

Tencent Kubernetes Engine

Notice

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

Notice

Instructions on Cluster Resource Quota Adjustment

Instructions on Stopping Delivering the Kubeconfig File to Nodes

Update Notes of TKE Kubernetes Major Versions

TKE Kubernetes Revision Version History

Notice

Instructions on Cluster Resource Quota Adjustment

Last updated : 2021-01-29 18:02:50

From January 13, 2021, TKE will automatically apply a set of resource quotas to the namespace on the cluster with less than 5 nodes ($0 < \text{nodeNum} \leq 5$), or with more than 5 and less than 20 nodes ($5 < \text{nodeNum} < 20$). You cannot remove these quotas, which are used to protect the cluster control plane from being unstable due to potential bugs in the applications deployed to the cluster.

You can run the following command to check the quota:

```
kubectl get resourcequota tke-default-quota -o yaml
```

If you need to view the `tke-default-quota` object of a specified namespace, you can add the `--namespace` option to specify the namespace.

The specific quota limits are as follows:

Cluster Scale	Quota Limits
$0 < \text{nodeNum} \leq 5$	Total Pods: 4000, configMap: 3000, CustomResourceDefinition(CRD): 4000
$5 < \text{nodeNum} < 20$	Total Pods: 8000, configMap: 6000, CustomResourceDefinition(CRD): 8000
$20 \leq \text{nodeNum}$	No limit

You can [submit a ticket](#) to apply to increase the quota.

Instructions on Stopping Delivering the Kubeconfig File to Nodes

Last updated : 2021-04-19 17:14:04

Note :

We plan to carry out an operation from 23:00 September 21 (Monday) to 06:00 September 22 (Tuesday) UTC+8 to stop delivering the Kubeconfig file.

Background

Currently, TKE stores the Kubeconfig file with the admin token in nodes by default. By using this Kubeconfig file, users can easily operate on Kubernetes clusters. However, if users fail to conduct node login permission management carefully, clusters may face security risks. Therefore, we decided to stop delivering the Kubeconfig file.

Existing clusters may use the Kubeconfig file to perform cluster initialization operations in user-defined scripts. To solve this issue, we will provide a client certificate for node initialization with the same permissions as the Kubeconfig file, but with a validity period of only 12 hours. After the certificate expires, the Kubeconfig file will be invalidated. If you still need the file after the expiration, refer to [Issues and Solutions](#).

Note :

If you still require default long-term admin permissions instead of a Kubeconfig file whose validity period is only 12 hours for some special scenarios, or if you encounter any other issues, [submit a ticket](#) to contact us.

Issues and Solutions

Issues

If you prefer to use the following command to log in to a TKE cluster node for kubectl operations, you will be prompted with the following error message:

```
$ kubectl get node
The connection to the server localhost:8080 was refused - did you specify the right host or port?
```

```
$ kubectl get node  
error: You must be logged in to the server (Unauthorized)
```

Solutions

1. Log in to the [TKE console](#).
2. Obtain the credential Kubeconfig file of the current account. For more information, see [Obtaining credentials](#).
3. After obtaining the Kubeconfig file, you can enable private network access or directly use the service IP address of Kubernetes.
 - Enabling private network access: on the cluster details page, choose **Basic Information** in the left sidebar, enable **Private Network Access** in the **Cluster API Server information** section, and operate according to the prompt.
 - Using the service IP address of Kubernetes: on the cluster details page, choose **Services and Routes** > **Service** in the left sidebar to obtain the service IP address of Kubernetes in the default namespace. Replace the clusters.cluster.server field in the Kubeconfig file with https://< IP >:443.
4. Copy the content of the Kubeconfig file to `$HOME/.kube/config` on the new node.
5. Access a Kubeconfig cluster and use `kubectl get nodes` to test connectivity.

Handling Special Scenarios

Special scenarios

A workload has mounted the `/root/.kube/config` or `/home/ubuntu/.kube/config` file of the host for use.

Solutions

Use Kubernetes serviceaccount correctly to access clusters in incluster mode. For more information, see [Configure Service Accounts for Pods](#).

Update Notes of TKE Kubernetes Major Versions

Last updated : 2020-12-25 11:02:44

1.18 changes since 1.16

Major updates

The cloud provider tag feature goes to the GA stage

The following table lists the deprecated tags and new tags:

Deprecated Tag	New Tag
<code>beta.kubernetes.io/instance-type</code>	<code>node.kubernetes.io/instance-type</code>
<code>failure-domain.beta.kubernetes.io/region</code>	<code>topology.kubernetes.io/region</code>
<code>failure-domain.beta.kubernetes.io/zone</code>	<code>topology.kubernetes.io/zone</code>

The Volume Snapshot feature goes to the Beta stage

VolumeSnapshotDataSource is enabled by default. For more information, see [Kubernetes 1.17 Feature: Kubernetes Volume Snapshot Moves to Beta](#).

CSI Migration enters the Beta stage

CSI Migration is enabled by default. For more information, see [Kubernetes 1.17 Feature: Kubernetes In-Tree to CSI Volume Migration Moves to Beta](#).

Kubernetes Topology Manager moves to Beta stage

Topology Manager enters the Beta stage in Kubernetes 1.18 and enables the CPU to implement NUMA alignment with other devices (such as SR-IOV-VF) so that workloads can support low-latency scenarios.

Before Topology Manager was introduced, the CPU and device manager could only make resource allocation decisions independently. As a result, a multi-socket CPU system might fail to allocate resources appropriately, affecting the performance of latency-sensitive applications.

Server-side Apply enters the Beta 2 stage

Server-side Apply entered the Beta stage in Kubernetes 1.16. In Kubernetes 1.18, its second beta version (ServerSideApply) was introduced. This version can record and manage changes to all new Kubernetes object fields, ensuring that users are informed of resource statuses.

IngressClass resources

`IngressClass` describes an Ingress controller type in a Kubernetes cluster. The `ingressClassName` field is added for `Ingress` resources to set the name of the controller that uses `IngressClass`. This replaced the deprecated `kubernetes.io/ingress.class` annotation.

Other updates

- Node Local DNSCache enters the GA stage.
- IPv6 enters the Beta stage.
- `kubectl debug` supports the Alpha feature.
- `Windows CSI support` supports the Alpha feature.
- `ImmutableEphemeralVolumes` supports the Alpha feature (which supports non-changeable ConfigMap and Secret and does not refresh the corresponding volume).
- The following features enter the GA stage:
 - `ScheduleDaemonSetPods`
 - `TaintNodesByCondition`
 - `WatchBookmark`
 - `NodeLease`
 - `CSINodeInfo`
 - `VolumeSubpathEnvExpansion`
 - `AttachVolumeLimit`
 - `ResourceQuotaScopeSelectors`
 - `VolumePVCDataSource`
 - `TaintBasedEvictions`
 - `BlockVolume` , `CSIBlockVolume`
 - `Windows RunAsUserName`
- The following features enter the Beta stage:
 - `EndpointSlices` : disabled by default
 - `CSIMigrationAWS` : disabled by default
 - `StartupProbe`
 - `EvenPodsSpread`

Deprecations and removals

Removed features

The following features, which were enabled by default and could not be configured, were removed:

- `GCERegionalPersistentDisk`
- `EnableAggregatedDiscoveryTimeout`
- `PersistentLocalVolumes`
- `CustomResourceValidation`
- `CustomResourceSubresources`
- `CustomResourceWebhookConversion`
- `CustomResourcePublishOpenAPI`
- `CustomResourceDefaulting`

Other removals

The following built-in cluster roles were removed:

- `system:csi-external-provisioner`
- `system:csi-external-attacher`

Deprecated features and parameters

- The default service IP CIDR block (`10.0.0.0/24`) was deprecated. You must set the service IP CIDR block by using the `--service-cluster-ip-range` parameter of kube-apiserver,
- The `rbac.authorization.k8s.io/v1alpha1` and `rbac.authorization.k8s.io/v1beta1` API groups were deprecated, and we plan to remove them in Kubernetes 1.20. Therefore, migrate your resources to `rbac.authorization.k8s.io/v1` .
- The `CSINodeInfo` feature, which has entered the GA stage and was enabled by default, was deprecated.

Parameters and other changes

kube-apiserver

- `--encryption-provider-config` : if `cacheSize: 0` is specified in the configuration file, versions earlier than 1.18 are configured to cache 1,000 Keys automatically, while version 1.18 will report a configuration verification error. You can disable the cache by setting `cacheSize` to a negative value.
- `--feature-gates` : the following features are enabled by default and no longer support CLI-based configuration.
 - `PodPriority`
 - `TaintNodesByCondition`
 - `ResourceQuotaScopeSelectors`
 - `ScheduleDaemonSetPods`

- The following resource versions (group versions) are no longer supported:
 - `apps/v1beta1` and `apps/v1beta2` . Use `apps/v1` instead.
 - Under `extensions/v1beta1` :
 - `daemonsets` , `deployments` , and `replicasets` . Use `apps/v1` instead.
 - `networkpolicies` . Use `networking.k8s.io/v1` instead.
 - `podsecuritypolicies` . Use `policy/v1beta1` instead.

kubelet

- `--enable-cadvisor-endpoints` : this parameter is disabled by default. For access to the cAdvisor v1 JSON API, you must enable this parameter.
- The `--redirect-container-streaming` parameter was deprecated and will be removed in later versions. Kubernetes 1.18 supports only the default action (kubelet proxy for streaming requests). If `--redirect-container-streaming=true` is set, it must be removed.
- The `/metrics/resource/v1alpha1` metrics endpoint was deprecated. Use `/metrics/resource` instead.

kube-proxy

- The following parameters were deprecated:
 - `--healthz-port` was deprecated. Use `--healthz-bind-address` instead.
 - `--metrics-port` was deprecated. Use `--metrics-bind-address` instead.
- The `EndpointSliceProxying` feature switch (disabled by default) was added to control whether to enable EndpointSlices in kube-proxy. The `EndpointSlice` feature switch does not affect the actions of kube-proxy any more.
- Added the following timeout settings for ipvs connection configuration:
 - `--ipvs-tcp-timeout`
 - `--ipvs-tcpfin-timeout`
 - `--ipvs-udp-timeout`
- The iptables mode added IPv4/IPv6 dual-protocol stack support.

kube-scheduler

- Deprecated the `scheduling_duration_seconds` metric:
 - Deprecated `scheduling_algorithm_predicate_evaluation_seconds` . Use `framework_extension_point_duration_seconds[extension_point="Filter"]` instead.
 - Deprecated `scheduling_algorithm_priority_evaluation_seconds` . Use `framework_extension_point_duration_seconds[extension_point="Score"]` instead.
- Deprecated `AlwaysCheckAllPredicates` in the scheduler policy API.

-enable-profiling

For the alignment of `kube-apiserver` , `kube-controller-manager` , and `kube-scheduler` , [profiling is enabled by default](#). To disable profiling, specify `--enable-profiling=false` .

kubectI

- Removed the deprecated `--include-uninitialized` parameter.
- `kubectI` and `k8s.io/client-go` no longer use `http://localhost:8080` as the default apiserver address.
- `kubectI run` supports pod creation and no longer supports using the deprecated generator to create other types of resources.
- Removed the deprecated `kubectI rolling-update` command. Use the `rollout` command instead.
- `-dry-run` supports three parameter values: `client` , `server` , and `none` .
- `-dry-run=server` supports the following commands: `apply` , `patch` , `create` , `run` , `annotate` , `label` , `set` , `autoscale` , `drain` , `rollout undo` , and `expose` .
- Added the `kubectI alpha debug` command, which can be used for [debugging and troubleshooting on ephemeral containers in pods](#) (the `EphemeralContainers` feature introduced in version 1.16 needs to be enabled.)

hyperkube

The implementation of hyperkube was changed from Go code to a bash script.

Changelogs

[kubernetes 1.18 changelog](#)

[kubernetes 1.17 changelog](#)

1.16 changes since 1.14

Major updates

Improved cluster stability and availability

Improved and enhanced the production-ready features like BM cluster tool and HA.

The kubeadm support for HA has entered the Beta stage. Users can run the `kubeadm init` and `kubeadm join` commands to deploy HA control planes. Certificate management is more stable and robust. During cluster update, kubeadm can seamlessly update all certificates prior to certificate expiry. For more information, see [pr357](#) and [pr970](#).

Continuous CSI improvement

Storage SIG continues to migrate built-in storage plug-ins to CSI APIs and support features such as built-in storage plug-in resizing and inline storage volumes. It also introduced some Alpha features

that are not available in the Kubernetes storage subsystem, such as storage volume cloning. Storage volume cloning allows users to specify another PVC as "DataSource" when configuring a new storage volume. If the underlying storage system supports this feature and has implemented the "CLONE_VOLUME" feature in its CSI driver, the new storage volume will become a clone of the source storage volume. For more information, see [pr625](#).

Features

- The following features enter the GA stage:
 - `CRD`
 - `Admission Webhook`
 - `GCERegionalPersistentDisk`
 - `CustomResourcePublishOpenAPI`
 - `CustomResourceSubresources`
 - `CustomResourceValidation`
 - `CustomResourceWebhookConversion`
- CSI support for volume resizing enters the Beta stage.

General updates

- The Kubernetes core code supports the Go module.
- Preparation for cloud-provider code extraction and organization continues. cloud-provider code has been moved to `kubernetes/legacy-cloud-providers` to facilitate future deletion and external use.
- [Kubectrl get and describe commands support expansion](#).
- [Nodes support third-party monitoring plug-ins](#).
- Launched a new Alpha scheduling framework for developing and managing plug-ins and expanding the features of the scheduler. For more information, see [pr624](#).
- The APIs of `extensions/v1beta1`, `apps/v1beta1`, and `apps/v1beta2` are still deprecated and will be completely removed in version 1.16.
- Adds Topology Manager component to the kubelet to coordinate resource allocation decisions and optimize resource allocation.
- Supports IPv4/IPv6 dual stacks, so that you can assign both IPv4 and IPv6 addresses to pods and services.
- (Alpha feature) Adds the API server network proxy.
- More **expansion options** are provided for cloud controller manager migration.
- Deprecates `extensions/v1beta1`, `apps/v1beta1`, and `apps/v1beta2` APIs.

Known issues

In version 1.15, when the `--log-file` parameter was used, a log might be written to the same file multiple times. For more information, see [pr78734](#).

Update notes

• Clusters

- The following tags are no longer configured on nodes: `beta.kubernetes.io/metadata-proxy-ready` , `beta.kubernetes.io/metadata-proxy-ready` , and `beta.kubernetes.io/kube-proxy-ds-ready` .
 - `ip-mask-agent` uses `node.kubernetes.io/masq-agent-ds-ready` as the node selector and no longer uses `beta.kubernetes.io/masq-agent-ds-ready` .
 - `kube-proxy` uses `node.kubernetes.io/kube-proxy-ds-ready` as the node selector and no longer uses `beta.kubernetes.io/kube-proxy-ds-ready` .
 - `metadata-proxy` uses `cloud.google.com/metadata-proxy-ready` as the node selector and no longer uses `beta.kubernetes.io/metadata-proxy-ready` .

• API Machinery

`k8s.io/kubernetes` and other released components, including `k8s.io/client-go` and `k8s.io/api`, now contain the Go module file including the dependent library version information. When using `k8s.io/client-go` through the Go module, you can refer to [go-modules](#) and [pr74877](#).

• Apps

Remove hyperkube short aliases from source code, as hyperkube docker image currently create these aliases. For more information, see [pr76953](#).

• Lifecycle

- Removes the deprecated `kubeadm v1alpha3` configurations.
- `kube-up.sh` no longer supports `centos` and `local` .

• Storage

- The CSI volume no longer configures the `Node.Status.Volumes.Attached.DevicePath` field. The external controller of this field needs to be updated.
- The Alpha CRD was removed.
- The `StorageObjectInUseProtection` admission [plug-in](#) is enabled by default. Previously, if this plug-in is not enabled, cluster actions may change.
- After `PodInfoOnMount` is enabled for the CSI driver, the new `csi.storage.k8s.io/ephemeral` parameter will be added in the volume context. During `NodePublishVolume` implementation, this parameter allows the driver to determine, one by one, whether the current volume is ephemeral storage or persistent storage. For more information, see [pr79983](#).
- `VolumePVCDDataSource` (the storage volume cloning feature) entered the Beta phase. For more information, see [pr81792](#).
- The built-in and CSI volume limits were combined into the predicate scheduler. For more information, see [pr77595](#).

• kube-apiserver

- Deprecated the `--enable-logs-handler` parameter. We plan to remove it in v1.19.
- Deprecated `--basic-auth-file` and the corresponding authentication mode. Both of them will be removed in the future.
- Deprecated the default service IP CIDR block (10.0.0.0/24), which will be removed in six months or after two future releases. The `--service-cluster-ip-range` parameter is required for configuring the service IP CIDR block.

• **kube-scheduler**

The `v1beta1` Events API is used. Tools that consume scheduler events need to use the `v1beta1` Events API.

• **kube-proxy**

- Removed the `--conntrack-max` parameter (which can be replaced with `--conntrack-min` and `--conntrack-max-per-core`).
- Removed the `--cleanup-iptables` parameter.
- Removed `--resource-container`.

• **kubelet**

- Removed the `--allow-privileged`, `--host-ipc-sources`, `--host-pid-sources`, and `--host-network-sources` parameters (which can be replaced with the admission controller of `PodSecurityPolicy`).
- Deprecated the cAdvisor JSON API.
- Removed `--containerized`.
- The `--node-labels` parameter can no longer be used to configure tags with the forbidden prefixes of `kubernetes.io-` or `k8s.io-`.

• **kubect**

- Removed `kubectl scale job`.
- Removed the `--pod/-p` parameter in the `kubectl exec` command.
- Removed the `kubectl convert` command.
- Removed `--include-uninitialized`.
- `kubectl cp` no longer supports copying symbol links in a container. This feature can be replaced with the following commands:
 - `local to pod` : `tar cf - /tmp/foo | kubectl exec -i -n <some-namespace> <some-pod> -- tar xf - -C /tmp/bar`
 - `pod to local` : `kubectl exec -n <some-namespace> <some-pod> -- tar cf - /tmp/foo | tar xf - -C /tmp/bar`

• **kubeadm**

- Deprecated the commands `kubeadm upgrade node config` and `kubeadm upgrade node experimental-control-plane` and replaced them with `kubeadm upgrade node`.
- Deprecated the `--experimental-control-plane` parameter and replaced it with `--control-plane`.
- Deprecated the `--experimental-upload-certs` parameter and replaced it with `--upload-certs`.

- Deprecated the `kubeadm config upload` command and replaced it with `kubeadm init phase upload-config`.
- CoreDNS uses the ready plug-in to check readiness.
- Deprecated the `proxy` plug-in and replaced it with `forward`.
- Removed the `resyncperiod` option from the `kubernetes` plug-in.
- Deprecated the `upstream` option. If it is specified, it will be ignored.

Changelogs

[kubernetes 1.16 changelog](#)

[kubernetes 1.15 changelog](#)

1.14 changes since 1.12

Major updates

- [Container Storage Interface entered the GA phase.](#)
- [CoreDNS replaced kube-dns and became the default DNS server.](#)
- kubeadm is used to simplify cluster management.
- [Support for Windows Nodes entered the stable phase.](#)
- [Local storage entered the GA phase.](#)
- [Pid Limiting entered the Beta phase.](#)
- [Pod Priority and Preemption are supported.](#)

General updates

- dry-run entered the Beta phase (dry-run enables users to simulate real API requests without actually changing the cluster status.)
- kubectl diff entered the Beta phase.
- kubectl plug-in registration entered the stable phase.
- The kubelet plug-in mechanism entered the Beta phase.
- CSIPersistentVolume entered the GA phase.
- TaintBasedEviction entered the Beta phase.
- kube-scheduler perception of volume topology entered the stable phase.
- Support for out-of-tree CSI Volume plug-ins entered the stable phase.
- Supports third-party device monitoring plug-ins.
- kube-scheduler subnet feasibility entered the Beta phase.
- Pod Ready supports custom probe conditions.
- Node memory supports HugePage.

- RuntimeClass entered the Beta phase.
- Node OS/Arch labels entered the GA phase.
- node-leases entered the Beta phase.
- kubelet resource metrics endpoint entered the Alpha phase and supports data collection through Prometheus.
- runAsGroup entered the Beta phase.
- kubectl apply server-side entered the Alpha phase, allowing apply operations to run on the server side.
- kubectl supports kustomize.
- Supports the configuration of resolv.conf in pods.
- CSI volumes support resizing.
- CSI supports topology.
- volume mount supports the configuration of sub-path parameters.
- CSI supports bare devices.
- CSI supports local ephemeral volumes.

Update notes

• kube-apiserver

- etcd2 is no longer supported. Default: `--storage-backend=etcd3`.
- Deprecated the `--etcd-quorum-read` parameter.
- Deprecated the `--storage-versions` parameter.
- Deprecated the `--repair-malformed-updates` parameter.

• kube-controller-manager

Deprecated the `--insecure-experimental-approve-all-kubelet-csrs-for-group` parameter.

• kubelet

- Deprecated the `--google-json-key` parameter.
- Deprecated the `--experimental-fail-swap-on` parameter.

• kube-scheduler

`componentconfig/v1alpha1` is no longer supported.

• kubectl

The `run-container` command is no longer supported.

• taints

`node.alpha.kubernetes.io/notReady` and `node.alpha.kubernetes.io/unreachable` are no longer supported and have been replaced with `node.kubernetes.io/not-ready` and `node.kubernetes.io/unreachable`, respectively.

Changelogs

[kubernetes 1.14 changelog](#)[kubernetes 1.13 changelog](#)

1.12 changes since 1.10

Major updates

API

- The CustomResources sub-resources have entered the Beta phase and are enabled by default, allowing users to update the `/status` sub-resources, except for the `.status` field (previously, only `.spec` and `.metadata` could be updated.) When the `/status` sub-resources are enabled, `required` and `rescription` can be used for CRD OpenAPI verification schemas. In addition, users can create multiple versions of CustomResourceDefinitions, without the need for automatic switching. The `spec.additionalPrinterColumns` field of CustomResourceDefinitions can be used to allow the output of `kubectll get` to contain additional columns.
- The `dry run` feature is supported. It allows users to view the execution results of some commands without having to submit relevant modifications.

Authentication authorization

- RBAC aggregation of ClusterRoles entered the GA phase. The client-go credentials plug-in entered the Beta phase, allowing users to obtain TLS authentication information from external plug-ins.
- Added the following annotations to audit events, so that users can be better informed of the audit decision-making process:
 - The Authorization component can configure `authorization.k8s.io/decision` (the allow or forbid authorization decision) and `authorization.k8s.io/reason` (the reason for this decision).
 - The PodSecurityPolicy admission controller can configure `podsecuritypolicy.admission.k8s.io/admit-policy` and `podsecuritypolicy.admission.k8s.io/validate-policy`, containing the names of the policies allowing pod admission (PodSecurityPolicy can also restrict mount targets of the hostPath type to the read-only mode.)
- The NodeRestriction admission controller can prohibit nodes from modifying the taint information of their corresponding node objects so that users can control and track the taint settings of nodes more easily.

CLI

The CLI implemented a new plug-in mechanism and provided a development library that contains common CLI tools to facilitate plug-in development by plug-in developers.

Network

- The ipvs mode entered the GA phase.
- CoreDNS entered the GA phase and replaced kube-dns.

Nodes

- DynamicKubeletConfig entered the Beta phase.
- cri-tools entered the GA phase.
- PodShareProcessNamespace entered the Beta phase.
- Added the RuntimeClass and CustomCFSQuotaPeriod Alpha features.

Scheduler

- Pod Priority and Preemption entered the Beta phase.
- DaemonSet Pod scheduling is no longer managed by the DaemonSet controller but by the default scheduler.
- TaintNodeByCondition entered the Beta phase.
- The local image preferential selection feature is enabled by default. During pod scheduling, nodes that have locally pulled the images required by all or some pods will have a higher priority. This accelerates the launch of pods.

General updates

- The ClusterRole and StorageObjectInUseProtection features entered the GA phase.
- The external Cloud Provider feature entered the Beta phase.

Update notes

• kube-apiserver

- The `--storage-version` parameter was removed and replaced by `--storage-versions` . Meanwhile, `--storage-versions` was also deprecated.
- The default value of `--endpoint-reconciler-type` was changed to `Lease` .
- When `--enable-admission-plugins` is used, it is contained by default. When the `--admission-control` parameter is used, it must be explicitly specified.

• kubelet

- Deprecated the `--rotate-certificates` parameter and replaced it with the `.RotateCertificates` field in the configuration file.

• kubectl

- Exception `run-pod/v1` , other `kubectl run` generators have been deprecated.
- Removed the `--interactive` parameter from `kubectl logs` .
- `--use-openapi-print-columns` was deprecated and replaced with `--server-print` .

Changelogs

[kubernetes 1.12 changelog](#)

[kubernetes 1.11 changelog](#)

TKE Kubernetes Revision Version History

Last updated : 2021-07-13 10:38:37

TKE kubernetes 1.18.4 revisions

Date	Version	Updates
2021-05-14	v1.18.4-tke.9	<ul style="list-style-type: none">Ported pr93370 to support CronJobControllerV2. (kube-controller-manager)Merged pr100376 to enable HTTP/2 health check, which prevented the issue that the underlying layer connection is closed but can still be used incorrectly. (kube-apiserver, kube-controller-manager, kube-scheduler, kubelet, kube-proxy, kubectl)Merged pr100317, which fixed the issue where CVE-2021-25735 node updates might bypass the Validating Admission Webhook. (kube-apiserver)When TKE cluster adds the virtual node, ComputeResource, EKS ClusterIP, and HPA are supported. (kube-controller-manager, kube-scheduler)

Date	Version	Updates
2021-04-02	v1.18.4-tke.8	<ul style="list-style-type: none"> Merged pr97752, which fixed the issue where NewReplicaSet is displayed as <code>&dxlt;none&dxgt;</code> when describing deployment (kubectl). Merged pr93808, which fixed the issue where unnecessary information is returned when <code>kube-scheduler --version</code> is executed (kube-scheduler). Merged pr91590, which fixed the issue of warning that the port has been allocated when using the multiprotocol service of NodePort type (kube-apiserver). Merged pr98262, which allows kube-controller-manager to dynamically adjust the log level (kube-controller-manager). Merged pr95154, which fixed the issue where kube-scheduler snapshot contains the nodes being deleted (kube-scheduler). Merged pr95711, which fixed the issue where kubectl drain command occupies too much CPU (kubectl). Merged pr96602, which fixed the issue where apiserver memory leaks before or after the time gaps (kube-apiserver). Merged pr97023, which deletes the related metadata directory when unmounting an emptyDir type volume (kubelet). Merged pr97527, which fixed the issue where map access operations are not synchronized in cpumanager (kubelet). Merged pr100190, which automatically deletes the volume directory left by orphaned Pod (kubelet). Merged pr92614, when all containers of the Pod whose restart policy is RestartPolicyOnFailure exit successfully, no new sandbox will be created (kubelet). Merged pr94833, which fixed the issue where the image tag does not match in status when Pod image has multiple tags (kubelet).
2020-12-28	v1.18.4-tke.6 (ARM clusters are supported starting from this version)	<ul style="list-style-type: none"> Added metrics to QcloudCbs (kube-controller-manager). Fixed the issue where extra space exists in the value of serial when mounting CBS disk (Kubelet).

Date	Version	Updates
2020-12-21	v1.18.4-tke.5	<ul style="list-style-type: none"> Merged pr94712, which fixed CVE-2020-8564 - fixed the issue when the file format was incorrect and logLevel >= 4, Docker configuration leaked (kubelet). Merged pr95316, which fixed CVE-2020-8565 - fixed the issue where incomplete fix for CVE-2019-11250 resulting in log token leak (logLevel >= 9) (kube-apiserver, kubect!). Merged pr95245, which fixed CVE-2020-8566 - fixed the issue where Ceph RBD adminSecrets was exposed in the log when loglevel >= 4 (kube-controller-manager). Fixed the issue where restarting kubelet caused Pod readiness check failed (kubelet). Merged pr90825, which fixed the issue where the pop operation of the fifo queue in client-go might be stuck due to race condition, which caused the pod to remain in the pending state (kubelet). The scheduler supports virtual nodes (kube-scheduler). kube-controller-manager supports virtual nodes (kube-controller-manager). Set the instance-type label based on the actual model of the node, instead of being fixed as QCLOUD (kubelet). Added the CBS to OpenAPI (kube-apiserver). Merged pr91126, which fixed the issue where the scheduler cache was inconsistent when Pod had the same name but different UID (kube-scheduler). Merged pr93387, which fixed the issue where the daemonset pod could not be scheduled to nodes due to the disorder of node cache information in the scheduler (kube-scheduler). Merged pr89465, which fixed the issue where the HPA based on Pod metrics incorrectly calculated the number of instances during rolling updates (kube-controller-manager).
2020-10-13	v1.18.4-tke.3	<ul style="list-style-type: none"> Merged pr89629, which fixed the issue where the container that mounted the subpath would fail to restart after the configmap is changed (kubelet). QcloudCbs supports BulkVolumeVerification (kube-controller-manager). Merged pr94430, which fixed the issue where the client-go reflector could not detect the "Too large resource version" error (kubelet).
2020-08-12	v1.18.4-tke.2	<ul style="list-style-type: none"> Merged pr93403, which removed the printed error information of pod condition irrelevant to the kubelet during kubelet update (kubelet).
2020-08-04	v1.18.4-tke.1	<ul style="list-style-type: none"> revert pr63066 Fixed the LB health check and IPVS issues (kube-proxy).

Date	Version	Updates
		<ul style="list-style-type: none"> Merged pr72914, which fixed the issue where mounting might fail if you deleted a pod, created a new one, and scheduled it to the same node (kube-controller-manager). Fixed the issue where creating containers in CentOS resulted in cgroup leakage (kubelet). Fixed the issue where upgrading lxcfs in Ubuntu 16 caused pods to exit (kubelet). metadata added cache and timeout. cloud-provider now supports using node names as hostnames (kubelet). metadata added local cache (kubelet). Incorporated CBS and relevant fixing code (kubelet). Merged pr90260, which fixed the issue of missing monitoring records for containerd cluster networks (kubelet). TKE can perceive the maximum number of qcloudcbs that can be mounted to a single node. In 1.12 and later versions, the value is maxAttachCount-2. In version 1.10, the value is 18 by default (kube-scheduler). Fixed the issue where CBS intree continued to unmount a non-existent disk, causing a large number of invalid requests (kubelet). Merged pr2359, which fixed the issue with missing monitoring records when the system was unable to obtain docker root (kubelet). kube-scheduler now supports dynamic logging level configuration (kube-scheduler). Produced a workaround for the missing CBS device path (/dev/disk/by-id/virtio-xxx/...) issue that prevents some users from accessing CBS properly (kubelet). TKE can perceive the maximum number of qcloudcbs that can be mounted to a single node. The kubelet side will not patch node (kubelet). Merged pr89296, so that the log will not record whether the iptables random-fully parameter is enabled (kube-proxy). Fixed the aws issue pr92162(kubelet). Merged pr91277, which prevents the issue of large numbers of TLS handshake error logs generated by kube-apiserver as a result of CLB health checks (kube-apiserver). Merged pr91500, which fixed the issue of missing environmental variables of KUBERNETES_SERVICE_HOST (kubelet). Merged 92537, which fixed the issue where client-go reflector could not recover from the error "Too large resource version" (kube-apiserver, kube-controller-manager, kube-scheduler, kubelet, and kube-proxy). Merged pr92969, which fixed the issue where CVE-2020-8559 privilege escalation from an invaded node resulted in invasion into

Date	Version	Updates
		<p>other nodes (kube-apiserver).</p> <ul style="list-style-type: none"> Merged pr92921, which fixed the DOS attack issue where CVE-2020-8557 exhausted the disk space by writing into “/etc/hosts” (kubelet).

TKE kubernetes 1.16.3 revisions

Date	Version	Updates
2021-05-24	v1.16.3-tke.17	<ul style="list-style-type: none"> Ported pr93370 to support CronJobControllerV2. (kube-controller-manager) When the TKE cluster adds the virtual node, the local replicas can be retained. (kube-scheduler)
2021-05-06	v1.16.3-tke.16	<ul style="list-style-type: none"> Updated the launch method of running kube-proxy as an image, and automatically adapted to the iptables running mode of the node to support the operating system that uses the NF_TABLES mode to run iptables by default.
2021-04-14	v1.16.3-tke.15	<ul style="list-style-type: none"> Merged pr97752, which fixed the issue where NewReplicaSet is displayed as <code>&dxlt;none&dxgt;</code> when describing deployment (kubectl). Merged pr92614, when all containers of the Pod whose restart policy is RestartPolicyOnFailure exit successfully, no new sandbox will be created (kubelet). Merged pr91590, which fixed the issue of warning that the port has been allocated when using the multiprotocol service of NodePort type (kube-apiserver). Merged pr98262, which allows you to use kube-controller-manager to dynamically adjust the log level (kube-controller-manager). Merged pr95301, which automatically deletes the volume directory left by orphaned Pod (kubelet).
2020-12-28	v1.16.3-tke.14	<ul style="list-style-type: none"> Added metrics to QcloudCbs (kube-controller-manager). Fixed the issue where extra space exists in the value of serial when mounting CBS disk (Kubelet).

Date	Version	Updates
2020-12-21	v1.16.3-tke.13	<ul style="list-style-type: none"> • Merged pr94712, which fixed CVE-2020-8564 - fixed the issue when the file format was incorrect and logLevel >= 4, Docker configuration leaked (kubelet). • Merged pr95316, which fixed CVE-2020-8565 - fixed the issue where incomplete fix for CVE-2019-11250 resulting in log token leak (logLevel >= 9) (kube-apiserver, kubectI). • Merged pr95245, which fixed CVE-2020-8566 - fixed the issue where Ceph RBD adminSecrets was exposed in the log when loglevel >= 4 (kube-controller-manager). • Merged pr86191, which fixed the issue where Pod might be in the wrong state when the node was restarted (kubelet). • Merged pr86140, which fixed the issue where the Controller Manager did not handle the timeout error correctly, so that the expanded pod could not be created (kube-controller-manager). • Merged pr90825, which fixed the issue where the pop operation of the fifo queue in client-go might be stuck due to race condition, which caused the Pod to remain in the pending state (kubelet). • The scheduler supports virtual nodes (kube-scheduler). • kube-controller-manager supports virtual nodes (kube-controller-manager). • Set the instance-type label based on the actual model of the node, instead of being fixed as QCLOUD (kubelet). • Added the CBS to OpenAPI (kube-apiserver). • Merged pr81344, which fixed the issue where the CPU Manager did not support SourcesReady (kubelet). • Merged pr91126, which fixed the issue where the scheduler cache was inconsistent when Pod had the same name but different UID (kube-scheduler). • Merged pr89224, which fixed the issue where kube-scheduler restarted abnormally because NodeInfo did not check (kube-scheduler). • Merged pr89465, which fixed the issue where the HPA based on Pod metrics incorrectly calculated the number of instances during rolling updates (kube-controller-manager).

Date	Version	Updates
2020-10-13	v1.16.3-tke.11	<ul style="list-style-type: none"> Merged pr92971, which fixed the issue where CVE-2020-8559 privilege escalation from an invaded node resulted in invasion into other nodes (kube-apiserver). Merged pr92924, which fixed the DOS attack issue where CVE-2020-8557 exhausted the disk space by writing into /etc/hosts (kubelet). Merged pr93403, which removed the printed error information of pod condition irrelevant to the kubelet during kubelet update (kubelet). Merged pr89629, which fixed the issue where the container that mounted the subpath would fail to restart after the configmap is changed (kubelet). QcloudCbs supports BulkVolumeVerification (kube-controller-manager). Merged pr84998, which resolved the issue where the corresponding node lease object might be rebuilt after the node was deleted and caused junk data (kubelet).
2020-07-28	v1.16.3-tke.10	<ul style="list-style-type: none"> Incorporated pr91277, which prevents the issue of large numbers of TLS handshake error logs generated by kube-apiserver as a result of CLB health checks (kube-apiserver). Incorporated pr91500, which fixed the issue of missing environmental variables of KUBERNETES_SERVICE_HOST (kubelet).
2020-06-17	v1.16.3-tke.9	Temporarily fixes the AWS issue pr92162 . AWS Credential Provider is no longer registered to prevent this issue from causing slow node launches.
2020-06-11	v1.16.3-tke.8	Merges pr85993 , which allows you to use CNI results to set kubernetes gateway addresses.
2020-06-10	v1.16.3-tke.7	<ul style="list-style-type: none"> Merges pr90260, which fixes the issue of missing monitoring records for containerd cluster networks. Merges pr89515, which fixes the issue where HPA miscalculates the number of pods during rolling updates. Merges pr91252, which ignores Pod Condition updates generated by other components to avoid unnecessary scheduling. Merges pr89794, which clears kube-controller-manager error logs to avoid CVE-2020-8555 Half-Blind SSRF attacks.
2020-05-18	v1.16.3-tke.6	TKE can perceive the maximum number of qcloudcbs that can be mounted to a single node. The max value cannot be dynamically obtained.

Date	Version	Updates
2020-04-20	v1.16.3-tke.5	Merges pr69047 , which fixes the <code>node.Spec.Unschedulable</code> backward compatibility issue. (This fix is overwritten when the in-tree cbs code is incorporated).
2020-04-14	v1.16.3-tke.4	<ul style="list-style-type: none"> Merges pr87913, which fixes the CVE-2020-8551: Kubelet DoS attack issue. Merges pr87669, which fixes the CVE-2020-8552: apiserver DoS attack issue. TKE can perceive the maximum number of qcloudcbs that can be mounted to a single node. (In 1.12 and later versions, the value is <code>maxAttachCount-2</code>. In version 1.10, the value is 18 by default). Merges pr87467, which fixes the issue of excessive CPU consumption by kubectrl in parsing YAML files when an authorized user sends a malicious YAML file.
2020-03-11	v1.16.3-tke.3	<ul style="list-style-type: none"> Fixed the issue where CBS intree continued to unmount a non-existent disk, which caused a large number of invalid requests. Added a local metadata cache.
2020-02-14	v1.16.3-tke.2	<ul style="list-style-type: none"> Merges pr2359, which fixes the issue of missing monitoring records when the system is unable to obtain docker root. Merges pr86583, which increases the logging level to reduce the amount of logs caused by the lack of support for random-fully in earlier versions of iptables. kube-scheduler now supports dynamic logging level configuration. Produces a workaround for the missing CBS device path (<code>/dev/disk/by-id/virtio-xxx/...</code>) issue that prevents some users from accessing CBS properly. Merges pr86230, which skips assumed pod updates when pods are scheduled.

Date	Version	Updates
2020-01-06	v1.16.3-tke.1	<ul style="list-style-type: none">• Incorporates pr79036, which fixes the issue where upon being opened, the CPU Manager disables the CPU quota if the QoS setting of a pod is Guaranteed.• Incorporates pr84167, which fixes the issue where an incorrect Etcd key prefix causes an apiserver health check failure.• Reverts pr63066, which fixes the CLB health check and IPVS issues.• Incorporates pr72914, which fixes the issue where mounting may fail if you delete a pod, create a new one, and schedule it to the same node.• Fixes the issue where creating containers in CentOS results in cgroup leakage.• Fixes the issue where upgrading lxcfs in Ubuntu 16 causes pods to exit.• Adds metadata cache and timeout. cloud-provider now supports using node names as hostnames.• Reverts pr79036, which fixes the issue where upon being opened, the CPU Manager disables the CPU quota if the QoS setting of a pod is Guaranteed.• Produces a workaround for the missing CBS device path (/dev/disk/by-id/virtio-xxx/...) issue that prevents some users from accessing CBS properly.

TKE kubernetes 1.14.3 revisions

Date	Version	Revisions
2021-05-06	v1.14.3-tke.22	Updated the launch method of running kube-proxy as an image, and automatically adapted to the iptables running mode of the node to support the operating system that uses the NF_TABLES mode to run iptables by default.

Date	Version	Revisions
2021-04-14	v1.14.3-tke.21	<ul style="list-style-type: none"> Merged pr97752, which fixed the issue where NewReplicaSet is displayed as <code>&dxlt;none&dxgt;</code> when describing deployment (kubectl). Merged pr78999, which fixed the issue of judging the case of the protocol during graceful close (kube-proxy). Merged pr91590, which fixed the issue of warning that the port has been allocated when using the multiprotocol service of NodePort type (kube-apiserver). Merged pr98262, which allows kube-controller-manager to dynamically adjust the log level (kube-controller-manager). Merged pr95301, which automatically deletes the volume directory left by orphaned Pod (kubelet).
2020-12-28	v1.14.3-tke.19	<ul style="list-style-type: none"> Added metrics to QcloudCbs (kube-controller-manager). Fixed the issue where extra space exists in the value of serial when mounting CBS disk (Kubelet).
2020-12-21	v1.14.3-tke.18	<ul style="list-style-type: none"> Merged pr94712, which fixed CVE-2020-8564 - fixed the issue when the file format was incorrect and logLevel ≥ 4, Docker configuration leaked (kubelet). Merged pr95316, which fixed CVE-2020-8565 - fixed the issue where incomplete fix for CVE-2019-11250 resulting in log token leak (logLevel ≥ 9) (kube-apiserver, kubectl). Merged pr95245, which fixed CVE-2020-8566 - fixed the issue where Ceph RBD adminSecrets was exposed in the log when loglevel ≥ 4 (kube-controller-manager). Merged pr86140, which fixed the issue where the Controller Manager did not handle the timeout error correctly, so that the expanded Pod could not be created (kube-controller-manager). The scheduler supports virtual nodes (kube-scheduler). kube-controller-manager supports virtual nodes (kube-controller-manager). Set the instance-type label based on the actual model of the node, instead of being fixed as QCLOUD (kubelet). Merged pr79338, when both SupportPodPidsLimit and SupportNodePidsLimit are not enabled, the pids cgroup subsystem will not be enabled (kubelet). Merged pr89224, which fixed the issue where kube-scheduler restarted abnormally because NodeInfo is not checked (kube-scheduler). Merged pr89465, which fixed the issue where the HPA based on Pod metrics incorrectly calculated the number of instances during rolling updates (kube-controller-manager).

Date	Version	Revisions
2020-10-13	v1.14.3-tke.17	<ul style="list-style-type: none"> Merged pr74781, which changed the default update strategy of ConfigMap and Secret from Cache to Watch (kubelet). Merged pr93403, which removed the printed error information of pod condition irrelevant to the kubelet during kubelet update (kubelet). Merged pr89629, which fixed the issue where the container that mounted the subpath would fail to restart after the configmap is changed (kubelet). Merged pr80942, which fixed the issue where rules were not deleted after the service was deleted in ipvs mode (kube-proxy). QcloudCbs supports BulkVolumeVerification (kube-controller-manager).
2020-08-04	v1.14.3-tke.16	Merged pr78883 , which fixed the bug where the default value for pod.spec.container.SecurityContext.ProcMount was added by default.
2020-07-28	v1.14.3-tke.15	<ul style="list-style-type: none"> Incorporated pr76518 and pr82514, which limits the return size of http and exec probe to prevent occupation of large amounts of node memory (kubelet). Incorporated pr91277, which prevents the issue of large numbers of TLS handshake error logs generated by kube-apiserver as a result of CLB health checks (kube-apiserver). Incorporated pr91500, which fixed the issue of missing environmental variables of KUBERNETES_SERVICE_HOST (kubelet). Incorporated pr77475, which fixed the issue of Cronjob scheduling failure when the number of jobs exceeded 500 (kube-controller-manager).
2020-06-10	v1.14.3-tke.14	<ul style="list-style-type: none"> Merges pr85027, which fixes the issue where HPA miscalculates of the number of pods during rolling updates. Merges pr79708, which uses spec.replicas to calculate the current number of replicas of HPA. Merges pr91252, which ignores Pod Condition updates generated by other components to avoid unnecessary scheduling. Merges pr89794, which clears kube-controller-manager error logs to avoid CVE-2020-8555 Half-Blind SSRF attacks.
2020-06-04	v1.14.3-tke.13	<ul style="list-style-type: none"> Merges pr90260, which fixes the issue of missing monitoring records for containerd cluster networks. Merges pr79451, which fixes the issue where if restartPolicy is set to Never, kubelet does not try to create SandBox again after the first attempt fails.

Date	Version	Revisions
2020-05-18	v1.14.3-tke.12	TKE can perceive the maximum number of qcloudcbs that can be mounted to a single node. The max value cannot be dynamically obtained.
2020-04-14	v1.14.3-tke.11	<ul style="list-style-type: none"> Merges pr75442, which changes the bandwidth unit from Kb to b. Merges pr87669, which fixes the CVE-2020-8552: apiserver DoS attack issue. TKE can perceive the maximum number of qcloudcbs that can be mounted to a single node. (In 1.12 and later versions, the value is maxAttachCount-2. In version 1.10, the value is 18 by default).
2020-04-14	v1.14.3-tke.10	Fixes the issue where CBSintree continues to unmount a non-existent disk, which causes a large number of invalid requests.
2020-01-13	v1.14.3-tke.9	<ul style="list-style-type: none"> Merges pr2359, which fixes the issue of missing monitoring records when the system is unable to obtain docker root. Merges pr86583, which increases the logging level to reduce the amount of logs caused by the lack of support for random-fully in earlier versions of iptables. kube-scheduler now supports dynamic logging level configuration. Produces a workaround for the missing CBS device path (/dev/disk/by-id/virtio-xxx/...) issue that prevents some users from accessing CBS properly. Merges pr86230, which skips assumed pod updates when pods are scheduled.
2019-12-23	v1.14.3-tke.8	Reverts pr79036 , which fixes an issue where upon being opened, the CPU Manager disables the CPU quota if the QoS setting of a pod is Guaranteed.
2019-12-17	v1.14.3-tke.7	<ul style="list-style-type: none"> Added metadata cache and timeout. Fixed the issue where upgrading lxcfs in Ubuntu 16 caused pods to exit. Avoided the readiness state of "pod not ready" when kubelet was restarted.
2019-11-28	v1.14.3-tke.6	cloud-provider supports using node names as hostnames.

Date	Version	Revisions
2019-11-18	v1.14.3-tke.5	<ul style="list-style-type: none"> Merges pr83435, which fixes an issue that allows DoS attacks that use malicious YAML or JSON files to exhaust kube-apiserver CPU or memory resources, resulting in a loss of service. Merges pr84167, which fixes an issue where an incorrect ETCD prefix causes apiserver health checks to fail. Merges pr75622, which fixes an issue where, when there is a high sts (>2000) workload in a cluster, it takes too long to sync sts changes to pod (about 20s).
2019-10-23	v1.14.3-tke.4	Merges pr79036 , which fixes an issue where upon being opened, the CPU Manager disables the CPU quota if the QoS setting of a pod is Guaranteed.
2019-09-10	v1.14.3-tke.3	Incorporated pr63066 , which fixed the issue where CLB health checks failed in IPVS mode.
2019-09-06	v1.14.3-tke.2	<ul style="list-style-type: none"> Fixes the cve-2019-9512&cve-2019-9514 HTTP/2 DDoS security issue. Merges pr72914, which fixes an issue where deleting a Pod and then creating a new one and scheduling it to the same node could cause mounting a volume to fail. Resolves the issue where creating containers in CentOS results in cgroup leakage.

TKE kubernetes 1.12.4 revisions

Date	Version	Revisions
2021-05-06	v1.12.4-tke.28	Updated the launch method of running kube-proxy as an image, and automatically adapted to the iptables running mode of the node to support the operating system that uses the NF_TABLES mode to run iptables by default.
2020-12-28	v1.12.4-tke.27	<ul style="list-style-type: none"> Added metrics to QcloudCbs (kube-controller-manager). Fixed the issue where extra space exists in the value of serial when mounting CBS disk (Kubelet).
2020-12-15	v1.12.4-tke.26	QcloudCbs supports BulkVolumeVerification (kube-controller-manager).
2020-11-17	v1.12.4-tke.25	Merged pr79495 , which fixed the issue where the webhook call failed when there were multiple versions of CRD (kube-apiserver).

Date	Version	Revisions
2020-10-13	v1.12.4-tke.24	Merged pr93403 , which removed the printed error information of pod condition irrelevant to the kubelet during kubelet update (kubelet).
2020-08-04	v1.12.4-tke.23	Merged pr78881 , which fixed the bug where the default value for pod.spec.container.SecurityContext.ProcMount was added by default.
2020-07-28	v1.12.4-tke.22	<ul style="list-style-type: none"> Incorporated pr91277, which prevents the issue of large numbers of TLS handshake error logs generated by kube-apiserver as a result of CLB health checks (kube-apiserver). Incorporated pr91500, which fixed the issue of missing environmental variables of KUBERNETES_SERVICE_HOST (kubelet).
2020-06-10	v1.12.4-tke.21	<ul style="list-style-type: none"> Merges pr73915, which prevents the watcher from receiving events before the watch is started. Merges pr91252, which ignores Pod Condition updates generated by other components to avoid unnecessary scheduling. Merges pr73915, which clears kube-controller-manager error logs to avoid CVE-2020-8555 Half-Blind SSRF attacks.
2020-06-04	v1.12.4-tke.20	<ul style="list-style-type: none"> Merges pr90260, which fixes the issue of missing monitoring records for containerd cluster networks. Merges pr79451, which fixes the issue where if restartPolicy is set to Never, kubelet does not try to create SandBox again after the first attempt fails.
2020-05-18	v1.12.4-tke.19	<ul style="list-style-type: none"> Merges pr77802, which disables graceful termination for UDP traffic. Merges pr68741, which fixes the issue of when the soft link /var/lib/kubelet and subpath are used, the host fails to unmount after pod deletion, resulting in mount target leakage and the pod being stuck in terminating. TKE can perceive the maximum number of qcloudcbs that can be mounted to a single node. The max value cannot be dynamically obtained.
2020-04-14	v1.12.4-tke.18	<ul style="list-style-type: none"> Merges pr73401, pr73606, and pr76060, which deletes DaemonSet pods allocated to non-existent nodes. Merges pr68619, which fixes the CPU Manager dirty data issue. Merges pr87669, which fixes the CVE-2020-8552: apiserver DoS attack issue. TKE can perceive the maximum number of qcloudcbs that can be mounted to a single node. (In 1.12 and later versions, the value is maxAttachCount-2. In version 1.10, the value is 18 by default).

Date	Version	Revisions
2020-02-14	v1.12.4-tke.17	<ul style="list-style-type: none"> Upgrades the CBS V2 interface to V3. Fixes the issue where CBS intree continues to unmount a non-existent disk, which causes a large number of invalid requests.
2020-01-13	v1.12.4-tke.16	<ul style="list-style-type: none"> Merges pr2359, which fixes the issue of missing monitoring records when docker root fails to be obtained. Merges pr86583, which increases the logging level to prevent excessive logs from being generated when iptables does not support random-fully. kube-scheduler supports dynamic logging level configuration. Produces a workaround for the missing CBS device path (/dev/disk/by-id/virtio-xxx/...) issue that prevents some users from accessing CBS properly. Merges pr86230, which skips assumed pod updates when pods are scheduled.
2019-12-23	v1.12.4-tke.15	Reverted pr79036 , which fixed the issue where the enabled CPU Manager disabled `cpu quota` if `QoS` of a pod was set to `Guaranteed`.
2019-12-17	v1.12.4-tke.14	<ul style="list-style-type: none"> Added metadata cache and timeout. Fixed the issue where upgrading lxcfs in Ubuntu 16 caused pods to exit. Avoided the readiness state of "pod not ready" when kubelet was restarted.
2019-11-28	v1.12.4-tke.13	cloud-provider supports using node names as hostnames.
2019-11-18	v1.12.4-tke.12	Merges pr75622 , which fixes an issue where, when there is a high sts (>2000) workload, it takes too long to sync sts changes to pod (about 20s).
2019-10-23	v1.12.4-tke.11	<ul style="list-style-type: none"> Merges pr79036, which fixes an issue where upon being opened, the CPU Manager disables the CPU quota if the QoS setting of a pod is Guaranteed. Merges pr72868, which adds a new parameter <code>--metrics-port</code> to kube-proxy and addresses the issue where <code>--metrics-bind-address</code> does not recognize port numbers.

Date	Version	Revisions
2019-09-06	v1.12.4-tke.10	<ul style="list-style-type: none"> Fixes the cve-2019-9512&cve-2019-9514 HTTP/2 DDoS security issue. Merges pr72914, which fixes an issue where deleting a Pod and then creating a new one and scheduling it to the same node could cause mounting a volume to fail. Merges pr71834, which fixes an issue with IPVS load balancing where, if sessionAffinity is set to ClientIP, traffic is routed to an invalid real server.
2019-08-09	v1.12.4-tke.9	Fixed the issue where creating containers in CentOS resulted in cgroup leakage.
2019-08-08	v1.12.4-tke.8	Incorporated pr72118 , which fixed the issue where mounting failed if a CBS StatefulSet was rescheduled to the same node.
2019-07-17	v1.12.4-tke.7	Incorporated pr75037 , which resolved the security risks of the cp command in kubectl.
2019-07-16	v1.12.4-tke.6	Fixed the compatibility issue between the TLinux kernel and IPVS and fixed CLB health check failures in IPVS mode.
2019-07-09	v1.12.4-tke.5	Incorporated pr72361 , which fixed the kube-proxy deadlock issue.
2019-06-25	v1.12.4-tke.4	Fixes the compatibility issue between the TLinux kernel and IPVS.
2019-06-17	v1.12.4-tke.3	Incorporated pr71114 , which fixed the IPVS throughput issue.
2019-06-04	v1.12.4-tke.2	<ul style="list-style-type: none"> Merges pr74755, which fixes a hang/timeout issue when running large numbers of pods with unique configmap/secret references. Merges pr69047, which fixes a backward compatibility issue with <code>node.Spec.Unschedulable</code>.

TKE kubernetes 1.10.5 revisions

Date	Version	Revisions
------	---------	-----------

Date	Version	Revisions
2021-05-06	v1.10.5-tke.20	Updated the launch method of running kube-proxy as an image, and automatically adapted to the iptables running mode of the node to support the operating system that uses the NF_TABLES mode to run iptables by default.
2020-06-10	v1.10.5-tke.19	<ul style="list-style-type: none"> Merges pr90260, which fixes the issue of missing monitoring records for containerd cluster networks. Merges pr91252, which ignores Pod Condition updates generated by other components to avoid unnecessary scheduling. Merges pr89794, which clears kube-controller-manager error logs to avoid CVE-2020-8555 Half-Blind SSRF attacks.
2020-05-18	v1.12.4-tke.19	Merges pr61549 , which adds volumeSpec data for mountedPods cache and fixes the issue of deletion failure when multiple pods use the same volume.
2020-04-29	v1.10.5-tke.17	Merges pr75622 , which fixes the issue where, when a large number (>2000) of sts workloads exist in a cluster, it takes too long (about 20s) to synchronize sts changes to a pod.
2020-04-14	v1.10.5-tke.16	<ul style="list-style-type: none"> Merges pr68619, which fixes the CPU Manager dirty data issue. Merges pr87669, which fixes the CVE-2020-8552: apiserver DoS attack issue. TKE can perceive the maximum number of qcloudcbs that can be mounted to a single node. (In 1.12 and later versions, the value is maxAttachCount-2. In version 1.10, the value is 18 by default).
2020-02-14	v1.10.5-tke.15	<ul style="list-style-type: none"> Upgrades the CBS V2 interface to V3. Fixes the issue where CBS intree continues to unmount a non-existent disk, which causes a large number of invalid requests.
2020-01-13	v1.10.5-tke.14	<ul style="list-style-type: none"> Merges pr2359, which fixes the issue of missing monitoring records when docker root fails to be obtained. Merges pr86583, which increases the logging level to prevent excessive logs from being generated when iptables does not support random-fully. kube-scheduler supports dynamic logging level configuration. Produces a workaround for the missing CBS device path (/dev/disk/by-id/virtio-xxx/...) issue that prevents some users from accessing CBS properly. Merges pr86230, which skips assumed pod updates when pods are scheduled.

Date	Version	Revisions
2019-12-23	v1.10.5-tke.13	Reverted pr79036 , which fixes an issue where cpumanager disables cpu quota when it opens if the QoS setting of a pod is Guaranteed.
2019-12-13	v1.10.5-tke.12	<ul style="list-style-type: none"> kubelet does not delete nodes when checking externalID. Added metadata cache and timeout. Fixed the issue where upgrading lxcfs in Ubuntu 16 caused pods to exit. Avoided the readiness state of “pod not ready” when kubelet was restarted.
2019-11-18	v1.10.5-tke.11	Removes the kube-controller-manager probe that sends heartbeats to kubelet.
2019-10-23	v1.10.5-tke.10	<ul style="list-style-type: none"> Merges pr79036, which fixes an issue where upon being opened, the CPU Manager disables the CPU quota if the QoS setting of a pod is Guaranteed. Merges pr72868, which adds a new parameter <code>--metrics-port</code> to kube-proxy and addresses the issue where <code>--metrics-bind-address</code> does not recognize port numbers.
2019-09-06	v1.10.5-tke.9	<ul style="list-style-type: none"> Fixes the cve-2019-9512&cve-2019-9514 HTTP/2 DDoS security issue. Merges pr72914, which fixes an issue where deleting a Pod and then creating a new one and scheduling it to the same node could cause mounting a volume to fail. Merges 67430 to rollback the state if updateContainerCPUSet fails.
2019-08-08	v1.10.5-tke.8	Merges pr72118 , which fixes an issue where, if kubelet mounts a device immediately after unmounting it, an error occurs with the message `resource name may not be empty`.
2019-07-17	v1.10.5-tke.7	Merges pr75037 , which fixes a security issue affecting the cp command in kubectl.
2019-06-25	v1.10.5-tke.6	Fixed a Linux kernel and IPVS load balancing issue.
2019-06-17	v1.10.5-tke.5	Merges pr71114 , which fixes an issue with IPVS throughput.
2019-03-19	v1.10.5-tke.4	Incorporated pr65092 , which fixed the issue where apiserver would panic when handling specific requests.
2019-02-19	v1.10.5-tke.3	Incorporated pr67288 , which fixed the issue where apiserver did not close the other side of the connection immediately when proxying.

Date	Version	Revisions
2018-09-28	v1.10.5-tke.2	Moves the CLB creation logic from controller-manager to an independent service controller.
2018-09-27	v1.10.5-tke.1	Backports pr63321 , which fixes an issue where termination takes too long when there are multiple service containers in a pod.
2018-09-21	v1.10.5-qcloud-rev1	If a kubelet status update times out, controller-manager probes the kubelet port.

TKE kubernetes 1.8.13 revisions

Date	Version	Revisions
2020-01-13	v1.8.13-tke.7	<ul style="list-style-type: none"> Merges pr2359, which fixes the issue of missing monitoring records when the system is unable to obtain docker root. Produces a workaround for the missing CBS device path (/dev/disk/by-id/virtio-xxx/...) issue that prevents some users from accessing CBS properly.
2019-12-13	v1.8.13-tke.6	<ul style="list-style-type: none"> kubelet does not delete nodes when checking externalID. Adds metadata cache and timeout. Fixes an issue where upgrading lxcfs in Ubuntu 16 causes pods to exit. Adds the ability to reboot kubelet to avoid pod not ready.
2019-11-18	v1.8.13-tke.5	<ul style="list-style-type: none"> Removes the kube-controller-manager probe that sends heartbeats to kubelet. Adds metrics to CBS PVC.
2018-09-28	v1.8.13-tke.2	Moved load balancing logic from controller-manager to independent service controllers.
2018-09-27	v1.8.13-tke.1	<ul style="list-style-type: none"> Disables kmem statistics to prevent cgroup numbers from leaking. Reduces resourcequota conflicts caused by creating pods.
2018-09-21	v1.8.13-qcloud-rev1	If a kubelet state update times out, controller-manager probes kubelet port.

TKE kubernetes 1.7.8 revisions

Date	Version	Revisions
2019-12-17	v1.7.8-tke.4	<ul style="list-style-type: none">• kubelet does not delete nodes when checking externalID.• Added metadata cache and timeout.• Fixed the issue where upgrading lxcfs in Ubuntu 16 caused pods to exit.• Avoided the readiness state of “pod not ready” when kubelet was restarted.
2018-09-28	v1.7.8-tke.2	Fixes a conflict between controller-manager and an external service controller.
2018-09-27	v1.7.8-tke.1	Moves load balancing logic from controller-manager to an independent service controllers.
2018-09-21	v1.7.8-qcloud-rev1	If a kubelet status update times out, controller-manager probes the kubelet port.