


Tencent Kubernetes Engine Release Notes and Announcements 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

Release Notes and Announcements

Release Notes

TKE Release Notes

Elastic Kubernetes Service

Tencent Kubernetes Engine for Edge

Announcements

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

Release Notes and Announcements

Release Notes

TKE Release Notes

Last updated : 2021-09-23 10:52:00

May 2021

Update	Description	Date	Related Documents
The features of cloud native monitoring were enhanced.	<ul style="list-style-type: none"> The data collection configuration is optimized. The status information of the collection target is added. The interaction process is optimized. The detection of the collection target is supported. 	2021-05-28	Log Collection
The OLM add-on was launched.	The OLM add-on helps users install, update, and manage the lifecycle of Operators.	2021-05-28	OLM
The HPC add-on was launched.	HPC is an add-on to periodically modify the number of replicas of K8s workload. Used in conjunction with HPC CRD resources, it can support scheduled actions in seconds.	2021-05-28	HPC
The feature of TKE console was enhanced.	<ul style="list-style-type: none"> Users can select the operating system for when creating node. Users can modify the desired number of nodes during the node pool scaling out. Users can search for workload by tag. 	2021-05-20	Adding a Node

April 2021

Update	Description	Date	Related Documents
The feature of	<ul style="list-style-type: none"> StatefulSet and DaemonSet can be 	2021-04-	•

TKE console was enhanced.	redeployed with one-click. <ul style="list-style-type: none"> • Secret supports TLS certificate. You can import it from file or paste multiple key-value pairs to enter in a batch. • The container health check support the use of port name. • Namespace supports the selection of “All namespaces” and can be searched by keyword. 	30	StatefulSet Management <ul style="list-style-type: none"> • DaemonSet Management • Secret Management • Setting the Health Check for a Workload • Namespaces
The log collection capability is enhanced.	<ul style="list-style-type: none"> • It supports the extraction mode of multiple lines - full regex, which is suitable for searching for log by key-value pair in multi-line logs such as java program. • It supports automatic generation of regular expressions based on user’s log samples, and automatic extraction of corresponding key-value pairs. 	2021-04-15	Log Collection

March 2021

Update	Description	Date	Related Documents
The beta ARM cluster starts.	The beta of ARM cluster starts. To join the beta, please submit a ticket .	2021-03-31	-
The feature of TKE console was enhanced.	<ul style="list-style-type: none"> • Kubernetes objects support batch input by pasting multiple key-value pairs. • When configMap is used as an environment variable, all keys can be selected with one-click. • Secret can be modified through console. • ConfigMap can be managed through console. You can manually add or upload file to add key-value pairs. • subPathExpr can be configured through console. • CronJob supports displaying the generated active Jobs. User can customize the 	2021-03-31	<ul style="list-style-type: none"> • Kubernetes Object Management • Secret Management • ConfigMap Management • Instructions for Other

	“Completed Jobs Retained” and “Failed Jobs Retained”, pause the generation of a scheduled Job, resume the paused Job, and manually generate a new Job.		Storage Volumes <ul style="list-style-type: none"> • CronJob Management
The node search feature was enhanced.	The nodes can be searched in batch by label or IP address.	2021-03-19	Node Overview
The nodes of the scaling group can be removed from the node pool.	The nodes of the scaling group can be removed from the node pool. The monthly-subscribed nodes can only be removed, but cannot be terminated.	2021-03-15	Node Pool Overview
Monthly-subscribed node pools support auto scale-out.	Monthly-subscribed node pools only support auto scale-out. The Auto Scale-in is not supported currently.	2021-03-02	Creating a Node Pool

February 2021

Update	Description	Date	Related Documents
The operations for nodes were enhanced.	The node in the node pool can be cordoned and drained, and the node management ability of the node pool is improved.	2021-02-25	Node Pool Overview
Users can mount data disk partition when adding the existing node.	When adding existing node to a cluster or node pool, users can select partition or logical volume name for the data disk mounting, which provides more flexible mount Settings.	2021-02-20	Adding an existing node

January 2021

Update	Description	Date	Related Documents
Master version	This feature provides a more flexible version	2021-01-	Upgrading

upgrade supports minor version.	upgrade mechanism.	14	the Master Kubernetes Version
---------------------------------	--------------------	----	---

December 2020

Update	Description	Date	Related Documents
Descheduler addon was launched.	Based on the actual node loads, this addon supports automatic rescheduling of marked services on high-load nodes to maintain the cluster load balance.	2020-12-25	DeScheduler
Nginx-ingress addon was fully launched.	<ul style="list-style-type: none"> The issue of nginx-ingress-controller toleration scheduling was fixed. Nginx-Ingress UI experience was improved, including the regular matching of forwarding rule, configuration of backend ClusterIP mode Service, and certificate supporting kubernetes.io/tls type Secret. 	2020-12-24	Nginx-ingress
CBS-CSI addon was launched.	CBS-CSI addon supports: <ul style="list-style-type: none"> Creating static volume/dynamic volume Storage topology awareness Scheduler awareness of node maxAttachLimit Online volume expansion Volume snapshot and restoration 	2020-12-22	CBS-CSI
TKE node pool was fully launched.	Basic node pool features allow you to conveniently and quickly create, manage, and terminate nodes and dynamically scale nodes in or out.	2020-12-21	Node Pool Overview
The feature of TKE console was enhanced.	<ul style="list-style-type: none"> The console supports Resource Quota, which can be used to configure the resource quotas and default resource request values for the namespace. The ConfigMap can be generated based on the imported file. 	2020-12-09	-
NetworkPolicy addon was	This addon supports automatic synchronization of NetworkPolicy to make the	2020-12-03	Network Policy

launched.	network isolation policy effective.		
-----------	-------------------------------------	--	--

November 2020

Update	Description	Date	Related Documents
The beta of new TKE network solution was launched.	Leveraging intelligent ENI, TKE launches a new container network solution. In this solution, each pod is assigned a dedicated ENI. Pod-to-pod communications can be implemented without traveling through the node protocol stack (default namespace), so as to shorten the container access link and the access latency.	2020-11-27	-
The beta of Nginx-Ingress feature was launched.	TKE is fully compatible with the native Nginx-ingress, and provides more advanced features to help users quickly deploy and build production-level traffic access gateways. It provides comprehensive Nginx-ingress full lifecycle management, automatic cloud native monitoring, CLS, and OPS capabilities.	2020-11-26	-
Event dashboard was launched.	This feature implements the aggregation search and trend observation of top events and exception events.	2020-11-26	Event Dashboard
Auditing dashboard was launched.	This feature implements the aggregation search and direct observation of cluster global, nodes, K8s objects and other important operations.	2020-11-26	Auditing Dashboard
The node pool and cluster operating system can be changed.	Users can create node pools of different operating systems as needed to facilitate the standardized management of nodes.	2020-11-23	Creating a Node Pool
DynamicScheduler addon is added.	This addon performs scheduling based on actual node loads, so as to prevent traffic hotspots.	2020-11-21	-

October 2020

Update	Description	Date	Related Documents
TPS supports monitoring of edge clusters.	TPS supports monitoring of edge clusters, and provides the capability of cross-VPCs multi-cluster management.	2020-10-30	PROM Instance Management
TPS supports WebHook in alarm policies.	TPS supports configuring WebHook, which can better help users find and solve service exceptions in time and is more conducive to the stable operation of services.	2020-10-30	Alarm Configurations
TKE node pool adds the ability to view scaling records.	This feature can help users view the changes in the number of nodes in the node pool and the causes and results for scaling activities.	2020-10-13	Viewing Node Pool Scaling Logs

September 2020

Update	Description	Date	Related Documents
TKE ServiceConfig was optimized	You can configure service/ingress to create tkeserviceconfig automatically.	2020-09-23	Using TKEServiceConfig to Configure CLBs
The DNSAutoscaler add-on was launched.	This add-on can obtain number of nodes and cores of the cluster via Deployment, and auto-scaling the number of DNS replicas according to the preset scaling policy, so as to improve DNS availability.	2020-09-23	DNSAutoscaler
The beta cloud native ETCD was launched.	This feature enables you to one-click deploy the high-reliability and high-performance ETCD cluster, which is profusely verified through Tencent's internal services. It also provides cross-AZ disaster recovery capabilities and optimal performance configuration.	2020-09-16	-

One-click addon configuration was available when creating the cluster.	You can easily and quickly configure the required addons for the cluster.	2020-09-15	Creating a Cluster
--	---	------------	------------------------------------

August 2020

- It can be associated with a local Alertmanager component.
- Supports CRD management of Prometheus rules.

Update	Description	Date	Related Documents
The monitoring capability of the cloud native monitoring service was optimized.	<ul style="list-style-type: none"> • The cluster monitoring collection items are preset, and a diverse Grafana dashboard is available. • The Targets list page is added to allow users to view the real-time status of monitoring tasks. 	2020-08-31	-
The alarm module of the cloud native monitoring service was upgraded.	2020-08-31	Alarm Configurations	
The NodeProblemDetectorPlus add-on was launched.	It supports configuring node self-healing policy on the basis of existing detection feature.	2020-08-25	Node-Problem-Detector-Plus Description
TKE launched in-place major-version upgrade capabilities.	The in-place major-version upgrade feature supports major-version upgrade without node reinstallation.	2020-08-25	-
TKE add-ons were fully launched.	The add-on feature enables users to install or uninstall multiple advanced add-ons for clusters.	2020-08-25	Add-on Overview
TKE Kubernetes 1.18 version was fully	Allows users to create clusters of the Kubernetes 1.18 version and	2020-08-24	-

launched.	upgrade clusters to the 1.18 version.		
-----------	---------------------------------------	--	--

July 2020

Update	Description	Date	Related Documents
The capabilities of storage plug-ins are optimized.	<ul style="list-style-type: none"> The TKE console supports PV creation without specifying StorageClass. Users can set and mount COS subdirectories. 	2020-07-28	<ul style="list-style-type: none"> PV and PVC binding rules Using COS
Cluster creation supports setting node configuration placement groups.	This feature enables disaster recovery and high availability for nodes when they are launched.	2020-07-15	Creating a Cluster
Cloud native monitoring is launched for beta testing.	It supports one-click deployment of the high-availability monitoring architecture and quick association with TKE clusters and EKS clusters.	2020-07-15	Cloud Native Monitoring
The collection configuration and alarm configuration of cloud native monitoring are implemented through products.	<ul style="list-style-type: none"> Three configuration modes are supported: service monitor, pod monitor, and raw job. Alarm history rewinding is supported. 	2020-07-15	-
RBAC-based permission control with finer granularity is launched for beta testing.	<ul style="list-style-type: none"> It allows cluster admins to configure management permissions for different roles regarding different resources in the cluster. It supports certificate revocation. It is suitable for enterprises' compliance permission management scenarios. 	2020-07-10	TKE Kubernetes Object-level Permission Control

June 2020

Update	Description	Date	Related Documents
The IPVS-bpf mode is launched for beta testing.	TKE uses eBPF to bypass conntrack and optimize the Kubernetes Service, improving the non-persistent connection performance by over 40% and reducing the p99 latency by over 31%.	2020-06-19	-
TKE supports the creation of services in CLB-to-Pod direct access mode.	The forwarding performance of pods with LoadBalancer directly connected to ENI can be improved by over 10%.	2020-06-18	Using services with LoadBalancer directly connected to pods
TKE supports balanced forwarding and local binding.	TKE has strengthened the Loadbalancer Service and the LoadBalancer Ingress backend binding with the RS feature. TKE supports balanced forwarding and local binding.	2020-06-18	Service backend selection
The TKE app market was comprehensively upgraded.	The app market provides an output window for Tencent Cloud's practical cloud-native technologies and also provides a variety of great community apps that users can easily and quickly use.	2020-06-10	Application Market

May 2020

Update	Description	Date	Related Documents
TKE launches the ContainerNative network LoadBalancer (supports CLB-to-Pod direct access).	In TKE, you can use services and ingresses with LoadBalancer directly connected to pods, which provides higher performance and more robust product capabilities. This feature can resolve issues such as imbalanced load for persistent connections, health check session persistence configuration issues, and IPVS jitter.	2020-05-12	-

The cluster deletion feature is optimized.	<ul style="list-style-type: none"> When deleting a cluster, you can view the existing nodes, security groups, cloud disks, and other resources in the cluster. A deletion risk reminder is added to prevent accidental deletion that may interrupt your business. You can delete the nodes, cloud disks, and other resources in the cluster at the same time. 	2020-05-12	Deleting Clusters
TKE launches the open-source KMS plug-in.	<ul style="list-style-type: none"> The Tencent Cloud TKE-KMS plug-in integrates the rich key management features of the Key Management Service (KMS) to provide robust encryption/decryption capabilities for secrets in Kubernetes clusters. By using the TKE-KMS plug-in, you can perform KMS encryption on your business credential information stored in clusters to enhance your security. 	2020-05-08	-

April 2020

Update	Description	Date	Related Documents
The TKE console supports multidimensional node filtering and node list export.	<ul style="list-style-type: none"> Cluster nodes can be filtered based on lock status. Cluster nodes can be filtered based on CVM attributes such as node status and IP address. Cluster nodes can be exported in batches. 	2020-04-22	Node Overview
TKE Image Registry can configure a global image lifecycle management policy.	TKE Image Registry adds the image lifecycle management feature, which allows users to configure a global image version clearing policy for the main account and supports independent version clearing policies retained for individual repositories.	2020-04-16	Image Lifecycle Management
The TKE beta	The node pool feature can be used in the	2020-04-	Node Pool

version supports the node pool feature.	<p>following scenarios:</p> <ul style="list-style-type: none"> When a cluster contains multiple heterogeneous nodes (different models), node pools can standardize node group management. If a cluster needs to scale nodes in or out frequently, node pools can reduce the operation costs. If application scheduling rules in a cluster are complex, node pool tags can quickly specify business scheduling rules. During routine cluster node maintenance, node pools can conveniently manage Kubernetes and Docker version upgrades. 	10	Management
TKE removes Kubernetes 1.8 as an option.	TKE no longer supports creating clusters using Kubernetes 1.8.	2020-04-03	Creating a Cluster
Self-deployed cluster master update.	You can now use the TKE console to perform rolling updates of Kubernetes masters on self-deployed clusters.	2020-04-02	Updating a Cluster

March 2020

Update	Description	Date	Related Documents
TKE now supports both GlobalRouter and VPC-CNI network modes.	TKE now supports GlobalRouter and VPC-CNI network modes for your business needs. Choose the one that fits your needs.	2020-03-30	How to Choose TKE Network Mode
TKE has stopped providing features related to TencentHub.	We plan to discontinue support for TencentHub this month, so TKE has officially stopped providing features related to TencentHub and no longer supports related APIs.	2020-03-25	-
TKE supports enabling "Local Disk Formatting" for	TKE now allows you to enable "Local Disk Formatting" for BM and big data model nodes and also allows you to mount and set container directories.	2020-03-02	Creating Clusters

BM and big data models.

February 2020

Update	Description	Date	Related Documents
TKE cluster scaling groups support node shutdown when scaling in.	When scaling in, cluster scaling groups now support shutting nodes down instead of terminating or draining them . To enable this feature, you need to submit a ticket .	2020-02-17	Cluster Scaling
TKE fully launches Kubernetes 1.16 and passes conformance verification .	<ul style="list-style-type: none"> Users can create self-deployed clusters and managed clusters of the Kubernetes 1.16 version. Users can update a cluster from Kubernetes 1.14 to 1.16. 	2020-02-14	<ul style="list-style-type: none"> Creating a Cluster Upgrading a Cluster

January 2020

Update	Description	Date	Related Documents
TKE allows users to create clusters using a cluster template.	The template-based cluster creation feature provides multiple templates for creating managed clusters, self-deployed clusters, and elastic clusters, simplifying the current cluster creation process and improving the cluster creation experience . It applies to various business scenarios such as HA clusters and GPU clusters.	2020-01-12	Creating a Cluster

December 2019

Update	Description	Date	Related
--------	-------------	------	---------

			Documents
TKE supports the PVs and the PVCs of the Cloud File Storage (CFS) and Cloud Object Storage (COS) types.	TKE supports the PVCs and the PVs of the CFS and COS types connecting storage resources with Kubernetes , which makes it convenient for users to use basic Tencent Cloud products through the native Kubernetes mode and allows users to manage file storage and object storage via PVs and PVCs .	2019-12-27	-
TKE Kubernetes 1.16 beta is launched.	<ul style="list-style-type: none"> This allows users to create Kubernetes 1.16 self-deployed clusters and managed clusters via the console. It also allows users to upgrade the Kubernetes version of a cluster from 1.14 to 1.16. 	2019-12-18	<ul style="list-style-type: none"> Creating a Cluster Upgrading a Cluster
TKE supports purchasing multiple data disks during node initialization as well as custom formatting.	TKE allows users to purchase multiple data disks during node initialization and supports custom data disk formatting, allowing users to isolate data and format settings flexibly based on their actual needs .	2019-12-12	Adding Nodes
TKE nodes support the in-place rolling updates of minor Kubernetes versions.	Nodes in in-place updates support the rolling update mode. <ul style="list-style-type: none"> Only one node is updated at a time, and the next node will be updated only after the current node is successfully updated. Currently, in-place updates only support updating different minor versions of the same major version. 	2019-12-03	Updating a Cluster

November 2019

Update	Description	Date	Related Documents
The beta custom Hostname	The TKE custom Hostname feature provides the following advantages:	2019-11-15	Adding a Node

supported by TKE is launched.	<ul style="list-style-type: none"> • Helps clusters interwork with enterprises' internal domain name service systems. • Makes it easier for users to quickly create nodes with a specified Hostname in batches. 		
TKE Ingress performance optimization is released.	<p>TKE Ingress performance is optimized to better serve users.</p> <ul style="list-style-type: none"> • CLB changes are optimized to allow batch calling APIs to process backend binding. • CVM backend query is optimized to help users avoid unnecessary repeated queries. 	2019-11-07	Ingress Management

October 2019

Update	Description	Date	Related Documents
Cluster worker nodes support configuring several security groups simultaneously and using the default security group.	TKE allows a cluster worker node to bind multiple security groups and provides a default security group, helping users quickly configure available security groups.	2019-10-22	TKE Security Group Settings
Node labels can be added in batches during creation of clusters/nodes.	When a cluster is created or new nodes are added to an existing cluster, TKE allows users to add labels for nodes that run the same business or have the same configurations. The labels help users divide resources, label resource attributes, and filter and batch process massive resource volumes.	2019-10-21	Cluster Management
Runtime component Containerd	The TKE runtime component Containerd supports the GPU model. When users need to create a GPU application in a cluster,	2019-10-17	How to Choose Containerd and Docker

supports the GPU model.	they can choose Containerd as the runtime component.		
The beta for rolling Kubernetes reinstallation and upgrade of TKE nodes is launched.	TKE supports the batch update of nodes in a cluster from an earlier version to a later version. This feature applies to clusters whose Kubernetes version is outdated and clusters whose nodes do not have relevant custom configurations. Custom configurations will become invalid after the rolling reinstallation and upgrade of nodes.	2019-10-15	Upgrading a Cluster
TKE supports GPU monitoring metrics.	TKE supports GPU monitoring metrics, enabling users to monitor GPU-related resources. By checking monitoring data, users can precisely identify specific problems, shorten troubleshooting time, and reduce OPS costs, ensuring the continuous and stable running of businesses.	2019-10-15	List of Monitoring and Alarm Metrics

September 2019

Update	Description	Release Time	Related Document
Related APIs of the TKE cluster scaling group have been updated to API 3.0.	TKE APIs have been updated to 3.0 and support all-region access. The new API documentation is more standardized and comprehensive, with unified parameter styles and common error codes. The SDK/CLI version is consistent with the API documentation, providing a simple and convenient user experience.	2019-09-12	Related APIs of the Scaling Group
TKE Kubernetes 1.14 is fully launched and has passed conformance verification.	TKE Kubernetes 1.14 is fully launched and has passed conformance verification to ensure that the latest Kubernetes version is available.	2019-09-07	Conformance Verification
TKE supports	If the Tencent Cloud tag is added to a cluster	2019-09-	-

the Tencent Cloud tag, allowing authorization by tag.	when the cluster is created, the Tencent Cloud services, cloud disks, CLBs, and other resources created in the cluster will automatically inherit the cluster's tag , allowing users to clearly view resource categories.	06	
The default instance type for created LoadBalancer-type services is CLB.	When TKE creates a LoadBalancer-type service, the default instance type is CLB. This instance type covers all features of a conventional CLB. <ul style="list-style-type: none"> The CLB instance type supports the TCP, UDP, HTTP, and HTTPS protocols. It provides flexible forwarding capabilities based on domain names and URLs. 	2019-09-06	Instance Types
TKE self-deployed clusters support the separate viewing of Master and Etcd nodes.	This feature allows users to intuitively view the list of all Master and Etcd nodes of a self-deployed cluster and the details of such nodes . Users no longer have trouble distinguishing Master and Etcd nodes in self-deployed clusters.	2019-09-05	Node Management

August 2019

Update	Description	Date	Related Documents
When a “self-deployed cluster” is created, a security group is automatically bound to the Master node.	This feature can automatically bind an applicable security group to the Master node in a self-deployed cluster . This prevents the Master node from being associated with a security group with communication problems and improves the success rate of creating self-deployed clusters.	2019-08-27	Creating a Cluster
TKE supports the visualized display of the cluster creation progress.	The visualized display of the cluster creation progress enables users to see the waiting time for cluster creation and troubleshoot the steps with exceptions . This improves the success rate of cluster	2019-08-23	Creating a Cluster

	creation and ensures the continuous and stable running of businesses.		
Open source components: TencentCloud-controller-manager and cbs-csi support Kubernetes 1.14.	The open source components Tencent Cloud-controller-manager and cbs-csi support Kubernetes 1.14.	2019-08-12	Open Source Components
Use existing CLB instances to create Ingress.	Users no longer have to create new CLB instances in order to create a new Ingress. They can now avoid additional costs by using existing CLB instances to create a new Ingress.	2019-08-08	Basic Ingress Features
TKE Kubernetes 1.14 beta is launched.	Users can now use the TKE console to create clusters based on Kubernetes 1.14.	2019-08-04	-
Related APIs of TKE cluster nodes have been updated to API 3.0.	TKE APIs have been updated to 3.0 and support all-region access. The new API documentation is more standardized and comprehensive, with unified parameter styles and common error codes. The SDK/CLI version is consistent with the API documentation, providing a simple and convenient user experience.	2019-08-04	API 3.0
TKE now supports application-level log collection.	By checking the collected file logs in the container, users can view the running status of applications in the container , precisely identify specific problems, shorten the troubleshooting time, and reduce OPS costs to ensure the continuous and stable running of businesses.	2019-08-01	Log Collection

July 2019

Update	Description	Date	Related Documents

The CLB health check failure issue in IPVS mode is fixed.	Fixes the compatibility issue between the TLinux kernel and IPVS and fixes the CLB health check failures in IPVS mode.	2019-07-16	-
TKE scaling groups support spot models.	When TKE creates a scaling group, users can choose spot instances and purchase pods at a certain discount . However, the system may automatically recall these pods that are sold at a discount.	2019-07-10	Spot Instances
TKE supports choosing Containerd as the container runtime component.	When Containerd serves as the container runtime component, it only runs necessary features to manage images and the container lifecycle , providing users with more stable and more resource-efficient container running infrastructures.	2019-07-05	How to Choose Containerd and Docker

June 2019

Update	Description	Date	Related Documents
The beta VPC-CNI network mode is launched.	TKE provides the VPC-CNI extended network mode, which can assign intra-VPC IP addresses to Pods in a cluster . In the VPC-CNI mode, clusters can create StatefulSet that supports fixed IP address types, and the Pod IP addresses will not change because of restart or migration.	2019-06-29	Enabling VPC-CNI When Creating the Cluster
The beta StatefulSet with fixed IP addresses is launched.	The StatefulSet with fixed IP addresses help resolve issues related to IP address changes caused by Pod restart or migration . Users can create the StatefulSet with fixed IP addresses for source IP address authorization, IP-based process review, log query based on Pod IP addresses, and other business needs to ensure the continuous and stable running of businesses.	2019-06-29	Managing StatefulSets with Static Pod IP Addresses
TKE uses the new console	In order to provide a better product user experience, TKE now uses the new Kubernetes-compatible console .	2019-06-17	The New TKE Console

version by default.			
Fixes an issue where cordoning a node while it is being created causes the process to freeze.	Fixes an issue where cordoning a node while it is being created causes the process to freeze.	2019-06-13	pr69047
Fixes an issue where too many secrets results in a pod creation failure.	Fixes an issue where too many secrets results in a pod creation failure.	2019-06-13	pr74755
The new version of the TKE international console is launched.	The new version of the TKE international console adjusts a series of functional modules and provides a native, easier-to-use platform , which helps users resolve environmental issues in development, testing, and OPS, reduce costs, and improve efficiency.	2019-06-05	TKE international console.
Managed clusters support configuring ACLs for public network access.	Users can set security group rules for managed clusters that enable public network access.	2019-06-05	TKE Security Group Settings

May 2019

Update	Description	Date	Related Documents
Nodes in a scaling group tolerate drain failures during automatic scaling in.	When scale-in conditions such as the number of idle nodes are met, the cluster automatically scales in. However, only when all pods of a node are successfully scheduled to other nodes can the pods be drained successfully and scale-in be performed successfully.	2019-05-20	Cluster Scaling

Supports registering the TKE network to CCN.	TKE allows users to register existing clusters to CCN, which can manage the container's network . After the container's network is registered, you can enable or disable its IP range routing on the CCN side to achieve interconnection between the container's cluster and the resources in CCN.	2019-05-17	Supports registering the TKE network to CCN.
TKE supports GPU virtualization.	<ul style="list-style-type: none"> Extension components support the installation and the deployment of GPU visualization components. Clusters that have deployed GPU nodes and gpu_manager can extend GPU-related settings during workload creation. 	2019-05-17	Using a GPU Node

April 2019

Update	Description	Date	Related Documents
Kubelet applies CNI mode by default	TKE Kubelet uses the VPC-CNI network mode by default.	2019-04-24	-
Docker 18.06 is launched for beta testing.	Runtime components that use Docker 18.06 can create clusters.	2019-04-22	-
The new alarm version is launched and supports all regions.	Alarms enable users to discover exceptions in TKE in a timely manner to ensure business stability and reliability. The new alarm version provides more alarm metrics . We recommend that you configure necessary alarms for all production clusters.	2019-04-22	Setting Alarms
Cluster management - Kubernetes online updates - managed master nodes	In the managed cluster mode, the Master and Etcd nodes of your Kubernetes cluster will be centrally managed and maintained by the Tencent Cloud technical team. The online updates of the Kubernetes version ensure business stability .	2019-04-11	Cluster Hosting Modes
Self-deployed clusters support	Users can query monitoring information about Master and Etcd nodes on the **Node	2019-04-11	Viewing Monitoring

Master and Etcd monitoring.	Management** page of self-deployed clusters.		Data
-----------------------------	--	--	----------------------

March 2019

Update	Description	Date	Related Documents
TKE supports Bare Metal (BM 2.0) nodes.	BM physical servers are a type of on-demand pay-as-you-go physical server rental service that provides high-performance and securely isolated physical server clusters for cloud users.	2019-03-28	-
Users can use a purchased CVM to create clusters.	Using existing CVMs to create clusters helps users reuse existing resources and reduce costs .	2019-03-28	Create a cluster using existing CVMs
Cluster auto-scaling (CA) supports disabling pod draining.	When there are multiple idle nodes in a cluster, scale-in will be triggered. CA supports disabling pod draining.	2019-03-16	Cluster Scaling
Cluster scaling groups support the scale-out of GPU nodes.	When a pod in a cluster cannot be scheduled due to a lack of available resources in the cluster, the previously set auto scale-out policy will be triggered. GPU nodes can be added during scale-out .	2019-03-12	Cluster Scaling

February 2019

Update	Description	Date	Related Documents
A new monitoring system is released.	A good monitoring environment ensures the high reliability, high availability, and high performance of Tencent Cloud TKE. You can collect monitoring data in different dimensions for different resources to	2019-02-18	Overview of Monitoring and Alarms

	quickly understand the resource usage situation and easily locate errors.		
Self-deployed clusters support Kubernetes 1.12.	Users can now create Kubernetes 1.12 self-deployed clusters in the TKE console.	2019-02-15	-
Fixes the runC vulnerability CVE-2019-5736.	The lightweight container runtime environment runc was found to have a container escape vulnerability, which allowed attackers to overwrite the host runc file (and consequently obtain host root access). This vulnerability has been fixed.	2019-02-13	[WARNING] runC Container Escape Vulnerability

January 2019

Update	Description	Date	Related Documents
Existing CLBs can be used to create Service.	Using existing CLBs to create Service can save resources and help users reduce costs.	2019-01-24	Service Management
Custom images can be used to create clusters.	TKE allows users to create custom images based on the basic image provided by TKE and use these custom images to create clusters. To enable this feature, submit a ticket to apply.	2019-01-24	Custom Images
Affinity scheduling can be set during workload creation.	YAML is delivered to the Kubernetes cluster to schedule pods in a workload. The affinity and anti-affinity mechanism of Kubernetes ensures that pods are scheduled according to specific rules.	2019-01-24	Setting the Scheduling Rule for a Workload
TKE allows multiple Services to use the same CLB instance.	Multiple Services can now use the same CLB instance to avoid additional resource costs.	2019-01-10	Service Management

December 2018

Update	Description	Date	Related Documents
TencentHub supports Helm Chart management.	Helm is a package management tool of Kubernetes. Chart is a collection of files describing Kubernetes resources. Tencent Hub provides an address for users to store Helm Charts.	2018-12-26	Overview of Helm Charts
TKE supports Helm application installation.	Helm is a packaging tool for managing Kubernetes applications. TKE has integrated Helm-related features to visually add, delete, modify, and query Helm Charts in a specified cluster.	2018-12-26	Helm Application Management
The privilege escalation vulnerability in Kubernetes is fixed.	Tencent Cloud Security Center detected that a severe privilege escalation vulnerability existed in Kubernetes (vulnerability ID: CVE-2018-1002105). This vulnerability has been fixed. Now, TKE can effectively prevent attackers from using the vulnerability to illegally access Kubernetes cluster resources, inducing privilege escalation and initiating malicious requests that ultimately jeopardize the security of the business system.	2018-12-04	[WARNING] Privilege Escalation Vulnerability in Kubernetes
Removes Kubernetes 1.7.8 as an option for creating clusters.	Users can disable the entry for creating clusters of Kubernetes 1.7.8 in the console. To enable this feature, submit a ticket to apply.	2018-12-04	-
pr71415 is merged to fix CVE-2018-1002105.	CVE-2018-1002105 is fixed and backend error responses are processed.	2018-12-04	pr71415
Kubelet disables kmem accounting to avoid kernel cgroup leakage.	Kernel cgroup leakage has an adverse impact on the system. Kubelet disables kmem accounting to avoid kernel cgroup leakage.	2018-12-04	-

November 2018

Update	Description	Date	Related Documents
The kubelet inotify leakage issue is fixed.	The kubelet inotify leakage problem is fixed.	2018-11-12	-

October 2018

Update	Description	Date	Related Documents
The beta TKE console is launched.	The new TKE console adjusts a series of feature modules to provide you with a native and easy-to-use platform. The new and old consoles are fully compatible in terms of features. Switching consoles will not affect your business. You can use the new console to continue to operate existing clusters.	2018-10-31	Notes on the New Console
Service CLB can be bound to specified nodes.	If your cluster is large, you will need to set affinity for entry-type applications to schedule them to certain nodes. You can configure the Service CLB to be bound only to specified nodes.	2018-10-31	-
Conflicts and Pod creation failures caused by the frequent updates of quota statuses by the quota controller are resolved.	Previously, if the quota controller frequently updated the quota status, conflicts and even Pod creation failures would occur. This problem has been resolved.	2018-10-22	-

September 2018

Update	Description	Date	Related Documents
The default Kubernetes version in TKE is 1.10.	When a new cluster is created, the default Kubernetes version is 1.10. However, you can change the version based on your actual needs.	2018-09-10	Creating a Cluster
BM clusters support Kubernetes 1.10.	TKE allows users to create BM clusters with Kubernetes 1.10.	2018-09-10	-
BM clusters support Ubuntu 16.04.	When TKE creates a BM cluster, the default operating system is Ubuntu 16.04.	2018-09-10	-

July 2018

Update	Description	Date	Related Documents
TKE supports the Russia and India regions.	The TKE console supports the Russia and India regions. You can go to the console to switch to and use these regions.	2018-07-30	-
TKE supports private network access to the Master node.	After the private network access entry is enabled, TKE allows private network access to the Master node.	2018-07-30	-
The open source component tencentcloud-cloud-controller-manager is released.	<p>This component is the Cloud Controller Manager implementation for TKE and allows the following features to be implemented on the Kubernetes clusters built by Tencent Cloud CVMs:</p> <ul style="list-style-type: none"> • Updates the relevant addresses information of the Kubernetes nodes. • routecontroller: creates routes within pod IP ranges in a VPC. • servicecontroller: creates a corresponding CLB when a load balancer-type service is created in a cluster. 	2018-07-30	Open Source Components

The open source component kubernetes-csi-tencentcloud is released.	This component is a plug-in for the Tencent Cloud CBS service and complies with CSI standards. It allows users to use CBS on Kubernetes clusters built by Tencent Cloud CVMs.	2018-07-30	Open Source Components
The BM cluster ingress plug-in is released.	ingress-tke-bm is the ingress controller for Tencent Cloud TKE BM clusters. This controller monitors ingress resources, creates BM CLBs, and binds them to the corresponding services.	2018-07-30	-

June 2018

Update	Description	Date	Related Documents
CCS is renamed TKE.	Tencent Kubernetes Engine (TKE) is a highly scalable and high-performance container management service. It allows you to easily run applications on a managed CVM instance cluster.	2018-06-22	Tencent Kubernetes Engine
Cluster auto scaling supports custom configurations.	TKE allows users to customize cluster scaling settings based on their actual needs, making it easier for them to configure businesses flexibly.	2018-06-22	Cluster Scaling
Node initialization supports the import of scripts.	This feature allows users to configure a node using custom data. As long as the script can be re-inputted and has a clear retry pattern, it will be used to configure the node after startup.	2018-06-22	Adding a Node

May 2018

Update	Description	Date	Related Documents
TKE supports BM clusters.	BM container clusters extend Tencent Cloud's CPM, BM Load Balancer, and other Kubernetes	2018-05-01	-

	plug-ins, providing a complete set of features such as high-efficient deployment and resource scheduling for containerized applications. This helps industries such as gaming and AI easily cope with the challenges of high-performance computing business scenarios.		
TKE supports GPU clusters.	If your business involves scenarios such as deep learning and high-performance computing , you can use the GPU feature supported by TKE, which can help you quickly use a GPU container.	2018-05-01	Enabling GPU Scheduling for a Cluster

April 2018

Update	Description	Date	Related Documents
TKE integrates the new Tencent Cloud UI version.	The new Tencent Cloud UI is elegant and easy to use, offering a better container service experience.	2018-04-01	TKE Console
TKE now supports all CVM models.	During cluster creation or node addition, the models available for selection on the TKE console are consistent with those on the CVM platform.	2018-04-01	Creating a Cluster

March 2018

Update	Description	Date	Related Documents
TKE supports the auto-scaling of services.	Horizontal Pod Autoscaler (HPA) can automatically scale the number of pods for services according to the average CPU utilization and other metrics of target pods .	2018-03-01	Basic Operations of Automatic Scaling
The TKE console interface is	The feature modules of the TKE console are adjusted.	2018-03-01	-

updated.			
----------	--	--	--

February 2018

Update	Description	Date	Related Documents
TKE supports the auto-scaling of clusters.	Cluster auto scaling adjusts the number of nodes dynamically according to resource demand : <ul style="list-style-type: none"> If pods become unschedulable due to a lack of resources, the cluster will automatically scale out. If there are enough idle nodes, the cluster will automatically scale in to reduce costs. 	2018-02-08	Cluster Scaling
TKE supports log collection.	This feature allows log files from services or specific node paths to be sent to Kafka, Elasticsearch, or CLS so that users can store and analyze them.	2018-02-06	Log Collection
TKE supports application management.	TKE supports the group management of services via applications, which significantly simplifies service management.	2018-02-06	-

December 2017

Update	Description	Date	Related Documents
Vouchers can be used to purchase cluster nodes.	TKE allows users to use vouchers in their accounts to purchase nodes.	2017-12-20	-
Empty clusters can be created.	This feature allows users to create clusters that do not contain nodes.	2017-12-20	-
Users can set the container directory and the project of	<ul style="list-style-type: none"> Container directory: users can set the directory for storing containers and images. We recommend that they be stored in data disks. 	2017-12-20	Adding an Existing Node

the resources when adding existing nodes.	<ul style="list-style-type: none"> Project: newly added resources will be automatically assigned to this project. 		
---	--	--	--

November 2017

Update	Description	Date	Related Documents
Cluster reservation policy.	Reserves system process resources such as dockerd and kubelet: when a cluster runs the retention policy, certain resources are reserved to ensure the proper running of system processes such as dockerd and kubelet.	2017-11-30	-
Cluster draining policy.	To ensure that there are sufficient resources for system processes, pods will be drained when necessary.	2017-11-30	Draining or Cordoning a Node
Dockerd log rollback.	Logs are recycled automatically to ensure that there is sufficient disk space: when log files occupy a certain amount of memory, the log rollback feature will be triggered to automatically recycle logs to ensure that there is sufficient disk space.	2017-11-30	-
Ingress forwarding rules support wildcards.	Ingress forwarding rules must comply with both the rules for the public network load balancing domain names and the Kubernetes rules for the Ingress domain names. <ul style="list-style-type: none"> They support regular expressions with a length of 1-80 characters. Other than regular expressions, they also support `a - z, 0 - 9, and -`. For domain names with wildcards, currently, only one `*` can be used in a domain name, such as `*.example.com`. 	2017-11-30	-

October 2017

Update	Description	Date	Related Documents
The beta TKE application management feature is launched.	With the rise of micro-service and Devops, users need to deploy and manage multiple services in multiple environments. TKE supports the group management of services via applications , which significantly simplifies service management.	2017-10-31	-
The multi-region deployment of Image Registry supports the new Hong Kong (China) region.	Image Registry is used to store Docker images, which are used to deploy TKE. Each image has a unique ID (the image's repository address + the image name + the image Tag). Image Registry can be deployed in multiple regions, including the Hong Kong (China) region that is now also supported.	2017-10-31	Image Registry Overview
The Tencent Cloud international console supports TKE.	The TKE international console is launched, which helps users solve environmental issues in development, testing, and OPS, reduce costs, and improve efficiency.	2017-10-31	CCS International Console

September 2017

Update	Description	Date	Related Documents
TKE Image Registry integrates access permission management.	The address format of a TKE image is as follows: <code>ccr.ccs.tencentyun.com/\${namespace}/\${name}:\${tag}</code> . The following fields are required for configuring the permissions of Image Registry: <ul style="list-style-type: none"> <code>\${namespace}</code>: the namespace of the image repository. <code>\${name}</code>: the name of the image repository. 	2017-09-26	TKE Image Registry Resource-level Permission Settings
TKE supports setting labels for services.	TKE supports setting labels for service pods. When searching services, you can filter them by label.	2017-09-26	-
Configuration	When deploying a container in a pod, users can	2017-	Configuration

items can be imported to environment variables.	import the configuration items ConfigMap and Secret to environment variables.	09-26	
Clusters support the Project attribute.	<ul style="list-style-type: none"> Clusters are not project-specific, but CVMs, CLBs, and other resources in a cluster are project-specific. Project: new resources added to the cluster will be allocated to the project. 	2017-09-26	Projects of New Resources
TKE supports the Singapore region.	TKE now supports purchasing resources and deploying businesses in the Singapore region.	2017-09-26	TKE Console

August 2017

Update	Description	Date	Related Documents
TKE integrates the alarm platform.	TKE allows users to set multi-dimensional alarms for clusters to discover cluster exceptions quickly and reduce business risks.	2017-08-23	Setting Alarms
TKE clusters support Kubernetes 1.7.	TKE allows users to create clusters with Kubernetes 1.7.	2017-08-23	-
Continuous integration and deployment based on TencentHub.	TencentHub is a management platform created by Tencent Cloud for storing R&D process files and creating DevOps workflows. TencentHub allows users to quickly and conveniently perform operations such as storage, query, and calls for files generated during the full project cycle.	2017-08-23	TencentHub Product Overview
Image Registry adds the trigger feature.	The Image Registry trigger feature allows users to trigger actions such as service update, webhook, and message push after creating an image. The trigger feature can be combined with continuous integration for continuous deployment.	2017-08-23	Trigger Overview
Image Registry	Operation logs allow users to view image	2017-08-	-

supports operation logs.	uploads and download records, which helps troubleshoot problems.	23	
Kubectl is used to operate clusters on public networks.	Kubectl is a CLI tool for Kubernetes cluster operations. You can use Kubectl to connect a local client to a TKE cluster.	2017-08-04	Connecting to a Cluster
TKE clusters integrate access permission management.	Access management is mainly used to help you securely manage and control access to resources under your Tencent Cloud accounts. Using CAM, you can create, manage, and terminate users (or user groups) and manage the use of Tencent Cloud resources through identity management and policies.	2017-08-04	TKE Cluster-level Permission

July 2017

Update	Description	Date	Related Documents
TKE supports configuration file management.	<ul style="list-style-type: none"> The configuration file management feature can help you manage the configurations of different businesses under different environments. It supports multiple versions and the YAML format. The configuration file supports multiple versions, allowing you to update and roll back applications. It also allows you to quickly import configurations, in the form of files, into containers. 	2017-07-19	-
TKE supports CI source code building.	Continuous container integration enables the automatic and manual building of container images on the Tencent TKE Platform.	2017-07-18	Image Registry Overview
Image Registry adds TencentHub images.	Image Registry allows users to view and use TencentHub images.	2017-07-18	TencentHub Product Overview

Image Registry adds "Favorite Public Images".	"Favorite Public Images" will display the images bookmarked by users, allowing users to query and use specific images.	2017-07-18	Image Registry Overview
Image Registry supports multiple namespaces.	Image Registry supports the creation of multiple namespaces. The names of namespaces are globally unique. If the namespace name you want to use is already being used by another user, try using another appropriate name.	2017-07-18	Creating a Namespace

June 2017

Update	Description	Date	Related Documents
TKE supports NFS volumes.	NFS volumes are used for the persistent storage of data that is read and written many times. They can also be used in scenarios such as big data analysis, media processing, and content management.	2017-06-24	Volume Management
TKE supports privileged containers and working directory configurations.	<ul style="list-style-type: none"> A privileged container has a certain priority. WorkingDir: specifies the current working directory. If it does not exist, one will be automatically created. If no directory is specified, the default directory when the container runs is used. If workingDir is not specified in the image or through the console, the default workingDir is `/`. 	2017-06-24	-
TKE supports cluster capacity.	A cluster is a collection of cloud resources required for running a container, including several CVMs and CLBs. You can run your applications in your cluster.	2017-06-07	Cluster Overview
TKE supports auto-formatting data disks and specifying container directories while	If the system disk capacity is small or a server with a data disk needs to format the data disk, you can set the storage directory of the containers and images.	2017-06-07	<ul style="list-style-type: none"> Creating a Cluster Adding a Node

creating/adding CVMs in container clusters.			
TKE supports service re-deployment.	Re-deployment means to re-deploy containers under a service and re-fetch images.	2017-06-07	<ul style="list-style-type: none"> • Service Management • Basic Ingress Features

April 2017

Update	Description	Date	Related Documents
TKE supports adding existing CVMs to container clusters.	TKE allows users to add existing CVMs to container clusters, which helps users reuse existing resources and effectively reduce costs.	2017-04-27	Adding an Existing Node
TKE supports the query of monitoring metrics for instances, services, and clusters.	A good monitoring environment ensures the high reliability, high availability, and high performance of Tencent Cloud TKE. You can collect monitoring data in different dimensions for different resources to quickly understand the resource usage situation and easily locate errors.	2017-04-27	Overview of Monitoring and Alarms
TKE supports viewing container logs.	By creating log collection rules, TKE can provide users with log information from within a cluster, making it easier for them to maintain and troubleshoot containers.	2017-04-27	Log Collection
The TKE remote terminal allows you to upload and download files remotely.	<ul style="list-style-type: none"> • When uploading files, you need to specify the file directory. • When downloading files, you need to specify the file path. 	2017-04-19	Basic Remote Terminal Operations
TKE supports the creation of monthly	Monthly subscription is a prepaid mode that requires customers to pay for CVMs for a period of one or multiple	2017-04-19	Adding a Node

subscription CVMs to clusters.	months/years in advance. It is cheaper than the pay-as-you-go mode and is suitable for scenarios where device demands can be predicted in advance.		
TKE supports custom security groups when creating a cluster.	If the current default security group cannot meet your business requirements, you can customize cluster security groups. For details, see Managing Security Group Rules .	2017-04-19	Creating a Cluster

March 2017

Update	Description	Date	Related Documents
TKE allows remote web terminals to log in to containers.	Remote terminals help you debug containers quickly and connect to the containers for troubleshooting. It supports file copy, paste, upload, and download operations, and helps solve the problems of long container login paths and difficult debugging.	2017-03-15	Remote Terminal Basic Operations
TKE supports third-party image creation services.	The third-party image creation service helps users deploy applications flexibly based on their actual business needs.	2017-03-15	-
TKE now supports 7-layer load balancing.	An Ingress is a collection of rules that allow access to services within a cluster. You can configure different forwarding rules to allow different URLs to access different services within the cluster.	2017-03-06	Ingress Management
Users can query monitoring information about clusters, services, and pods.	A good monitoring environment ensures the high reliability, high availability, and high performance of Tencent Cloud TKE. You can collect monitoring data in different dimensions for different resources to quickly understand the resource usage situation and easily locate errors.	2017-03-06	Overview of Monitoring and Alarms
TKE supports	TKE makes it easy for you to build, operate,	2017-03-	Overview of

native Kubernetes APIs, requesting Kubernetes certificates via Tencent Cloud APIs, and all Kubernetes features.	and manage container clusters by seamlessly utilizing Tencent Cloud computing, networking, storage, monitoring, and security capabilities. You can refer to corresponding examples in the API documentation to perform operations such as adding, deleting, modifying, and querying scaling groups, networks, nodes, and clusters.	06	APIs
---	---	----	----------------------

December 2016

Update	Description	Date	Related Documents
Cluster management.	Cluster management supports cluster addition, deletion, modification, and query, VPC-based container clusters, cross-AZ clusters, and open-source native Kubernetes APIs.	2016-12-26	Cluster Management
Service management.	Service management supports service addition, deletion, modification, and query, the creation of services via private images and official Docker images, and cross-AZ scheduling of services.	2016-12-26	Service Management
Image management.	Image management supports official Docker images, My Images, uploading and downloading private images, and official Docker image acceleration.	2016-12-26	-
Cluster monitoring and container monitoring.	TKE provides the basic monitoring feature for all clusters by default.	2016-12-26	Viewing Monitoring Data
Service creation, event updates, and rolling updates for services.	Rolling updates indicate that pods are updated one by one, which allows you to update the service without interrupting your business.	2016-12-26	-

Elastic Kubernetes Service

Last updated : 2021-06-29 17:47:29

May 2021

Update	Description	Date	Related Documents
Supports automatic allocation of EIP for Pods when they are created.	The EIP can be automatically allocated for Pod when it is created. Pod no longer strongly relies on NAT gateway for public network communication.	2021-05-28	-
Supports modifying the custom DNS of the virtual node.	Users can modify the custom DNS of the virtual node. After modification, the Pods scheduled to this virtual node will adopts this DNS configuration by default.	2021-05-26	Virtual Node Overview
Elastic cluster supports log collection via CRD configuration.	Users can use the Custom Resource Definitions (CRD) to configure log collection for the Elastic cluster. CRD is non-intrusive to Pod and supports a variety of log parsing methods. It sends standard output and file logs in the container to Tencent Cloud CLS, which provides search and analysis, visual application, log download and consumption, and other services. It is recommended to use CRD to configure log collection.	2021-05-25	Using a CRD to Configure Log Collection

March 2021

Update	Description	Date	Related Documents
EKS has connected to	Users can add Tencent Cloud tag to EKS cluster, and manage bills through tags.	2021-03-20	-

Tencent Cloud Tags.			
---------------------	--	--	--

December 2020

Update	Description	Date	Related Documents
Supports spot instance.	The spot instance costs are 20% of the original cost, which is expected to reduce business costs by 65%.	2020-12-25	-
Event dashboard was launched.	This feature supports the multi-dimensional statistics of top events, exception events, etc. and supports aggregation search and trend observation.	2020-12-08	Event Dashboard

November 2020

Update	Description	Date	Related Documents
The event storage feature was added.	Users can observe resource change and locate the problem in time.	2020-11-26	Event Storage

August 2020

Update	Description	Date	Related Documents
Pod Event completion	The Pod Event is aligned with the native K8S, making the K8S cluster running events more abundant and locating problems in Pod operation more convenient.	2020-08-15	-

July 2020

Update	Description	Date	Related Documents
Supports binding Pods with CAM roles.	Users can bind Pods with CAM roles to obtain the permission policies owned by the roles.	2020-07-22	Permission Management
Supports static IP addresses of Pods.	The IP addresses of Pods can remain unchanged when the StatefulSet/Bare Pod updates its workload.	2020-07-15	-
Supports Pod login.	Users can use the console or run <code>kubect exec -it</code> to remotely log in to a Pod.	2020-07-01	-
Supports third-party image repositories.	When creating a workload, users can select images from third-party image repositories and set the image repository access credential.	2020-07-01	-

June 2020

Update	Description	Date	Related Documents
The EKS console provides a command line window for interaction with containers.	This feature improves the user experience and helps you quickly identify issues.	2020-06-30	-
Supports updates of StatefulSets and Pods without changing their IP addresses.	This feature enhances service stability and simplifies service network management.	2020-06-30	-
LoadBalancer supports IPv6.	The service IP address supports the IPv6 network.	2020-06-30	-
EKS supports	Tencent's self-developed Star Lake servers	2020-06-	Resource

the purchase of Tencent's self-developed Star Lake servers.	provide reliable, secure, and stable high performance at low costs.	18	Specifications
EKS was fully launched.	EKS is a service mode launched by Tencent Cloud TKE that allows users to deploy workloads without having to purchase nodes.	2020-06-01	Elastic Kubernetes Service

December 2019

Update	Description	Date	Related Documents
Elastic Kubernetes Service (EKS) beta is launched.	EKS allows users to deploy workloads without having to purchase nodes. It is fully compatible with native Kubernetes and supports resource purchase and management in the native mode. Resources are billed based on the amount of resources used by the containers.	2019-12-27	Elastic Kubernetes Service

Tencent Kubernetes Engine for Edge

Last updated : 2021-02-03 11:26:21

December 2020

Update	Description	Date	Related Documents
TKE Edge opens source for SuperEdge.	SuperEdge is an edge container management system based on the native Kubernetes. Tencent Cloud has provided the edge-related source code in the TKE Edge for the SuperEdge open source project.	2020-12-19	SuperEdge GitHub

November 2020

Update	Description	Date	Related Documents
ServiceGroup feature was launched.	You can find ServiceGroup in the cluster details page.	2020-11-27	ECK Overview
The node installation script supports "check" and "clear" parameters.	<ul style="list-style-type: none"> The "check" parameter is convenient for users to use scripts to manually check where the installation requirements are not met in the node environment. The "clear" parameter is convenient for one-click cleaning of dirty data in the node, turning off the firewall, etc. 	2020-11-13	-
Edge DNS solution was launched.	The edge DNS solution will no longer occupy 53 port of the nodes.	2020-11-4	-

October 2020

Update	Description	Date	Related
--------	-------------	------	---------

			Documents
ECK supports multi-architecture hybrid management.	Users can manage the nodes in both ARM and X86 CPU architectures within a cluster at the same time.	2020-10-28	-
ECK supports edge Pod HPA.	The feature of Edge Pod HPA was launched, while the native Kubernetes HPA feature is also available on the edge.	2020-10-23	Utilizing HPA to Implement Auto Scaling of Business on TKE
Upgraded feature of using script to add node.	Users can use the same script to add self-owned nodes to the cluster multiple times (the script validity is 1 hour), making it convenient to add self-owned nodes in batches.	2020-10-22	-

September 2020

Update	Description	Date	Related Documents
Users can enable edge health feature on the console.	The "Enable Edge Health" switch is added to the basic information page of the edge cluster. Users can enable or disable this feature based on actual needs.	2020-09-28	-
ECK supports ECM security groups.	When purchasing ECM resources on the ECK console, users can select the node security group for security management.	2020-09-24	-
The permission convergence of edge node is launched.	This feature is automatically enabled, and can effectively prevent malicious users from disrupting the normal operation of the system through edge nodes.	2020-09-15	-

August 2020

--	--	--	--

Update	Description	Date	Related Documents
Edge cluster is available in Beijing region.	Users can create edge clusters in Beijing region.	2020-08-28	Creating a Cluster
The node installation script is optimized.	The node installation script can automatically obtain the default ENI.	2020-08-12	-
The Pod access mode is added.	Pods can access Apiserver in incluster mode.	2020-08-05	-

July 2020

Update	Description	Date	Related Documents
The application market, Helm Chart, and assembly line now support ECK.	Users can create apps directly or with application market, and use assembly line with ECK.	2020-07-06	<ul style="list-style-type: none"> Application Market Application Overview
Users can customize the node initialization script.	<ul style="list-style-type: none"> Node initialization operations include mounting a data disk, creating directories, and so on. The script is run only once during node initialization. 	2020-07-01	-
Users can obtain the metrics of all pods in the cluster via the apiserver.	<ul style="list-style-type: none"> Users can obtain the metrics of all cluster pods in the cluster(if any) by requesting the api-server. In such cases, a monitoring component should be deployed in the cluster. 	2020-07-01	-

June 2020

--	--	--	--

Update	Description	Date	Related Documents
ECK supports GPU.	Currently, ECK supports the NVIDIA Tesla (T4, P40, M40, P4, and V100) GPU models.	2020-06-30	-
The ECK image acceleration feature is launched for beta testing.	The launch time of big-image pods is shortened by 30%, and the public traffic consumption for pulling images is reduced to 1/n (n: the number of nodes in the same LAN) of the original traffic consumption.	2020-06-30	-
ECK supports custom parameters.	<ul style="list-style-type: none"> • Supports custom node initialization scripts. • Supports custom container directories. • Supports custom node max-pod. 	2020-06-30	-
ECK supports Kubernetes v1.18.2.	Supports the creation of Kubernetes v1.18.2 clusters.	2020-06-01	Creating Edge Clusters

March 2020

Update	Description	Date	Related Documents
ECK is launched.	ECK is a container system that manages edge cloud resources from the central cloud. You can use it to manage distributed nodes in the same cluster across multiple regions. ECK is fully compatible with native Kubernetes, supports one-click app delivery, and comes with edge autonomy and distributed health checks.	2020-03-25	Edge Cloud Kubernetes Engine

Announcements

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 Kueconfig 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` .

kubectl

- Removed the deprecated `--include-uninitialized` parameter.
- `kubectl` and `k8s.io/client-go` no longer use `http://localhost:8080` as the default apiserver address.
- `kubectl run` supports pod creation and no longer supports using the deprecated generator to create other types of resources.
- Removed the deprecated `kubectl 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 `kubectl 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.
- [Kubectl 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-`.
- **kubectl**
 - 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 `kubectl 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><none></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, 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 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 kubectl 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 CBS intree 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 (&dxgt;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.