

c. For the **Specify Availability Zones**, choose where to place the management cluster nodes, and then fill in the specifics:

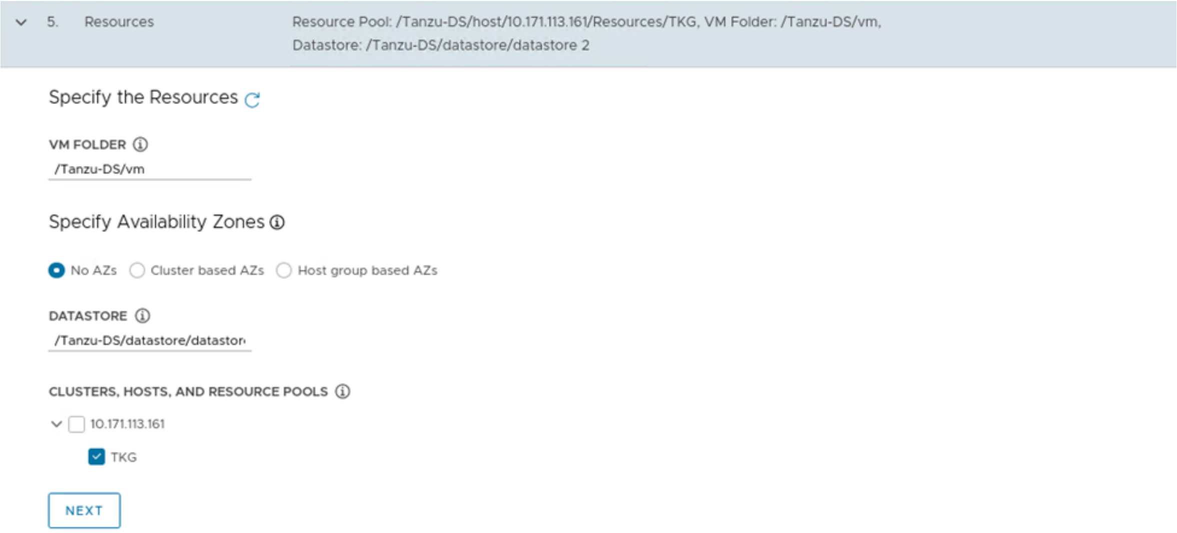

5. Resources — Resource Pool: /Tanzu-DS/host/10.171.113.161/Resources/TKG, VM Folder: /Tanzu-DS/vm, Datastore: /Tanzu-DS/datastore/datastore 2

Specify the Resources

VM FOLDER: /Tanzu-DS/vm

Specify Availability Zones

● No AZs ○ Cluster based AZs ○ Host group based AZs

DATASTORE: /Tanzu-DS/datastore/datastor

CLUSTERS, HOSTS, AND RESOURCE POOLS

□ 10.171.113.161

☑ TKG

11. Configure the Kubernetes network. Use the IP addresses from the Kubernetes CIDR block. Click **Next** when finished.

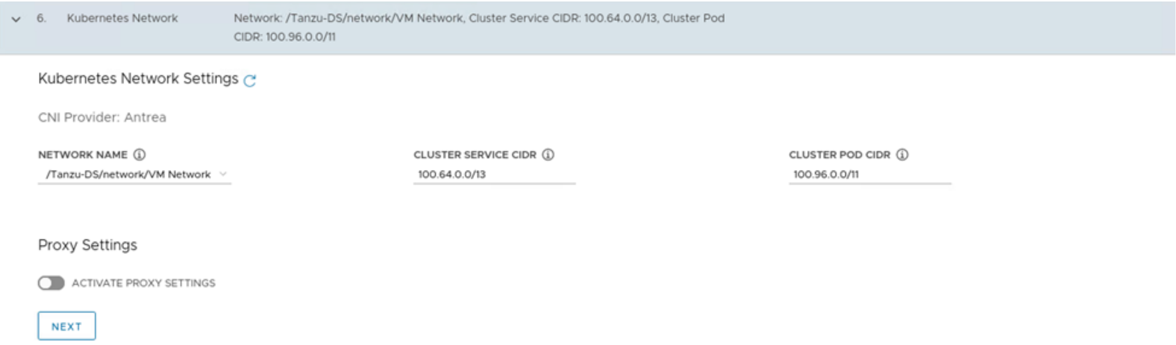

6. Kubernetes Network — Network: /Tanzu-DS/network/VM Network, Cluster Service CIDR: 100.64.0.0/13, Cluster Pod CIDR: 100.96.0.0/11

Kubernetes Network Settings

CNI Provider: Antrea

NETWORK NAME: /Tanzu-DS/network/VM Network

CLUSTER SERVICE CIDR: 100.64.0.0/13

CLUSTER POD CIDR: 100.96.0.0/11

Proxy Settings

○ ACTIVATE PROXY SETTINGS

12. Select the base OS image.

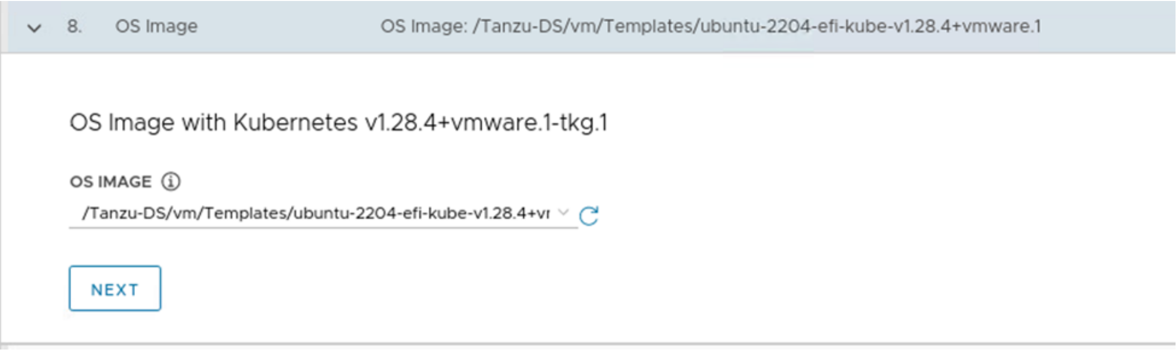

8. OS Image — OS Image: /Tanzu-DS/vm/Templates/ubuntu-2204-efi-kube-v1.28.4+vmware.1

OS Image with Kubernetes v1.28.4+vmware.1-tkg.1

OS IMAGE: /Tanzu-DS/vm/Templates/ubuntu-2204-efi-kube-v1.28.4+vı

13. Finalize the deployment, click **Review configuration**.

14. Click **Deploy Management Cluster** to deploy the cluster.

Post Deployment

After the deployment completes, you can see one control plane and one worker node.

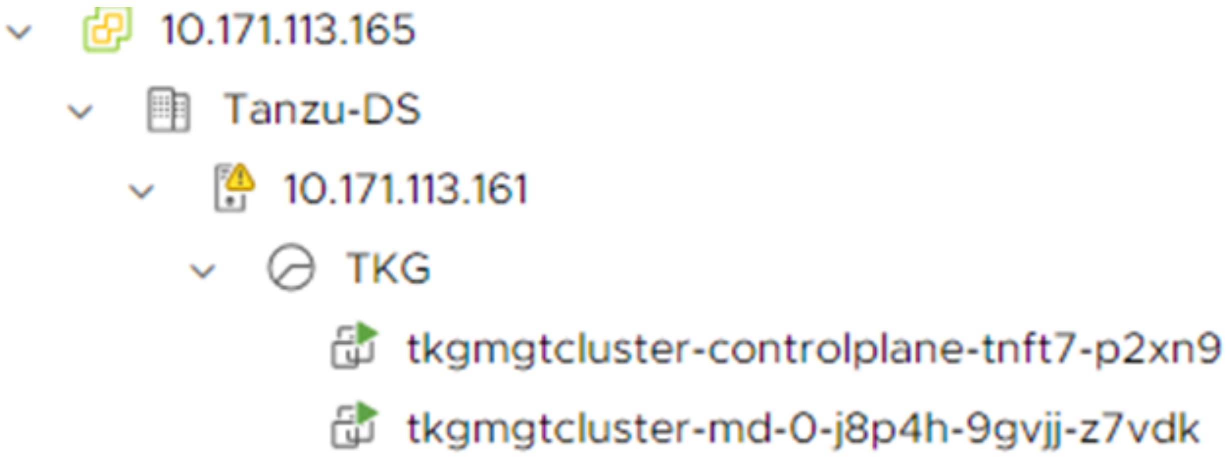

- 10.171.113.165
  - Tanzu-DS
    - 10.171.113.161
      - TKG
        - tkgmgtcluster-controlplane-tnft7-p2xn9
        - tkgmgtcluster-md-0-j8p4h-9gvjj-z7vdk

1. After the deployment completes, a configuration `yaml` file is created in `.../.config/tanzu/tkg/clusterconfigs/<random_name>.yaml`.

2. Create a copy of this `yaml` file and name it: `<newName>.yaml`.

3. Edit the `yaml` file. Provide a static IP address for the **VSPHERE_CONTROL_PLANE_ENDPOINT**.

4. Save the file.

Create a Workload Cluster

1. Create workload cluster, type:

   ```
   tanzu cluster create workernew -f <newName>.yaml
   ```

2. Scale the worker node count to three, type:

   ```
   tanzu cluster scale tkgwork01 --controlplane-machine-count 1 --
   worker-machine-count 3
   ```

3. Create a Tanzu cluster list, type:

   ```
   tanzu cluster list
   ```

4. Get the Kubernetes node list, type:

   ```
   kubectl get nodes
   ```

5. Retrieve the information about the current Kubernetes context, type:

   ```
   kubectl config get-contexts
   ```

Installation

See Install CTE for Kubernetes for more information.

Policy and Policy Elements

Use a policy with the following access:

CTE for Kubernetes Policy

Name	Access
	all_ops with Permit, Apply Key, Audit
	all_ops with Deny, Audit
Default	No Access

Support Contacts

If you encounter a problem while installing, registering, or operating the product, please refer to the documentation before contacting support. If you cannot resolve the issue, contact your supplier or Thales Customer Support.

Thales Customer Support operates 24 hours a day, 7 days a week. Your level of access to this service is governed by the support plan arrangements made between Thales and your organization. Please consult this support plan for further information about your entitlements, including the hours when telephone support is available to you.

Customer Support Portal

The Customer Support Portal, at Thales Customer Support, is where you can find solutions for most common problems. The Customer Support Portal is a comprehensive, fully searchable database of support resources, including software and firmware downloads, release notes listing known problems and workarounds, a knowledge base, FAQs, product documentation, technical notes, and more. You can also use the portal to create and manage support cases.

Tip

You require an account to access the Customer Support Portal. To create a new account, go to the portal and click on the REGISTER link.

Telephone Support

If you have an urgent problem, or cannot access the Customer Support Portal, you can contact Thales Customer Support by telephone at +1 410-931-7520. Additional local telephone support numbers are listed on the support portal.

Email Support

You can also contact technical support by email at technical.support@Thales.com.