


Tencent Kubernetes Engine

FAQs

Product Documentation



Copyright Notice

©2013-2019 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

FAQs

FAQ on Service

Clusters-related

Auto-scaling related

How to Choose Containerd and Docker

Cluster FAQ

How to Choose TKE Network Mode

About Services

Image Repositories

About Remote Terminals

Event FAQs

FAQs

FAQ on Service

Last updated : 2020-07-20 10:38:27

FAQ on Service Creation

Why must a service name be unique?

A service name is a unique service identifier under the current cluster. Services access each other by entering service names + access ports.

Can I use a third-party image instead of Tencent Cloud or Docker Hub image when creating services?

Yes. You can log in to your CVM and execute the command `docker login` to log in to the third-party image repository and pull the image.

What are the prerequisites for using internet services?

Ensure that CVM's in the cluster have internet bandwidth. Otherwise the creation of internet services may fail.

How do I configure memory and CPU limits?

For more information, see [Setting Limit on Service Resources](#).

What does "Privilege" mean?

If this option is enabled, applications in the container can have true root permission. We recommend that enable this option when you need to perform higher-level system operations on the applications in the container; for example, building NFS server.

FAQ on Updating Number of Service Containers

What should I take note of when I update the number of containers?

Confirm if CPU and memory resources are sufficient. Otherwise container creation will fail.

Can I set the number of containers to 0?

Yes. You can set the number to 0. In this case, the service configuration is retained but the resources will be freed up.

FAQ on Service Configuration Updating

Is rolling update supported?

Both rolling update and quick update are supported.

FAQ on Service Deletion

When I delete a service, will the load balancer created by the service also be terminated?

All containers and internet load balancers under a service are terminated upon service deletion. Back up the data in advance.

FAQ on Service Operation

How do I set the container system time to Beijing time?

Containers use UTC by default. If the problem of an 8-hour difference between the container system time and Beijing time always occurs when you use a container, create a time zone file in Dockerfile.

```
RUN echo "Asia/shanghai" > /etc/timezone;
```

What can I do when some Docker Hub images, such as ubuntu, php and busybox, encounter exceptions in TKE?

If the startup command is not set or the default startup command is `bash`, the container will exit after the startup procedure is completed. To keep the container running, the process in the container with PID set to 1 must be the permanent process. Otherwise, the container exits when this process ends. For some images like CentOS, you can create services by using `/bin/bash` as the running command and `-c sleep 800000` as the running parameter. `-c` and `sleep 800000` must be entered in different rows in the console.

Currently, the images that cannot be started using default parameters include: clearlinux, ros, mageia, amazonlinux, ubuntu, clojure, crux, gcc, photon, java, debian, oraclelinux, mono, bash,

buildpack-deps, golang, sourcemage, Swift, openjdk, centos, busybox, docker, alpine, ibmjava, php, and python.

Clusters-related

Auto-scaling related

Last updated : 2020-10-10 11:29:25

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

Cluster Autoscaler (CA) ensures that all pods in the cluster can be scheduled regardless of the actual load, while node auto scaling based on monitoring metrics do not take pods into consideration during auto scaling. Therefore, nodes without pods might be added, or nodes with system-critical pods such as kube-dns might be deleted during auto scaling. Kubernetes discourages the latter auto scaling mechanism. In conclusion, these two modes conflict and should not be enabled at the same time.

How does CA work with auto scaling groups?

A CA-enabled cluster will, according to the configuration of the selected node, create a launch configuration and bind an auto scaling group to it. The cluster will then perform scale-in/out in this bound auto scaling group. CVM instances scaled out are automatically added to the cluster. Nodes that are automatically scaled in/out are billed on a pay-as-you-go basis. For more information about auto scaling group, see [Auto Scaling \(AS\)](#).

Can a node manually added in the TKE Console be scaled in by CA?

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

Can a CVM instance be added or removed in the AS Console?

No. We do not recommend making any modifications on the [AS Console](#).

What configurations of the selected node will be inherited during scaling?

When creating an auto 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

How do I use multiple auto scaling groups?

Based on the level and type of the service, you can create multiple auto scaling groups, set different labels for them, and specify the label for the nodes scaled out in the auto scaling groups to classify the service.

What is the maximum quota for scaling?

Each Tencent Cloud user is provided with a quota of 30 pay-as-you go CVM instances in each availability zone. You can [submit a ticket](#) to apply for more instances for your auto scaling group. For more information about the quotas, see [CVM Instance Quantity and Quota](#) for your current availability zone. In addition, there is a maximum limit of 200 instances for Auto Scaling. You can [submit a ticket](#) to apply for a higher quota.

Is scale-in safe for the cluster?

Since pods will be rescheduled when a node is scaled in, scale-in can be performed only if the service can tolerate rescheduling and short-term interruption. We recommend using [PDB](#). PDB can specify the minimum number/percentage of replicas for a pod set that remains available at all times. With PodDisruptionBudget, application deployers can make sure that the cluster operations that actively remove pods will not terminate too many pods at a time, which helps prevent data loss, service interruption, or unacceptable service degradation.

What types of pods in a node will not be scaled in?

- If you set strict PodDisruptionBudget for a pod, and the PDB is not met, it will not be scaled in.
- Pods under Kube-system.
- Pods on a node that are not created by controllers such as Deployment, ReplicaSet, Job, or StatefulSet.
- Pods with local storage.
- Pods that cannot be scheduled to another node.

How long does it take to trigger scale-in after the condition is met?

10 minutes.

How long does it take to trigger scale-in when the node is marked as Not Ready?

20 minutes.

How often is a node scanned for scaling?

Every 10 seconds.

How long does it take to scale out a CVM instance?

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

Why is a node with an unschedulable pod not scaled out?

Please check the following:

- Whether the requested resource of the pod is too large.
- Whether a node selector is set.
- Whether the maximum number of nodes in the auto scaling group has been reached.
- Whether the account balance is sufficient (scale-out cannot be triggered if the account balance is insufficient) or the quota is insufficient. For more information, see [other reasons](#).

How can I prevent CA from scaling in 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
```

Where can I find details on the scaling events?

You can query the scaling events of an auto scaling group and view K8s events in the AS Console. Events can be found in the following three resources:

- kube-system/cluster-autoscaler-status config map
 - **ScaledUpGroup** - CA triggers scale-out.
 - **ScaleDownEmpty** - CA deletes a node with no running pod.
 - **ScaleDown** - CA triggers scale-in.
- node
 - **ScaleDown** - CA triggers scale-in.
 - **ScaleDownFailed** - CA failed to trigger scale-in.
- pod
 - **TriggeredScaleUp** - CA triggers scale-out because of this pod.
 - **NotTriggerScaleUp** - CA cannot find an available auto scaling group to schedule this pod.
 - **ScaleDown** - CA tries to drain this pod to scale in the node.

How to Choose Containerd and Docker

Last updated : 2020-02-24 16:37:46

How do you select a container runtime?

As one of the most important components of Kubernetes (K8s), a container runtime manages the lifecycle of images and containers. Kubelet interacts with the container runtime through the `Container Runtime Interface (CRI)` to manage images and containers.

TKE supports containerd and Docker as a container runtime:

- (Recommended) Containerd has a shorter calling chain and fewer components and features higher stability and lower node resource consumption.
- Docker should be used as the runtime component if you need to use the following:
 - Docker in docker;
 - Commands such as `docker build/push/save/load`;
 - Docker API;
 - Docker compose or docker swarm.

What are the commands commonly used in containerd and Docker?

containerd does not support docker API or docker CLI. However, you can get these features with cri-tool commands.

Image-related Features	Docker	containerd
Display the local image list	<code>docker images</code>	<code>crictl images</code>
Download an image	<code>docker pull</code>	<code>crictl pull</code>
Upload an image	<code>docker push</code>	None
Delete a local image	<code>docker rmi</code>	<code>crictl rmi</code>
View image details	<code>docker inspect IMAGE-ID</code>	<code>crictl inspecti IMAGE-ID</code>

Container-related Features	Docker	containerd
Display the container list	<code>docker ps</code>	<code>crictl ps</code>
Create a container	<code>docker create</code>	<code>crictl create</code>
Start a container	<code>docker start</code>	<code>crictl start</code>

Container-related Features	Docker	containerd
Stop a container	docker stop	crictl stop
Delete a container	docker rm	crictl rm
View container details	docker inspect	crictl inspect
attach	docker attach	crictl attach
exec	docker exec	crictl exec
logs	docker logs	crictl logs
stats	docker stats	crictl stats

Pod-related Features	Docker	containerd
Display the Pod list	None	crictl pods
View Pod details	None	crictl inspectp
Run a Pod	None	crictl runp
Stop a Pod	None	crictl stopp

What are the differences between the calling chains of containerd and Docker?

- When Docker is used as the K8s container runtime, the calling chain is as follows:
`kubelet --> docker shim (in the kubelet process) --> dockerd --> containerd`
- When containerd is used as the K8s container runtime, the calling chain is as follows:
`kubelet --> cri plugin (in the containerd process) --> containerd`

Although Docker offers more features such as swarm cluster, docker build, and docker API, it may also introduce some bugs and requires one more calling step than containerd.

Stream Service

Commands such as `kubectl exec` and `kubectl logs` require the establishment of a stream forwarding channel between the apiserver and the container runtime.

How are stream services used and configured in containerd?

The docker API itself provides a stream service, and the docker-shim inside the Kubelet forwards streams through the docker API.

The stream service of containerd needs to be configured separately:

```
[plugins.cri]
stream_server_address = "127.0.0.1"
stream_server_port = "0"
enable_tls_streaming = false
```

What are the differences between versions before and after k8s 1.11?

The stream service of containerd has different configurations for different versions of K8s.

- Before K8s v1.11:
Kubelet performs redirection but not stream proxying. That is, Kubelet sends the stream server address opened by containerd to the apiserver which then directly accesses the stream service of containerd. You need to authenticate the stream service forwarder for security purposes.
- K8s v1.11 and later:
K8s v1.11 introduced [kubelet stream proxy](#), so that the stream service of containerd only needs to listen to the local address.

Other Differences

Container Log and Related Parameters

Item	Docker	containerd
Storage path	Docker saves container logs to a directory such as <code>/var/lib/docker/containers/\$CONTAINERID</code> . Kubelet will create a soft link under <code>/var/log/pods</code> and <code>/var/log/containers</code> pointing to the container log files in the <code>/var/lib/docker/containers/\$CONTAINERID</code> directory.	Kubelet saves container logs to the <code>/var/log/pods/\$CONTAINER_NAME</code> directory, and creates a soft link under <code>/var/log/containers</code> pointing to the log files.
Configuration parameters	Specify in the Docker configuration files: <pre>"log-driver": "json-file", "log-opts": {"max-size": "100m", "max-file": "5"}</pre>	<ul style="list-style-type: none"> • Method 1: Specify in the Kubelet parameters: <pre>--container-log-max-files=5 --container-log-max-size="100Mi"</pre>

		<ul style="list-style-type: none"> Method 2: Specify in KubeletConfiguration: <pre> "containerLogMaxSize": "100Mi", "containerLogMaxFiles": 5, </pre>
Save container logs to the data disk	Mount the data disk to "data-root" (/var/lib/docker by default).	<p>Create a soft link /var/log/pods to point to a directory under the data disk mounting point.</p> <p>Selecting "Store containers and images in the data disk" in TKE will automatically create the soft link /var/log/pods .</p>

CNI Network

Item	Docker	containerd
Component responsible for calling CNI	docker-shim inside Kubelet	containerd's built-in cri-plugin (in containerd v1.1 or later)
How to configure CNI	Kubelet parameters <code>--cni-bin-dir</code> and <code>--cni-conf-dir</code>	containerd configuration file (toml): <pre> [plugins.cri.cni] bin_dir = "/opt/cni/bin" conf_dir = "/etc/cni/net.d" </pre>

Cluster FAQ

Last updated : 2020-04-29 17:39:35

Cluster Creation

Can I choose not to set a public IP address for the CVM when creating a cluster?

Yes. A CVM without a public IP address can only pull images from "My Images" in Image Repository but not Docker Hub images or third-party images.

A CVM without a public IP address but with Internet bandwidth can access the Internet by binding an EIP.

Why do I need to choose a network when creating a cluster?

The selected network and subnet are where the CVM resides. You can add different CVMs to subnets in different AZs for cross-AZ disaster recovery.

What CVM models are supported when creating a cluster?

All pay-as-you-go models that use cloud disks as system disks are supported.

What operating systems are supported for TKE hosts?

Currently, Ubuntu 16.04 and CentOS 7.2 are supported. If you need other operating systems, [submit a ticket](#) to contact us.

Adding CVMs

What are the limitations on adding CVMs to a cluster?

You can only choose the CVMs in the same region as the cluster. But you may choose a different AZ to implement cross-AZ deployment of the cluster.

Is there a limit on the number of CVMs?

Yes. The number of pay-as-you-go CVMs cannot exceed the usage quota of the current account. For details, see [CVM Overview](#).

CVM Termination

After a CVM is terminated, what will happen to the containers deployed in it?

When a CVM instance is terminated, the resources it contains, such as containers, will also be terminated. If the number of containers for a service drops below the expected number of running containers, the cluster will launch more containers in other CVM instances until the desired number is met.

How to Choose TKE Network Mode

Last updated : 2020-09-10 10:22:47

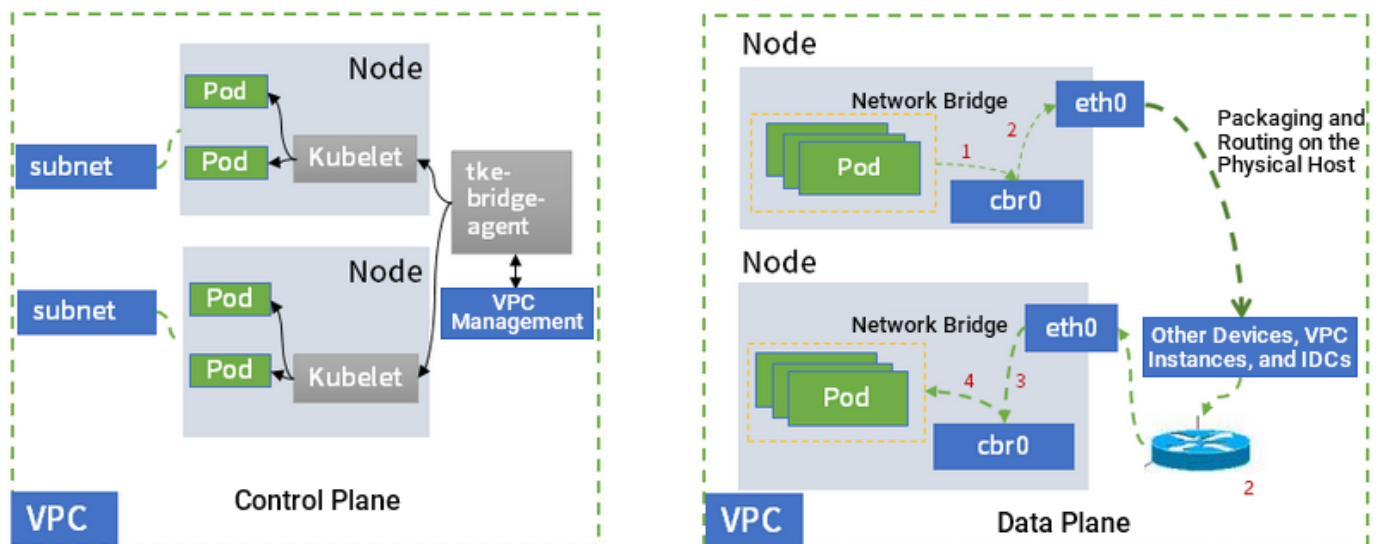
Tencent Kubernetes Engine (TKE) provides different network modes for different scenarios. This document gives a detailed description of the GlobalRouter and VPC-CNI network modes, as well as comparisons on use cases, advantages, and use limits. You can select a network mode based on your business needs.

GlobalRouter Mode

GlobalRouter is a global routing capability provided by TKE based on the underlying VPC instance. It implements a routing policy for mutual access between the container network and the VPC instance. This network mode has the following characteristics:

- Container routing is performed directly through the VPC instance.
- Containers and nodes are located on the same network plane.
- Container IP ranges are dynamically assigned without occupying other IP ranges in the VPC instance.

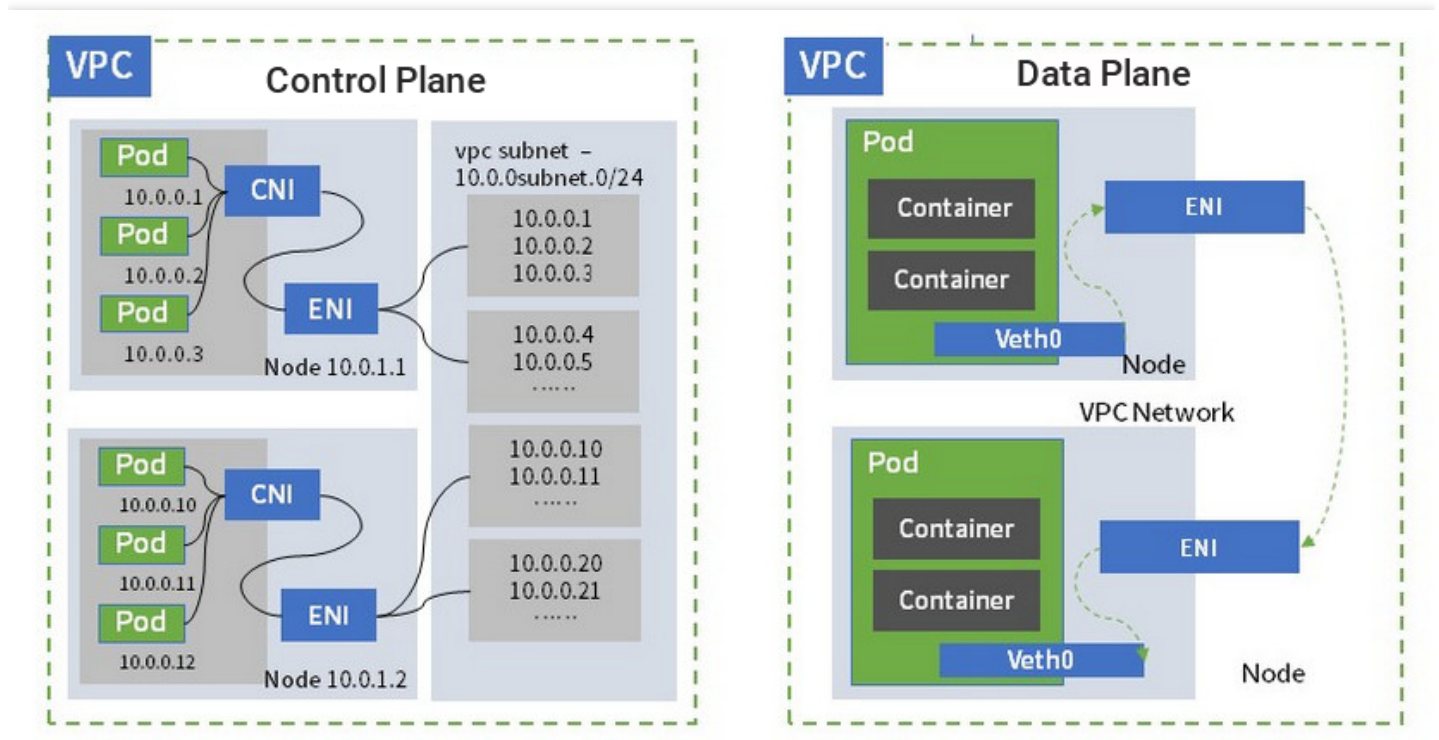
The GlobalRouter mode is suitable for general use cases and can be seamlessly used with standard Kubernetes features. The following diagram illustrates how it works.



VPC-CNI Mode

VPC-CNI is a container network capability provided by TKE based on CNIs and VPC ENIs. It is suitable for scenarios with high latency requirements. In this network mode, containers and nodes are located on the same network plane, and container IP addresses are ENI IP addresses assigned by the IPAMD component.

The following diagram illustrates how the VPC-CNI mode works.



Enabling support for static pod IP addresses

⚠ Note :

- By default, the VPC-CNI mode **does not support static pod IP addresses**. You can set this capability only when [creating a cluster](#) and cannot modify it after the cluster is created.
- When support for static pod IP addresses is enabled, you can only choose an empty subnet to set up the cluster network.
- Pods with static IP addresses cannot be migrated across subnets.

In the step where you configure "Cluster Information", select **VPC-CNI** for "Container Network Add-on" and select "Enable Support" for "Static Pod IP", as shown below.

Container Network Add-on

Global Router

VPC-CNI

[How to select](#)

VPC-CNI mode is a container network plug-in implemented based on ENI. The container network and CVM network are in the same VPC.

Static Pod IP

 Enable Support

By default, VPC-CNI mode does not support static pod IP. You need to enable it manually. If static pod IP is enabled, the subnet must be used by the container exclusively. For more information, [Learn more](#)

For information on how to use the VPC-CNI mode with static pod IP addresses, please see [Managing StatefulSets with Static IP Addresses](#).

Choosing a Network Mode

This section compares the GlobalRouter and VPC-CNI network modes in terms of the use cases, advantages, and use limits. You can choose the network mode that best fits your needs.

Dimension	GlobalRouter	VPC-CNI
Use cases	<ul style="list-style-type: none"> General container businesses Offline computing 	<ul style="list-style-type: none"> Scenarios with high network latency requirements Scenarios where static container IP addresses are required after a traditional architecture is migrated to a container platform
Advantages	<ul style="list-style-type: none"> Container routing is performed directly through the VPC instance, and containers and nodes are located on the same network plane. Container IP ranges are dynamically assigned without occupying other IP ranges in the VPC instance. Therefore, available IP addresses are abundant. 	<ul style="list-style-type: none"> The container network of the ENI is a VPC subnet and can be managed as a VPC product. Static IP addresses and LB-to-Pod direct access are supported. The network performance is better than that of GlobalRouter.
Use limits	<ul style="list-style-type: none"> Additional configuration is required for interconnection scenarios such as Direct Connect, Peering Connection, and Cloud Connect Network. Static pod IP addresses are not supported. 	<ul style="list-style-type: none"> The container network and node network belong to the same VPC instance. Therefore, IP addresses are limited. The number of pods on a node are limited by the ENI and the available IP addresses of the ENI. The static IP address mode does not allow subnets to be shared by containers and other services.

		<ul style="list-style-type: none">• The static IP address mode does not allow pods to be scheduled across availability zones.• Network planning must be done in advance and is difficult to adjust later.
Additional capabilities	Standard Kubernetes features	<ul style="list-style-type: none">• TKE supports static pod IP addresses.• The container network is managed in the VPC console.• LBs directly forward requests to pods, and the pods can obtain source IP addresses.

About Services

Last updated : 2020-08-27 10:58:13

FAQs on Service Creation

Why must a service name be unique?

A service name uniquely identifies a service in a cluster. Services can use service name + access port to access each other.

Can I use a third-party image instead of a Tencent Cloud or Docker Hub image to create services?

Yes. You can log in to your CVM and run the `docker login` command to log in to the third-party image repository and pull the image.

What are the prerequisites for using Internet services?

CVM instances in the cluster must be configured with Internet bandwidth. Otherwise, Internet services will fail to be created.

How do I configure memory and CPU limits?

For more information, please see [Setting the Resource Limit of Workload](#).

What does the "Privileged" option mean when I create a service?

If this option is enabled, applications in a container will have true root permissions. We recommend that you enable this option when you need to perform high-level system operations on the applications in the container, for example, when building an NFS server.

Can I specify a security group when a CLB is created?

Yes. You can use either of the following solutions to specify a security group when a service uses a CLB:

- Use existing CLBs. First, create a CLB and configure a security group and then mount them to a service. For more information, see "Using Existing CLBs".
- Use `TkeServiceConfig` in a service to configure a security group. When a CLB is created, the CLB will use the corresponding security group based on the configuration. To use this feature, [submit a ticket](#).

Do not access a service in a cluster using the CLB IP address to prevent an access failure. Usually a Layer-4 CLB will bind multiple nodes as real servers (RSs). Ensure that the client and RS are not on the same CVM. Otherwise, packets may fail to be sent out due to the loopback. When a pod accesses a CLB, the pod IP address is the source IP address. When the packet is routed to the private network, the CLB will not convert the source IP address to a node IP address using SNAT. As a result, the CLB cannot identify the source node of the packet, the CLB's loopback avoidance policy will not take effect, and the packet may be forwarded to any RS. When the packet is forwarded to the node where the client is located, the CLB cannot receive the response, leading to an access failure.

FAQs on Updating the Number of Service Containers

What should I take note of when I update the number of containers?

Confirm whether the CPU and memory resources are sufficient. Otherwise, container creation will fail.

Can I set the number of containers to 0?

Yes. When you set the number of containers to 0, the service configuration is retained and the resources will be freed up.

FAQs on Updating Service Configuration

Is rolling update supported?

Yes. Both rolling update and quick update are supported.

Can I change a public CLB to a private CLB?

Yes. Currently, you can change the service access mode from Via Internet to Via VPC or from Via VPC to Via Internet or switch the VPC subnet. For more information, see "Service lifecycle management".

- If the service manages the CLB lifecycle, the CLB and its public IP address will be released.
- It takes time to switch from the Internet to a VPC. This means that it takes some time from when the public CLB goes offline to when the private CLB starts to provide service. We recommend that you first configure a private service in the cluster and test it. Then, delete the original public service after traffic is switched over to the new service.

FAQs on Service Deletion

Will the CLB created by a service be destroyed after the service is deleted?

If a CLB is automatically created when you create a service, the CLB will be deleted when the service is deleted. If you select an existing CLB when you create a service, the CLB will not be deleted.

Will service data be affected when a service is deleted?

No. When a service is deleted, the service container will not be deleted and data will not be affected. You do not need to back up data before deleting a service.

FAQs on Service Operation

How do I set the container system time to Beijing time?

The default system time of containers is Universal Time Coordinated (UTC). If the container system time is different from the local time zone when a container is used, you can create a time zone file in Dockerfile to solve this problem. For more information, see [Solve the inconsistent time zone problem in the container](#).

What can I do when some Docker Hub images, such as ubuntu, php, and busybox, encounter exceptions in TKE?

If the startup command is not set or the default startup command is `bash`, the container will exit after the startup procedure is completed. To keep the container running, the process in the container with PID set to 1 must be the permanent process. Otherwise, the container exits when this process ends. For some images like CentOS, you can create services by using `/bin/bash` as the running command and `-c sleep 800000` as the running parameter. `-c` and `sleep 800000` must be entered in different rows in the console.

Currently, the images that cannot be started using default parameters include clearlinux, ros, mageia, amazonlinux, ubuntu, clojure, crux, gcc, photon, java, debian, oraclelinux, mono, bash, buildpack-deps, golang, sourcemage, Swift, openjdk, centos, busybox, docker, alpine, ibmjava, php, and python.

Image Repositories

Last updated : 2020-04-21 12:36:34

Image Registry Activation FAQs

What is a namespace used for?

A namespace is an address prefix that identifies a user's private images.

What is an Image Registry account?

By default, this account is your Tencent Cloud account.

What should I do if I forgot the password set during activation?

You can reset the password in the console.

Cluster Creation FAQs

Is there a quota limit on the number of created images?

Yes. The default quota is 500 image repositories per region and 100 image tags per image repository. If you need a higher quota, [submit a ticket](#).

Can I share my created images with other users?

No, this feature is currently not supported.

How do I use the created images?

First, upload available image tags and then create services by using specific image tags.

Image Deletion FAQs

How do I delete a certain tag of an image?

You can specify and delete a certain tag directly from the console.

When I delete an image from the image list, will all tags of the image be deleted too?

Yes. Be sure to back up your data in advance.

Image Building FAQs

How do I specify the Dockerfile path and building directory when building an image by using source code?

- If you do not enter a path, the system uses the following default values:
 - Default Dockerfile path: Dockerfile under the root directory of the repository (`Dockerfile`).
 - Default building directory: root directory of the repository (`./`).
- To specify a path, **enter a relative path with the project path as the root path**, as shown in the following figure:

The screenshot shows the 'Build Config' interface with the following fields and options:

- Image Address:** `ccr.ccs.tencentyun.com/test-yunx/helloworld`
- Code source:** Buttons for **Github** and **Gitlab**.
- Image Tag Naming Rules:** A text input field with the placeholder 'Please enter the naming rule' and three checkboxes: Branch/label, Update Time, and Commit No. Below it, the text reads 'Custom prefix, supports variables in the format of \$(Foo)'.
- Overwrite the image tag:** A text input field with the placeholder 'Please enter the tag'. Below it, the text reads 'The generated image also contains the tag'.
- Dockerfile path:** A text input field containing `qcloud/Dockerfile`. Below it, the text reads 'Path of the Dockerfile in the code source'.
- Building Directory:** A text input field containing `qcloud`. Below it, the text reads 'The working directory for building, should be a relative path'.
- Building Parameters:** A link labeled 'Add a variable'.
- Complete:** A blue button at the bottom.

The 'Dockerfile path' and 'Building Directory' fields are highlighted with a red border in the original image.

How does the source code building feature use the Dockerfile path and building directory?

Clone the user-specified repository, then switch to the corresponding branch or tag, and lastly run the command `docker build -f $DOCKERFILE_PATH $WORKDIR` in the root directory of the code repository to build a container image.

How do I specify the source path in Dockerfile?

For `COPY` , `ADD` , and other commands related to the source path, specify the source path as a path relative to the building directory.

About Remote Terminals

Last updated : 2020-04-26 16:18:05

What should I do if there is no bash in the container?

If there is no bash in the container, enter the command you want to run in the command line. The result of the command is then displayed on the screen. You can consider the command line as a simplified bash without certain features such as autocomplete. We recommend that you run the bash installation command before proceeding.

Why is software installation using `apt-get` so slow?

When software installation using `apt-get` takes up too much time, this may be due to slow server access to software sources outside China. We provide acceleration solutions for containers with different operating systems. You can select a solution when required.

Precautions

Before selecting the acceleration solution, check the operating system of the container as instructed below:

- Ubuntu: run `cat /etc/lsb-release` to check if the `/etc/lsb-release` file exists.
- CentOS: run `cat /etc/redhat-release` to check if the `/etc/redhat-release` file exists.
- Debian: run `cat /etc/debian_version` to check if the `/etc/debian_version` file exists.

">

Solutions

Ubuntu 16.04

For containers running Ubuntu 16.04, run the following commands on your terminal to set the apt source to a Tencent Cloud source server:

```
cat << EOF > /etc/apt/sources.list
deb http://mirrors.tencentyun.com/ubuntu/ xenial main restricted universe multiverse
deb http://mirrors.tencentyun.com/ubuntu/ xenial-security main restricted universe multiverse
deb http://mirrors.tencentyun.com/ubuntu/ xenial-updates main restricted universe multiverse
deb-src http://mirrors.tencentyun.com/ubuntu/ xenial main restricted universe multiverse
deb-src http://mirrors.tencentyun.com/ubuntu/ xenial-security main restricted universe multiverse
deb-src http://mirrors.tencentyun.com/ubuntu/ xenial-updates main restricted universe multiverse
EOF
```

CentOS 7

For containers running CentOS 7, directly change the source address as instructed below to speed up the installation:

1. Copy and run the following code in the container:

```
cat << EOF > /etc/yum.repos.d/CentOS-Base.repo
[os]
name=Qcloud centos os - ¥$basearch
baseurl=http://mirrors.tencentyun.com/centos1/¥$releasever/os/¥$basearch/
enabled=1
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7
[updates]
name=Qcloud centos updates - ¥$basearch
baseurl=http://mirrors.tencentyun.com/centos1/¥$releasever/updates/¥$basearch/
enabled=1
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7
#[centosplus]
#name=Qcloud centosplus - ¥$basearch
#baseurl=http://mirrors.tencentyun.com/centos1/¥$releasever/centosplus/¥$basearch/
#enabled=1
#gpgcheck=1
#gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7
#[cloud]
#name=Qcloud centos contrib - ¥$basearch
#baseurl=http://mirrors.tencentyun.com/centos1/¥$releasever/cloud/$basearch/openstack-kilo/
#enabled=1
#gpgcheck=1
#gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7
#[cr]
#name=Qcloud centos cr - ¥$basearch
#baseurl=http://mirrors.tencentyun.com/centos1/¥$releasever/cr/¥$basearch/
#enabled=1
#gpgcheck=1
#gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7
[extras]
name=Qcloud centos extras - ¥$basearch
baseurl=http://mirrors.tencentyun.com/centos1/¥$releasever/extras/¥$basearch/
enabled=1
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7
#[fasttrack]
#name=Qcloud centos fasttrack - ¥basearch
#baseurl=http://mirrors.tencentyun.com/centos1/¥$releasever/fasttrack/¥$basearch/
#enabled=1
#gpgcheck=1
```

```
#gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7
ENDOF
```

2. Run the following command to clear and rebuild the YUM cache.

```
yum clean all && yum clean metadata && yum clean dbcache && yum makecache
```

Debian

For containers running Debian, run the following commands on your terminal to set the apt source to a Tencent Cloud source server:

```
cat << ENDOF > /etc/apt/sources.list
deb http://mirrors.tencentyun.com/debian stretch main contrib non-free
deb http://mirrors.tencentyun.com/debian stretch-updates main contrib non-free
deb http://mirrors.tencentyun.com/debian-security stretch/updates main

deb-src http://mirrors.tencentyun.com/debian stretch main contrib non-free
deb-src http://mirrors.tencentyun.com/debian stretch-updates main contrib non-free
deb-src http://mirrors.tencentyun.com/debian-security stretch/updates main
ENDOF
```

Conclusion

Changing the source address in the container directly is an interim solution. When the container is rescheduled, your change becomes invalid. Therefore, we recommend that you solve this problem with the following method when creating the image:

Add the source address provided in [Solutions](#) for corresponding operating systems to the `RUN` field of the Dockerfile file for creating the container image. For example, append the following code to the Dockerfile file for an image running the Ubuntu operating system:

```
RUN cat << ENDOF > /etc/apt/sources.list
deb http://mirrors.tencentyun.com/ubuntu/ xenial main restricted universe multiverse
deb http://mirrors.tencentyun.com/ubuntu/ xenial-security main restricted universe multiverse
deb http://mirrors.tencentyun.com/ubuntu/ xenial-updates main restricted universe multiverse
#deb http://mirrors.tencentyun.com/ubuntu/ xenial-proposed main restricted universe multiverse
#deb http://mirrors.tencentyun.com/ubuntu/ xenial-backports main restricted universe multiverse
deb-src http://mirrors.tencentyun.com/ubuntu/ xenial main restricted universe multiverse
deb-src http://mirrors.tencentyun.com/ubuntu/ xenial-security main restricted universe multiverse
deb-src http://mirrors.tencentyun.com/ubuntu/ xenial-updates main restricted universe multiverse
#deb-src http://mirrors.tencentyun.com/ubuntu/ xenial-proposed main restricted universe multivers
e
#deb-src http://mirrors.tencentyun.com/ubuntu/ xenial-backports main restricted universe multiver
```

```
se
ENDOF
```

What should I do if I cannot find tools such as vim and netstat after logging in to a container?

Download the required tools by running commands such as `apt-get install vim` and `apt-get install net-tools` (run `yum install vim` on CentOS.)

What should I do if a tool cannot be found when I run the `apt-get install` command?

Install the software program as follows:

1. Run the following command to update the software list.

```
apt-get update
```

2. Run the following command to install the software program (run `yum updateinfo` on CentOS.)

```
apt-get install
```

How do I use an in-house tool in a container?

Go to the remote terminal page, click "File Assistant" in the lower-right corner, and upload your tool.

How do I upload a live file such as dump or log for local analysis?

Go to the remote terminal page, click "File Assistant" in the lower-right corner, and upload your file.

Why can't I upload a file to a container or download a file to the local system?

This issue occurs if the tar program is not included in your container image. To correct it, you can install the tar program by running `apt-get install vim` or `apt-get install net-tools` (run `yum install vim` on CentOS) and try again.

Why can't I find the tools that I installed previously?

When Kubernetes re-schedules your container, it pulls an image to create a new container, and tools that are not in the image will not be installed in the new container. Therefore, we recommend that you install some common troubleshooting tools when creating the image.

How do I copy text in the console?

You can copy text by highlighting and copying it.

How do I paste the copied text?

Press Shift+Insert to paste the copied text.

Why is the connection lost?

This happens when you perform some operations on the container or CVM on another page that modify the container status, or when the current page remains idle for more than 3 minutes. In both cases, the server disconnects the connection.

What should I do if the "TERM environment variable not set" error occurs when I run commands such as `top` ?

Run `export TERM linux` .

Why does the bash prompt display only "<" and part of the path when I enter a directory with a long absolute path?

The bash prompt is set to display "@ " by default. If the current path is too long, bash displays only "<" and the last part of the path by default.

Event FAQs

Last updated : 2019-09-18 18:07:09

Back-off restarting failed docker container

Description: An exceptional Docker container is being restarted.

Solution: Check if the Docker process running in the image has exited due to an exception. If there is no process running in the image, you can add an execution script in the service creation page.

fit failure on node: Insufficient cpu

Description: The cluster CPU is insufficient.

Solution: The node cannot provide enough computing cores. Modify the CPU limit on the service page or scale out the cluster.

no nodes available to schedule pods

Description: Cluster resources are insufficient.

Solution: The reason for this error is that the number of nodes is not sufficient to carry the Pods. From the service page, modify the number of service Pods, the number of Pods or the CPU limit.

pod failed to fit in any node

Description: There is no proper node for the Pod to use.

Solution: This error is caused when the service configures an inappropriate resource limit, which leads to a situation where there are no proper nodes to carry the Pods. Modify the number of service Pods or the CPU limit from the service page.

Liveness probe failed

Description: The container health check failed

Solution: Check if the container processes in the image are normal, and whether the health check port is correctly configured.

Error syncing pod, skipping

Error syncing pod, skipping failed to "StartContainer" for with CrashLoopBackOff: "Back-off 5m0s restarting failed container

Description: The container process crashed or exited.

Solution: Check whether there is a foreground process that is persistently running. If so, check whether it is exceptional. For more information, see [Building a Docker Image Guide](#).

Contact customer service if the solutions offered above failed to resolve your issues.