

Tencent Kubernetes Engine Clusters Product Introduction



Copyright Notice

©2013-2018 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

Contents

Clusters

- Overview

- Basic Operations of Clusters

- Cluster Autoscaler

- Usage of Nodes

Clusters

Overview

Last updated : 2018-06-21 14:41:28

A cluster is a collection of cloud resources required for containers to run, including several CVMs, load balancers and other Tencent Cloud resources.

Cluster Information

Cluster Type: Container clusters in VPCs are supported.

Cluster Configuration: When creating a cluster, you can configure CVM model, operating system, system disk and data disk size, login password, etc.

Cluster Structure: Prepaid and postpaid CVMs are supported.

Cluster Management

You can create clusters, scale up/down nodes, delete clusters and work with clusters directly by using Kubernetes APIs.

Retaining Resources for Cluster

Some resources are retained for each node in a cluster to keep Kubernetes running normally. Specific rules are as follows:

CPU

Total CPU capacity on node (in core)	1	2	4	8	16	32
Retained CPU capacity on node (in core)	0.06	0.07	0.08	0.09	0.11	0.14

Memory

Total memory capacity on node (in Gib)	1	2	4	8	12	16	24	32	48
--	---	---	---	---	----	----	----	----	----

Total memory capacity on node (in Gib)	1	2	4	8	12	16	24	32	48
Retained memory capacity on node (in Mib)	160	320	420	830	1200	1300	1660	1830	2420

Moreover, an extra 100 Mib of memory is retained for each node to avoid the OOM error.

After log collection feature is enabled for a cluster, about 0.3 core of CPU and 250 Mib of memory are used to run log collection plug-in.

Help Documentations

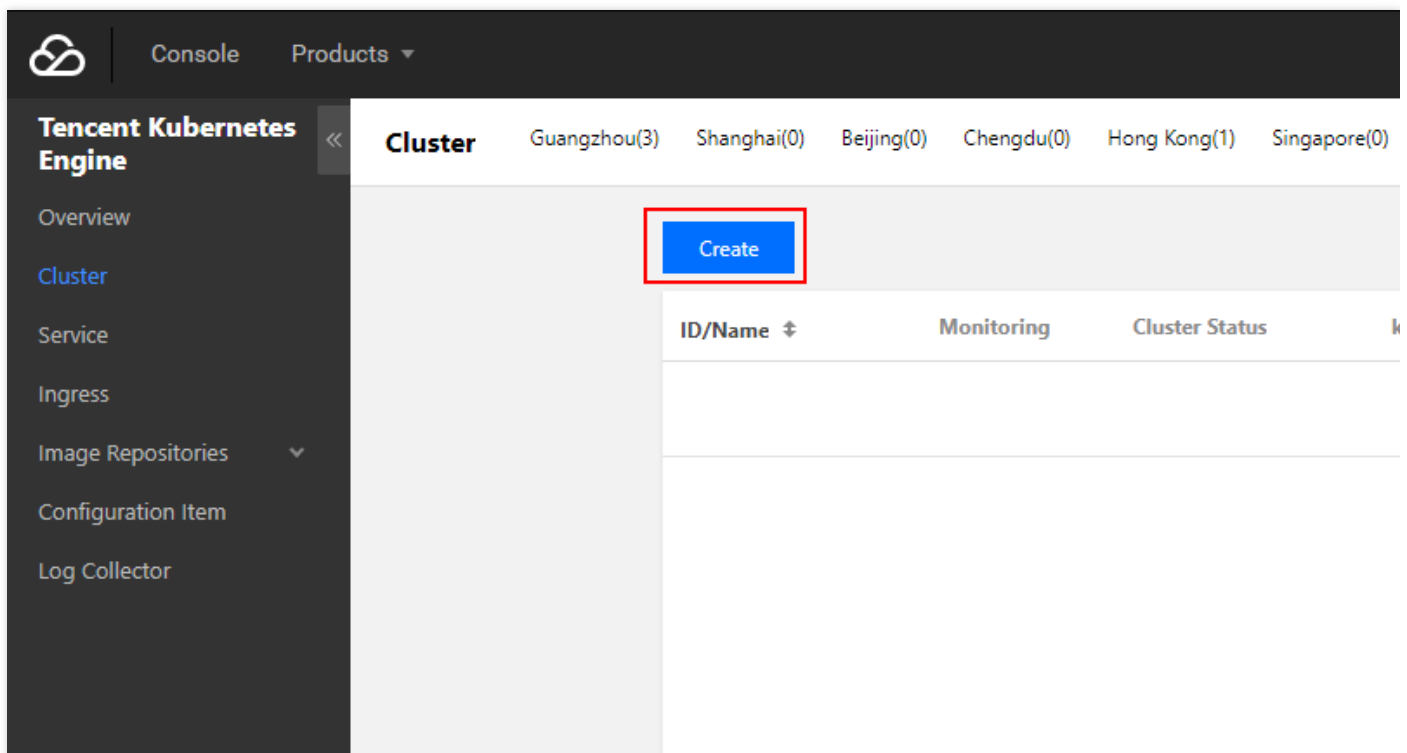
- [Basic Operations of Cluster](#)
- [Cluster Lifecycle](#)
- [Cluster Quota Limit](#)
- [Network Configuration of Cluster Nodes and Containers](#)
- [Disk Configuration of Cluster Nodes](#)
- [Public IP Configuration of Cluster Nodes](#)
- [Security Group Configuration of Cluster Nodes](#)

Basic Operations of Clusters

Last updated : 2018-10-09 18:02:03

Creating a Cluster

1. Log in to the [Tencent Cloud TKE console](#).
2. Click **Cluster** in the left navigation bar, and click **New** in the cluster list page.



3. Configure the basic cluster information.
 - **Cluster name:** The name of a cluster to be created, with a length limited to 60 characters.
 - **Billing method:** Both Prepaid and Postpaid are supported. For more information, please see [Billing Method](#).
 - **Region:** Select a closest region based on your location. This helps minimize access latency and improve download speed.
 - **Availability zone:** Clusters in the same region are interconnected with each other through private network, but this is not the case for those distributed in different regions. The same region must be chosen for communication over private networks.
 - **Node network:** The system assigns the IP addresses within the node network address range to the CVMs in the cluster. For more information, please see [Network Configuration of Containers and](#)

Nodes.

- **Container network:** The system assigns the IP addresses within the container network address range to the containers in the cluster. For more information, please see [Network Configuration of Containers and Nodes](#).
- **Cluster description:** Information about cluster creation, which is displayed on the **Cluster Information** page.

The screenshot displays the 'Create Cluster' interface in the Tencent Cloud console. The left sidebar shows the 'Tencent Kubernetes Engine' menu with options like Overview, Cluster, Service, Ingress, Image Repositories, Configuration Item, and Log Collector. The main area is titled 'Create Cluster' and has a progress bar with four steps: 1. Cluster Information (active), 2. Select the model, 3. CVM Configuration, and 4. Confirm Info.

A blue banner at the top of the form states: 'To use TKE, you need to create a cluster. A cluster consists of several nodes (CVMs) on which services are running. To learn more, please see [Cluster Overview](#)'.

The form fields include:

- Cluster Name:** A text input field with a hint 'Up to 60 characters'.
- Project:** A dropdown menu currently set to 'Default Project'. A note below states: 'New resources such as CVMs and load balancers in cluster are automatically allocated to this project. [Instruction](#)'.
- Node settings:** Two buttons: 'Add node' (highlighted in blue) and 'Empty cluster'.
- Kubernetes version:** A dropdown menu set to '1.10.5'.
- Region:** A grid of buttons for different regions: Guangzhou, Shanghai, Beijing, Chengdu, Hong Kong, Singapore, Mumbai, Silicon Valley (highlighted in blue), Virginia, Moscow, and Frankfurt. A note below states: 'Tencent Cloud products in different regions CANNOT communicate via private network and CANNOT be changed once purchased. Please choose a region close to your customers to minimize access latency and improve download speed.'
- Cluster network:** A dropdown menu set to 'VINT' and a text field showing 'CIDR: 192.168.0.0/19'. A note below states: 'If no suitable network is found, you can [Create a VPC](#)'.
- Container Network:** A section with a title and an information icon. It contains:
 - CIDR:** A field with a dropdown set to '172', a text input '16', a dropdown set to '0', a text input '0', and a dropdown set to '16', followed by a slash and a dropdown set to '16'. An 'Instruction' link is to the right.
 - Maximum number of pods/node:** A dropdown menu set to '256'.
 - Maximum number of services/cluster:** A dropdown menu set to '256'.A note at the bottom of this section states: 'Under the current container network configuration, the maximum number of nodes for each cluster is: 255 nodes'.

At the bottom of the form are 'Cancel' and 'Next' buttons.

4. Select a model (all models with cloud disks as system disks are supported).

- **Series:** **Series 1** and **Series 2** are available. For more information, please see [Pod Types](#).
- **Model:** For more information on how to select a model, please see [Select CVM Configuration Solution](#).

Selected Configuration

Cluster Name: virgilttest

Kubernetes version: 1.10.5

Region: West US(Silicon Valley)

Container Network: 172.16.0.0/16

Billing Mode: ☒ Postpaid

Region: West US(Silicon Valley)

Availability Zone: ☒ Silicon Valley Zone 1 ☐ Silicon Valley Zone 2

Node Network: Total subnet IPs: 253; Available subnet IPs: 253

CIDR: 192.168.0.0/19

If no suitable network is found, you can [Create a VPC](#) or [Create a Subnet](#)

Pod family: ☒ All pod families ☐ Standard ☐ High IO ☐ MEM-optimized ☐ Compute ☐ GPU-based

Instance type: ☒ All pod types ☐ Standard S2 ☐ High IO I2 ☐ MEM-optimized M2 ☐ Compute C2 ☐ GPU Compute GN8

Model	Specifications	CPU	MEM	Configuration Fee
<input checked="" type="radio"/> Standard S2	S2.SMALL1	1-core	1GB	\$ 0.02 USD/hourup
<input type="radio"/> Standard S2	S2.SMALL2	1-core	2GB	\$ 0.03 USD/hourup

5. Enter CVM configuration.

- **System disk:** Always 50 GB.
- **Data disk:** The increment is 10 GB. The maximum value is 4,000 GB.
- **Public network bandwidth:** Two billing methods are available. For more information, please see [Purchase Network Bandwidth](#).
- **Bandwidth:** If you select **Assign public IP for free**, the system assigns a public IP at no cost. If there is no need, set bandwidth value to 0.
- **Login method:** Three login methods are provided.
 - i. **Set password:** Set a password as instructed.
 - ii - **Immediately associate key:** A key pair is a pair of parameters generated by an algorithm, which makes it more secure to log in to a CVM than using a password. For more information, please see [SSH Key](#).
 - iii. **Automatically generate password:** The automatically generated password is sent to you via the internal message.

- **Security group:** It functions as a firewall and is used to set the network access control of CVMs. For more information, please see [Configure TKE Security Group](#).
- **Number of CVMs:** Select the number of CVMs.

Selected Configuration

Cluster Name	virgilttest
Kubernetes version	1.10.5
Region	West US(Silicon Valley)
Container Network	172.16.0.0/16
Billing Mode	Postpaid
Model	Standard S2
Specifications	S2.SMALL1(1-core 1 GB)

Operating system: Ubuntu Server 16.04.1 LTS 64位

System disk: HDD Cloud Storage (0GB to 500GB slider, 50GB selected)

Data disk: Purchase Later

Data Disk Type: HDD Cloud Storage (0GB to 16000GB slider, 0GB selected)

Public IP: Purchase Now

Back Next

6. The created cluster is displayed in the cluster list.

Add a CVM

1. On the cluster list page, click **Add Node** on the right.
2. Set the **Network** to which the new CVM belongs, its **Model** and **Configuration information**.
You may create CVMs in different subnets under different availability zones in the same region.
3. The new CVM can be found in the **ID/Node Name** column.

Terminating a CVM

1. Click the **ID/Name** of a cluster on the cluster list page to enter the following page. Select a CVM to be terminated, and click **Remove** on the right.
2. When a prompt page that displays details of the node appears, click **OK** to remove the node.

Viewing Node Information

1. Click **ID/Name** of the cluster in the cluster list.
2. Click **Node List** to view the cluster node list information.

Logging in to a Node

Tencent Cloud CVM can be used as a node. For more information on login, please see [Log in to CVM](#).

Creating a Namespace for a Cluster

1. Select **ID/Name** of a cluster on the cluster list page.
2. Click **Namespace List**, and then click **New Namespace**.
3. Enter the information and click **Submit**.

Deleting a Namespace from a Cluster

1. Select **ID/Name** of a cluster on the cluster list page.
2. Click **Namespace List**, select a Namespace to be deleted, and click **Delete** on the right.
3. When a prompt page that displays details of the namespace appears, click **OK** to delete it.

Note:

Deleting a namespace will terminate all the resources under this namespace. All the data will be cleared and cannot be recovered after termination. Please back up your data in advance.

Cluster Autoscaler

Last updated : 2018-10-09 18:02:55

Cluster Autoscaler

1. Overview

Cluster Autoscaler (CA) is an independent program that dynamically adjusts the number of nodes in a cluster to meet your needs. When the pods in the cluster cannot be scheduled due to insufficient resources, scaling up is automatically triggered to reduce the labor costs. When conditions such as idle nodes are met, scaling down is automatically triggered to save your resource costs.

2. How to Use It

2.1 Enable CA

Create an auto scaling group before using auto scaling. You can specify the minimum/maximum number and label.

- Minimum/Maximum number - Limit the number of nodes in the scaling group.
- Label - Nodes automatically scaled up are configured with a label that is set for the scaling group to implement a flexible scheduling policy for the service.

Notes:

1. You need to configure `request` of the container under the service: Auto scaling up is triggered when some pods in the cluster cannot be scheduled due to insufficient resources, and the pod request is the basis for determining whether the resources are sufficient or not.
2. Do not directly modify the nodes belonging to the scaling group.
3. All nodes in the same scaling group should have the same configuration (model, label, etc.).
4. You can use PodDisruptionBudget (UI support will be available soon) to prevent pods from being deleted during scaling down.
5. Before specifying the minimum/maximum number of nodes in a scaling group, check whether the quota for the available zone where these nodes reside in is large enough.
6. It is not recommended to enable node auto scaling based on monitoring metrics.

7. Deleting a scaling group will terminate the CVMs in the scaling group. Please proceed with caution.

2.2 Triggering Conditions for Scaling Up/Down

Scaling Up

When pods in the cluster cannot be scheduled due to a lack of available resources, the auto scaling up policy is triggered to scale up the node to run these pods.

When the Kubernetes scheduler cannot find a place to run a pod, it sets the pod's PodCondition to false and the reason to "Unschedulable". The CA program performs scaling up by scanning whether there are unschedulable pods at regular intervals. If so, it will scale up the node to run these pods.

Scaling Down

When the proportion of both CPU and memory requests of all the pods on a node is less than 50%, the node is used as an alternative for scaling down. If the following conditions for scaling down are met, the node can be scaled down only when all the pods on it can be scheduled to other nodes.

The node containing the following types of pods cannot be scaled down:

- Pods with a strict PodDisruptionBudget. The node cannot be scaled down when PDB is not met
- Kube-system pods
- Pods that are not created by such controllers as deployment, replica set, job, and stateful set.
- Pods with local storage
- Pods that cannot be scheduled to other nodes

3. FAQ

3.1. What is the difference between Cluster Autoscaler and the node auto scaling based on monitoring metrics?

Cluster Autoscaler ensures that all the pods in the cluster can be scheduled, regardless of load. It also makes sure that the cluster contains no unnecessary nodes.

Node auto scaling based on monitoring metrics pays no attention to the pods during auto scaling. It may add a node with no pods, or delete a node with some system-critical pods, such as kube-dns. This auto scaling mechanism is not recommended by Kubernetes. Do not enable them at the same time because they conflict with each other.

3.2. What is the relationship between CA and scaling group?

A CA-enabled cluster will create a launch configuration and a scaling group that is bound with this launch configuration based on the selected node configuration. It will then perform scaling up/down in this bound scaling group. CVMs scaled up are automatically added to the cluster. Nodes that are automatically scaled up/down are billed on a postpaid basis. For more information on scaling group, please see [Auto Scaling](#).

3.3. Will CA scale down the nodes that I added manually on the TKE console?

No. CA only scales down the nodes within the scaling group. Nodes that are added on the [TKE console](#) are not added to the scaling group.

3.4. Can I add or remove CVMs on the AS console?

Yes but NOT RECOMMENDED. ed making any modifications on the [AS console](#).

3.5. Which configuration of the selected node will be inherited?

When creating a scaling group, you need to select a node in the cluster as a reference to create a [launch configuration](#). The node configuration for reference includes:

- vCPU
- Memory
- System disk size
- Data disk size
- Disk type
- Bandwidth
- Bandwidth billing method
- Whether to assign public IP
- Security group
- VPC
- Subnet

3.6. How to use multiple scaling groups?

According to the level and type of the service, you can create multiple scaling groups and set different labels for them to specify the label for the nodes scaled up in scaling groups, so as to classify the service.

3.7. Is the maximum number limited?

Each Tencent Cloud user is provided with a quota of 30 postpaid CVMs in each availability zone. Submit a ticket to apply for more CVMs for your scaling group.

For more information, please see [here](#). Besides, the maximum number is limited to 200 for auto scaling. Submit a ticket to apply for more quota.

3.8. Will the availability of cluster be affected during scaling down?

Since pods are rescheduled when a node is scaled down, the service must tolerate the rescheduling and short-term interruption before re-enabling scaling down. It is recommended that you set [PDB](#) for your service. PDB specifies the minimum number or percentage of replicas of a collection of pods that must be up at any time. With PodDisruptionBudget, an application deployer can ensure that cluster operations that voluntarily evict pods will never take down so many simultaneously as to cause data loss, an outage, or an unacceptable service degradation.

3.9. What types of pods will not trigger scaling down of a node?

- Pods with a strict PodDisruptionBudget. The node cannot be scaled down when PDB is not met
- Kube-system pods
- Pods on the node are created by such controller as non-deployment, replica set, job, statefulset, etc.
- Pods with local storage
- Pods that cannot be scheduled to other nodes

3.10. How long before a node is scaled down when the condition for scaling down is met?

10 minutes

3.11. How long before a node is scaled down when it is marked as Not Ready?

20 minutes

3.12. How often a node is scanned for scaling?

10 seconds

3.13. How long will it take to complete the scaling up of CVMs?

It generally takes less than 10 minutes. For more information, please see [Auto Scaling](#).

3.14. Why is a node with unschedulable pods not scaled up?

The reasons may include the requested resource of a pod is too large, a node selector is set, the maximum number of nodes in the scaling group is reached, account balance is sufficient (AS cannot trigger scaling up if the account balance is insufficient), quota is insufficient. See [here](#).

3.15. How to prevent Cluster Autoscaler from scaling down a specific node?

You can set the following information in the annotations of the node:

`kubectl annotate node <nodename> cluster-autoscaler.kubernetes.io/scale-down-disabled=true`

3.16. How to provide feedback on the scaling events to users?

You can query the scaling events of a scaling group and view K8S events on the AS console. Events can be found on the following three resources:

1. kube-system/cluster-autoscaler-status

2. pod

3. node

4. kube-system/cluster-autoscaler-status config map:

- **ScaledUpGroup** - CA triggers scaling up.
- **ScaleDownEmpty** - CA deletes a node with no pod running on it.
- **ScaleDown** - CA triggers scaling down.

5. node:

- **ScaleDown** - CA triggers scaling down.
- **ScaleDownFailed** - CA triggers scaling down but failed.

6. pod:

- **TriggeredScaleUp** - CA triggers scaling up because of this pod.
- **NotTriggerScaleUp** - CA cannot find a scalable scaling group to schedule this pod.
- **ScaleDown** - CA tries to evict this pod to scale down the node.

Usage of Nodes

Last updated : 2018-08-08 20:47:15

A node is a CVM registered in a cluster. Multiple nodes compose a cluster. Tencent TKE supports adding both new and existing nodes to the cluster.

Note:

Currently, existing nodes can only be added to the cluster within the same VPC. The reuse of CVM resources in the basic network and different VPCs will be available in the future.

Preconditions

You need to create a cluster first. For more information on how to create a cluster, please see [New Cluster](#).

Adding a Node

1. Log in to [TKE Console](#).
2. Click **Cluster** on the left navigation bar, and click **Add Node** on the right of the cluster list.

Cluster							
Guangzhou(1) Shanghai(0) Beijing(0) Hong Kong(0)							
Cluster Operation Documentation							
Enter the cluster name							
ID/Name	Monitoring	Cluster Status	Node Status	Number of Nodes	CPU Usage	MEM Usage	Operation
cls-696wkkse yunxiaoc_t1		Running	Normal	1 units	1.12/4	0.21/8	Add Nodes Add Existing Node Manage Namespaces Delete

3. Set **Network**, **Model** and **Configuration** for the new node.

[< Back](#) | **Create Cluster**

1

 Cluster Information

To use CCS, you need to create a cluster. A cluster consists of several nodes (CVMs) on which services are running. To learn more, please s

Cluster Name

yunyxiao_c2

Project of new resource

Select Projects

New resources such as CVMs and load balancers in cluster are automatically allocated to this project.[Instruction](#)

Kubernetes version

1.7.8

Billing Mode ⓘ

Postpaid

Region

Guangzhou

Shanghai

Beijing

Hong Kong

Cloud products in different regions CANNOT communicate via private network and CANNOT be changed once p

4. Then the new node will display on the node list.

[< Back](#) | **cls-696wkkse**

Node List

Namespace list

Cluster Information

+ Create a Node

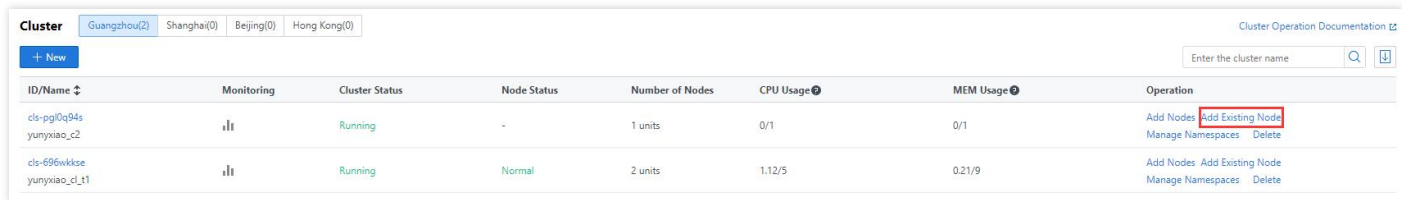
Add Existing Node

Remove

<input type="checkbox"/>	ID/Node Name ↕	Status	IP address
<input type="checkbox"/>	ins-e9acgqqu ccs_cls-696wkkse_node	Healthy	
<input type="checkbox"/>	ins-7pjxeb5s	Creating	

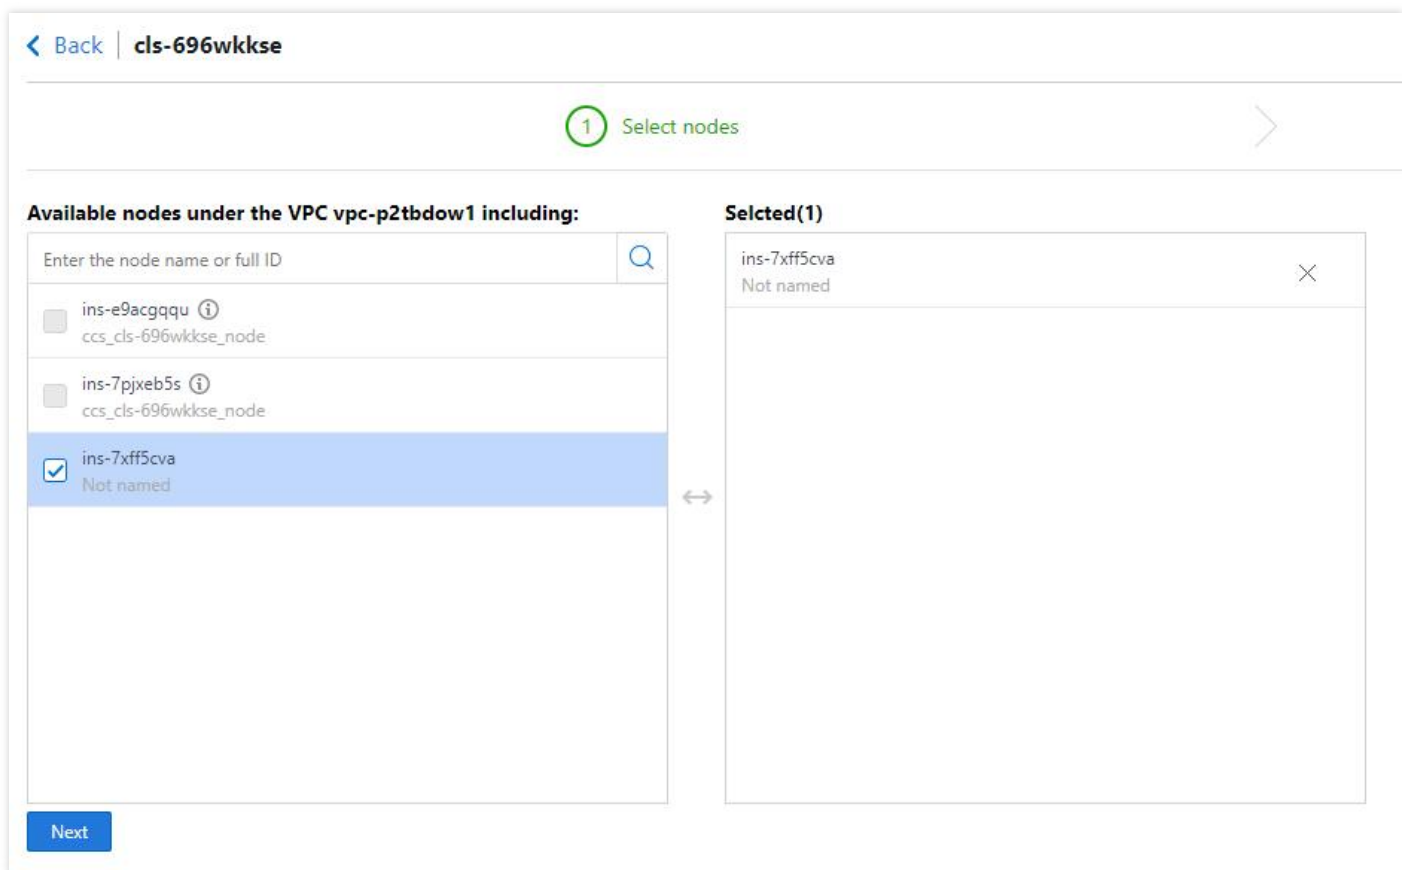
Adding an Existing Node

1. Log in to [TKE Console](#).
2. Click **Cluster** on the left navigation bar, and click **Add Existing Node** on the right of the cluster list.



ID/Name	Monitoring	Cluster Status	Node Status	Number of Nodes	CPU Usage	MEM Usage	Operation
cls-pgl0q94s yunyixiao_c2		Running	-	1 units	0/1	0/1	Add Nodes Add Existing Node Manage Namespaces Delete
cls-696wkkse yunyixiao_c11		Running	Normal	2 units	1.12/5	0.21/9	Add Nodes Add Existing Node Manage Namespaces Delete

3. Select a node on the left available node list, and then the ID of this node will display on the right column.



< Back | **cls-696wkkse**

1 Select nodes

Available nodes under the VPC vpc-p2tbdown1 including:

- ☐ ins-e9acgqqu
ccs_cls-696wkkse_node
- ☐ ins-7pjxeb5s
ccs_cls-696wkkse_node
- ☒ ins-7xff5cva
Not named

Selected(1)

ins-7xff5cva
Not named

Next

4. Enter CVM configuration. Three login methods are provided.
 - **Set Password:** Set a password according to instructions.
 - **Associate with key immediately:** Key pair is a pair of parameters generated by an algorithm. It is a more secure method to log in to CVM than a password. For more information, please see [SSH Key](#).
 - **Automatically generated password:** The automatically generated password is sent to you through the internal message.

< Back | cls-696wkkse

✓ Select nodes > 2 CVM Configuration

Nodes selected ins-7xff5cva ,

Note: you need to reinstall system for the above nodes

Note: After reinstallation, data in the system disk is cleared while data in the data disk is NOT cleared. However you need to mount the data disk manually again.

Operating system ⓘ Ubuntu 16.04 64-bit ▾

Login Methods Set Password SSH Key Pair Automatic password generation

Note: Please keep your password in mind. If you forgot your password, please reset it on CVM Console.

5. Click **Finish**, and then the new node will be on the node list.

< Back | cls-696wkkse

Node List Namespace list Cluster Information

+ Create a Node Add Existing Node Remove



<input type="checkbox"/>	ID/Node Name ↕	Status
<input type="checkbox"/>	ins-7pjxeb5s ccs_cls-696wkkse_node	Healthy
<input type="checkbox"/>	ins-e9acgqqu ccs_cls-696wkkse_node	Healthy
<input checked="" type="checkbox"/>	ins-7xff5cva Not named	Creating

Note:

- i. Currently, only the CVM under the same VPC can be added.
- ii. If you add an existing CVM to the cluster, the operating system of CVM should be reinstalled.

Viewing Node Information

1. Click **ID/Name** of the cluster (such as cls-696wkkse) in the cluster list.

Cluster			
<div>Guangzhou(2) Shanghai(0) Beijing(0) Hong Kong(0)</div>			
<div>+ New</div>			
ID/Name ↕	Monitoring	Cluster Status	Node Status
cls-pgl0q94s yunyxiao_c2		Running	Normal
cls-696wkkse yunyxiao_cl_t1		Running	Normal

2. Go to **Node List** to view the cluster node information.

< Back cls-696wkkse	
<div>Node List Namespace list Cluster Information</div>	
<div><div>+ Create a Node</div><div>Add Existing Node</div><div>Remove</div></div>	
<input type="checkbox"/> ID/Node Name ↕	Status
<input type="checkbox"/> ins-7pjxeb5s ccs_cls-696wkkse_node	Healthy
<input type="checkbox"/> ins-e9acgqqu ccs_cls-696wkkse_node	Healthy
<input checked="" type="checkbox"/> ins-7xff5cva Not named	Creating

Removing a Node

1. Click **ID/Name** of the cluster (such as cls-696wkkse) in the cluster list.

Cluster			
<div>Guangzhou(2) Shanghai(0) Beijing(0) Hong Kong(0)</div>			
<div>+ New</div>			
ID/Name ↕	Monitoring	Cluster Status	Node Status
cls-pgl0q94s yunyxiao_c2		Running	Normal
cls-696wkkse yunyxiao_cl_t1		Running	Normal

2. Go to **Node List** page, and click **Remove** on the right side.

<div>Back cls-696wkkse</div>							
<div>Node List Namespace list Cluster Information</div>							
<div>+ Create a Node Add Existing Node Remove</div>							
<div>Enter the IP or node name/ID</div>							
<input type="checkbox"/> ID/Node Name ↕	Status	IP address	CPU Usage	MEM Usage	Billing Mode	Operation	
<input type="checkbox"/> ins-7pjreb5s ccs_cls-696wkkse_node	Healthy		0 / 1	0 / 1	Postpaid Created by 2017-12-12	Remove	
<input type="checkbox"/> ins-e9acqqqu ccs_cls-696wkkse_node	Healthy		1.12 / 4	0.21 / 8	Postpaid Created by 2017-12-01	Remove	
<input type="checkbox"/> ins-7xft5cva Not named	Creating		0 / 1	0 / 2		Remove	

3. When a prompt page that displays details of the node appears, click **OK** to remove it.

Note

Confirm to remove the following nodes?

Selected 1 node. [Learn More](#)

ID	Status	description
ins-7pjxeb5s ccs_cls-696wkkse_node	Healthy	Can be Remove and Terminate

Data is not terminated as you remove the node. Prepaid nodes CANNOT be terminated.

☒ Terminate postpaid nodes as well (Data terminated CANNOT be restored. Please back up your data in advance)

OK

Cancel