VMware Integrated OpenStack with Kubernetes Getting Started Guide

VMware Integrated OpenStack 4.0

vmware®

You can find the most up-to-date technical documentation on the VMware Web site at:

https://docs.vmware.com/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

About This Book 1

The *VMware Integrated OpenStack with Kubernetes Getting Started Guide* provides information about how to install, deploy, and use VMware Integrated OpenStack with Kubernetes.

Intended Audience

As a system administrator, you can use VMware Integrated OpenStack with Kubernetes to manage containers in your Kubernetes cluster.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to http://www.vmware.com/support/pubs.

About VMware Integrated OpenStack with Kubernetes 2

With VMware Integrated OpenStack with Kubernetes you can deploy and maintain enterprise class Kubernetes clusters in an OpenStack environment.

The Kubernetes clusters are configured to use VMware Integrated OpenStack enterprise-grade services such as Keystone authentication for your cluster, Block Storage Cinder to provide persistent storage for your stateful applications, and Neutron Load Balancing as a Service (LBaaS) for your application services.

You deploy Kubernetes clusters through the VMware Integrated OpenStack with Kubernetes vApp in vCenter. The vApp provides a workflow that guides you through and completes the Kubernetes deployment process.

This chapter includes the following topics:

VMware Integrated OpenStack with Kubernetes Architecture

Kubernetes is an open-source platform for automating deployment, scaling, and operations of application containers across clusters of hosts, providing container-centric infrastructure. By combining Kubernetes with VMware Integrated OpenStack, you can use a common infrastructure management layer to provision both VMs and containers.

VMware Integrated OpenStack with Kubernetes builds high-availability Kubernetes clusters that support scalability and multi-tenancy.

Figure 2-1. VMware Integrated OpenStack with Kubernetes Built Cluster

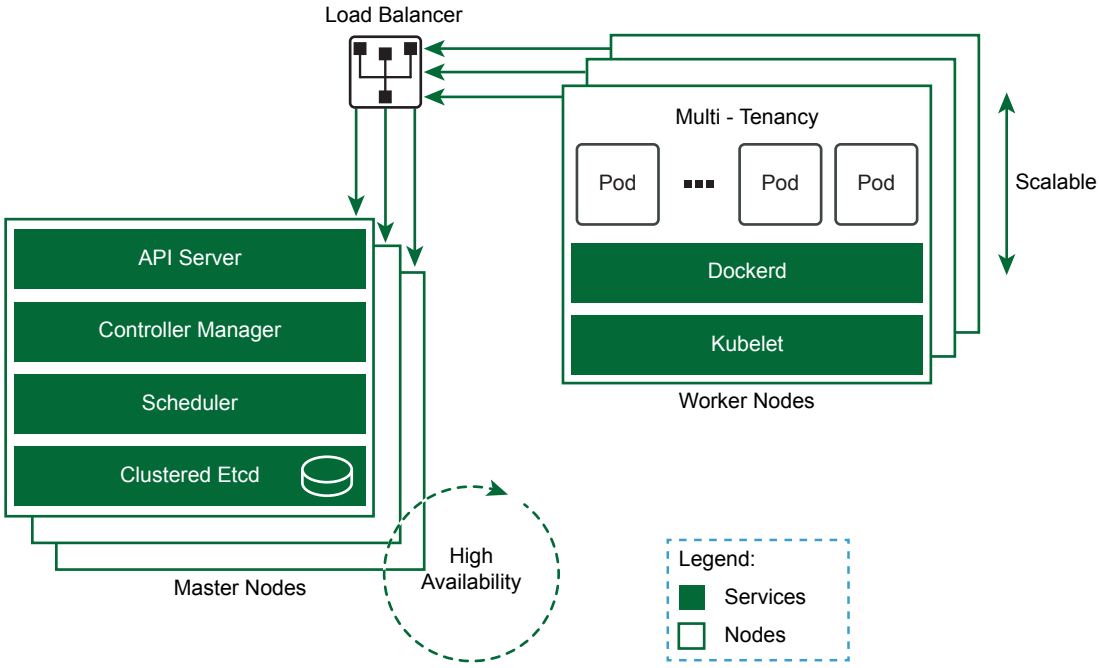

The high-availability Kubernetes cluster consists of load-balanced master nodes, replicated API servers, and clustered etcd services. In addition, you can scale out or scale in the worker nodes in a Kubernetes cluster to meet changing demands for capacity.

Using the concept of a namespace, a single Kubernetes cluster can be shared among multiple users or groups, or partitioned into multiple virtual clusters. With the namespace management feature, you can configure multi-tenancy on a shared Kubernetes cluster. Or you can create a Kubernetes cluster in exclusive mode, where any authorized user or group has privileges to manage the namespace.

Backend Networking

VMware Integrated OpenStack with Kubernetes supports VDS, NSX-V, and NSX-T backend networking.

Networking Support

Container network and load balancer support for Kubernetes Services is dependent on the backend networking.

Backend Networking	Container Network	Load Balancer
VDS	Flannel	Kubernetes Nginx Ingress Controller
NSX-V	Flannel	NSX Edge
NSX-T	NSX Container Plugin	Kubernetes Nginx Ingress Controller

Where:

- Flannel is a network fabric for containers and is the default for VDS and NSX-V networking.
- NSX Container Plugin (NCP) is a software component that sits between NSX manager and the Kubernetes API server. It monitors changes on Kubernetes objects and creates networking constructs based on changes reported by the Kubernetes API. NCP includes native support for containers. It is optimized for NSX-T networking and is the default.
- The NSX Edge load balancer distributes network traffic across multiple servers to achieve optimal resource use, provide redundancy, and distribute resource utilization. It is the default for NSX-V.

VDS Backend

VDS or vSphere Distributed Switch supports virtual networking across multiple hosts in vSphere.

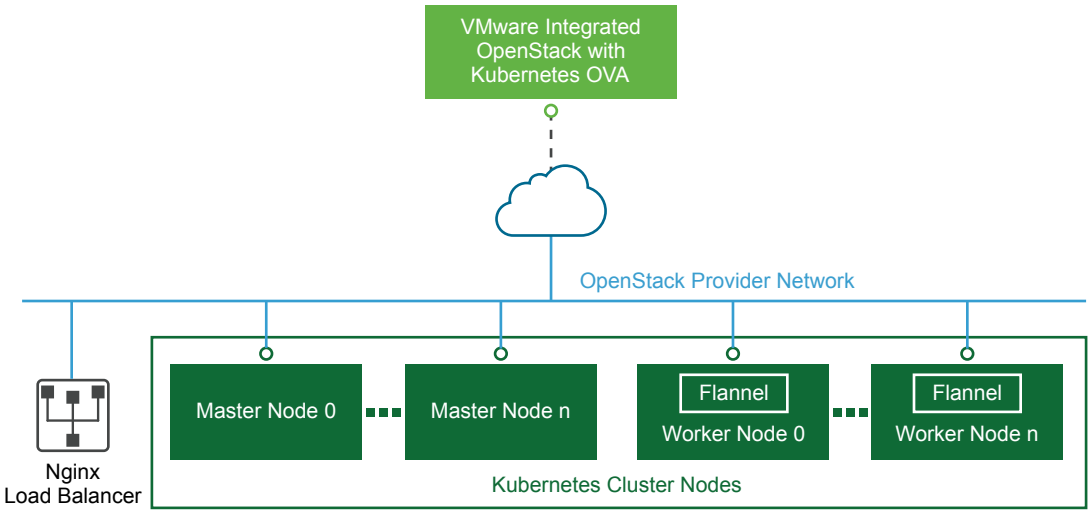

With the VDS backend, VMware Integrated OpenStack with Kubernetes deploys Kubernetes cluster nodes directly on the OpenStack provider network. The OpenStack cloud administrator must verify that the provider network is accessible from outside the vSphere environment. VDS networking does not include native load balancing functionality for the cluster nodes, so VMware Integrated OpenStack with Kubernetes deploys HAProxy nodes outside the Kubernetes cluster to provide load balancing.

NSX-V Backend

NSX-V is the VMware NSX network virtualization and security platform for vSphere.

Figure 2-2. NSX-V backend networking

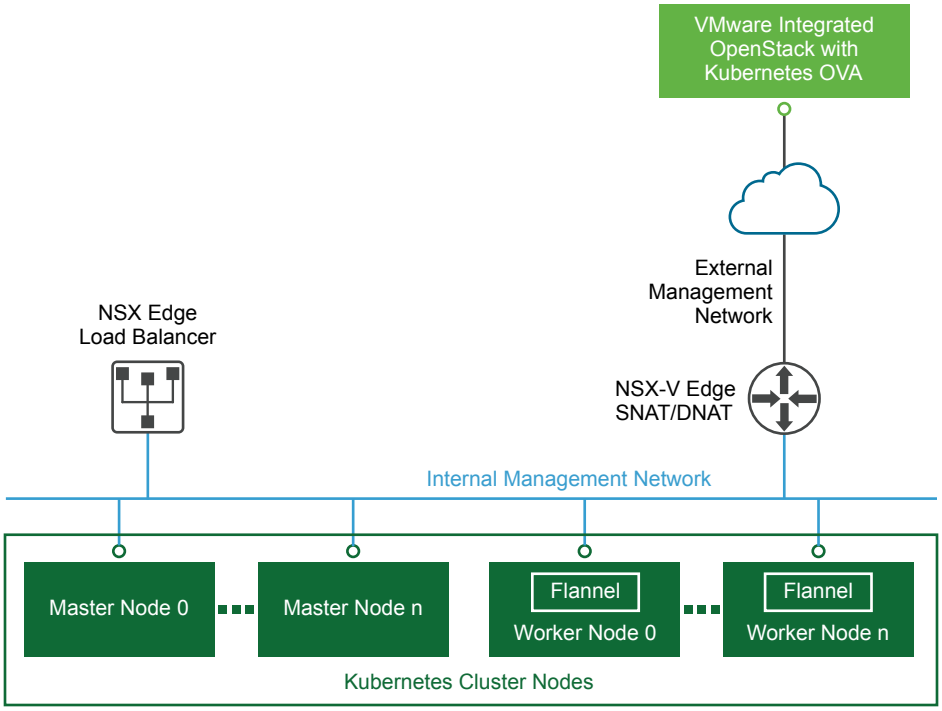

With the NSX-V backend, VMware Integrated OpenStack with Kubernetes deploys multiple nodes within a single cluster behind the native NSX Edge load balancer. The NSX Edge load balancer manages up to 32 worker nodes. Every node within the Kubernetes cluster is attached to an internal network and the internal network is attached to a router with a default gateway set to the external management network.

NSX-T Backend

NSX-T or NSX Transformer supports networking on a variety of compute platforms including KVM.

Figure 2-3. NSX-T backend networking

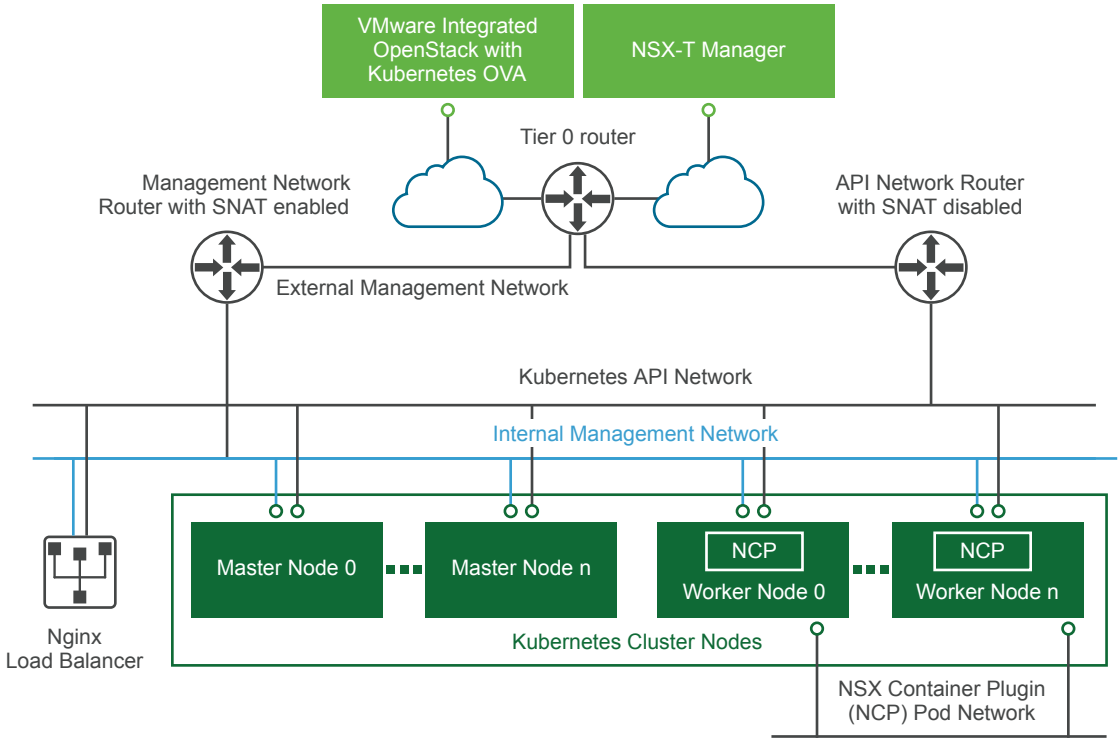

With the NSX-T backend, the VMware Integrated OpenStack with Kubernetes deploys worker nodes each with three NICs.

- One NIC connects to the NSX Container Plugin (NCP) Pod network, an internal network with no routing outside the vSphere environment. When the NSX Container Plugin is enabled, this NIC is dedicated to Pod traffic.
- One NIC connects to the Kubernetes API network which is attached to a router with SNAT disabled. Special Pods such as KubeDNS can access the API server using this network.
- One NIC connects to the internal management network. This NIC is accessible from outside the vSphere environment through a floating IP that is assigned by the external management network.

NSX-T networking does not include native load balancing functionality, so VMware Integrated OpenStack with Kubernetes creates two separate load balancer nodes. The nodes connect to the management network and API network.

Installing VMware Integrated OpenStack with Kubernetes 3

You install VMware Integrated OpenStack with Kubernetes as a virtual appliance in vSphere.

This chapter includes the following topics:

System Requirements

Before you begin the deployment tasks, your system must comply with all hardware, software, networking, and storage requirements.

System requirements specific to VMware Integrated OpenStack with Kubernetes are required in addition to system requirements for VMware Integrated OpenStack. See the VMware Integrated OpenStack Installation and Configuration Guide.

Hardware Requirements for VMware Integrated OpenStack with Kubernetes

The hardware requirements are based on the number of VMs used for each component.

Core VMware Integrated OpenStack with Kubernetes Components

Component	VMs	vCPU	vRAM (GB)	vDisk Space (GB)
VMware Integrated OpenStack with Kubernetes	1	4	16	60

Software Requirements for VMware Integrated OpenStack with Kubernetes

Before you begin installing VMware Integrated OpenStack with Kubernetes, verify that the software components meet all of the version prerequisites for vSphere, ESXi hosts, and the NSX product.

Requirement	Description
VMware Integrated OpenStack	■ Deployment on an OpenStack provider requires VMware Integrated OpenStack version 3.1 or 4.0
vSphere version	■ vSphere 6.0 or later with Virtual Distributed Switch. ■ Deployment with NSX-T backend networking requires vSphere 6.5.
ESXi hosts	■ Deployment on an SDDC provider requires ESXi version 6.0 or later

Requirement	Description
NSX	■ Deployment with NSX backend networking requires NSX version 6.2 or 6.3.
NSX-T	■ Deployment with NSX-T backend networking requires NSX-T version 2.0

Deploy the VMware Integrated OpenStack with Kubernetes OVA in the vSphere Web Client

VMware Integrated OpenStack with Kubernetes is a vApp that you deploy using a wizard in the vSphere Web Client.

Prerequisites

- Verify that vSphere is installed and correctly configured. See "System Requirements," on page 13.
- Obtain the VMware Integrated OpenStack with Kubernetes OVA from VMware.

Procedure

1. Download the VMware Integrated OpenStack with Kubernetes OVA file from the VMware Integrated OpenStack download page.
2. Login to the vSphere Web Client.
3. Go to the vCenter Hosts and Cluster view.
4. Right click the cluster where you want to deploy VMware Integrated OpenStack with Kubernetes and select **Deploy OVF Template**.
5. Access the downloaded VMware Integrated OpenStack with Kubernetes OVA.
6. Specify the destination and configure the OVA deployment.
   a. Select the target datacenter created specifically for the VMware Integrated OpenStack with Kubernetes OVA, and click **Next**.
   b. Select a target host, cluster, resource pool, or vApp and click **Next**.
   c. Review the product, version, vendor, publisher, size and description details and click **Next**.
   d. Select your storage options and click **Next**.
   e. To set up your networks, select the destination network for your source and click **Next**.

      The source is your management network and provides access to infrastructure components such as vCenter server, ESX, NSX-V, and NSX-T. The destination network you select must have access to the infrastructure components.
   f. Customize the deployment properties by configuring the networking properties and root user properties.
      - Networking properties are required to configure a static IP for the VM. Leave the properties blank if you want to the DHCP server to provide the IP address.
      - The root user uses the initial password to log in to VMware Integrated OpenStack with Kubernetes.
7. Click **Next**.
8. Review the deployment settings and click **Finish** to deploy VMware Integrated OpenStack with Kubernetes.
9. After the file finishes deploying into vCenter Server, power on the VMware Integrated OpenStack with Kubernetes appliance.

10 Wait for the machine to start, and right-click to obtain the IP address for the VM.

Adding a Cloud Provider 4

Before you deploy a Kubernetes cluster, you must configure a cloud provider. VMware Integrated OpenStack with Kubernetes uses the cloud provider to create the infrastructure required to deploy all your Kubernetes clusters. User management on the provider provides the basis for VMware Integrated OpenStack with Kubernetes user management on the cluster.

When choosing the type of provider to create, consider the following:

- With an existing VMware Integrated OpenStack deployment, you can create an OpenStack provider.
- Without an existing VMware Integrated OpenStack deployment, you can create an SDDC provider if you do not want to deploy a standalone VMware Integrated OpenStack instance.

This chapter includes the following topics:

OpenStack Cloud Provider

With an OpenStack provider, VMware Integrated OpenStack with Kubernetes deploys and configures Kubernetes clusters on an existing VMware Integrated OpenStack deployment.

An OpenStack provider supports NSX-V or NSX-T networking.

Network Requirements for OpenStack Provider

In addition to system requirements required for installation, your NSX-V or NSX-T network must satisfy requirements for an OpenStack provider.

NSX-V Requirements

The following requirements apply to NSX-V networking:

- Ubuntu 16.04 image installed in your VMware Integrated OpenStack cloud.
- An external network used to communicate with the VMware Integrated OpenStack with Kubernetes OVA. The external network must have at least one floating IP available. Additional floating IPs are required if you expose Kubernetes applications via Service with Load Balancer.
- An internal network for Kubernetes nodes. DNS must be available in your internal network.
- A centralized, exclusive router with a gateway configured for the external network.
- Security group for the ingress traffic. The security group must have ports 22 and 443 opened from outside.

NSX-T Requirements

In addition to the requirements for NSX-V networking, verify that NSX-T is version 2.0.

Input Parameters for an OpenStack Provider

This section describes the input parameters required to add an OpenStack provider. In addition, NSX-T backend networking requires specific configuration parameters.

An OpenStack provider requires the following information.

Table 4-1. OpenStack Authentication

Variable	Description
Keystone Public URL	Full Keystone public endpoint URL including protocol (http or https), port and API version. For example, https://openstack.cloud:5000/v3.
Username	OpenStack username
Password	OpenStack password
Project name	OpenStack project name
Region name (Default: nova)	OpenStack region name
Domain name (optional)	OpenStack domain name. Leave blank when using version 2 of authentication API. Must be set for v3.
CA Certificate	Certificate for authentication with the OpenStack Keystone service. See "Update a VMware Integrated OpenStack Load Balancer Certificate for an OpenStack Provider," on page 30

Table 4-2. Image and Flavor

Variable	Description
Image username	Used to establish SSH connection with cluster nodes. This user must be able run sudo without a password. For example, the default user for Ubuntu cloud images is **`ubuntu`**.
Image ID of the Ubuntu image	OpenStack image ID
Flavor ID	OpenStack flavor ID

Table 4-3. Networking and Security

Variable	Description
NSX-T Networking	See "Configuration Information for NSX-T Networking," on page 18.
Security Group ID	Security group ID to be applied to all VMs
Internal Network ID of Kubernetes cluster network	Internal network ID used for nodes IPs
Internal network Subnet ID	Subnet ID of the internal network used for allocating the IPs
External Network ID used for floating IPs	External network ID used to assign floating IPs

Configuration Information for NSX-T Networking

NSX-T networking requires specific input parameters.

Variable	Description
Manager address	NSX-T manager FQDN or IP
Username	NSX-T manager username
Password	NSX-T manager password
Tier 0 Router	Tier 0 router ID configured for OpenStack
Transport zone	Transport zone ID configured for OpenStack

Add an OpenStack Provider

Use the deployment wizard to add an OpenStack provider.

Prerequisites

Verify that you have the data for provider configuration. See "Input Parameters for an OpenStack Provider," on page 18.

Procedure

1 Login to VMware Integrated OpenStack with Kubernetes.

2 On the Cloud Providers home page, click **Deploy a New Provider**.

3 On the Intro page, click **Next**.

Alternatively, you can click **Choose File** to upload a JSON file containing provider information. The information automatically populates fields in the subsequent wizard screens.

4 Specify the provider name and select the **OpenStack** provider type. Click **Next**.

5 Configure OpenStack authentication.

- Specify the Keystone Public URL.
- Specify the username, the password, and the Project name.
- The Region name and the Domain name are optional.

6 Specify the image username. VMware Integrated OpenStack with Kubernetes displays a list of image IDs. Select an ID and click **Next**.

To search for a particular image name, click the filter icon and type a few letters of the image name.

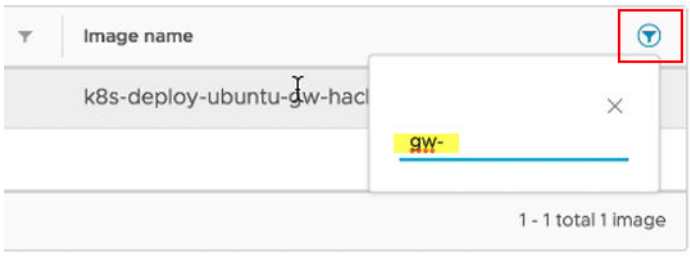

7 Select a flavor for the Kubernetes cluster nodes and click **Next**.

8 Configure the Neutron networking.

- NSX-V networking is the default. No special networking information is required.
- If NSX-T networking is being used, click the **NSX-T Networking** box and add NSX-T networking information.

9 Click **Next**.

10 Select the Security group and click **Next**.

11 Select the External network and click **Next**.

12 Select the Internal network and scroll down to select the Subnet ID. Click **Next**.

13 Review the Configuration Summary and click **Finish** to add the provider.

VMware SDDC Provider

With an SDDC provider, VMware Integrated OpenStack with Kubernetes creates an embedded VMware Integrated OpenStack deployment on an existing vSphere infrastructure. If you configure the embedded VMware Integrated OpenStack deployment for authentication with the local user database, you must add a cluster user on the provider.

An SDDC provider supports VDS, NSX-V, or NSX-T networking.

Network Requirements for SDDC Provider

In addition to system requirements required for installation, your VDS, NSX-V, or NSX-T network must satisfy requirements for an SDDC provider.

VDS Requirements

The following requirements apply to VDS networking:

- One or more vSphere clusters with at least one ESXi host configured.
- Uplinks from all hosts in the active clusters.
- Port group that can access the management network.
- Each cluster requires a port group with IP addresses for:
  - two load balancers and one virtual IP
  - each master node
  - each worker node

A network with multiple clusters requires IP addresses for all port groups.

NSX-V Requirements

The following requirements apply to NSX-V networking:

- vSphere 6.0 or later installed.
- NSX-V 6.2 or later installed and configured.
- One or more vSphere clusters with at least one ESXi host configured in an NSX transport zone.
- The Kubernetes cluster must be in a single NSX transport zone.
- Each cluster requires one VDS-based port group with at least five static IP addresses. A network with multiple clusters requires IP addresses for all port groups.
- Datastore connected to the vSphere cluster

NSX-T Requirements

In addition to the requirements for NSX-V networking, the following requirements apply to NSX-T networking:

- NSX edge installed and configured in NSX-T.

- A range of IP addresses available to avoid conflict in the datacenter.

Input Parameters for an SDDC Provider

This section describes the input parameters required to add an SDDC provider. In addition, NSX-V or NSX-T backend networking require specific configuration parameters. Authentication also requires specific configuration parameters.

An SDDC provider requires the following information.

Table 4-4. vSphere Authentication

Variable	Description
vSphere hostname	FQDN or IP of vCenter server
vSphere username	vCenter server username
vSphere password	vCenter server password
Ignore the vCenter Server certificate validation?	If checked, VMware Integrated OpenStack with Kubernetes does not verify the vCenter Server certificate when connecting to the vCenter.

Table 4-5. vSphere Cluster and Datastore Configuration

Variable	Description
Compute cluster	vSphere compute cluster used to deploy Kubernetes cluster nodes
Datastores	vSphere datastores used to store Kubernetes cluster nodes, images, and volumes

Table 4-6. Management Network Setting for Kubernetes Cluster Nodes

Variable	Description
Port Group	Distributed port group that Kubernetes cluster nodes connect to. Not applicable for NSX-T networking.
VLAN ID (optional)	VLAN ID of the management portgroup. Leave blank if not using VLAN.
Network Address	Management network address in CIDR format such as 192.168.0.0/24.
IP Range	Start and end IP addresses of the management network allocation IP range.
Gateway	Gateway IP for the management network
DNS (optional)	DNS servers to be used if DNS for the management network is unavailable. To specify multiple servers, use comma separated values.

Networking Parameters

NSX-V or NSX-T networking requires specific input parameters.

Table 4-7. Configuration Information for NSX-V Networking with SDDC provider

Variable	Description
Manager address	FQDN or IP of the NSX-V manager
Username	NSX-V manager username
Password	NSX-V manager password

Table 4-7. Configuration Information for NSX-V Networking with SDDC provider (Continued)

Variable	Description
Ignore the NSX-V SSL certificate validation?	If checked, VMware Integrated OpenStack with Kubernetes does not verify the NSX-V SSL certificate when connecting to the NSX-V server.
Transport zone	Transport zone configured for NSX-V networking
Edge resource pool	vSphere resource pool for the NSX Edge VMs
Edge datastore	vSphere datastore for NSX Edge VMs
Virtual Distributed Switch	vSphere Distributed Switch configured for NSX-V networking
External network	vSphere distributed port group on the distributed switch

Table 4-8. Configuration Information for NSX-T Networking with SDDC provider

Variable	Description
Manager address	FQDN or IP of the NSX-T manager
Username	NSX-T manager username
Password	NSX-T manager password
Ignore the NSX-T SSL certificate validation?	If checked, VMware Integrated OpenStack with Kubernetes does not verify the NSX-T SSL certificate when connecting to the NSX-T server.
Tier 0 Router	Tier 0 router pre-configured for NSX-T networking
Default overlay transport zone	Overlay transport zone pre-configured for NSX-T networking
Default VLAN transport zone	VLAN transport zone pre-configured for NSX-T networking

Authentication Source Parameters

If you create a standalone user database, VMware Integrated OpenStack with Kubernetes creates a Kubernetes cluster admin user in the database to start. VMware Integrated OpenStack with Kubernetes also supports both Active Directory as an LDAP server on Windows and LDAP server for Unix and Linux.

Table 4-9. Local Admin User Authentication Source

Variable	Description
Kubernetes cluster admin user	Admin user for authentication with the local user database
Kubernetes cluster admin password	Password for authentication with the local user database

Table 4-10. Active Directory as LDAP Backend Authentication Source

Variable	Description	Default
Encryption	SSL or None	None
Hostname	FQDN or IP of the LDAP or AD server	None
Port	Port	636 for SSL 389 for non-SSL
Bind user	LDAP bind user.. Same as Kubernetes cluster admin user.	None

Table 4-10. Active Directory as LDAP Backend Authentication Source (Continued)

Variable	Description	Default
Bind Password	Password for LDAP bind user. Same as Kubernetes cluster admin user.	None
User Tree DN	Search base for users	None
Group Tree DN	Search base for groups	None
User object/class	LDAP objectclass for users	organizationalPerson
User ID attribute	LDAP attribute mapped to user ID. This must not be a multivalued attribute.	cn
User name attribute	LDAP attribute mapped to user name.	userPrincipalName
User mail attribute	LDAP attribute mapped to user e-mail	mail
User password attribute	LDAP attribute mapped to password	userPassword
User enabled attribute	LDAP attribute mapped to user enabled flag	userAccountControl
Group object/class	LDAP objectclass for groups	group
Group ID attribute	LDAP attribute mapped to group ID	cn
Group name attribute	LDAP attribute mapped to group name	sAMAccountName
Group member attribute	LDAP attribute mapped to group member	memberOf
Group description attribute	LDAP attribute mapped to group description	description

Add a VMware SDDC Cloud Provider

Use the deployment wizard to add an SDDC cloud provider.

Prerequisites

Verify that you have the data for provider configuration. See "Input Parameters for an SDDC Provider," on page 21.

Procedure

1 Login to VMware Integrated OpenStack with Kubernetes.

2 On the Cloud Providers home page, click **Deploy a New Provider**.

3 On the Intro page, click **Next**.

Alternatively, you can click **Choose File** to upload a JSON file containing provider information. The information automatically populates fields in the subsequent wizard screens.

4 Specify the provider name and select the **SDDC** provider type. Click **Next**.

5 Specify the vSphere vCenter hostname, username, and password. Leave the **Ignore the vCenter Server certificate validation** option checked. Click **Next**.

6 With the vSphere vCenter information provided, VMware Integrated OpenStack with Kubernetes displays a list of available vSphere clusters. Select a cluster and click **Next**.

7 Select at least one vSphere datastore and click **Next**.

8 Configure networking.

Select VDS, NSX-V, or NSX-T networking and select the network.

9 Configure the management network.

- Select a Port Group.
- VLAN ID is optional.
- Provide the Network Address.
- Provide the IP range.
- Provide the Gateway.
- DNS is optional.

10 Click **Next**

11 Configure the authentication source.

Select **Local Admin User** or **Active Directory as LDAP Backend**.

- For Local Admin User, specify the Kubernetes cluster admin username and password.
- For Active Directory as LDAP backend, see "Authentication Source Parameters," on page 22.

12 Click **Next**

13 Review the Configuration Summary and click **Finish** to add the provider.

Manage Users and Groups on Your SDDC Provider from the UI

If you selected Local Admin User as the authentication source for your SDDC provider, you must add a user for the provider. Later when you create a Kubernetes cluster, you will select this user for the cluster.

Procedure

1 Login to VMware Integrated OpenStack with Kubernetes.

2 On the **Cloud Providers** home page, select an SDDC cloud provider.

3 To manage or add users, click **Users**.

A list of the existing users appears.

- Click **+New** to add a new user with username and password.
- Click the three dots to the right of a user name to delete or edit an existing user.

4 To manage or add groups, click Groups.

A list of the existing groups appears.

- Click **+New** to add a new group and select the users to add to the group.
- Click the three dots to the right of a group name to delete or edit an existing group.

Creating Your First Cluster 5

You create a Kubernetes cluster based on a provider and populate it with master and worker nodes. VMware Integrated OpenStack with Kubernetes supports exclusive and shared cluster types.

This chapter includes the following topics:

Understanding Cluster Configuration Settings

After creating a provider, you create a Kubernetes cluster. To configure your cluster correctly, review the cluster configuration settings.

Node Types

A Kubernetes cluster is comprised of two types of nodes. Each node in the VMware Integrated OpenStack with Kubernetes is a VM. Node settings can be changed after cluster deployment.

Master Nodes A master node provides the Kubernetes API service, scheduler, replicator, and so on. It manages the worker nodes. A cluster with a single master node is valid but has no redundancy.

Worker Nodes A worker node hosts your containers. A cluster with a single worker node is valid but has no redundancy.

Cluster Types

VMware Integrated OpenStack with Kubernetes supports two types of clusters. The cluster type cannot be changed after cluster deployment.

Exclusive Cluster In an exclusive cluster, multi-tenancy is not supported. Any authorized Kubernetes user using the Kubernetes CLI of APIs has namespace management privileges.

The exclusive cluster provides a familiar environment for developers who deploy Kubernetes themselves.

Shared Cluster

In a shared cluster, multi-tenancy is supported and enforced by the Kubernetes namespace. Only a VMware Integrated OpenStack with Kubernetes administrator using the VMware Integrated OpenStack with Kubernetes interface or CLI has namespace management privileges.

The shared cluster is an environment where the administrator can manage resource isolation among users.

Add a New Kubernetes Cluster

You deploy a Kubernetes cluster on an OpenStack or SDDC provider. VMware Integrated OpenStack with Kubernetes supports exclusive and shared cluster types.

Prerequisites

- Verify that the cloud provider you deployed is active. See Chapter 4, "Adding a Cloud Provider," on page 17.
- Determine the type of cluster you want to add.

Procedure

1. Login to VMware Integrated OpenStack with Kubernetes.
2. On the Clusters home page, click **Deploy a New Cluster** or click **+NEW** to add a cluster to a set of existing clusters.
3. On the Intro page, click **Next**.

   Alternatively, you can click **Choose File** to upload a JSON file containing cluster configuration information. The information automatically populates fields in the subsequent wizard screens.
4. Highlight the infrastructure provider and click **Next**.
5. If no node profiles are listed, click the **Ignore node profile** box and click **Next**.
6. Configure the cluster.
   - Provide a cluster name.
   - Specify the number of master nodes in the cluster. Because etcd servers coexist on master nodes, you must specify an odd number of nodes to support high availability and fault tolerance.
   - Specify the number of worker nodes in the cluster.
   - The DNS server IP is optional.
   - Select the **Shared Cluster** or **Exclusive Cluster** type.
     - For a shared cluster, specify a namespace.
     - For an exclusive cluster, specify a user and group.
7. Click **Next**.
8. Review the Configuration Summary and click **Finish** to add the cluster

Configure User and Group Access for an Exclusive Cluster

Once a Kubernetes cluster is created, you can authorize users or groups for the cluster. The users and groups belong to the SDDC or OpenStack provider where the cluster was created.

Procedure

1 Login to VMware Integrated OpenStack with Kubernetes.

2 On the Clusters home page, click the three dots to the right of an existing exclusive cluster and select **+Configure user & group**.

3 In the Configure user and group for cluster dialogue box, check the boxes for users or groups that you want to authorize for access to the cluster. Or check off the boxes for users or groups that you no longer want to authorize for access to the cluster.

4 Click **OK**.

Create a Namespace for Users and Groups on a Shared Cluster

The shared cluster has restricted access with multi-tenancy support based on the Kubernetes namespace. An administrator can create a namespace and authorize access to users or groups specified in a namespace policy.

Procedure

1 Login to VMware Integrated OpenStack with Kubernetes.

2 On the Clusters home page, click the three dots to the right of an existing shared cluster and select **+Add namespace**.

3 In the Add a new namespace dialogue box, type a name for the namespace and check the boxes for users or groups that you want to authorize for access to the namespace.

4 Click **OK**.

Managing Your Deployment 6

Once you create your provider and cluster, you can use VMware Integrated OpenStack with Kubernetes to manage your deployment.

This chapter includes the following topics:

Cluster Management

After creating your Kubernetes cluster, you may want to expand the cluster.

Scale Your Cluster

If the cluster you added is not large enough, you can increase the number of Kubernetes worker nodes to increase capacity. You can also decrease the number of worker nodes.

Prerequisites

Verify that the state of the Kubernetes cluster you want to scale is ACTIVE.

Procedure

1 Login to VMware Integrated OpenStack with Kubernetes.

2 On the Clusters home page, click the three dots to the right of a cluster name and select **Scale Cluster**.

3 In the Scale cluster dialogue box, enter the desired number of worker nodes for the cluster.

4 Click **OK**

Delete Your Cluster

If a cluster is no longer needed, you can delete it to free resources for another use.

Procedure

1 Login to VMware Integrated OpenStack with Kubernetes.

2 On the Clusters home page, click the three dots to the right of a cluster name and select **Delete Cluster**.

3 Click **OK** to confirm.

Certificate Management Using the CLI

You use self-signed certificates in VMware Integrated OpenStack with Kubernetes for authentication. In some cases, certificates may need to be modified and regenerated.

Update a VMware Integrated OpenStack Load Balancer Certificate for an OpenStack Provider

If you deploy an OpenStack provider on a VMware Integrated OpenStack deployment with a self-signed certificate, you must replace the self-signed certificate with a certificate that the Kubernetes cluster recognizes.

Procedure

1 On the VMware Integrated OpenStack VM, perform the following steps:

   a In the `/tmp/vio.cnf` file, edit the value `basicConstraints=CA:TRUE`

   b On the command line, type the commands:

```
openssl req –new –key /var/lib/vio/jarvis/{uuid}/credentials/etc/ssl/private/vio.key –
out /tmp/vio.csr –config /tmp/vio.cnf
openssl x509 –req –days 3650 –extensions v3_req –extfile /tmp/vio.cnf –in /tmp/vio.csr –
signkey /var/lib/vio/jarvis/{uuid}/credentials/etc/ssl/private/vio.key –out /tmp/vio.crt
cat /var/lib/vio/jarvis/{uuid}/site–req–{uuid}–hosts.ini | grep deployment_name
viocli deployment cert–update –d deployment_name –f /tmp/vio.crt
```

   c From the OpenStack Horizon web portal, verify that the certificate is working.

2 On the VMware Integrated OpenStack with Kubernetes VM, perform the following steps:

   a Copy the `vio.crt` file to the local machine.

   b When adding a provider, choose the file as the Root CA file for certificate validation. See "Input Parameters for an OpenStack Provider," on page 18.

Manage SSL Certificates for Kubernetes API Servers

If Kubernetes API servers are accessed over a public internet, you may want to use a certificate signed by a trusted certificate authority (CA) to further secure your Kubernetes deployment.

You can use VMware Integrated OpenStack with Kubernetes CLI to prepare a certificate signing request. After the CA generates a signed certificate, you can use the CLI to upload it to a target Kubernetes cluster.

Prerequisites

If application pods are deployed already and these pods use the Kubernetes secret tokens, back up application data and remove the pods before updating the API server certificate. The certificate update invalidates the secret tokens, so you must re-create the pods following the update.

Procedure

1 Login as root to the VMware Integrated OpenStack with Kubernetes VM. Provide the root password set during OVA deployment.

```
vkube login ––insecure
```

2 Generate a certificate signing request.

```
vkube cluster list --insecure
vkube cluster csr <cluster_id>
   --country-name <value1>
   --locality-name <value2>
   --organization name <value3>
   --organizatio-unit-name <value4>
   --state-name <value5>
   --insecure
```

3 Copy the existing Kubernetes `/etc/kubernetes/openssl.conf` file from the Kubernetes Master0 node and send it with the Certificate Signing Request (CSR) file to your company's CA administrator.

4 Using the CSR file and the extfile from openssl.conf, the CA administrator generates a signed certificate. Upload the API server's certificate and corresponding CA certificate to the Kubernetes cluster.

```
vkube cluster crt <cluster-id> --insecure --ca-file-name ca.pem --crt-file-name apiserver.pem
```

5 Login to the Master0 node and type:

```
kubectl get pod --namespace=kube-system
```

When all the pods change to status running, the cluster is ready to use.

Optimizing Kubernetes Cluster Performance

7

To ensure optimal Kubernetes cluster performance, you should follow certain best practices. This section highlights some of the key best practices.

Setting Adequate Quotas in OpenStack

For an OpenStack provider, set quotas that are large enough to accommodate a large cluster.

Table 7-1. Sample Commands

Command	Description
`nova quota-update --key-pairs 500 --instances 500 --cores 4000 --ram 12288000 <tenant_ID>`	Set quotas for a 500-node cluster, where each node has 8 vCPUs and 24G RAM
`neutron quota-update --tenant-id <tenant_ID> --pool 300 --port 1000 --loadbalancer 300 --floatingip 150`	Neutron command to allocate quota according to your network. Port number should be greater than instance plus load balancer number.
`cinder quota-update --volumes 500 --gigabytes 5000 <tenant_ID>`	Cinder command to allocate quota according to the number of persistent volumes that you want to create.

Best Practices for Creating Large Clusters

To create a large cluster, a best practice is to first create a small cluster, then scale it out. For example, to create a stable 500-node cluster, start by creating a 30-node cluster, then scale it out with a maximum of 30 nodes at a time until you reach 500 nodes.

Tips:

- If your cluster is larger than 200 nodes, you might see RPC timeouts in the OpenStack service logs. If that occurs, increase the RPC timeout setting for those services. For example for a Nova service, increase the value of the rpc_response_timeout configuration option in the nova.conf file.
- It may take time to refresh the status of created resources when scaling out a cluster. Add the --skip-refresh option to the `vkube cluster scaleout` command to decrease the deployment time. With this option, the scale out operation does not check the state of existing resources such as VMs or load balancers, and assumes that the resources are successfully created.

Managing High CPU Usage with an OpenStack Provider

If you are using VMware Integrated OpenStack deployed in compact mode as your OpenStack provider, you may notice high CPU usage on the controller or compute service VM's. If so, increase the number of vCPU's to 16 per VM.

Alternatives to Load Balancing with NSX-V Backing

When you create services in Kubernetes and you specify the type as LoadBalancer, NSX Edge load balancers are deployed for every service. The load balancer distributes the traffic to all Kubernetes worker nodes up to 32 members. If your Kubernetes cluster includes more that 32 worker nodes, use the Kubernetes Ingress resource instead.

Persistent Volume Claim Management

If you create many persistent volume claims and associated pods in parallel, dynamic provisioning of persistent volumes may fail. If you check the OpenStack service logs and see that the failures are due to RPC timeouts, increase the RPC timeout setting for those services. For example for a Nova service, increase the value of the rpc_response_timeout configuration option in the nova.conf file.

Index