


Tencent Kubernetes Engine

Release Notes and Announcements

Product Documentation



Copyright Notice

©2013-2025 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 General Cluster Updates (2024)
- TKE General Cluster Updates (2023)
- TKE General Cluster Updates
- TKE Serverless Cluster Updates (2024)
- TKE Serverless Cluster Updates (2023)
- TKE Serverless Cluster Updates

Announcements

- Security Vulnerability Fix Description
 - CVE-2024-21626 Vulnerability Fix Description
 - OOM Causes docker to Stop and Not Restart for Repair
 - Runc Vulnerability (CVE-2021-30465) Fix Description
- Host Operation System Release for Super Node Pods (Mitigated NodeLost Issue)
- TKE Native Node Sub-product Name Change Notice
- Announcement on Authentication Upgrade of Some TKE APIs
- Discontinuing Update of NginxIngress Addon
- qGPU Service Adjustment
- Version Upgrade of Master Add-On of TKE Managed Cluster
- Upgrading tke-monitor-agent
- Instructions on Cluster Resource Quota Adjustment
- Decommissioning Kubernetes Version
- Deactivation of Scaling Group Feature
- Notice on TPS Discontinuation on May 16, 2022 at 10:00 (UTC +8)
- Basic Monitoring Architecture Upgrade
- Starting Charging on Managed Clusters
- Instructions on Stopping Delivering the Kubeconfig File to Nodes

Release Notes

- Kubernetes Version Maintenance
 - TKE Kubernetes Version Maintenance Mechanism
 - Update Notes of TKE Kubernetes Major Versions
 - TKE Kubernetes Revision Version History
- Add-On Version Maintenance Description
- CLS Add-on Version Description
- Change Records of VPC-CNI Component

Runtime Version Maintenance Description
TKE Runtime Revision Version History

Release Notes and Announcements

Release Notes

TKE General Cluster Updates (2024)

Last updated : 2024-05-27 15:44:20

January 2024

Update	Description	Reference
Support for hostname customization for native nodes	Hostname can be customized at the node pool level. The node hostname will be consistent with the host name.	Creating Native Nodes
Launch of the policy management capability for TKE	Protection against cascade deletion of clusters, nodes, namespaces, and CRD is supported, addressing instability caused by accidental resource deletion.	Policy Management
Launch of the 2022 version of Cloud API for TKE	Parameter information of different types of nodes (general/native/super/registered nodes) can be queried.	Container Service Cloud API Update History
New models for general nodes	Users can create 4090 consumer-grade card models GC49 and high-performance models HCCPNV5x.	Supported CVM Models for General Nodes

TKE General Cluster Updates (2023)

Last updated : 2024-05-27 15:44:41

December 2023

Update	Description	Reference
New model for TKE native nodes	TKE native nodes support the M3 models.	Native Node Parameters
Introduction of the industry's first memory compression capability for container scenarios	"Cold" memory is recycled and compressed based on the "WuJing" kernel. This realizes control of the pod-level compression scope, helping customers solve cluster memory bottlenecks.	-
Support for users to use Nginx-ingress-controller to create ingress after TKE registered nodes are enabled	After registered nodes are enabled in a TKE cluster, ingress of the Nginx-Ingress-controller type is no longer blocked. Users can follow the documentation to create Nginx-ingress-controller instances through standard scheduling policies, and then use them according to the standard process.	Installing Nginx-ingress Instance
Optimization of the node lifecycle management experience	Users can adjust the expected number of instances when deleting a node.	Deleting Native Nodes

November 2023

Update	Description	Reference
Support for importing Tencent Cloud SSM credentials using ExternalSecretOperator	ExternalSecretOperator can import key credentials, which are uniformly stored and managed in Tencent Cloud Secrets Manager (SSM), into a cluster in the form of Kubernetes native secret objects and achieve automatic synchronization of key data. This enables the unified storage and management of key lifecycles by SSM.	-
Support for TKE registered	Registered nodes are fully launched with a new	Registered Node

nodes to access to public network nodes	public network version, allowing users to easily manage IDC nodes in a public cloud environment (without dedicated lines), quickly build Kubernetes clusters, and reduce costs.	Overview
---	---	--------------------------

October 2023

Update	Description	Reference
Support for CFS Turbo in cluster storage	CFS Turbo is supported for cluster storage, suitable for business with high throughput and hybrid workloads.	CFSTURBO-CSI Instructions
Upgrade of the native node dedicated scheduler	The native node dedicated scheduler is upgraded to enhance stability with high availability deployment, support configuring independent drain stop watermarks, and support scheduling event logs.	Instructions
New models for TKE native nodes	TKE native nodes support the M6 and GT4 models.	Native Node Overview

September 2023

Update	Description	Reference
GooseFS-Lite supported as a COS mounting method	GooseFS-Lite is supported as a COS mounting method, enabling higher large file read and write speeds that are not limited by local disk performance.	Using COS
Verification for deprecated API versions in pre-upgrade cluster check	Pre-upgrade cluster check includes verification for deprecated API versions to avoid compatibility issues during upgrade, ensuring cluster stability and reliability.	Upgrading Clusters
Support for updating Management parameters of legacy nodes	Native node pools support in-place updates of Kubelet, Nameservers, Hosts, and KernelArgs (kernel) parameters through Management.	Management Parameters
Support for wildcard domain name configuration for	Forwarding rules and certificates of the wildcard domain name type can be configured for CLB-type	Ingress Certificate Configuration

Ingress	Ingress.	
Local binding and weighting enabled by default in the externalTrafficPolicy local mode of TKE services	In the externalTrafficPolicy local mode of TKE services, CLB by default binds to the node where the pod is located and weights according to the number of pods on the node.	Basic Features
Graceful Shutdown and Graceful Deletion enabled by default in the CLB Direct Connect mode	Graceful Shutdown and Graceful Deletion are enabled by default in the CLB Direct Connect mode, enhancing network stability.	Graceful Ingress Shutdown

August 2023

Update	Description	Reference
Optimization of TKE native node scheduler capabilities	The native node dedicated scheduler supports configuring independent drain stop watermarks.	Instructions
New log collection capabilities	Based on log collection rules, users can configure multiple paths as container file paths and node file paths.	-
Premium SSD CBS, supporting setting additional performance	Users can create a Premium SSD CBS StorageClass, which supports setting additional performance for higher performance.	Managing CBS Templates by Using a StorageClass
Optimization of Nginx Ingress capabilities	Nginx Ingress supports configuring forwarding rules of the HTTP and HTTPS types through the console.	Basic Ingress Features

July 2023

Update	Description	Reference
Optimization of general node experience	Error messages in the initialization process of general nodes are more precise for error locating, supporting the display of specific reasons for initialization errors.	Creating a Node Pool
	Cluster upgrade prompts are more readable, displaying the version limit details of dependent	View a Node Pool

	components required for the upgrade.	
New models for TKE general nodes	Premium CVM models of the A800/H800 type are supported.	Supported CVM Models for General Nodes
New capabilities for TKE native nodes	Enhanced stability: The native node dedicated scheduler is fixed for higher stability.	Native Node Overview
	New capability: Modification of security groups in native node pools takes effect on legacy nodes.	
	New capability: Native node pools support modifying subnets.	
	New capability: Native nodes support setting Tencent Cloud tags.	
	Experience optimization: Monitoring button display is added for native node pools.	

June 2023

Update	Description	Reference
KMS encryption for Secrets in managed/independent clusters	TKE clusters can use KMS to encrypt and decrypt various sensitive data in Secrets, such as database usernames, passwords, certificates, OAuth Tokens, and SSH KEYS, enhancing data security.	Using KMS for Kubernetes Data Source Encryption
Cluster upgrade capability iteration	TKE Kubernetes version supports upgrading from 1.22 to 1.24.	Update Notes of TKE Kubernetes Major Versions
Update on node/container login method for TKE	New capability: Node/Container login for TKE supports OrcaTerm.	-
New capabilities for TKE native nodes	New capability: Native nodes support console-initiated restart.	Native Node Overview
	New capability: Trace scenario detection is added for RuntimeUnhealthy, supporting runtime	

	restart for kubelet status exceptions caused by this scenario.	
	Experience optimization: Pod running details are displayed during node draining.	
	Experience optimization: For prepaid node pools, users can search for target nodes and check node cordon status when adjusting the node quantity.	
	Experience optimization: Prepaid nodes support bulk renewal on the billing platform.	

May 2023

Update	Description	Reference
Connection to a cluster through Cloud Shell	TKE has integrated Tencent Cloud Cloud Shell , allowing users to connect to a cluster with one click from the Tencent Cloud console and flexibly manage the cluster via kubectl.	Connect to a Cluster
Ingress certificate modification and upgrade	All existing certificates for the TKE service can be modified with one click on the certificate platform.	-
New capabilities for TKE Insight	Users can view the cost trend and forecasting of TKE clusters, including the cost distribution of workloads and namespaces.	-
New capabilities for TKE native nodes	New models: The S6, GN10X, and GN10Xp models are supported for the China site.	Native Node Overview
	New drivers: GPU models support the installation of 450/470/515 drivers.	
	New capability: Enhanced security capability can.	
	New capability: Native node pools support modifying security groups and data disks.	
Strong verification for service activation provided for using CFS	Before the CFS component is installed in TKE, strong verification and guidance for service activation is provided to avoid scenarios where	CFS Instructions

	services have not been activated but resources are in use through weak API validation, resulting in billing. This enhances the linkage and billing accuracy between TKE and CFS.	
Support for search within the namespace of log collection rules	When configuring a log source for log collection rules, users can select a namespace in a drop-down menu and search one or more namespaces. In this way, users can quickly find and select namespaces from which logs need to be collected.	-

April 2023

Update	Description	Reference
Support for cross-region shipment of cluster audit logs and event logs	To facilitate multi-region centralized management, cluster audit logs and event logs can be shipped to regions outside of the cluster's location for centralized and unified management.	Collecting Container Logs to CLS
Upgrade of access layer components	The mode of Service component lifecycle management is upgraded to reduce anomalies during Service modification. For certain misconfigurations of Service/Ingress, TKE Service/Ingress components will opt to skip rather than stop running, preventing the risk of traffic interruption due to Service/Ingress anomalies.	-
Launch of TKE Backup Center	TKE Backup Center provides an integrated solution for backup, recovery, and migration of containerized applications.	Backup Center

March 2023

Update	Description	Reference
Nginx Ingress independent of EIP	The Nginx Ingress controller no longer depends on EIP, reducing user costs.	Installing Nginx-ingress Instance
Support for formatted mounting of data disks and	Data disk: Native node pools support formatted mounting of Premium SSD CBS, with a	Creating Native Nodes

addition of alternative models with the same specifications during the creation of native node pools

recommended data disk size of at least 50 GB. Alternative model: Native node pools support configuring multiple alternative models with the same specifications, which can be set on the node pool details page.

February 2023

Update	Description	Reference
Launch of the runtime upgrade capability	The cluster node upgrade entry also supports the minor version updates of docker and containerd components.	-
Support for metadata customization by using cluster log collection rules	When configuring rules for collecting logs and shipping them to Tencent Cloud Log Service (CLS) in the TKE console, you can customize metadata to reduce unnecessary storage resources.	Using CRD to Configure Log Collection

References

For more information on TKE general cluster updates, see [TKE General Cluster Updates \(Before 2023\)](#).

TKE General Cluster Updates

Last updated : 2025-01-09 17:58:44

Note :

For more information on TKE General cluster updates, see [TKE General Cluster Updates \(2023\)](#).

December 2022

Update	Description	Date	Related Documents
Supports upgrading a Kubernetes version to 1.22	You can upgrade a cluster Kubernetes version from 1.20 to 1.22.	2022-12-08	Update Notes of TKE Kubernetes Major Versions

November 2022

Update	Description	Date	Related Documents
Launches the imc-operator component	You can cache super nodes in clusters on images with CRD.	2022-11-10	Image Cache
Enhanced the dedicated scheduler for native nodes	Supports virtually expanding the capacity of native nodes and configuring scheduling and running thresholds for the nodes.	2022-11-01	-

September 2022

Update	Description	Date	Related Documents
Launched the Ceberus enhanced add-on	Signature images in the TCR repository can be verified for trustworthiness, reducing the security risks of invalid images being deployed in the container environment.	2022-09-29	-
Enabled all custom kubelet parameters	It provides an entry for you to modify parameters.	2022-09-29	Custom Kubernetes Component Launch Parameters
Defined the error codes of Service/Ingress events	It helps quickly troubleshoot business exceptions and provides solutions.	2022-09-28	-
Supported scheduling Pods in kube-system to monthly subscribed super nodes	It helps reduce resource use costs.	2022-09-26	Pod Schedulable to Super Node
Integrated the basic monitoring add-on `tke-monitor-agent` into the add-on management module for maintenance	You can upgrade the add-on in the console.	2022-09-22	Upgrading tke-monitor-agent
Launched the `SecurityGroupPolicy` enhanced add-on	You can bind a Pod not matching a policy to a security group to control the inbound and outbound network traffic of the Pod.	2022-09-08	-

Launched the in-place Pod configuration adjustment capabilities for TKE native nodes	You can directly modify the Request/Limit values of CPU and memory without restarting a Pod, which helps handle traffic surges and reduce business costs.	2022-09-08	-
--	---	------------	---

August 2022

Update	Description	Date	Related Documents
Launched Request smart recommendation for native nodes	It recommends the Request/Limit values of resources at the container level for Kubernetes workloads to reduce resource waste.	2022-08-22	-
Launched the dedicated scheduler for native nodes	It virtually expands the capacity of native nodes, solving the problem of resource occupation and low utilization.	2022-08-19	-
Supported automatic update for the image cache feature	Automatic update is supported for the image cache feature. After it is enabled, a cache update will be triggered when the updated image is uploaded to TCR, eliminating the need to recreate an image cache.	2022-08-08	Image Cache
Supported edit operations for the image cache feature	The image cache feature allows you to edit security groups, images, cache sizes, and expiration policies.	2022-08-08	Image Cache

Released the brand-new native node	A native node is a brand-new node type of TKE. It supports declarative APIs and adopts the FinOps concept, facilitating cost reduction and efficiency improvement for Tencent Cloud resources through the "visual resource management dashboard", "smart request recommendation", and "dynamic dedicated scheduler".	2022-08-04	-
Supported managing image caches through CRD	Image caches can be managed through CRD on super nodes. Currently, management can be performed in the console and through TencentCloud API. With the image caching feature, Pods can be started within seconds.	2022-08-04	Image Cache
Supported declaring the size of resources required by a Pod.	You can declare the size of the system disk required by a Pod through an annotation on super nodes. If the size exceeds 20 GiB, the excessive part will be billed based on the pay-as-you-go published price of CBS Premium Cloud Storage, meeting your requirements for system disk resources in special scenarios.	2022-08-03	Annotations
Added third-party nodes	The TencentOS Server 2.4 (TK4) operating system is supported, which is compatible with the CentOS 7 user mode. The entry to create Cilium-Overlay clusters is added in the TKE console for easier use. In Cilium-Overlay network mode, the check for a conflict between the IP range of a container	2022-08-01	Cilium-Overlay Mode

and that of another container in another cluster within the same VPC can be ignored, thereby solving your problem of IP resource insufficiency.

July 2022

Update	Description	Date	Related Documents
Released all the super nodes	A super node is a brand-new node form of TKE, which supports custom node size and flexible configuration adjustment. The monthly subscription mode is more cost-effective, while the pay-as-you-go mode eliminates the need to reserve resource buffer, contributing to your cost reduction and efficiency improvement. Node management and resource operations are simplified, so that you can concentrate on the business.	2022-07-26	Super Node Overview
Supported productized Nginx-ingress	The Nginx-ingress add-on can be quickly deployed for out-of-the-box Nginx services. By default, CM and CLS are integrated to meet the observability requirement.	2022-07-25	Installing Nginx-ingress Instance
Quickly cleared invalid RBAC resources	RBAC resource objects can be quickly cleared. De-registered Tencent Cloud accounts can be automatically detected, which is suitable for scenarios where cluster permissions need to be reclaimed	2022-07-15	Clearing De-registered Tencent Cloud Account Resources

	after personnel turnover.		
Released the deletion protection feature for the node pool	This feature is suitable for node resource protection, fixing the issue where node resources were batched released due to maloperations.	2022-07-13	Node Pool Overview
Supported associating with a TMP instance	A cluster can be associated with a TMP instance during creation, simplifying the association process.	2022-07-06	Quick Migration from TPS to TMP

June 2022

Update	Description	Date	Related Documents
TMP provides metrics recommended by experts	Simplifies the metric selection.	2022-06-21	-
qGPU supports viewing GPU card resources with CRD	Unifies and streamlines the GPU resource view.	2022-06-07	-
qGPU supports scheduling at both the node and GPU card levels	Realizes refined scheduling of GPU resources and solves fragmentation problems.	2022-06-07	-
Unschedulable nodes	Solves the issue of unbalanced traffic load.	2022-06-01	Service Backend Selection

can be bound to the CLB backend			
---------------------------------	--	--	--

May 2022

Update	Description	Date	Related Documents
TKE provides audit/log/event APIs	You can install relevant add-ons in TKE clusters by calling the audit/log/event APIs.	2022-05-18	-

April 2022

Update	Description	Date	Related Documents
Kubernetes versions are discontinued	Kubernetes v1.14 and earlier versions cannot be created from May 20, 2022 (UTC +8).	2022-04-26	Discontinuing Kubernetes v1.14 and Earlier Versions
The scaling group service is discontinued	Starting from June 6, 2022 (UTC +8), no more TKE scaling groups can be created in the console and by using the API. The feature will be officially discontinued from June 13, 2022 (UTC +8). Existing scaling groups will be migrated to node pools automatically.	2022-04-26	Deactivation of Scaling Group Feature
Updates CFS dynamic persistence storage	TKE supports sharing a CFS storage instance when dynamically	2022-04-18	CFS-CSI

	create a PVC.		
Launches TMP	Supports cross-region monitoring in multiple VPCs and supports checking multiple monitoring instances on a single Grafana instance.	2022-04-01	TMP Overview
TMP Overview	You can have enhanced cloud native monitoring capabilities with quick migration.	2022-04-01	Quick Migration of TPS Instances to TMP

March 2022

Update	Description	Date	Related Documents
Starts charging on managed clusters	For TKE managed clusters created before March 21, 2022 10:00 (UTC +8), select a specification for them in the console before April 1, 2022 10:00 (UTC +8). If you do not select a specification, they will be billed on the basis of the recommended specification.	2022-03-21	Starting Charging on Managed Clusters
Backend architecture upgrade of TKE basic monitoring	Backend architecture upgrade of TKE basic monitoring	2022-03-21	Description of tke-monitor-agent
Supports upgrading of log collection add-on in	Supports upgrading of log collection add-on in the Ops center. Please upgrade the add-on to the latest	2022-03-15	Log Add-On Version Upgrade

the Ops center	version.		
Optimized user experience in the TKE console	<p>On the Pods management details page, you can search by Pod status in the search box.</p> <p>In the Log column of cluster workload details page, you can search by name in the drop-down list.</p>	2022-03-14	Cluster Management

February 2022

Update	Description	Date	Related Documents
Optimized the search feature in the console	You can filter clusters by status, and prompts are available when you search by node/node pool label.	2022-02-27	Cluster Overview
Upgraded the add-on architecture	Add-ons can be installed via cloud APIs after the add-on architecture upgrade.	2022-02-26	Overview of APIs
Optimized CFS Turbo add-on	Fixed known mounting issues.	2022-02-25	Static Mounting of CFS-Turbo for TKE Clusters
Optimized operations of log collection rules in the console	When Pod Labels are specified as the log source, you can add != label filter and optimize selection for namespaces in container standard output (multiple namespaces can be chosen and excluded).	2022-02-15	-

January 2022

Update	Description	Date	Related Documents
Configures the httpHeaders for health check in the TKE console	Health Check is a service provided by TKE to check the status and running of resources in the cluster.	2022-01-17	Health Check
Optimized operations of node pool in the console	You can set cloud tags and configure instance creation policies in the node pool.	2022-01-15	Node Pool Overview
Updates CFS add-on	CFS supports using V3 and V4 protocols for mounting.	2022-01-13	CFS Instructions
Updates CBS-CSI add-on	Cloud disks dynamically created by CBS-CSI inherit the cluster cloud tags.	2022-01-12	CBS-CSI
Updates Request Recommendation	Supports downloading of recommended Request values.	2022-01-10	Resource Utilization Analysis and Optimization Suggestions
Supports connecting to the Pods in TKE clusters via the API gateway	Supports exposing services via API gateways. It can be specified when you create an Ingress.	2022-01-03	API Gateway TKE Tunnel Configuration

December 2021

Update	Description	Date	Related Documents

Updates CFS-CSI add-on	Cloud tags can be added when CFS instances are created dynamically through StorageClass.	2021-12-25	Managing CFS Templates by Using a StorageClass
------------------------	--	------------	--

November 2021

Update	Description	Date	Related Documents
Enhances qGPU features	qGPU supports multiple GPU cards in one container. A container can be bound with multiple entire cards or a scorecard. In addition, GPU monitoring is supported. qGPU supports GPU monitoring. It supports GPU card/Pod/container utilization monitoring. qGPU supports TKE BM clusters and supports Ampere GPU cards. qGPU supports the online/offline hybrid deployment feature and supports the native priority scheduling capabilities of online inference and offline training.	2021-11-05	-
Feature updates of CFS-CSI add-on	The CFS-CSI add-on in clusters of v1.20 supports reading the fsGroup configuration of security context in workloads.	2021-11-02	CFS-CSI

October 2021

Update	Description	Date	Related Documents
Launches Request Recommendation	Request Recommendation is a resource recommendation feature launched by TKE. It recommends Request and Limit values for resources at the container level in TKE and TKE Serverless clusters.	2021-10-27	-
Launches the registered	Registered cluster is a new cluster	2021-10-26	-

cluster feature	type of TKE. It allows you to register Kubernetes clusters in your local infrastructure or those of other cloud vendors with TKE for unified management.		
Launches the cloud-native AI service	Cloud-native AI is a complete solution for AI scenarios launched by TKE, which has the characteristics of modularity, low coupling, and high scalability.	2021-10-18	-

September 2021

Update	Description	Date	Related Documents
Launches qGPU add-on	qGPU is a GPU sharing technique launched by Tencent Cloud. It supports sharing of GPU cards among multiple containers and provides the capacity to isolate the vRAM and computing resource among containers.	2021-09-16	qGPU

August 2021

Update	Description	Date	Related Documents
Uses pre-request script to add existing nodes	When adding existing nodes, you can use pre-request scripts via API. The script is executed before the K8s node is initialized. You can customize the configuration such as data disk mounting.	2021-08-01	-

July 2021

Update	Description	Date	Related Documents

Supports specifying the model when creating a GPU CVM in the console	Chooses the GPU driver version, CUDA version and cuDNN version in the console.	2021-07-13	Using a GPU Node
Supports MIG	Once enabled, you can improve the GPU utilization in scenarios where multiple jobs are running in parallel.	2021-07-13	Using a GPU Node
Enhances TKE console	Supports formatting and mounting of multiple data disks when adding existing nodes.	2021-07-05	Adding an Existing Node

June 2021

Update	Description	Date	Related Documents
Introduced the feature of adding external nodes	External nodes allow users to add non-Tencent Cloud CVMs to a TKE cluster. Users provide computing resources and TKE manages cluster lifecycle.	2021-06-28	-

May 2021

Update	Description	Date	Related Documents
Improves cloud native monitoring	Optimizes data collecting configuration. Displays status of collection targets. Optimizes interaction process. Tests the collection target.	2021-05-28	Log Collection
Launches OLM add-on	The OLM add-on helps users automatically install and upgrade Operator, and manage its lifecycle.	2021-05-28	OLM
Launches HPC add-on	Developed by Tencent Cloud, HPC can help 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

Enhances TKE console	Selects OS upon node creation. Modifies desired number of nodes while scaling up a node pool. Searches workloads by labels.	2021-05-20	Adding a Node
----------------------	---	------------	-------------------------------

April 2021

Update	Description	Date	Related Documents
Enhances TKE console	StatefulSet and DaemonSet can be redeployed with a few steps. 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 readinessProbe supports configuration and naming of a port. Namespace supports the selection of "All namespaces" and can be searched by keyword.	2021-04-30	StatefulSet Management DaemonSet Management Secret Management Setting the Health Check for a Workload Namespaces
Enhances log collection capability	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
		2021-03-31	-

The beta ARM cluster starts	The beta of ARM cluster starts. To join the beta, submit a ticket .		
Enhances TKE console	<p>Kubernetes objects support batch input by pasting multiple key-value pairs.</p> <p>When configMap is used as an environment variable, all keys can be selected with one-click.</p> <p>Secret can be modified through console.</p> <p>ConfigMap can be managed through console. You can manually add or upload file to add key-value pairs.</p> <p>SubPathExpr can be configured through console.</p> <p>CronJob supports displaying the generated active Jobs. User can customize the “Completed Jobs Retained” and “Failed Jobs Retained”, pause the generation of a scheduled Job, resume the paused Job, and manually generate a new Job.</p>	2021-03-31	Kubernetes Object Secret Management ConfigMap Management Instructions for Other Storage Volumes CronJob Management
Enhances node search capability	The nodes can be searched in batch by label or IP address.	2021-03-19	Node Overview

February 2021

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

which provides more flexible mount settings.

January 2021

Update	Description	Date	Related Documents
Master version upgrade supports minor version	This feature provides a more flexible version upgrade mechanism.	2021-01-14	Upgrading the Master Kubernetes Version

December 2020

Update	Description	Date	Related Documents
Node pool supports configuration of removal protection	Removal protection is used to prevent the nodes of important application from being scaled in or being removed from the node pool during manually size adjustment.	2020-12-30	-
Node pool supports configuration of multiple alternative models that have the same specification	This feature can reduce the risk of scale-out failure caused by sold out single model resource.	2020-12-29	Adjusting a Node Pool
Launches Descheduler add-on	Based on the actual node loads, this add-on supports automatic rescheduling of marked services on high-load nodes to maintain the cluster load balance.	2020-12-25	DeScheduler
Fully launches Nginx-ingress add-on	The issue of nginx-ingress-controller toleration scheduling is fixed. Nginx-Ingress UI experience is 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

Launches CBS-CSI add-on	<p>CBS-CSI add-on supports:</p> <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
Enhances TKE console	<p>The console supports Resource Quota, which can be used to configure the resource quotas and default resource request values for the namespace.</p> <p>The ConfigMap can be generated based on the imported file.</p>	2020-12-09	-
Launches NetworkPolicy add-on	This add-on supports automatic synchronization of NetworkPolicy to make the network isolation policy effective.	2020-12-03	Network Policy

November 2020

Update	Description	Date	Related Documents
Launches beta for new TKE network solution	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	-
Launches beta for productized Nginx-Ingress	TKE extends and maintains native Nginx-ingress to help you quickly deploy and build a traffic access gateway in the production environment, and provide complete Nginx-ingress lifecycle management, automated cloud	2020-11-26	Nginx-ingress

	native monitoring, CLS and supporting OPS capabilities.		
Launches event dashboard	This feature implements the aggregation search and trend observation of top events and exception events.	2020-11-26	Event Dashboard
Launches audit dashboard	This feature implements the aggregation search and direct observation of cluster global, nodes, K8s objects and other important operations.	2020-11-26	Audit Dashboard
The node pool and cluster operating system can be modified	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
Adds DynamicScheduler add-on	This add-on performs scheduling based on actual node loads, so as to prevent traffic hotspots.	2020-11-21	DynamicScheduler

October 2020

Update	Description	Date	Related Documents
TPS supports using edge cluster as monitoring object to access the monitoring instance.	TPS supports the monitoring of edge clusters and the management of multiple clusters across VPCs.	2020-10-30	PROM Instance Management
TPS alarm policy supports webhook configuration.	The alarm policy supports webhook configuration, which enables users to troubleshoot abnormal services in time and improve service stability.	2020-10-30	Alarm Configurations
TKE node pool adds the capability of viewing scaling log	This feature helps users to more easily observe the change of node number in the node pool as well as the trigger cause and result of scaling, improving the node pool observability.	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
Launches DNSAutoscaler add-on	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
Launches beta for cloud native ETCD	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	-
Quick add-on configuration is available when creating the cluster	You can easily and quickly configure the required add-ons for the cluster.	2020-09-15	Creating a Cluster

August 2020

Update	Description	Date	Related Documents
Optimizes the monitoring capability of TPS	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	Data Collection Configurations Alarm Configurations
The alarm module of the cloud native monitoring service was upgraded.	It can be associated with a local Alertmanager add-on. It supports managing PROM instance rules with CRD.	2020-08-31	Alarm Configurations

TKE launches NodeProblemDetectorPlus add-on	It supports configuring node self-healing policy on the basis of existing detection feature.	2020-08-25	Node-Problem-Detector-Plus
TKE launches in-place major-version upgrade capabilities	The in-place major-version upgrade feature supports major-version upgrade without node reinstallation.	2020-08-25	-
Fully launches TKE add-ons	The add-on feature enables users to install or uninstall multiple advanced add-ons for clusters.	2020-08-25	Add-on Overview
Fully launches TKE Kubernetes 1.18 version	Allows users to create clusters of the Kubernetes 1.18 version and upgrade clusters to the 1.18 version.	2020-08-24	-

July 2020

Update	Description	Date	Related Documents
Optimizes capabilities of storage plug-ins	The TKE console supports PV creation without specifying StorageClass. Users can set and mount COS sub-directories.	2020-07-28	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 TKE Serverless clusters.	2020-07-15	-
The collection configuration and alarm configuration of TPS are implemented through products	Three configuration modes are supported: service monitor, pod monitor, and raw job. Alarm history rewinding is supported.	2020-07-15	-

Launches beta for RBAC-based permission control with finer granularity	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
Launches beta for IPVS-bpf mode	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 CLB-to-Pod Direct Access Mode
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
Comprehensively upgrades The TKE app market	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	-
Optimizes the cluster deletion feature	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	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	Cluster nodes can be filtered based on lock status.	2020-04-22	Node Overview

multidimensional node filtering and node list export	Cluster nodes can be filtered based on CVM attributes such as node status and IP address. Cluster nodes can be exported in batches.		
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	-
The TKE beta version supports the node pool feature	The node pool feature can be used in the following scenarios: 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.	2020-04-10	-
TKE removes Kubernetes 1.8 as an option	TKE no longer supports creating clusters using Kubernetes 1.8.	2020-04-03	Creating a Cluster
Upgrades self-deployed cluster master	You can now use the TKE console to perform rolling updates of Kubernetes masters on self-deployed clusters.	2020-04-02	Upgrading 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 BM and big data models	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 a Cluster

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, submit a ticket .	2020-02-17	Cluster Scaling
TKE fully launches Kubernetes 1.16 and passes conformance verification	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	Creating a Cluster Upgrading a Cluster

January 2020

Update	Description	Date	Related Documents
TKE allows users to create clusters using a	The template-based cluster creation feature provides multiple templates	2020-01-12	Creating a Cluster

cluster template	for creating managed clusters, self-deployed clusters, and serverless 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.		
------------------	--	--	--

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	-
Launches beta for TKE Kubernetes 1.16	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	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 a Node
TKE nodes support the in-place rolling updates	Nodes in in-place updates support the rolling update mode.	2019-12-03	Upgrading a Cluster

of minor Kubernetes versions	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 .		
------------------------------	---	--	--

November 2019

Update	Description	Date	Related Documents
Launches beta for TKE custom Hostname	The TKE custom Hostname feature provides the following advantages: 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.	2019-11-15	Adding a Node
Releases TKE Ingress performance optimization	TKE Ingress performance is optimized to better serve users. 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	Basic Ingress Features

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	-
Runtime component Containerd supports the GPU model	The TKE runtime component Containerd supports the GPU model. When users need to create a GPU application in a cluster, they can choose Containerd as the runtime component .	2019-10-17	How to Choose Containerd and Docker
Launches beta for rolling Kubernetes reinstallation and upgrade of TKE nodes	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	Date	Related Documents

<p>Related APIs of the TKE cluster scaling group have been updated to API 3.0</p>	<p>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.</p>	2019-09-12	CreateClusterAsGroup
<p>TKE Kubernetes 1.14 is fully launched and has passed conformance verification</p>	<p>TKE Kubernetes 1.14 is fully launched and has passed conformance verification to ensure that the latest Kubernetes version is available.</p>	2019-09-07	Conformance Verification
<p>TKE supports the Tencent Cloud tag, allowing authorization by tag</p>	<p>If the Tencent Cloud tag is added to a cluster 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.</p>	2019-09-06	-
<p>The default instance type for created LoadBalancer-type services is CLB</p>	<p>When TKE creates a LoadBalancer-type service, the default instance type is CLB. This instance type covers all features of a conventional CLB. The CLB instance type supports the TCP, UDP, HTTP, and HTTPS protocols. It provides flexible forwarding capabilities based on domain names and URLs.</p>	2019-09-06	Instance Types
<p>TKE self-deployed clusters support the separate viewing of Master and Etcd nodes</p>	<p>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.</p>	2019-09-05	-

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 creation and ensures the continuous and stable running of businesses.	2019-08-23	Creating a Cluster
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.	August 8, 2019	Basic Ingress Features
Launches beta for TKE Kubernetes 1.14	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	2019-08-04	API 3.0

	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.		
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	Fixed the compatibility issue between the TLinux kernel and IPVS and fixed the CLB health check failures in IPVS mode.	2019-07-16	-
Fixed the compatibility issue between the TLinux kernel and IPVS and fixed the CLB health check failures in IPVS mode.	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	When Containerd serves as the container runtime component, it only	2019-07-05	How to Choose Containerd and Docker

container runtime component	runs necessary features to manage images and the container lifecycle , providing users with more stable and more resource-efficient container running infrastructures.		
-----------------------------	---	--	--

June 2019

Update	Description	Date	Related Documents
Launches beta for VPC-CNI network mode	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	-
Launches beta for StatefulSet with fixed IP addresses	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 version by default	In order to provide a better product user experience, TKE now uses the new Kubernetes-compatible console .	2019-06-17	The New TKE Console
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	Fixes an issue where too many	2019-06-13	pr74755

many secrets results in a pod creation failure	secrets results in a pod creation failure.		
Launches the new version of the TKE international console	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	Register Container Clusters with CCN
TKE supports GPU virtualization	Extension components support the installation and the deployment of GPU visualization components.	2019-05-17	Using a GPU Node

	Clusters that have deployed GPU nodes and gpu_manager can extend GPU-related settings during workload creation.		
--	---	--	--

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	-
Launches beta for Docker 18.06	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 Master and Etcd monitoring	Users can query monitoring information about Master and Etcd nodes on the Node Management page of self-deployed clusters.	2019-04-11	Viewing Monitoring Data

March 2019

Update	Description	Date	Related Documents
TKE supports Bare	BM physical servers are a type of on-	2019-03-28	-

Metal (BM 2.0) nodes	demand pay-as-you-go physical server rental service that provides high-performance and securely isolated physical server clusters for cloud users.		
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	Creating a Cluster
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
Release a new monitoring system	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.	2019-02-18	Overview of Monitoring and Alarms
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	2019-02-13	-

obtain host root access). This vulnerability has been fixed.		
--	--	--

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	-
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	-
TKE supports Helm application installation	Helm is a packaging tool for managing Kubernetes	2018-12-26	Helm Application Management

	applications. TKE has integrated Helm-related features to visually add, delete, modify, and query Helm Charts in a specified cluster.		
Fixes the privilege escalation vulnerability in Kubernetes	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	-
Removes Kubernetes 1.7.8 as an option for creating clusters	The entry for creating clusters of Kubernetes 1.7.8 in the console is disabled. To enable it, submit a ticket .	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
Fixes the kubelet inotify leakage	The kubelet inotify leakage problem is fixed.	2018-11-12	-

October 2018

--	--	--	--

Update	Description	Date	Related Documents
Launches beta for the new TKE console	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	-
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 India regions	The TKE console supports the 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	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: 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 plugins, 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.	2018-05-01	-
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 updated	The feature modules of the TKE console are adjusted.	2018-03-01	-

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: 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 the resources when adding existing nodes	Container directory: users can set the directory for storing containers and images. We recommend that they be stored in data disks. Project: newly added resources will be automatically assigned to this project.	2017-12-20	Adding an Existing Node

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. They support regular expressions with a length of 1-80 characters. Other than regular expressions, they also support <code>a - z</code> , <code>0 - 9</code> , and <code>-</code> For domain names with wildcards, currently, only one <code>*</code> can be used in a domain name, such as <code>*.example.com</code> .	2017-11-30	-

October 2017

Update	Description	Date	Related Documents
Launches beta for TKE application management feature	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	Image Registry is used to store Docker images, which are used to deploy TKE. Each image has a	2017-10-31	Image Registry Overview

new Hong Kong (China) region	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.		
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	TKE International Console

September 2017

Update	Description	Date	Re Do
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: \${namespace}: the namespace of the image repository. \${name}: the name of the image repository.	2017-09-26	TK Re Re lev Pe Se
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 items can be imported to environment variables	When deploying a container in a pod, users can import the configuration items ConfigMap and Secret to environment variables.	2017-09-26	-
Clusters support the Project attribute	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	-
TKE supports the	TKE now supports purchasing resources and deploying businesses in the Singapore region.	2017-09-26	TK Co

Singapore region			
------------------	--	--	--

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	-
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	-
Image Registry supports operation logs	Operation logs allow users to view image uploads and download records, which helps troubleshoot problems.	2017-08-23	-
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	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 adds TencentHub images	Image Registry allows users to view and use TencentHub images.	2017-07-18	-
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.	2017-07-18	Creating a Namespace

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.

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	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 <code>/</code> .	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 creating/adding CVMs in container clusters	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	Creating a Cluster Adding a Node
TKE supports service re-deployment	Re-deployment means to re-deploy containers under a service and re-fetch images.	2017-06-07	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 supports uploading and downloading files remotely	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 custom security groups when creating a cluster	If the current default security group cannot meet your business requirements, you can customize cluster security groups by referring to "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	Basic Remote Terminal 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 native Kubernetes APIs, requesting Kubernetes certificates via Tencent Cloud APIs, and all Kubernetes features	TKE makes it easy for you to build, operate, 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.	2017-03-06	Overview of 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	-
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	-
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	-

TKE Serverless Cluster Updates (2024)

Last updated : 2024-05-31 09:42:06

January 2024

Update	Description	Reference
Support for pods on super nodes to obtain eks-id	Obtaining eks-id in yaml and reconciling through eks-id are supported.	Annotation

TKE Serverless Cluster Updates (2023)

Last updated : 2024-05-31 09:42:26

September 2023

Update	Description	Reference
Release of the cloud API for reserved coupon	The cloud API for reserved coupon is supported, allowing users to call the Cloud API to achieve reserved coupon related features.	API Overview

August 2023

Update	Description	Reference
Support for image cache editing	Users can edit existing image caches.	Image Cache

July 2023

Update	Description	Reference
Support for high availability of pods on super nodes	Pods on super nodes support forced dispersion and soft dispersion by host/switch/rack.	Super Node Overview

June 2023

Update	Description	Reference
Support for SSD disks for pods on super nodes	New capabilities: Users can specify SSD disks for pods on super nodes.	Annotation
Image cache CRD update	Experience optimization: Users can specify subnet information for image caches.	Image Cache

April 2023

Update	Description	Reference
Support for hot update of security groups for pods on super nodes	Users do not need to update pods when modifying their security groups through annotation.	Annotation
Launch of the A10 vGPU model for super nodes	The A10 vGPU model is launched to support more granular A10 specifications.	Resource Specifications

February 2023

Update	Description	Reference
Support for downloading the deduction history of reserved instances for super nodes	Users can download the deduction history of reserved instances for super nodes in a specified period of time.	-
Support for creating and editing super node resources on Terraform	Users can use Terraform to manage super node resources: creating and editing super node resources.	-

References

For more information on TKE serverless cluster updates, see [TKE Serverless Cluster Updates \(Before 2023\)](#).

TKE Serverless Cluster Updates

Last updated : 2023-03-31 10:54:55

说明：

For more information on TKE Serverless cluster updates, see [TKE Serverless Cluster Updates \(2023\)](#)。

December 2022

Update	Description	Date	Related Documents
Supports running DaemonSet	Supports adding DaemonSet to super nodes to align with DS scenario product features.	2022-12-01	TKE Serverless Clusters-related

November 2022

Update	Description	Date	Related Documents
Launches Window containers	Supports the serverless Windows container service.	2022-11-01	TKE Serverless Cluster

October 2022

Update	Description	Date	Related Documents
Supports static IP addresses for Pods	Pods running on super nodes can use static IP addresses.	2022-10-01	Pod Schedulable to Super Node

August 2022

Update	Description	Date	Related Documents
Enhances Pod security	Supports configuring and managing Pod security groups by label to implement fine-grained secure access control.	2022-08-11	Scheduling Pod to Super Node

June 2022

Update	Description	Date	Related Documents
Launches disk cleanup capability	You can enable the capabilities such as cleaning up images, cleaning up exited containers and restarting the Pods with high disk usage through annotations.	2022-06-23	Annotations

April 2022

Update	Description	Date	Related Documents
Supports upgrading TKE clusters to TKE Serverless clusters through tools	Supports upgrading existing managed TKE clusters to TKE Serverless clusters through tools.	2022-04-20	Guide on Migrating Resources in a TKE Managed Cluster to a Serverless Cluster
Launches the feature of restarting Pods automatically when the system disk load reaches the threshold for TKE Serverless clusters	This feature was launched on April 27, 2022 (UTC +8). The Pods created before that date must be rebuilt to enable the feature.	2022-04-27	TKE Serverless Clusters-related
Launches the batch Job scheduling feature for TKE Serverless clusters	TKE Serverless supports batch scheduling of Jobs. You can submit a ticket to apply for it.	2022-04-18	Job Management
Supports restricting the number of schedulable GPU Pods on super nodes	Supports restricting the number of schedulable GPU Pods on super nodes. You can submit a ticket to apply for it.	2022-04-11	Notes on Pods Scheduled to a Super Node

March 2022

Update	Description	Date	Related Documents
Supports automatic specification upgrade and retry if the current	Dynamic specification upgrade can be enabled in TKE Serverless	2022-03-31	Creating a Cluster

specification has no resources available during TKE Serverless cluster resource creation	clusters.		
Adds reservation resource overview	Reservation resource overview is added for you to view the current numbers of Pods and corresponding reservations in real time.	2022-03-31	-
Supports restricting the number of schedulable GPU Pods on super nodes	You can set annotations for a super node to control the number of GPUs on it.	2022-03-31	Annotation
Supports the HPC component	TKE Serverless componentization supports scheduled scaling.	2022-03-21	HPC
Updates the image cache feature	Image cache allows you to use an annotation to customize the time for delayed CBS cloud disk termination.	2022-03-15	Image Cache
Supports global annotation settings in TKE Serverless clusters	You can use <code>ConfigMap</code> to set global annotations.	2022-03-12	Annotations

December 2021

Update	Description	Date	Related Documents
Supports the image cache feature for TKE Serverless	You can use image cache to accelerate image pull during instance creation so as to expedite instance startup.	2021-12-20	Image Cache

November 2021

Update	Description	Date	Related Documents
Supports discounts on resources with small specifications for TKE Serverless	TKE Serverless offers up to 65% discount for resources with small specifications such as 1-core 2 GiB and 2-core 4 GiB.	2021-11-01	Product Pricing

July 2021

Update	Description	Date	Related Documents
Supports creation of container instance	Container instance is a service model launched by Elastic Kubernetes Service that allows users to deploy container applications without the need to purchase servers or deploy K8s clusters.	2021-07-14	-

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 a super node	Users can modify the custom DNS of the super node. After modification, the Pods scheduled to this super node will adopt this DNS configuration by default.	2021-05-26	Supernodes Overview
TKE Serverless cluster supports log collection via CRD configuration	Users can use the Custom Resource Definitions (CRD) to configure log collection for the TKE Serverless 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
EKS supports the purchase of reserved instances	EKS reserved instances can be purchased by month. Purchased EKS reserved instances can be used to offset the cost of containers of the same location, resource type, and resource specifications.	2021-05-17	-

March 2021

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

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	-
Launches event dashboard	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 is available	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 the purchase of Tencent's self-developed Star Lake servers	Tencent's self-developed Star Lake servers provide reliable, secure, and stable high performance at low costs.	2020-06-18	Resource Specifications
Fully releases EKS	EKS is a service mode launched by Tencent Cloud TKE that allows users	2020-06-01	Elastic Kubernetes Service

	to deploy workloads without having to purchase nodes.		
--	---	--	--

December 2019

Update	Description	Date	Related Documents
Launches beta for Elastic Kubernetes Service (EKS)	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

Announcements

Security Vulnerability Fix Description

CVE-2024-21626 Vulnerability Fix Description

Last updated : 2024-05-27 16:04:09

Vulnerability Details

Agent: runc

Vulnerability Name: runc Container Escape Vulnerability

CVE No.: CVE-2024-21626

Impact: This vulnerability could compromise the isolation layer between the container and the host operating system, allowing attackers to access host files or execute binary programs without authorization. For details, see [Community Explanation](#) and [Fix Suggestions](#).

Scope of Impact

Runtime engines that use runc versions between 1.0.0-rc.93 and 1.1.11.

Note:

Preliminary verification indicates that exploiting this vulnerability requires kernel support for the openat2 system call (kernel version 5.6 and later). The affected node operating system distributions currently identified include Ubuntu 22.04 LTS and Red Hat Enterprise Linux 8.6. This vulnerability has not been reproduced on other operating system distributions. We are continuously following up.

Fix Method

1. Incremental TKE clusters and nodes created after February 3, 2024, are not affected by this vulnerability.
2. For legacy nodes, you can fix the vulnerability by executing the following command on the machine or replace the nodes:

```
wget http://static.ccs.tencentyun.com/fix-cve-2024-21626.tar.gz && tar -zxf fix-cve
```


OOM Causes docker to Stop and Not Restart for Repair

Last updated : 2024-05-27 16:08:53

Problem Description

In Docker 19 and later versions, when excessive system memory usage causes containerd to encounter an Out of Memory (OOM) situation, it may result in Docker stopping and not restarting automatically. This issue can be reproduced by executing the `pkill -9 containerd; systemctl is-active dockerd containerd` command. At this point, dockerd will be stopped by systemd.

The most severe impact could be general nodes becoming NotReady after OOM, and issues with the primary node in an independent cluster could trigger an avalanche effect.

Problem Analysis

Initially, the Docker community set the relationship between docker and containerd as dockerd.service BindsTo containerd.service. This causes systemd to actively stop dockerd when containerd is forcibly terminated by the `kill -9` command. Even if Restart is set in Docker, recovery is not possible. For more information, see:

<https://github.com/moby/moby/commit/a985655ac4eb6c5b60b5eab8d8d09a487e353e1d>

<https://github.com/moby/moby/commit/009639e0f46d270309821998836fb664046b3e07>

Fixing Incremental Nodes

Incremental nodes were fixed on April 20, 2023.

Fixing Legacy Nodes

For legacy nodes, you can fix the problem with the following script:

```
#!/bin/bash
insert_if_absent() {
    line="${1}"
    lead="$(echo "${line}" | cut -f1 -d=)""="
```

```
if ! grep "^${lead}" /usr/lib/systemd/system/containerd.service >
/dev/null 2>&1; then
    sed -i "/^ExecStart=/a${line}"
/usr/lib/systemd/system/containerd.service
fi
}

insert_if_absent OOMScoreAdjust=-999
insert_if_absent RestartSec=5
insert_if_absent Restart=always

sed -i '/BindsTo/d' /usr/lib/systemd/system/dockerd.service
sed -i 's/^Wants.*Wants\\=network-online.target containerd.service/'
/usr/lib/systemd/system/dockerd.service

systemctl daemon-reload
```

You can verify whether the issue of Docker not being able to restart after containerd is forcibly terminated has been successfully fixed by executing the command below. Additionally, you can further verify by executing the `docker run` command.

```
pkill -9 containerd;systemctl is-active dockerd containerd
```

Runc Vulnerability (CVE-2021-30465) Fix

Description

Last updated : 2024-12-13 15:46:01

Vulnerability Details

Add-on: runC

Vulnerability Name: runC path traversal vulnerability

CVE No.: CVE-2021-30465

Fix Policy: Upgrade runC to 1.0.0-rc95 or later.

Fix Progress

1. The vulnerability was fixed for **incremental nodes** in September 2021 in TKE.
2. For **legacy nodes**, see the following upgrade script and fix the vulnerability during off-peak hours to avoid affecting the business stability.

Note:

Upgrading the runC add-on will not restart the business Pod.

```
#!/bin/bash

util::is_docker() {
  if command -v docker 1>/dev/null 2>&1; then
    RUNTIME="docker"
    return 0
  else
    return 1
  fi
}

wget http://static.ccs.tencentyun.com/docker-19.03.9-install-1.2.tgz

tar -zxf docker-19.03.9-install-1.2.tgz

if ! docker-19.03/bin/runc --version; then
  echo "unmatch libseccomp version"
  # Get OS distribution
  OS_RELEASE="$(. /etc/os-release && echo "$ID")"
  OS_VERSION="$(. /etc/os-release && echo "$VERSION_ID")"

  if [ "ubuntu" = "${OS_RELEASE}" ]; then
    apt-get install libseccomp2
  else
```

```
    yum install -y libseccomp
fi
fi

if ! docker-19.03/bin/runc --version; then
    echo "bad libseccomp version"
    exit 1;
fi

if util::is_docker; then
    cp docker-19.03/bin/runc /usr/bin/docker-runc
    cp docker-19.03/bin/runc /usr/bin/runc
else
    cp docker-19.03/bin/runc /usr/local/sbin/runc
fi

rm -r    docker-19.03
rm      docker-19.03.9-install-1.2.tgz
```

Host Operation System Release for Super Node Pods (Mitigated NodeLost Issue)

Last updated : 2025-02-11 15:42:18

To enrich the features, Tencent Kubernetes Engine (TKE) plans to release a new host OS version for super nodes on March 10, 2025 (UTC+8). The incremental pods on super nodes will gradually use the new OS version.

Time

The operating system will gradually apply to the incremental pods starting from March 10, 2025 (UTC+8).

Impact Scope

This change will apply to the pods that are rebuilt or newly added on super nodes in TKE Standard Cluster and TKE Serverless Cluster. The pods currently running on super nodes will not be affected.

About the New Version of Host OS

The new version of the host OS is based on Tencent's commercial TencentOS distribution, with an upgraded kernel on the original OS, supporting a wider range of storage types. It greatly optimizes memory usage of system components, reducing their memory occupation.

Version Check Methods

After logging into the pod, execute the following command to view the host OS kernel version(5.4.241-1.0017):

```
# uname -a
Linux <pod-name> 5.4.241-1.0017.10.eks.8.t14
```

New Features

Mitigated unresponsiveness of the machine that is caused by memory inadequacy. It will trigger cgroup OOM kill timely to prevent the machine from entering the unresponsive state (nodelost state) when the memory in the pod is nearly reached.

Added the support for Cloud Object Storage (COS). For the operation guide, see [Using COS](#).

Added the support for Ceph File System (CephFS).

Optimized Features

The kernel version was upgraded from 5.4.119-1.0009 to 5.4.241-1.0017.

containerd was upgraded from v1.4.3 to v1.6.9.

The GPU driver was upgraded from 470 to 535, supporting CUDA version upgrade from up to 11.4 to 12.2.

The systemd software package version was updated from 239 to 255.

The dbus software package version was updated from 1.12.8 to 1.14.8.

Other Notes

To try the operating system in advance, you can [submit a ticket](#) for application.

To upgrade a running pod to the new OS version, you need to rebuild the pod first.

If you have specified a specific OS version through the annotation (`eks.tke.cloud.tencent.com/eklet-version`), you need to remove that annotation to make the new OS version take effect.

TKE Native Node Sub-product Name Change Notice

Last updated : 2024-11-27 10:44:51

Background

Native node is a new node type introduced for the Kubernetes environments by the Tencent Cloud TKE team, leveraging Tencent Cloud's extensive container operations experience to provide users with native, high-stability, and responsive Kubernetes node management capabilities. This product currently supports creation in TKE standard clusters and supports more than 20 models. For details, please refer to [Native Node Parameters](#).

During the initial launch of the product, some models were previously known as "HostingNode". To provide a consistent user experience, TKE "HostingNode" will be officially renamed as "NativeNode". For the names of sub-products affected by this change, please refer to the table below.

Billing Sub-product Name Before Renaming	Billing Sub-product Name After Renaming
TKE_HostingNode_C3	TKE_NativeNode_C3
TKE_HostingNode_C4	TKE_NativeNode_C4
TKE_HostingNode_GN7	TKE_NativeNode_GN7
TKE_HostingNode_IT5	TKE_NativeNode_IT5
TKE_HostingNode_M5	TKE_NativeNode_M5
TKE_HostingNode_MA3	TKE_NativeNode_MA3
TKE_HostingNode_SA3	TKE_NativeNode_SA3
TKE_HostingNode_SA2	TKE_NativeNode_SA2
TKE_HostingNode_S5	TKE_NativeNode_S5
TKE_HostingNode_S2	TKE_NativeNode_S2

Change Contents

This product renaming does not involve any changes to the product console, TencentCloud API, CAM, or other product features, nor does it affect billing amounts. It will not impact the actual use of users' businesses.

Billing Information

If you are using "TKE_HostingNode_xx", the product name field in the billing information will be changed from "TKE_HostingNode_xx" to "TKE_NativeNode_xx". This change does not affect the billing amounts or other information.

Announcement on Authentication Upgrade of Some TKE APIs

Last updated : 2024-10-12 16:04:32

To provide a more comprehensive authentication feature, Tencent Kubernetes Engine (TKE) plans to integrate Cloud Access Management (CAM) authentication for some APIs on October 16, 2024. If your sub-accounts still need access to the corresponding APIs, refer to [Creating Custom Policies via Policy Generator](#); otherwise, your sub-accounts will not be able to access the corresponding APIs. Thank you for your trust and support for Tencent Cloud. If you encounter any problems while using our cloud products, please [contact us](#).

Common Authorization Methods

Method 1: Creating a Custom Policy

Implementation Method

Create policies for different sub-accounts by following the principle of least privilege and bind them to sub-users.

Applicable Scenario

Strict permission control is required, and the operation scope of each sub-user needs to be refined as needed.

Directions

1. On the [Policy](#) page in the CAM console, click **Create Custom Policy** in the upper left corner.
2. In the pop-up Select Creation Method window, click **Create by Policy Syntax** to enter the Select Policy Template page.
3. On the Select Policy Template page, you can enter keywords to search. For example, if the template type is set to All Templates, you can enter the keyword 'a' to select the AdministratorAccess template.
4. Click **Next** to enter the Edit Policy page.
5. On the Edit Policy page, confirm the policy name and content, and then click **Complete** to finish creating a custom policy by policy syntax. The default policy name and content are automatically generated by the console. The default policy name is **policygen**, with a suffix number generated according to the creation date.
6. On the [Policy](#) page in the CAM console, locate the created policy and click Associate Users/Groups/Roles in the Operation column.
7. In the Associate Users/User Groups/Roles window, select the user, user group, or role you want to associate and click Confirm to complete the association.

Method 2: Binding a Preset Policy

Implementation Method

Bind the QcloudTKEInnerFullAccess preset policy, which contains all TKE API operation permissions, to sub-accounts.

Applicable Scenario

Business personnel have operation permissions for all feature modules.

Directions

1. On the [Policy](#) page in the CAM console, select the policy type TKE, locate QcloudTKEInnerFullAccess, and click Associate Users/Groups/Roles in the **Operation** column.
2. In the **Associate Users/User Groups/Roles** window, select the user, user group, or role you want to associate and click **Confirm** to complete the association.

List of APIs with Added Authentication

API Name	Description
DescribeInstancesVersion	Queries versions of worker nodes
ModifyClusterNodePool	Edits a node pool
DescribeInstanceTypesForDirectENI	Queries the models available in independent ENI mode
DescribeClusterStatus	Queries the cluster status list
DescribeClusterVirtualNodePools	Queries the list of virtual node pools
DescribeClusterVirtualNode	Queries the list of virtual nodes
GetPodSecurityGroups	Queries Pod security groups
GetPodById	Queries Pod information
DescribeNodePools	Queries the list of TKE node pools
DescribeSuperNodes	Queries the list of super nodes
DescribeRegions	Queries the list of regions
DescribeMachineConfiguration	Queries machine configuration information
DescribeVpcCniPodLimits	Queries the maximum number of Pods in the VPC-CNI mode supported by a model

DescribeClusterSecurityGroup	Queries cluster security group information
DescribeClusterSchedulerPolicy	Queries cluster scheduling policies
DescribeFlowIdStatus	Queries the status of enabling port flow in a cluster
DescribeClusterAvailableExtraArgs	Queries available custom parameters for a cluster
DescribeEksAlarmSetting	Queries whether monitoring and alarms are set in the cluster list
DescribeClusterRoutes	Queries cluster routes
DescribeClusterRouteTables	Queries cluster routing tables
GetPods	Queries information of all Pods within a cluster
DescribeClusterResourceLabels	Queries the list of cluster resource tags
DescribeClusterExtraArgs	Queries custom parameters for a cluster
DescribeClusterCIDRs	Queries the Classless Inter-Domain Routing (CIDR) used under the cluster VPC
DescribeHealthCheckTemplate	Queries health check policy templates
DescribeClusterNodePools	Queries the node pool list
DescribeClusterNodePoolDetail	Queries detailed information of a node pool
DescribeImageRegistryCredentials	Queries image repository credentials
DescribeImageCaches	Queries image cache information
DescribeOldStaticIPSubnets	Queries the list of subnets exclusive to legacy fixed IP clusters
DescribeEnableVpcCniProgress	Queries the async task progress of enabling VPC-CNI mode
DescribeAvailableInstanceConfigInfo	Queries available model information
GetZoneResource	Queries availability zone resources
DescribeMetaFeatureProgress	Queries the progress of enabling cross-tenant features
DescribeRouteTableConflicts	Queries the list of routing table conflicts

DescribeBatchModifyTagsStatus	Queries the status of modifying tags in batches
GetMostSuitableImageCache	Queries the matching image cache
DescribeLogSwitches	Queries log switches
DescribeEKSContainerInstanceEvent	Queries container instance events
DescribeEksContainerInstanceLog	Queries container instance logs
DescribeContainerLog	Queries container logs in a container group
DescribeExistedInstances	Queries whether existing nodes can be added to a cluster
DescribeProductVersions	Queries the list of application product versions
DescribeProductVersionDetails	Queries details of application product versions
DescribeProducts	Queries the list of application products
DescribeRIPodDetail	Queries reserved coupons and Pod quantity
DescribeZoneInstanceConfigInfos	Queries model configurations of native nodes
DescribeClusterMachines	Queries the native node list
DescribeCcnRoutes	Queries Cloud Connect Network (CCN) routes
DescribeCcnInstances	Queries CCN instances
DescribeZoneDiskQuota	Queries cloud disk quotas in each availability zone
GetAccountType	Queries account types
GetPodSpecQuota	Queries quotas for specified Pod specifications
DescribeClusterControllers	Queries Kubernetes controller status
DescribeMasterLog	Queries primary log switch information
GetPodChargeInfo	Queries the billing information of a Pod
CheckPodRetain	Queries whether a Pod is retained
GetPod	Queries Pod information
RunClusterInspections	Triggers cluster inspection
CreateUpdateNodeUnit	Creates or updates a NodeUnit

CreateCluster	Creates a cluster
CreateClusterRoute	Creates cluster routes
CreateClusterRouteTable	Creates a cluster routing table
CreateClusterAsGroup	Creates a cluster scaling group
CreateClusterNodePool	Creates a node pool
CreateImageRegistryCredential	Creates image repository credentials
CreateImageCache	Creates an image cache
CreateCLSLogConfig	Creates log collection rules
CreateClusterVirtualNode	Creates a virtual node
CreateClusterVirtualNodePool	Creates a virtual node pool
SwitchParameterCreateNativeNode	Creates native node pool parameter conversion
CreatePod	Creates a Pod
UninstallLogAgent	Uninstalls the Cloud Log Service (CLS) log collection component from a TKE cluster
AddClusterCIDR	Adds ClusterCIDR to a cluster
ModifyClusterInspection	Updates cluster inspection configurations
UpdateImageRegistryCredential	Updates image repository credentials
UpdateImageCache	Updates image cache
DisableVpcCniNetworkType	Disables additional VPC-CNI network capabilities
DisableMasterLog	Disables primary log collection
StopEks	Shuts down
DescribeClusterInspectionOverviews	Queries the cluster health check list
DescribeClusterInspectionReport	Queries cluster inspection report details
GetUpgradeInstanceProgress	Obtains the current progress of node upgrade
GetPodSpecification	Obtains Pod specifications

DescribeClusterCreateProgress	Queries the cluster creation progress
DescribeClusterMetricsData	Queries the monitoring data of a cluster
DescribeClusterResourceDetails	Queries detailed information of the resources corresponding to a cluster
DescribeClusterEndpoints	Queries cluster access addresses
DescribeClusterLevelAttribute	Queries the cluster scale
DescribeClusterStaticInstaller	Queries the static installation scripts of a cluster node
DescribeClusterInstancelds	Queries the cluster node ID list
GetQuota	Obtains the quota for creating Pods in a cluster
DescribeAvailableClusterVersion	Queries all versions that a cluster can upgrade to
DescribeClustersResourceStatus	Queries the resource status within a cluster
DescribeClusterNamespaces	Queries the namespace list of a cluster
ListExpiredClusterAuth	Obtains expired permission information within a cluster
DescribeQuota	Queries the cluster quota
DescribeClusterPods	Queries the container list of a cluster
DescribeClusterMasterDiff	Queries the differences between control plane components before and after cluster upgrade
DescribeClusterAuthorizationMode	Queries the cluster authorization mode
GetTkeAppUpgradeInfo	Obtains the upgrade information of components under a cluster
ListClusterCertificates	Obtains the cluster certificate list
DescribeClusterCommonNames	Queries the CommonNames for sub-accounts to access a cluster
DescribeClusterServices	Queries the cluster service list
DescribeInstanceCreateProgress	Queries the node creation progress
DescribeSpotPodDetail	Queries the spot Pod details

DescribeImages	Queries image information
DescribeOSImageId	Queries the image ID
DescribeEdgeRegion	Queries the region list
DescribeClusterGlobalStatistics	Queries the number of clusters in all regions
DescribeNodeParamUpdateProcess	Queries the parameters for rolling updates of native nodes
GetSubnetVip	Obtains the reserved VIP within a subnet
GetSubnetResource	Obtains subnet resources
DescribeECMInstances	Queries the Edge Computing Machine (ECM) instance information
DescribeIPAMD	Queries information about eniipamd components
DescribeGPUInfo	Queries GPU information
DescribeNodeGroup	Queries NodeGroup information
DescribeNodeUnit	Queries NodeUnit information
DescribeOSImages	Queries OS aggregation information
ListEKSPods	Obtains detailed Pod information
DescribeClusterKubeconfig	Queries the cluster kubeconfig file
GetTkeAppChartList	Obtains the list of applications supported by TKE
GetTkeAppDiff	Obtains Diff information for TKEApp
DescribeTKEEdgeAppChartList	Queries the list of applications supported by TKEEdge
DescribeSupportedRuntime	Queries available runtime versions based on Kubernetes versions
DescribeVersions	Queries cluster version information
CreateClusterRelease	Creates applications for a cluster
DescribeClusterHealthyStatus	Queries cluster health status
UpgradeClusterInstances	Upgrades cluster nodes
DescribeUpgradeClusterProgress	Queries cluster upgrade progress

ModifyClusterUpgradingState	Modifies cluster upgrade status
DescribeClusterAuthStatus	Queries whether a cluster automatically cleans up expired sub-accounts
DescribeClusterInspections	Queries cluster inspection overview
CheckClusterRuntimeConfig	Checks cluster runtime configurations
EnableMetaFeature	Enables cluster meta features
CheckClusterCIDR	Checks cluster CIDR for conflicts
CheckClusterHostName	Checks hostnames of cluster nodes
CheckClusterImage	Checks whether an image can be set as a cluster image
CheckUseTKE	Checks whether you are a new TKE user
CheckMigrateVm	Checks instance routing
AddNodeToNodePool	Moves nodes in a cluster to a node pool
SetNodePoolNodeProtection	Enables removal protection for the nodes in a node pool
OpUpgradeClusterInstances	Controls the node upgrade task status
CheckInstancesUpgradeAble	Checks whether the nodes can be upgraded
UpdateClusterAuthStatus	Enables or disables automatic cleanup of expired sub-accounts for a cluster
EnableMasterLog	Enables primary log collection
EnableVpcCniNetworkType	Enables the VPC-CNI network mode
RotateClusterToken	Rotates cluster static tokens
DescribeClusterRuntimeVersion	Queries cluster runtime version information
DrainClusterVirtualNode	Drains a virtual node
DeleteClusterInstances	Deletes nodes from a cluster
DeleteClusterRoute	Deletes a cluster route
DeleteClusterRouteTable	Deletes a cluster routing table

DeleteEKSNode	Deletes a node
DeleteClusterNodePool	Deletes a node pool
DeleteImageRegistryCredentials	Deletes image repository credentials
DeleteImageCaches	Deletes image caches
DeleteClusterVirtualNode	Deletes a virtual node
DeleteClusterVirtualNodePool	Deletes a virtual node pool
DeleteRetainPod	Deletes a reserved Pod
DeleteNodeUnit	Deletes a NodeUnit
DeletePod	Deletes a Pod
UpdateClusterVersion	Upgrades a cluster
UpgradeLogAgent	Upgrades a log component
UpgradeClusterRelease	Upgrades an application
GrantCodingClusterRoleBinding	Authorizes Coding service role permissions
InquirePriceRefundNativeNode	Inquires about refund for native nodes
ValidateClusterAddVirtualNodeConditions	Verifies the conditions for adding virtual nodes to a cluster
CheckComponentVersion	Verifies the cluster component versions
CheckSubaccountAuthority	Verifies sub-account permissions
ModifyPodSecurityGroups	Modifies Pod security groups
ModifySuperNodeTag	Modifies the tag of a super node
ModifyExternalNodePool	Modifies the third-party node pool
ModifyClusterTags	Modifies cluster tags
ModifyClusterImage	Modifies a cluster image
ModifyNodePoolInstanceTypes	Modifies the model of instances in a node pool
ModifyNodePoolDesiredCapacityAboutAsg	Modifies the expected number of instances for the scaling group associated with a node pool

ModifyClusterVirtualNodePool	Modifies a virtual node pool
ModifyReservedInstanceName	Modifies the name of a reserved coupon
SwitchParameterRenewNativeNode	Converts parameters for renewing native node pools
InquirePriceRenewNativeNode	Inquires about renewal for native nodes
ClearExpiredClusterAuth	Cleans up expired permission information within a cluster with one click
RemoveNodeFromNodePool	Removes a node from a node pool
InstallLogAgent	Installs the CLS log collection component in a TKE cluster
AddVpcCniSubnets	Adds cluster subnets
DescribeAggregationData	Queries aggregated data
ListHousekeeperRegion	Queries the Housekeeper region list
InquirePriceHousekeeper	Inquires about Housekeeper
GetPrice	Inquires about TKE Serverless
DisableEksEventPersistence	Disables the event persistence feature for a TKE Serverless cluster
EnableEksEventPersistence	Enables the event persistence feature for a TKE Serverless cluster
EnableMetaFeatureForEks	Enables the cross-tenant ENI for a TKE Serverless cluster
UpdateMetaFeatureForEks	Updates cross-tenant ENI configurations for a TKE Serverless cluster
UpdateMetaFeature	Updates cross-tenant ENI global configurations for a TKE Serverless cluster
AttachEKSDisks	Binds TKE Serverless cloud disks
DescribeEKSDisks	Queries the TKE Serverless cloud disk list
QueryEKSDiskTask	Queries TKE Serverless cloud disk tasks
GetEksClusterUsed	Queries elastic cluster quotas and the number of created clusters

DescribeEKSClusterStatus	Queries the elastic container cluster status
DescribeEKSInstances	Queries elastic container instances
ListEKSZone	Queries available availability zones for TKE Serverless in the corresponding region
DescribeEksMetaFeatureProgress	Queries the cross-tenant ENI enabling status for a TKE Serverless cluster
DescribeEksLogSwitches	Queries log switch information of a TKE Serverless cluster
GetEKSClusterResources	Queries TKE Serverless cluster resources
ListEKSRegion	Queries available regions for TKE Serverless
CreateEKSDisk	Creates a TKE Serverless CBS disk
CreateEKSInstances	Creates elastic container instances
CreateEksLogConfig	Creates log collection rules for a TKE Serverless cluster
UpdateEKSClusterKubeconfig	Updates kubeconfig information for a TKE Serverless cluster
DisableEksAudit	Disables elastic cluster audit
DescribeEKSClusterCredential	Queries authentication information of a TKE Serverless cluster
DescribeEKSClusterAuthorizationMode	Queries the authorization mode of a TKE Serverless cluster
GetEKSApplUpgradeInfo	Obtains component upgrade information under a TKE Serverless cluster
ListEKSClusterCertificates	Obtains the certificate list of a TKE Serverless cluster
DescribeEKSClusterCommonNames	Queries the CommonNames for sub-accounts to access a TKE Serverless cluster
GetEksSpecs	Obtains all supported specifications for TKE Serverless
ListEKSK8SVersion	Obtains supported Kubernetes versions for TKE Serverless

GetEksAppDiff	Obtains Diff information for TKE ServerlessApp
ActivateEKSCluster	Activates an elastic cluster
CheckEksClusterCIDR	Checks elastic cluster CIDR for conflicts
CheckUseEks	Checks whether a user is using elastic container service
DetachEKSDisks	Unbinds TKE Serverless cloud disks
EnableEksAudit	Enables elastic cluster audit
DrainEksNode	Drains a node from an elastic cluster
DeleteEKSDisk	Deletes a TKE Serverless cloud disk
DeleteEKSCluster	Deletes an elastic cluster
DeleteEKSInstances	Deletes elastic container instances
DeletePodsById	Deletes a TKE Serverless Pod
UpgradeEKSClusterAuthorizationMode	Upgrades the TKE Serverless cluster authorization mode to RBAC
NotifyResult	Notifies of TKE Serverless operation result
UninstallEksLogAgent	Uninstalls the log collection component in an elastic cluster
UpdateEKSCluster	Modifies elastic cluster TKE Serverless
InstallEksLogAgent	Installs the CLS log collection component in an elastic cluster

Discontinuing Update of NginxIngress Addon

Last updated : 2024-07-23 11:17:28

TKE's [NginxIngress Addon](#) has stopped iterating new versions and new features since June 2024, and it is planned to stop creating new NginxIngress instances from October 1, 2024. Existing NginxIngress instances continue to operate but will no longer be compatible with TKE version 1.30 and above. If instances of the NginxIngress addon are still in use in the cluster at that time, it will affect the cluster upgrade to version 1.30.

If you still have NginxIngress-related requirements, it is recommended to migrate to a self-built NginxIngress. For details, see [Ingress-Nginx Controller Installation Guide](#).

qGPU Service Adjustment

Last updated : 2023-11-22 15:44:33

Starting from November 30, 2023, Tencent Kubernetes Engine (TKE) will no longer provide qGPU feature support for normal nodes. If you want to continue using the qGPU to improve the utilization of GPU resources, we recommend migrating your business to [native nodes](#).

Please note that from November 30, 2023, Tencent Kubernetes Engine (TKE) will no longer provide technical support or guarantee functional validity for normal nodes. The specific instructions are as follows:

1. The platform will no longer provide a cluster dimension switch to use the qGPU feature for newly added clusters.
2. If your cluster has enabled the qGPU feature, the normal node business will remain unaffected. However, if you have disabled the cluster dimension qGPU switch, the platform will no longer support the reactivation of this feature for normal nodes.

We apologize for any inconvenience this change may cause and express our gratitude for your support and understanding of our product. If you encounter any problems during this process, or need our help, please do not hesitate to [contact us](#).

Version Upgrade of Master Add-On of TKE Managed Cluster

Last updated : 2023-05-18 10:40:07

Upgrade time

TKE plans to upgrade the minor Kubernetes version of the master add-on of the TKE managed cluster in batches from 10:00 PM to 5:00 AM on October 12–13, 17–18, and 24–27, 2022.

Upgrade scope

This upgrade involves the master add-on of managed clusters, with the major Kubernetes version between v1.16 and v1.22. After the upgrade, the master add-on will be on the latest minor version of the current major version. For example, if the current master add-on of your cluster is on v1.20.6-tke.15, and the latest version launched by TKE is v1.20.6-tke.27, the upgraded version will be v1.20.6-tke.27.

Upgrade content

This upgrade contains:

1. **Performance enhancement:** Optimize the kube-apiserver list performance of large clusters and improve etcd storage stability.
2. **Security enhancement:** Combine the fix PR of the key community vulnerability CVE-2022-3172.
3. **Stability enhancement:** Combine the fix PR of multiple community features and improves kubelet, kube-scheduler, and HPA capabilities.
4. **Feature support:** Support dedicated schedulers for native nodes, in-place Pod configuration adjustment, and super node creation.

Upgrade description

The minor version of the master add-on is fully compatible with earlier versions, and you don't need to reset the Kubernetes add-on startup parameters, as the original settings (if any) will be retained. This upgrade won't change the major Kubernetes version of the master add-on, nor the Kubernetes version of the node. For more information on TKE Kubernetes version revisions, see [TKE Kubernetes Revision Version History](#). If you have any questions, [contact us](#).

Note :

Kubernetes on v1.14 or earlier will no longer be iterated, and no more technical support is available. We recommend you upgrade your cluster as soon as possible. For detailed directions, see [Upgrading a Cluster](#).

Upgrading tke-monitor-agent

Last updated : 2025-01-23 17:52:40

September 2022

TKE's basic monitoring add-on `tke-monitor-agent` will be connected to the cluster add-on management module from 10:00 AM to 22:00 PM on September 27 and 28, 2022, after which you can maintain its lifecycle in the module. This change won't affect your businesses.

Change content

Connect the `tke-monitor-agent` add-on to the add-on management module to support version upgrades in the console. For detailed directions, see [Add-On Lifecycle Management](#).

Create a secret named `sh.helm.release.v1.monitoragent.v1` in the `kube-system` namespace.

July 2022

TKE's monitoring add-on will be upgraded to a new version from 10:00 AM to 23:00 PM on July 27–29 and August 1–4, 2022. This upgrade will fix some of the known issues and support the collection and reporting of some basic metrics. This change won't affect your businesses. If you don't want to have this change, [contact us](#) for assistance.

Change content

Fix the issue where the CPU utilization (usage/request) and CPU utilization (usage/limit) of the Pod are negative values.

Fix the data source configuration issue where the data source of the GPU metric is incorrectly configured, affecting GPU metric pull in GPU scenarios.

Fix the issue of reporting node CPU and memory utilization metrics.

Support the collection of some metrics:

Support the collection of the node CPU/memory packing rate.

Support the collection of metrics for calculating the recommended CPU/memory value and optimizable CPU/memory resources, which will be collected and reported only when the request recommendation add-on is installed.

Troubleshooting

If you encounter any problems during the upgrade, [contact us](#) for assistance.

Instructions on Cluster Resource Quota Adjustment

Last updated : 2024-12-13 15:46:01

TKE plans to apply a resource usage quota to managed clusters from **April 30, 2022 (UTC +8)**. This quota is automatically applied according to the cluster model, and can not be modified. It prevents the cluster from being unavailable due to excessive resource usage due to application bugs. Please check the resource usage of clusters under your account. Upgrade the cluster or delete unused resources before April 30 if necessary. For quota details, see [Quotas and Limits](#).

Decommissioning Kubernetes Version

Last updated : 2024-07-22 09:37:14

As per Tencent Cloud TKE's version support mechanism, Kubernetes Version 1.18 will be expired on July 20, 2023. TKE will no longer provide technical support for this version. After the version is decommissioned, you will not be able to create Kubernetes clusters of the decommissioned version.

Version	Decommissioning Time	Technical Support	Creating Cluster
v1.18	July 20, 2023	No longer provided	Unable to create
v1.16	December 20, 2022	No longer provided	Unable to create
v1.14 and earlier	May 20, 2022	No longer provided	Unable to create

To ensure normal operations of your business, upgrade your cluster version as soon as possible. For details on how to upgrade a cluster, see [Upgrading a Cluster](#). For details about TKE's version support mechanism, see [TKE Kubernetes Version Maintenance Mechanism](#).

Deactivation of Scaling Group Feature

Last updated : 2024-12-13 15:46:01

To ensure a consistent user experience, TKE plans to deactivate the scaling group service on June 13, 2022. If your cluster has a created scaling group, you can convert it into a node pool through the

CreateClusterNodePoolFromExistingAsg API or **in the console**. The conversion will not affect your business.

Through the API

Through the console

You can use the `CreateClusterNodePoolFromExistingAsg` API to convert the scaling group into a node pool. For detailed directions, see [CreateClusterNodePoolFromExistingAsg](#).

1. Log in to the [TKE console](#) and click **Cluster** on the left sidebar.
2. On the cluster management list page, click the target cluster ID.
3. On the cluster details page, select **Node Management > Scaling group**.
4. On the node pool list page, select **More > Create Node Pool** on the right of the scaling group.

A node pool provides more stable and comprehensive elastic scaling capabilities than a scaling group, facilitating node creation, management, and termination. After the deactivation, you can use the node pool to manage created scaling groups. For more information, see [Node Pool Overview](#).

Deactivation plan

Creation entry closure: Starting from June 6, 2022, no more scaling groups can be created in the console or through the API.

Feature deactivation and switch: The scaling group service will be deactivated on June 13, 2022. We recommend you convert your scaling group into the node pool as instructed in "Converting a scaling group into a node pool" before the deactivation date. Scaling groups that still exist after the date will be automatically converted into node pools for unified management.

Notice on TPS Discontinuation on May 16, 2022 at 10:00 (UTC +8)

Last updated : 2024-12-13 15:46:01

To provide better and more powerful product capabilities, TPS will be merged and upgraded into [Tencent Managed Service for Prometheus \(TMP\)](#). The new TMP service supports cross-region and cross-VPC monitoring and allows you to connect a unified Grafana dashboard to multiple TMP instances for unified data display. For TMP billing details, see [Pay-as-You-Go](#). For cloud resource usage details, see [Billing Mode and Resource Usage](#). [Free metrics](#) for basic monitoring will not be billed.

TMP has been officially launched. Click [here](#) to try it out. TPS instances can no longer be created, all the legacy instances will be deactivated on May 16, 2022, and relevant resources will be deleted at that time. You can use our quick [migration tool](#) to migrate your TPS instances to TMP. Before the migration, [streamline monitoring metrics](#) or reduce the collection frequency first; otherwise, higher costs may be incurred.

Basic Monitoring Architecture Upgrade

Last updated : 2024-12-13 15:46:01

Background

TKE will be upgraded to provide more stable basic monitoring services and improve the availability of monitoring data, alarming, and HPA scaling. The upgrade falls into three phases, that is, deploying the new version of the monitoring service add-on, switching the add-on version, and deactivating the earlier version, without affecting the services.

Operation Details

Phase 1. Deploying the monitoring service add-on

The `tke-monitor-agent` monitoring data collection add-on will be installed under the `kube-system` namespace of your cluster. The agent will use less than 70 MB memory and 0.01 CPU cores on each node. For add-on details, see [Description of tke-monitor-agent](#).

Phase 2. Switching the add-on version

This phase will be performed one week after the new version of the add-on is deployed to ensure the data source stability.

Note:

If a Kubernetes cluster on v1.8 or later is used, a Service and Endpoint named `metrics-service` will be created under the `kube-system` namespace and point to the metrics-server maintained on the container side.

The `hpa-metrics-service` in the cluster will point to the new data source. If `apiservice` `v1beta1.metrics.k8s.io` uses the default `kube-system/hpa-metrics-service` data source provided by the container, and the Kubernetes cluster is on v1.8 or later, it will switch to the more stable data source `kube-system/metrics-service`.

`apiservice` `v1alpha1.monitor.tencent.io` will be added to support TKE virtual nodes in reporting monitoring data and querying the monitoring data of Pods on each node.

Phase 3. Deactivating the earlier version of the monitoring add-on

This phase will be performed one week after the add-on version is switched.

Starting Charging on Managed Clusters

Last updated : 2024-12-13 15:46:01

Tencent Cloud plans to start charging on TKE managed clusters from **March 21, 2022 10:00 (UTC +8)**. Nine cluster models are available with different specifications, such as the maximum nodes and maximum Pods. SLA support is provided for all models.

The new Auto Cluster Upgrade feature is released, which upgrades the cluster model automatically when the load on control plane components reaches the threshold or the number of nodes reaches the upper limit.

Tencent Cloud provides a 10-day grace period for TKE managed clusters created before 10:00, March 21, 2022 (UTC +8), which means you do not need to pay for these clusters during this period. When the grace period ends, that is 10:00, April 1, 2022 (UTC +8), you will be charged based on the chosen models. If you do not choose the model for managed clusters, a system-recommended model will apply and incur charges accordingly.

For more information about the prices, see [TKE Billing Overview](#).

If you have any question, please [submit a ticket](#).

Instructions on Stopping Delivering the Kubeconfig File to Nodes

Last updated : 2024-12-13 15:46:01

Note

TKE plans to carry out an operation from 23:00 September 21 (Monday) to 06:00 September 22 (Tuesday), 2020 UTC+8 to stop delivering the Kubeconfig file.

Background

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

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

Issues and Solutions

Symptoms

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](#).

Release Notes

Kubernetes Version Maintenance

TKE Kubernetes Version Maintenance Mechanism

Last updated : 2024-10-15 10:48:42

Tencent Kubernetes Engine (TKE) provides container-centric solutions based on native Kubernetes. Since 3 to 4 minor versions of the community-supported open source edition of Kubernetes are released each year, TKE will also regularly release and maintain the TKE Kubernetes versions supported by the platform. This document mainly describes the support mechanism, release schedule, and expiration risks of the TKE Kubernetes versions.

Version Maintenance

Since September 24, 2018, TKE has only released versions that support even-numbered minor versions of Kubernetes, such as 1.30, 1.28, and 1.26.

Version Number

The TKE Kubernetes versions are expressed as x.y.z-tke.n. x.y.z represents the community Kubernetes version, where x represents the major version, y represents the minor version, z represents the patch version, and n represents the TKE patch version. For example, 1.28.3-tke.1 represents a patch version provided based on the Kubernetes version 1.28.3.

Maintenance Cycle

TKE provides up to 27 months (18+6+3) of service support for each minor version released. The support cycle is divided into the following four stages:

Version Stage	Description
GA (General Availability)	Indicates that the current version can be fully delivered to existing network customers. Typically, it will go to the GA stage 3 months after a new even-numbered minor version is released in the community. During this period, the platform provides patches and service support for the version in this stage.
EOM (End of Marketing)	Indicates that the creation of new clusters for the current version will be stopped across the entire network. Typically, it will go to the EOM stage 18 months after update to the GA version. During this period, the platform provides patches and service support for the version in this stage.

EOFS (End of Full Support)	Indicates that the platform will stop providing new patches for the current version and product features under this version may be limited, such as limiting the creation of new nodes (pools). Typically, it will go to the EOFS stage 6 months after update to the EOM version. During this period, the platform only provides service support for the version in this stage.
EOS (End of Service & Support)	Indicates the time to stop providing services for the current version when the version has expired. Typically, it will go to the EOS stage 3 months after update to the EOFS version.

Notes:

TKE will conduct a comprehensive risk assessment and strict compatibility verification for all open versions. Typically, the creation of new clusters will be gradually opened up 2 months after a new even-numbered minor version is first released in the community. The version is still in the grayscale stage at this time and is expected to enter the GA stage after grayscale for one month. In the grayscale stage, it is not recommended to deploy production business directly.

Patches and Service Support**Patch Scope**

New product features, feature defect repair, community feature merging, security risk remediation, and other features provided by the platform.

Service Support Scope

Cluster creation: TKE supports the creation of version clusters in the GA stage.

Upgrade and O&M assurance: TKE provides upgrade features for versions later than v1.10 and provides troubleshooting, fault recovery, and other support.

After-sales support: TKE provides answering questions, online guidance, and troubleshooting. However, for expired versions of Kubernetes clusters, the platform will not guarantee the quality and effectiveness of technical support, and the SLA may be affected due to failure to comply with best practices.

TKE Kubernetes Release Schedule

K8s Version	GA Date	EOM Date	EOFS Date	EOS Date
1.30	August 20, 2024	February 20, 2026	August 20, 2026	November 20, 2026
1.28	April 28, 2024	October 28, 2025	April 28, 2026	July 28, 2026
1.26 (upgrade)	April 25, 2023	October 25, 2024	April 25, 2025	July 25, 2025

supported)				
1.24 (upgrade supported)	January 6, 2023	July 6, 2024	January 6, 2025	April 6, 2025
1.22 (upgrade supported)	July 20, 2022	January 20, 2024	July 20, 2024	October 20, 2024
1.20 (upgrade supported)	September 9, 2021	March 9, 2023	September 9, 2023	December 9, 2023

Notes:

We have postponed the service support time for some versions. The EOM time for version 1.20/1.22/1.24 is postponed to October 25, 2024. The EOFS and EOS time for version 1.20/1.22 is postponed to December 20, 2024.

Expired Version Risks

1. The creation of new clusters for expired versions is not supported.
2. The platform no longer provides patch services for expired versions.
3. The quality and effectiveness of technical support cannot be guaranteed.
4. The platform reserves the right to forcibly upgrade expired versions of clusters. Before performing a forced upgrade, we will send relevant notifications at least one month in advance via SMS, email, and internal messages.

In summary, it is recommended that you should plan and upgrade the Kubernetes version in a timely manner. For details, see [Upgrading a Cluster](#).

Update Notes of TKE Kubernetes Major Versions

Last updated : 2024-05-28 10:25:38

1.26 changes since 1.24

Major Updates

PodSecurityPolicy has been removed, and Pod Security Admission becomes stable: PodSecurityPolicy was deprecated in version 1.21 and completely removed in version 1.25. The necessary updates to improve its usability would result in breaking changes, thus it was necessary to remove it and replace it with the more user-friendly Pod Security Admission. If you are currently using PodSecurityPolicy, please see [Migrating from PodSecurityPolicy to the built-in Pod Security admission](#).

Ephemeral Containers are now stable: Ephemeral containers were in beta in version 1.23 and have become stable in version 1.25. They can be used in a pod for inspection and troubleshooting when a container cannot be effectively kubect exec'd due to a crash or the lack of debugging tools in the image.

Cgroups v2 has become stable in version 1.25: Compared with cgroups v1, cgroups v2 offers many improvements. While cgroups v1 will continue to be supported, this cgroups v2 enhancement readies Kubernetes for eventually deprecating and switching to v2.

Windows continues to be optimized in version 1.25, such as support for unit tests, consistency tests, and [creating a new repository for Windows Operational Readiness](#).

The image repository has been migrated from k8s.gcr.io to registry.k8s.io.

SeccompDefault enters the beta phase in version 1.25.

NetWorkPolicy's endPort becomes stable in version 1.25. Previously, each NetworkPolicy only supported a single port, while the endPort field allows for specifying a range of ports.

Local Ephemeral Storage Capacity Isolation becomes stable in version 1.25, providing capacity isolation for local ephemeral container storage between pods, such as EmptyDir. If a pod's consumption of local ephemeral capacity storage exceeds the limit, eviction can be used to forcefully limit its consumption of shared resources.

CSI Migration becomes stable in version 1.25. The CSI migration is a continuous effort made by SIG Storage over multiple versions before, with the purpose of migrating in-tree volume plugins to out-of-tree CSI drivers.

CSI Ephemeral Volume becomes stable in version 1.25. The CSI ephemeral volume feature allows for directly specifying CSI volumes in the Pod specification for ephemeral use scenarios. They can be used to inject any state (such as configuration, secrets, identities, variables, or similar information) directly into a Pod through mounted volumes. This was initially introduced as an alpha feature in version 1.15 and has been upgraded to the official version. Some CSI drivers (such as the secret-store CSI driver) have used this feature.

CRD Validation Expression Language enters the beta phase in version 1.25.

Server Side Unknown Field Validation enters the beta phase in version 1.25, enabled by default. The apiserver supports validation for unknown fields, allowing for the planned, sequential removal of client-side validation features in kubectl.

KMS v2 alpha1 API is introduced in version 1.25 to enhance performance, implement rotation, and realize observable improvements. Use AES-GCM instead of AES-CBC, and implement static data encryption through DEK without requiring additional actions from users, while continuing to allow reading using both AES-GCM and AES-CBC.

CRI v1alpha2 API is removed in version 1.26, and CRI v1 version is recommended. Version 1.26 will not support containerd 1.5 and earlier versions. If you are using containerd, it is required to upgrade to containerd v1.6.0 or higher before upgrading the node to Kubernetes v1.26.

In the v1beta1 Traffic Control API Group in version 1.26: FlowSchema and PriorityLevelConfiguration's flowcontrol.apiserver.k8s.io/v1beta1 API version will no longer be available in v1.26. The flowcontrol.apiserver.k8s.io/v1beta2 API version is available since v1.23.

Starting with 1.25, the HorizontalPodAutoscaler of the autoscaling/v2beta1 API version will no longer be provided, and the HorizontalPodAutoscaler of the autoscaling/v2beta2 API version will no longer be available in version 1.26. It is recommended to use the autoscaling/v2 API version, available since v1.23.

The in-tree credential management code, originally as a part of Kubernetes and authentication code specific to vendors Azure and Google Cloud, will be removed from client-go and kubectl. As an alternative, Kubernetes has provided a vendor-neutral authentication plugin mechanism before the release of v1.26.

Deprecation and Removal

Parameter

kubectl Command-line Parameter: As part of the implementation of the Inclusive Naming Initiative, in version 1.26, the --prune-whitelist flag will be deprecated and replaced with --prune-allowlist.

kube-apiserver Command-line Parameter: The --master-service-namespace command-line parameter has no effect on kube-apiserver and has been informally deprecated. This command-line parameter will be officially marked as deprecated in v1.26, with plans for future removal. The Kubernetes project is not expected to be affected by this deprecation and removal.

kubectl run Command-line Parameters: In 1.26, several unused option parameters for the kubectl run sub-command will be marked as deprecated, including: cascade, filename, force, grace-period, kustomize, recursive, timeout, and wait.

The kube-proxy's userspace mode is removed in 1.26. Linux users shall use iptables or ipvs, while Windows users shall use . Now using --mode userspace will fail.

API Version

1.25

batch/v1beta1 has been removed, and batch/v1 is used instead.

discovery.k8s.io/v1beta1 has been removed, and discovery.k8s.io/v1 is used instead.

events.k8s.io/v1beta1 has been removed, and events.k8s.io/v1 is used instead.

policy/v1beta1 has been removed, and policy/v1 is used instead.

node.k8s.io/v1beta1 has been removed, and node.k8s.io/v1 is used instead.

autoscaling/v2beta1 has been removed, and autoscaling/v2 is used instead.

1.26

autoscaling/v2beta2 has been removed, and autoscaling/v2 is used instead.

flowcontrol.apiserver.k8s.io/v1beta1 has been removed, and flowcontrol.apiserver.k8s.io/v1beta2 and

flowcontrol.apiserver.k8s.io/v1beta3 can be used instead.

Others

1.25 deprecates the GlusterFS and Portworx in-tree volume plugins. Although a CSI driver was built for GlusterFS, it has not been maintained. The community discussed the possibility of migrating to a compatible CSI driver but ultimately decided to start deprecating the GlusterFS plugin from the in-tree drivers.

1.25 removes PodSecurityPolicy, Flocker, Quobyte, StorageOS.

1.25 change in vSphere version support: The in-tree vSphere volume driver will no longer support any vSphere versions earlier than 7.0u2.

1.25 cleaning ownership of IPTables chain: Kubelet will gradually migrate to not creating the following iptables chains in the nat table: KUBE-MARK-DROP, KUBE-MARK-MASQ and KUBE-POSTROUTING. This change will be implemented in phases through the IPTablesCleanup feature gate.

1.26 removes the GlusterFS in-tree driver, removal of the deprecated [OpenStack](#) in-tree storage integration (Cinder volume type).

Changelogs

[kubernetes 1.26 changelog](#)

[kubernetes 1.25 changelog](#)

1.24 changes since 1.22

Note:

1. Kubernetes 1.24 has removed the support for Docker through Dockershim. For new nodes in TKE, only Containerd 1.6.9 is supported upon the container runtime .
2. For clusters upgrading from 1.22 to 1.24, nodes whose runtime versions are docker or containerd 1.4 or later (excluding 1.4) can only be upgraded through reinstalling the rolling upgrade mode.

Major Update

Remove Dockershim from kubelet

The built-in dockershim of kubelet was deprecated in v1.20 and removed starting from v1.24. From v1.24 onwards, nodes using the Docker runtime need to migrate to its runtime, such as containerd or CRI-O. For more information, refer to [Is Your Cluster Ready for v1.24](#).

Upgrade EphemeralContainers to Beta

The ephemeral container feature entered Beta in 1.23. Ephemeral containers exist in a Pod for a limited time. They are particularly useful for troubleshooting when you need to inspect another container that has either crashed or cannot use `kubectl exec` because its image lacks the necessary debugging tools.

API with Beta disabled by default

New beta version APIs, starting from 1.24, are disabled by default. Existing beta version and upgraded Beta version APIs are not affected.

Storage capacity and volume expansion are upgraded to GA

This feature is GA in 1.24. [Storage capacity tracking](#) is supported through [CSIStorageCapacity Objects](#) to obtain the current available storage capacity, enhancing the scheduling of Pods with delayed binding CSI volumes.

[Volume Expansion](#) increased support for resizing existing persistent volumes.

NonPreemptingPriority Feature GA

NonPreemptingPriority feature is GA in 1.24, enabled by default and cannot be disabled.

GRPC-based detection upgraded to Beta

GRPCContainerProbe feature is upgraded to beta in 1.24 and enabled by default.

IdentifyPodOS feature upgraded to Beta

The IdentifyPodOS feature is introduced in 1.23 and upgraded to Beta in 1.24. Once enabled, the podSpec can specify an `os` field, indicating the desired OS type (linux/windows) for the Pod to run on. When creating a Pod, this field will be used to validate other Pod settings, such as the `SecurityContext` field. Note, it is still necessary to select nodes of the expected OS type for the Pod using the nodeSelector `kubernetes.io/os`. If the OS of the scheduled node does not match the OS specified in this field, the kubelet will refuse to run the Pod.

JobReadyPods feature enters Beta

The JobReadyPods feature introduced in 1.23 is upgraded to Beta in 1.24. Once enabled, the number of Ready Pods will be displayed in the Job status.

KubeletCredentialProviders feature upgraded to Beta

This feature is in 1.24 Beta and enabled by default, allowing kubelet to dynamically obtain credentials for container image repositories using the `exec` plugin, rather than storing credentials on the node's file system.

ServiceIPStaticSubrange alpha Feature

The ServiceIPStaticSubrange feature is introduced in 1.24. Once enabled, it can effectively prevent conflicts with the ClusterIPs allocated to services.

Dynamic Kubelet configuration has been removed

After being deprecated in 1.22, the Dynamic Kubelet configuration was removed in 1.24.

Incompatible changes related to CNI versions

The following runtime versions have been verified to be compatible with 1.24:

containerd v1.6.4+,v1.5.11+

CRI-O 1.24+

For containerd versions v1.6.0 to v1.6.3, it's necessary to upgrade the CNI plugin or declare the CNI configuration version; otherwise, errors such as "Incompatible CNI versions" or "Failed to destroy network for sandbox" may occur.

CSI snapshot v1beta1 was removed

VolumeSnapshot v1beta1 was deprecated in v1.20 and removed in 1.24.

Deprecate FlexVolume

Starting from 1.23, FlexVolume is deprecated, and the use of CSI for developing volume drivers is recommended.

Deprecate parameters related to klog

Starting from 1.23, all klog parameters are deprecated except for `--v<2>` and `--vmodule<3>:<4>:--log-dir, --log-file, --log-flush-frequency, --logtostderr, --alsologtostderr, --one-output, --stderrthreshold, --log-file-max-size, --skip-log-headers, --add-dir-header, --skip-headers, --log-backtrace-at<4>`

Upgrade of IPv4/IPv6 Dual Stack Support to Stable

Starting from k8s 1.21, dual stack support is enabled by default, entering GA in 1.23, and the IPv6DualStack feature has been removed. Although the cluster supports dual stack networks by default, pods and services are still single-stack by default. To use dual stack, nodes must have routable IPv4/IPv6 network interfaces and use a dual stack-supported CNI network plugin. Also, the Service must be configured to use dual stack: The Service's `.spec.ipFamilyPolicy` field should be set to `PreferDualStack` or `RequireDualStack`. For more information, refer to [IPv4/IPv6 dual-stack](#).

HorizontalPodAutoscaler v2 enters GA

`autoscaling/v2` API enters GA in 1.23, deprecating `autoscaling/v2beta2` API.

Generic Ephemeral Volume enters GA

The Ephemeral Volume feature enters GA in 1.23, allowing the use of all storage plugins that support dynamic creation to provide ephemeral volumes for Pods.

Ignore volume ownership change to enter GA

This feature enters GA in 1.23, allowing for setting the `.securityContext.fsGroupChangePolicy` field to `OnRootMismatch`. Owner and access permissions of the content inside are changed only if the root directory's Owner and access permissions do not match the expected permissions of the volume. Otherwise, it allows skipping the recursive setting of file permissions when mounting a volume, and speeding up Pod startup.

PodSecurity upgraded to Beta

PodSecurity, intended to replace the deprecated `PodSecurityPolicy` Admission Controller, enters the Beta phase in 1.23.

Default to using CRI v1

Starting from 1.23, kubelet supports and defaults to using the CRI v1 API. If v1 is not supported at runtime, it will fall back to v1alpha2.

Structured logging upgraded to Beta

Structured logging is upgraded to Beta in 1.23, with most logs from kubelet and kube-scheduler recorded in structured JSON format.

Simplified Scheduler Plugin Configuration

Starting from 1.23, the MultiPoint option is introduced in the Scheduler Plugins configuration. Plugins configured here are automatically registered to all supported extension point plugin collections, simplifying plugin configuration.

CustomResourceValidationExpressions Alpha Feature

The CustomResourceValidationExpressions feature was introduced in 1.23. Once enabled, it allows the use of rules written in [Common Expression Language \(CEL\)](#) to validate CRDs.

ServerSideFieldValidation Alpha Feature

The ServerSideFieldValidation feature is introduced in 1.23. Once enabled, the apiserver can return warning messages for unknown or duplicate fields. The fieldValidation parameter can be specified in the request:

Ignore: Before 1.23 and when the default behavior when the ServerSideFieldValidation feature is enabled, unknown or duplicate fields (except for the last one of the duplicate fields) will be ignored.

Warn: The default behavior when the ServerSideFieldValidation feature is enabled.

Strict: Request fails, returning an Invalid Request error.

OpenAPIV3 feature enters the Beta phase

The OpenAPIV3 feature introduced in 1.23 enters Beta in 1.24, allowing users to access the OpenAPI v3.0 specifications for all k8s types:

`$cluster/openapi/v3/apis/<group>/<version>` : For a specific type

`$cluster/openapi/v3` : All types

PodAndContainerStatsFromCRI Alpha Feature

The PodAndContainerStatsFromCRI Alpha feature is introduced in 1.23. Once enabled, Pod metrics will be obtained via the CRI interface, instead of from cAdvisor.

Other Updates

Features Entering GA

1.23:TTLAfterFinished,CSIVolumeFSGroupPolicy,IngressClassNamespacedParams,

1.24:DefaultPodTopologySpread,PodOverhead,IndexedJob,PodAffinityNamespaceSelector,SuspendJob,CSRDuration,ServiceLBNodePortControl,PreferNominatedNode,ControllerManagerLeaderMigration(For more information, refer to [Migrate Replicated Control Plane To Use Cloud Controller Manager](#))

Upgraded to Beta Features

1.23:JobTrackingWithFinalizers,StatefulSetMinReadySeconds.

1.24:AnyVolumeDataSource,MixedProtocolLBService,GracefulNodeShutdownBasedOnPodPriority.

Deprecation and Removal

Parameter

kube-apiserver:1.24 deprecates --master-count parameter --endpoint-reconciler-type=master-count option, by default using lease objects created by apiserver to synchronize the apiserver service's endpoint; 1.24 removes the following parameters:--address, --insecure-bind-address, --port,--insecure-port

1.24 removes the following kube-controller-manager parameters:--address,--port,--deployment-controller-sync-period. kube-controller-manager and kube-scheduler's --port and --address parameters, which were deprecated in 1.23 and no longer work, were in 1.24.

The legacy scheduler configuration methods and related parameters policy-config-file, policy-configmap, policy-configmap-namespace, and use-legacy-policy-config were removed in 1.23. You must use the KubeSchedulerConfiguration configuration file.

1.23 removes the kubelet experimental-bootstrap-kubeconfig parameter, replacing it with the --bootstrap-kubeconfig parameter.

1.23 removes the kubelet --seccomp-profile-root parameter and the seccompProfileRoot configuration item.

1.24 deprecates the kubelet --pod-infra-container-image parameter. The sandbox image specified by this parameter will not be reclaimed. Future image reclamation processes will obtain sandbox image information from the CRI interface.

API Version

1.23

rbac.authorization.k8s.io/v1alpha1 and scheduling.k8s.io/v1alpha1 have been removed, use rbac.authorization.k8s.io/v1 and scheduling.k8s.io/v1 instead.

The v1beta1 version of scheduler configuration has been removed, please use the v1beta2 or v1beta3 configuration file.

1.24

Remove `client.authentication.k8s.io/v1alpha1 ExecCredential`, and use `client.authentication.k8s.io/v1` instead.

Remove `node.k8s.io/v1alpha1 RuntimeClass`, and use `node.k8s.io/v1` instead.

Remove `audit.k8s.io/v1alpha1`, `audit.k8s.io/v1beta1`, and use `audit.k8s.io/v1` instead.

Deprecate `CSIStorageCapacity storage.k8s.io/v1beta1` version, and use `storage.k8s.io/v1` instead.

Others

1.24 removes the `Service tolerate-unready-endpoints` annotation deprecated since 1.11, and use `Service.spec.publishNotReadyAddresses` instead.

1.24 removes deprecated features: `ValidateProxyRedirects`, `StreamingProxyRedirects`.

`metadata.clusterName` field is deprecated in 1.24. This read-only field is always empty, leading to misunderstandings.

Changelogs

[kubernetes 1.24 changelog](#)

[kubernetes 1.23 changelog](#)

1.22 Changes Since 1.20

Major Updates

PodSecurityPolicy deprecated

`PodSecurityPolicy` is deprecated in 1.21 and will be removed in 1.25. You can evaluate and migrate it to Pod Security Admission or third-party admission plug-ins.

Immutable Secret and ConfigMap GA

After `Secret` and `ConfigMap` are set as immutable (`immutable: true`), kubelet no longer watches the changes of these objects and mounts them to the container again to reduce the load of apiserver. This feature enters GA in 1.21.

CronJobs GA

`CronJobs` enters GA (`batch/v1`) in 1.21, and the new version controller `CronJobControllerV2` with higher performance is enabled by default.

IPv4/IPv6 supports to enter Beta

Dual-stack networks allow Pods, services, and nodes to obtain IPv4 and IPv6 addresses. In 1.21, the dual-stack network is upgraded from alpha to beta and enabled by default.

Graceful Node Shutdown

This feature enters the beta phase in 1.21, allowing for notifying kubelet upon node shutdown events and gracefully terminates Pods on nodes.

Persistent Volume Health Monitoring

This alpha feature is introduced in 1.21, allowing for monitoring the running status of PV and marking when they become unhealthy. At this time, workloads can be adjusted accordingly to avoid data being written to or read from unhealthy PVs.

Server-side Apply GA

Server-side Apply helps users and controllers manage resources through declarative configuration, such as creating or modifying objects declaratively. Server-side Apply enters the GA phase in 1.22.

External Credential GA

External credentials enter the GA phase in 1.22, providing better support for interactive login process plug-ins. For more information, refer to [sample-exec-plugin](#).

Etcd Updated to 3.5.0

The Etcd 3.5.0 version is used by default in 1.22, which has improved security, performance, monitoring and developer experience, fixed multiple bugs, and added important new features such as structured log records and built-in log rotation.

MemoryQoS

The alpha MemoryQoS feature is supported starting from 1.22. Once enabled, the Cgroups v2 API will be used to manage and control memory allocation and isolation, ensuring memory usage for workloads and improving the availability of workloads and nodes in the case of memory resource competition. This feature was proposed by Tencent Cloud and contributed to the community.

Cluster's Default seccomp Configuration

Kubelet introduces the `SeccompDefault` alpha feature in 1.22. According to the `--seccomp-default` parameter and setting, kubelet will use the `RuntimeDefault` seccomp configuration instead of `Unconfined` to improve the security of workloads.

Other Updates

GA Features:

1.21: EndpointSlice, Sysctls, PodDisruptionBudget.

1.22: CSIServiceAccountToken

Features Entering to beta:

1.21: TTLAfterFinished

1.22: SuspendJob,PodDeletionCost,NetworkPolicyEndPort.

The new scheduler scoring plug-in `NodeResourcesFit` is introduced in 1.22, which is used to replace three plug-ins: `NodeResourcesLeastAllocated`, `NodeResourcesMostAllocated`, and `RequestedToCapacityRatio`.

After the alpha feature `APIServerTracing` is enabled since 1.22, the apiserver supports distributed tracing and allows users to use the `--service-account-issuer` parameter to set multiple issuers. In addition, it can provide uninterrupted service when issuers are changed.

Deprecation and Removal

Removed Parameters and Features

1. `Service TopologyKeys` is deprecated and replaced with `Topology Aware Hints`.

2. kube-proxy

Starting from 1.21, `net.ipv4.conf.all.route_localnet=1` will not be automatically set in ipvs mode. For upgraded nodes, `net.ipv4.conf.all.route_localnet=1` will be retained. But for new nodes, the default system value (usually `0`) is inherited.

The `--cleanup-ipvs` parameter is deleted and can be replaced with the `--cleanup` parameter.

3. kube-controller-manager

Starting from 1.22, the `--horizontal-pod-autoscaler-use-rest-clients` parameter is removed.

The `--port` and `--address` parameters become invalid and will be removed in 1.24.

4. kube-scheduler: The `--hard-pod-affinity-symmetric-weight` and `--scheduler-name` parameters are removed in 1.22, and instead, these information can be configured in the `config` file.

5. Kubelet: The `DynamicKubeletConfig` feature is deprecated and is disabled by default. If the `--dynamic-config-dir` parameter is set when kubelet is started, an alarm will be reported.

Removed or Deprecated Versions

1. CronJob batch/v2alpha1 is removed started from 1.21

2. Starting from 1.22, the following beta APIs are removed: (For more information, refer to [Kubernetes official documentation](#).)

rbac.authorization.k8s.io/v1beta1

admissionregistration.k8s.io/v1beta1

apiextensions.k8s.io/v1beta1

apiregistration.k8s.io/v1beta1

authentication.k8s.io/v1beta1

authorization.k8s.io/v1beta1

certificates.k8s.io/v1beta1
coordination.k8s.io/v1beta1
extensions/v1beta1 and networking.k8s.io/v1beta1 ingress API

Change logs

[kubernetes1.22changelog](#)

[kubernetes1.21changelog](#)

1.20 Changes Since 1.18

Major Updates

New Version of CronJob Controller

The new version of the CronJob controller is introduced in 1.20, which uses the informer mechanism to replace the original polling and optimize the performance. You can set `--feature-gates="CronJobControllerV2=true"` in `kube-controller-manager` to enable the new version. The new version will be enabled by default on later versions.

Dockershim Deprecation

Dockershim has been deprecated. The kubernetes' support for Docker is deprecated and will be removed from a future release. OCI-compliant images generated by Docker will continue to run in the CRI-compliant runtime. For more information, refer to [Don't Panic: Kubernetes and Docker](#) and [Dockershim Deprecation FAQ](#).

Structured Logs

The structures of log messages and k8s object reference are standardized to make log parsing, processing, storage, query, and analysis easier. Two methods are added to klog to support structured logs: `InfoS` and `ErrorS`. The `--logging-format` parameter is added to all components, and its default value is `text` in the previous format. You can set it to `json` to support structured logs, and the following parameters will become invalid: `--add_dir_header`, `--alsologtostderr`, `--log_backtrace_at`, `--log_dir`, `--log_file`, `--log_file_max_size`, `--logtostderr`, `--skip_headers`, `--skip_log_headers`, `--stderrthreshold`, `--vmodule`, and `--log-flush-frequency`.

Exec Probe Timeout Processing

A longstanding bug regarding exec probe timeouts that may impact existing Pod definitions has been fixed. Prior to this fix, the `timeoutSeconds` field was not respected for exec probes. Instead, probes would run indefinitely, even beyond their configured deadline, until a result was returned. With this change, the default value of `1 second` will be applied if no value is specified. If a probe takes longer than 1 second, existing Pod definitions may need to be

modified, displaying the specified `timeoutSeconds` field. A switch called `ExecProbeTimeout` has also been added with this fix, allowing for retaining previous behaviors (In later releases, this feature will be locked and removed). To retain previous behaviors, it is required to set `ExecProbeTimeout` to `false`.

For more information, refer to [Configure Liveness, Readiness and Startup Probes](#).

Volume Snapshot Operation Feature to GA

This feature provides a standard way to trigger volume snapshot operations and allows you to incorporate snapshot operations in a portable manner on any Kubernetes environment and supported storage providers.

Additionally, these Kubernetes snapshot primitives act as basic building blocks that unlock the ability to develop advanced and enterprise-grade storage administration features for Kubernetes, including application or cluster level backup solutions.

Note that snapshot support requires Kubernetes distributors to bundle and deploy the snapshot controller, snapshot CRDs, and validation webhook, as well as a CSI driver supporting the snapshot feature.

kubectl debug enters the beta phase

The `kubectl alpha debug` command enters the beta phase, and is replaced with `kubectl debug`. It supports common debugging work directly from `kubectl`, for example:

Troubleshoot workloads that crash on startup by creating a copy of the Pod with a different container image or command.

Troubleshoot distroless and other containers that do not contain debugging tools by adding a new container with debugging tools, either in a new copy of the Pod or using an ephemeral container. (`EphemeralContainers` are an alpha feature that are not enabled by default.)

Troubleshoot on a node by creating a container running in the host namespaces and with access to the host's file system.

Note that as a new built-in command, `kubectl debug` takes priority over any `kubectl` plugin named `debug`, and the affected plugins must be renamed.

`kubectl alpha debug` is now deprecated and will be removed in a subsequent version, and it is required to be updated to `kubectl debug`. For more information, refer to [Debug Running Pods](#).

API Priority and Fairness enters the beta phase

The API Priority and Fairness feature introduced in 1.18 will be enabled by default in 1.20, allowing `kube-apiserver` to categorize incoming requests by priority.

PID Limit Feature GA

`SupportNodePidsLimit` (node-to-Pod PID isolation) and `SupportPodPidsLimit` (ability to limit PIDs per Pod) have both moved to the GA phase.

Alpha Feature: Graceful Node Shutdown

Users and cluster admins expect that Pods will adhere to the expected Pod lifecycle, including Pod termination. Currently, when a node shuts down, Pods do not follow the expected Pod termination lifecycle and are not terminated normally, which may cause issues for some workloads. The `GracefulNodeShutdown` feature is now added for alpha in 1.20, making the kubelet be able to monitor the system shutdown events of nodes, thereby enabling graceful termination of Pods during a system shutdown.

CSIVolumeFSGroupPolicy enters the beta phase

CSIDrivers can use the `fsGroupPolicy` field to control whether ownership and permissions can be controlled upon mount. (`ReadWriteOnceWithFSType` , `File` , and `None`)

ConfigurableFSGroupPolicy enters the beta phase

The following can be set in a non-recursive manner: `fsgroup - PodFSGroupChangePolicy = OnRootMismatch`

Other Updates

The Cloud Controller Manager component is added.

Features graduating to GA:

[RuntimeClass](#) `node.k8s.io/v1beta1` is deprecated and replaced with `node.k8s.io/v1` .

[Built-in API types defaults](#)

[StartupProbe](#)

[Services and Endpoints](#) support `AppProtocol` field

[TokenRequest and TokenRequestProjection](#)

[SCTPSupport](#)

[Containerd support for Windows](#)

[Ingress](#)

`networking.k8s.io/v1beta1` is deprecated (it will be removed in 1.22) and replaced by `networking.k8s.io/v1` .

[seccomp](#)

`seccomp` annotations `seccomp.security.alpha.kubernetes.io/pod` and

`container.seccomp.security.alpha.kubernetes.io/...` are deprecated (they will be removed in 1.22). You can directly specify the following fields for Pods and container specs:

```
securityContext:
  seccompProfile:
    type: RuntimeDefault|Localhost|Unconfined ## choose one of the three
    localhostProfile: my-profiles/profile-allow.json ## only necessary if type ==
    Localhost
```

K8S converts annotations and fields automatically without additional operation.

Node Certificate Automatic Rotation

Limit Node Access to API

Node authentication mode features are all implemented.

Redesign Event API

To reduce the impact of events on the system performance and add more fields to provide more useful information, Event API is redesigned. This work is done in 1.19.

CertificateSigningRequest API

In addition to `certificates.k8s.io/v1beta1`, the `certificates.k8s.io/v1` version is added to `CertificateSigningRequest`. When using `certificates.k8s.io/v1`:

You must specify `spec.signerName` and stop using `kubernetes.io/legacy-unknown`.

You must specify `spec.usages`, which can not contain repeated value but only known usage.

You must specify `status.conditions[*].status`.

`status.certificate` must be PEM encoded and can contain only the `CERTIFICATE` block.

Enter the Beta feature:

The following features enter the beta phase and are enabled by default:

EndpointSliceProxying

kube-proxy reads information from EndpointSlices instead of Endpoints, which can greatly improve the cluster scalability and make it easier to add new features such as topology-aware routing.

KubeSchedulerConfiguration

HugePageStorageMediumSize

ImmutableEphemeralVolumes

The Secret and ConfigMap volumes can be marked as immutable. When there are many Secret and ConfigMap volumes, the pressure on apiserver can be greatly mitigated.

NodeDisruptionExclusion

NonPreemptingPriority

ServiceNodeExclusion

RootCAConfigMap

Pod Resource Metrics in Scheduler

ServiceAccountIssuerDiscovery

Deprecations and Removals

Deprecated Version

Deprecated Version	New Version
<code>apiextensions.k8s.io/v1beta1</code>	<code>apiextensions.k8s.io/v1</code>
<code>apiregistration.k8s.io/v1beta1</code>	<code>apiregistration.k8s.io/v1</code>

authentication.k8s.io/v1beta1	authentication.k8s.io/v1
authorization.k8s.io/v1beta1	authorization.k8s.io/v1
autoscaling/v2beta1	autoscaling/v2beta2
coordination.k8s.io/v1beta1	coordination.k8s.io/v1
storage.k8s.io/v1beta1	storage.k8s.io/v1

kube-apiserver

1. The `componentstatus` API is deprecated. This API is to provide the running status of etcd, kube-scheduler and kube-controller-manager components, but only worked when those components were local to apiserver, and when kube-scheduler and kube-controller-manager exposed unsecured health endpoints. After this API is deprecated, the etcd health check is included in the kube-apiserver health check and kube-scheduler/kube-controller-manager health checks can be made directly against those components' health endpoints.
2. The apiserver no longer listens on insecure ports.

The `--address` and `--insecure-bind-address` parameters can be set, but are invalid. The `--port` and `--insecure-port` parameters can be set to `0` only. These parameters will be removed in 1.24.

3. `TokenRequest` and `TokenRequestProjection` enter the GA phase. You need to set the following parameters for kube-apiserver:

- `--service-account-issuer` : Identify the fixed URL of the cluster API server.
- `--service-account-key-file` : One or multiple public keys for token verification.
- `--service-account-signing-key-file` : Private key for issuing service account, which can use the same file as the `--service-account-private-key-file` parameter of `kube-controller-manager` .

kubelet

1. The following parameters are removed:

- `--seccomp-profile-root`
- `--cloud-provider` and `--cloud-config` , which are replaced with `config`
- `--really-crash-for-testing` and `--chaos-chance`

2. The deprecated `metrics/resource/v1alpha1` endpoint is removed and please replace with `metrics/resource` .

Other Removals

The `failure-domain.beta.kubernetes.io/zone` and `failure-domain.beta.kubernetes.io/region` labels are deprecated and please replace with `topology.kubernetes.io/zone` and `topology.kubernetes.io/region` respectively. All labels

prefixed with `failure-domain.beta...` labels should be replaced with the labels prefixed with corresponding `topology...`.

PodPreset is removed, and you can use webhooks to implement this feature.

The `basic auth` authentication method is no longer supported.

Direct CBS inline mounting to workloads is no longer supported.

Note

When you upgrade from Kubernetes 1.18 to 1.20, successful mounting of [CSI ephemeral \(inline\) volumes](#) cannot be guaranteed. If your application uses a CSI ephemeral volume, we recommend you convert it to a persistent volume before upgrade.

Change logs

[Kubernetes 1.20 changelog](#), [Kubernetes 1.19 changelog](#)

1.18 Changes Since 1.16

Major Updates

Cloud Provider labels reach to the stable (GA) phase

Deprecated and new labels are as listed below:

Deprecated Labels	New Labels
<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>

Volume Snapshot enters the Beta phase

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

CSI Migration enters the Beta phase

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

Graduation of Kubernetes Topology Manager to Beta

The TopologyManager feature enters beta in 1.18. This feature enables NUMA alignment of CPU and devices (such as SR-IOV VFs) that will allow your workload to run in an environment optimized for low latency.

Prior to the introduction of the TopologyManager, the CPU and Device Manager would make resource allocation decisions independent of each other. This could result in undesirable allocations on multi-socket CPU systems, causing degraded performance on latency critical applications.

Graduation of Server-Side Apply to Beta 2

Server-Side Apply was upgraded to beta on Kubernetes 1.16, but is now introducing a second Beta (ServerSideApply) on Kubernetes 1.18. This new version will track and manage changes to fields of all new Kubernetes objects, allowing users to know resource changes.

IngressClass Resources

The `IngressClass` resource is used to describe a type of Ingress controller within a Kubernetes cluster.

`Ingresses` uses the new `ingressClassName` field to set up the controller name for `IngressClass`, and replaces the deprecated `kubernetes.io/ingress.class` annotation.

Other Updates

Graduation of NodeLocal DNSCache to GA.

Graduation of IPv6 to Beta.

`kubectl debug` : Alpha Feature.

Windows CSI support : Alpha Feature.

`ImmutableEphemeralVolumes` : Alpha Feature (supporting immutable ConfigMaps and Secrets without refreshing the corresponding volumes).

The following features graduate to GA:

`ScheduleDaemonSetPods`

`TaintNodesByCondition`

`WatchBookmark`

`NodeLease`

`CSINodeInfo`

`VolumeSubpathEnvExpansion`

`AttachVolumeLimit`

`ResourceQuotaScopeSelectors`

`VolumePVCDataSource`

`TaintBasedEvictions`

`BlockVolume`, `CSIBlockVolume`

`Windows RunAsUserName`

Features graduating to Beta:

`EndpointSlices` : Disabled by default

`CSIMigrationAWS` : Disabled by default

`StartupProbe``EvenPodsSpread`

Deprecations and Removals

Removed Features

The following features, which are enabled by default and cannot be configured, are removed:

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

Other Removals

The following built-in cluster roles are removed:

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

Deprecated Feature Switches and Parameters

The default service IP CIDR block (`10.0.0.0/24`) is deprecated, and must be set through the `--service-cluster-ip-range` parameter on kube-apiserver.

The `rbac.authorization.k8s.io/v1alpha1` and `rbac.authorization.k8s.io/v1beta1` API groups are deprecated and will be removed in 1.20. Therefore, migrate your resources to

`rbac.authorization.k8s.io/v1` .

The `CSINodeInfo` feature gate is deprecated. This feature has graduated to GA and is enabled by default.

Parameter and Other Changes

kube-apiserver

`--encryption-provider-config` : If `cacheSize: 0` is specified in the configuration file, versions earlier than 1.18 are automatically configured to cache 1,000 keys, 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 can no longer be configured through the command line.

`PodPriority`

TaintNodesByCondition

ResourceQuotaScopeSelectors

ScheduleDaemonSetPods

The following resource versions (group versions) are no longer supported:

apps/v1beta1 and apps/v1beta2 , which are replaced with apps/v1 .

Under extensions/v1beta1 :

daemonsets , deployments and replicaset , which are replaced with apps/v1 .

networkpolicies , which is replaced with networking.k8s.io/v1 .

podsecuritypolicies , which is replaced with policy/v1beta1 .

kubelet

`--enable-cadvisor-endpoints` : This parameter is disabled by default. To access the `cAdvisor v1` JSON API, you must enable it.

The `--redirect-container-streaming` parameter is deprecated and will be removed from later versions. 1.18 only supports the default behavior (kubelet proxy for streaming requests). If `--redirect-container-streaming=true` is set, it must be removed.

The `/metrics/resource/v1alpha1` endpoint is deprecated and replaced with `/metrics/resource` .

kube-proxy

The following parameters are deprecated:

The `--healthz-port` is deprecated and replaced with `--healthz-bind-address` .

The `--metrics-port` is deprecated and replaced with `--metrics-bind-address` .

The `EndpointSliceProxying` feature gate (disabled by default) is added to control whether to enable EndpointSlices in kube-proxy. The `EndpointSlice` feature gate no longer affects the behaviors of kube-proxy.

The following timeout settings for IPVS connection configuration are added:

`--ipvs-tcp-timeout`

`--ipvs-tcpfin-timeout`

`--ipvs-udp-timeout`

The iptables mode supports the IPv4/IPv6 dual-protocol stack.

kube-scheduler

The `scheduling_duration_seconds` metric is deprecated:

The `scheduling_algorithm_predicate_evaluation_seconds` is deprecated and replaced with `framework_extension_point_duration_seconds[extension_point="Filter"]`

The `scheduling_algorithm_priority_evaluation_seconds` is deprecated and replaced with `framework_extension_point_duration_seconds[extension_point="Score"]`

The `AlwaysCheckAllPredicates` is deprecated in the scheduler policy API.

-enable-profiling Parameter

To align `kube-apiserver` , `kube-controller-manager` and `kube-scheduler` , [profiling is enabled by default](#). To disable profiling, it is required to specify the `--enable-profiling=false` parameter.

kubect1

The deprecated `--include-uninitialized` parameter is removed.

The `kubect1` and `k8s.io/client-go` no longer use <http://localhost:8080> as the default apiserver address.

The `kubect1 run` supports Pod creation and no longer supports using the deprecated generator to create other types of resources.

The deprecated `kubect1 rolling-update` command is removed and please use the `rollout` command.

The `-dry-run` supports three parameter values: `client` , `server` , and `none` .

The `-dry-run=server` supports the following commands: `apply` , `patch` , `create` , `run` , `annotate` , `label` , `set` , `autoscale` , `drain` , `rollout undo` , and `expose` .

The `kubect1 alpha debug` command is added, which can be used for [debugging and troubleshooting on ephemeral containers in Pods](#) (the `EphemeralContainers` feature introduced on version 1.16 needs to be enabled).

hyperkube

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

Change Logs

[Kubernetes 1.18 changelog](#), [Kubernetes 1.17 changelog](#)

1.16 Changes Since 1.14

Major updates

Improved Cluster Stability and Availability

Production-ready features like bare metal cluster tool and high availability (HA) have been improved and enhanced. `kubeadm` support for HA capability enters the beta phase, allowing you to use the `kubeadm init` and `kubeadm join` commands to configure and deploy the highly available HA control plane. Certificate management has become more robust, with `kubeadm` now seamlessly updating all your certificates before they expire. For more information, see [pr357](#) and [pr970](#).

Continuous CSI Improvement

SIG Storage continues work to enable migration of in-tree volume plugins to Container Storage Interface (CSI). It works on bringing CSI to feature parity with in-tree functionality, including resizing and inline volumes. It introduces

new alpha functionality in CSI that doesn't exist in the Kubernetes Storage subsystem yet, like volume cloning. Volume cloning enables you to specify another PVC as a `DataSource` when configuring a new volume. If the underlying storage system supports this functionality and implements the `CLONE_VOLUME` capability in its CSI driver, the new volume becomes a clone of the source volume. For more information, see [pr625](#).

Features

Features Graduating to GA:

`CRD`

`Admission Webhook`

`GCERegionalPersistentDisk`

`CustomResourcePublishOpenAPI`

`CustomResourceSubresources`

`CustomResourceValidation`

`CustomResourceWebhookConversion`

CSI support for volume resizing graduates to Beta.

General Updates

Go modules are supported by the Kubernetes core.

Preparation on cloud provider code extraction and organization is continued. The cloud provider code has been moved to `kubernetes/legacy-cloud-providers` for easier removal later and external usage.

[kubectrl get and describe command support extension](#)

[Nodes support 3rd party device monitoring plugins.](#)

A new alpha version of scheduling framework for developing and managing plugins and extending the scheduler features is introduced. For more information, see [pr624](#).

`extensions/v1beta1` , `apps/v1beta1` and `apps/v1beta2` APIs continue to be depreciated. These extensions will be retired in 1.16.

The Topology Manager component is added to Kubelet, aiming to coordinate resource allocation decisions and optimize resource allocation.

IPv4/IPv6 dual stack is supported to assign both v4 and v6 addresses to Pods and Services.

The API Server Network Proxy in alpha Feature.

More extension options are provided for cloud controller manager migration.

`extensions/v1beta1` , `apps/v1beta1` and `apps/v1beta2` APIs are depreciated.

Known Issues

Using the `--log-file` parameter is known to be problematic on Kubernetes 1.15. This presents as things being logged multiple times to the same file. For more information, see [\[Failing Test\] timeouts in ci-kubernetes-e2e-gce-scale-performance](#).

Update notes

Cluster

The following labels can no longer be added to new nodes: `beta.kubernetes.io/metadata-proxy-ready` , `beta.kubernetes.io/metadata-proxy-ready` and `beta.kubernetes.io/kube-proxy-ds-ready` .

The `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` .

The `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` .

The `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 published components (such as `k8s.io/client-go` and `k8s.io/api`) now contain Go module files, including version information of the dependent library. For more information on consuming `k8s.io/client-go` in Go modules, see [go-modules](#) and [pr74877](#).

Apps: Hyperkube short aliases have been removed from source code, because these aliases will be created when compiling hyperkube docker image. For more information, see [pr76953](#).

Lifecycle

Deprecated kubeadm `v1alpha3` configurations have been totally removed.

`kube-up.sh` no longer supports `centos` and `local` .

Storage

The `Node.Status.Volumes.Attached.DevicePath` field is no longer set for CSI volumes. You must update any external controllers that depend on this field.

CRDs in alpha version are removed.

The `StorageObjectInUseProtection` admission [plugin](#) is enabled by default. If you previously had not enabled it, your cluster behavior may change.

When `PodInfoOnMount` is enabled for a CSI driver, the new `csi.storage.k8s.io/ephemeral` parameter in the volume context allows a driver's `NodePublishVolume` implementation to determine on a case-by-case basis whether the volume is ephemeral or a normal persistent volume. For more information, see [pr79983](#).

The `VolumePVCDDataSource` (Storage Volume Cloning feature) enters beta. For more information, see [pr81792](#).

Limits for in-tree and CSI volumes are integrated into one scheduler predicate. For more information, see [pr77595](#).

kube-apiserver

The `--enable-logs-handler` parameter is deprecated and will be removed in v1.19.

The `--basic-auth-file` flag and authentication mode are deprecated and will be removed from a future release.

The default service IP CIDR block (`10.0.0.0/24`) is deprecated and will be removed in six months/two releases.

The `--service-cluster-ip-range` parameter is required to configure the service IP range.

kube-scheduler

The `v1beta1` Event API is used. Any tool targeting scheduler events needs to use it.

kube-proxy

The `--conntrack-max` parameter is removed and replaced with `--conntrack-min` and `--conntrack-max-per-core`.

The `--cleanup-iptables` parameter is removed.

The `--resource-container` is removed.

kubelet

The `--allow-privileged`, `--host-ipc-sources`, `--host-pid-sources` and `--host-network-sources` parameters are removed and replaced with the admission controller of `PodSecurityPolicy`.

The cAdvisor JSON API is deprecated.

The `--containerized` is removed.

The `--node-labels` parameter can no longer be used to configure forbidden labels prefixed with `kubernetes.io-` or `k8s.io-`.

kubectx

The `kubectx scale job` is removed.

The `--pod/-p` parameter of the `kubectx exec` command is removed.

The `kubectx convert` command is removed.

The `--include-uninitialized` is removed.

The `kubectx cp` no longer supports copying symbolic links from containers. You can use the following commands instead:

```
local to pod : tar cf - /tmp/foo | kubectx exec -i -n <some-namespace> <some-pod> -  
- tar xf - -C /tmp/bar
```

```
pod to local : kubectx exec -n <some-namespace> <some-pod> -- tar cf - /tmp/foo |  
tar xf - -C /tmp/bar
```

kubeadm

The `kubeadm upgrade node config` and `kubeadm upgrade node experimental-control-plane` commands are deprecated and replaced with `kubeadm upgrade node`.

The `--experimental-control-plane` parameter is deprecated and replaced with `--control-plane`.

The `--experimental-upload-certs` parameter is deprecated and replaced with `--upload-certs`.

The `kubeadm config upload` command is deprecated and replaced with `kubeadm init phase upload-config`.

CoreDNS checks readiness via the `ready` plugin.

The `proxy` plugin is deprecated and replaced with the `forward` plugin.

The `resyncperiod` option is removed from the `kubernetes` plugin.

The `upstream` option is deprecated. If it is specified, it will be ignored.

Change Logs

[Kubernetes 1.16 changelog](#), [Kubernetes 1.15 changelog](#)

1.14 Changes Since 1.12

Major Updates

[Bump CSI Spec to 1.0.0 and gRPC to 1.13.0.](#)

[Make CoreDNS default in kubeup and update CoreDNS version/manifest in kubeup and kubeadm.](#)

kubeadm is used to simplify cluster management.

[Support Windows Nodes to enter stable.](#)

[Durable \(non-shared\) local storage management.](#)

[Support total process ID limiting for nodes.](#)

[Support Pod Priority and Preemption.](#)

General Updates

`dry-run` graduates to beta (`dry-run` enables you to simulate real API requests without actually changing the cluster status).

`kubectl diff` graduates to beta.

`kubectl` plugin registration becomes stable.

`kubelet` plugin mechanism graduates to beta.

`CSIPersistentVolume` graduates to GA.

`TaintBasedEviction` graduates to beta.

`kube-scheduler` perception of volume topology becomes stable.

Support for out-of-tree CSI volume plugins becomes stable.

Support third-party device monitoring plugins.

The `kube-scheduler` subnet feasibility graduates to beta.

Pod Ready supports customizing probe conditions.

Node memory supports HugePage.

The `RuntimeClass` graduates to beta.

Node OS/Arch labels graduate to GA.

Node-leases graduate to beta.

The `kubelet` resource metrics endpoint graduates to alpha and supports data collection through Prometheus.

The `runAsGroup` graduates to beta.

The `kubectl apply server-side` graduates to alpha, allowing you to perform apply operations on the server side.

The `kubectl` supports `kustomize`.

The `resolv.conf` can be configured in Pods.

CSI volumes support resizing.

CSI supports topology.

Volume mounting supports configuration of sub-path parameters.

CSI supports raw block devices.

CSI supports local ephemeral volumes.

Update Notes

kube-apiserver

The `etcd2` is no longer supported. The `--storage-backend=etcd3` is used by default.

The `--etcd-quorum-read` parameter is deprecated.

The `--storage-versions` parameter is deprecated.

The `--repair-malformed-updates` parameter is deprecated.

kube-controller-manager

The `--insecure-experimental-approve-all-kubelet-csrs-for-group` parameter is deprecated.

kubelet

The `--google-json-key` parameter is deprecated.

The `--experimental-fail-swap-on` parameter is deprecated.

kube-scheduler

`componentconfig/v1alpha1` is no longer supported.

kubectx

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

taints

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

Change Logs

[Kubernetes 1.14 changelog](#), [Kubernetes 1.13 changelog](#)

1.12 Changes Since 1.10

Major Updates

API

Subresources for `CustomResources` graduate to beta now and are enabled by default. With this, updates to the `/status` subresource will disallow updates to all fields other than `.status` (not just `.spec` and

`.metadata` as before). Also, `required` and `description` can be used at the root of the CRD OpenAPI validation schema when the `/status` subresource is enabled. In addition, you can now create multiple versions of CustomResourceDefinitions, but without any kind of automatic conversion, and CustomResourceDefinitions now allow specification of additional columns for `kubectl get` output via the `spec.additionalPrinterColumns` field.

The `dry run` feature is supported. It allows you to view the execution results of some commands without having to submit relevant modifications.

Authentication and Authorization

RBAC aggregation of ClusterRoles graduates to GA. The `client-go credentials` plugin graduates to beta, allowing you to get TLS authentication information from external plugins.

The following annotations are added to audit events, so that you can be better informed of the audit decision-making process:

The Authorization component sets `authorization.k8s.io/decision` (the `allow` or `forbid` authorization decision) and `authorization.k8s.io/reason` (the reason for this decision).

The PodSecurityPolicy admission controller sets `podsecuritypolicy.admission.k8s.io/admit-policy` and `podsecuritypolicy.admission.k8s.io/validate-policy` annotations containing the name of the policy that allows a Pod to be admitted. (`PodSecurityPolicy` also gains the ability to limit `hostPath` volume mounts to be read-only.)

The NodeRestriction admission controller prevents nodes from modifying taints on their node objects, making it easier to keep track of which nodes should be in use.

CLI Command Line

CLI implements a new plugin mechanism, providing a library with common CLI tooling for plugin authors and further refactorings of the code.

Internet

The IPVS mode graduates to GA.

CoreDNS graduates to GA to replace `kube-dns`.

Node

`DynamicKubeletConfig` graduates to the Beta phase.

`cri-tools` graduates to GA.

`PodShareProcessNamespace` graduates to the Beta phase.

Alpha features `RuntimeClass` and `CustomCFSQuotaPeriod` are added.

Scheduler

Pod Priority and Preemption graduate to the Beta phase.

DaemonSet Pod scheduling is no longer managed by the DaemonSet controller, but by the default scheduler.

`TaintNodeByCondition` graduates to the Beta phase.

The Use Local Image First 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

Features graduating to GA: `ClusterRole` and `StorageObjectInUseProtection`.

Features graduating to Beta: External Cloud Provider.

Update Notes

kube-apiserver

The `--storage-version` parameter is removed and replaced with `--storage-versions`. The `--storage-versions` parameter is also deprecated.

The default value of `--endpoint-reconciler-type` is 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

The `--rotate-certificates` parameter is deprecated and replaced with the `.RotateCertificates` field in the configuration file.

kubectrl

All `kubectrl run` generators except `run-pod/v1` are deprecated.

The `--interactive` parameter is removed from `kubectrl logs`.

The `--use-openapi-print-columns` is deprecated and replaced with `--server-print`.

Change Logs

[Kubernetes 1.12 changelog](#), [Kubernetes 1.11 changelog](#)

TKE Kubernetes Revision Version History

Last updated : 2024-05-28 10:24:39

TKE kubernetes 1.28.3 revisions

Time	Version	Update Contents
2024-04-24	v1.28.3-tke.3	lb vip binding issue fixed. (kube-proxy)
2024-04-10	v1.28.3-tke.2	Merge PR119561 to fix the issue where kubelet using gpu might cause panic. (kubelet) Fix issues with scheduling pods with eni-ip to super nodes. (kube-scheduler)
2024-03-06	v1.28.3-tke.1	Support specified Pods during scale-down. (kube-controller-manager) Avoid scheduling Pods using cbs to external CHC Nodes. (kube-scheduler) Pods with added annotation support mounting cgroupfs. (kubelet) Support native node scaling feature. (kube-scheduler) Correctly display disk and network metrics when runtime is containerd. (kubelet) Super nodes support DaemonSet. (kube-apiserver/kube-controller-manager) Optimize super node HPA. (kube-controller-manager) Support extended scheduler Prebind and Unreserve operations. (kube-scheduler) Extend the apiserver's self-signed certificate validity from 1 year to 5 years. (kube-apiserver) Super nodes allow scheduling Pods in TKE trunking ENI mode. (kube-scheduler)

TKE kubernetes 1.26.1 revisions

Time	Version	Update Contents
2024-04-10	v1.26.1-tke.6	Merge PR119561 to fix the issue where kubelet using gpu might cause panic. (kubelet) Merge PR113145 to fix the issue where forced pod deletion might leave behind a sandbox, preventing pod reconstruction. (kubelet)
2024-02-27	v1.26.1-tke.5	Super nodes allow scheduling Pods in TKE trunking ENI mode. (kube-scheduler)
2024-02-01	v1.26.1-tke.4	Super node supports Reserved Specification Awareness. (kube-scheduler) Merge PR119605 to fix incomplete pod information in kubelet prober_probe_total metric. (kubelet)

		<p>kube-proxy Optimization. (kube-proxy)</p> <p>job-controller supports settings for multi-worker parameters <code>--concurrent-job-syncs</code>. (kube-controller-manager)</p>
2023-12-04	v1.26.1-tke.3	<p>Pods with added annotation support mounting cgroupfs. (kubelet)</p> <p>Merge PR119706 to fix the issue of Pods using device plugin resources being reconstructed when kubelet restarts. (kube-apiserver)</p> <p>When deleting resources, add userAgent reporting and record detailed logs. (kube-apiserver)</p>
2023-10-12	v1.26.1-tke.2	<p>Support in-place updates for Pod resources. (kube-apiserver, kubelet)</p> <p>Merge PR117550 to add <code>--concurrent-cron-job-syncs</code> parameter to set the worker quantity of the cronjob controller. (kube-controller-manager)</p> <p>Error in fixing the scheduler iterator, avoiding score exceeding the range of a single plugin due to adding duplicate scorePlugin. (kube-scheduler)</p> <p>Read cache optimization. (kube-apiserver)</p> <p>Extend the self-signed certificate validity from 1 year to 5 years. (kube-apiserver)</p>
2023-03-30	v1.26.1-tke.1	<p>Support specified Pods during scale-down. (kube-controller-manager)</p> <p>Avoid scheduling Pods using cbs to external CHC Nodes. (kube-scheduler)</p> <p>Fix issues with LB health checks and IPVS. (kube-proxy)</p> <p>Optimize super node HPA. (kube-controller-manager)</p> <p>Support extended scheduler Prebind and Unreserve operations. (kube-scheduler)</p> <p>Support binding LB addresses to IPVS NICs. (kube-proxy)</p> <p>Support native node scaling feature. (kube-scheduler)</p> <p>Enhancements to kube-apiserver's priority and fairness features. (kube-apiserver)</p> <p>APF Enhancement: Sort by pl name upon <code>dump_priority_levels</code>, and add PL debug information, which is merged into community PR112393.</p> <p>Add metrics for list/watch. (kube-apiserver)</p> <p>Support <code>exec/logs</code> commands for super nodes and DaemonSet Pods. (kube-apiserver)</p> <p>Add metrics for network/disk, etc., for the Containerd runtime. (kubelet)</p> <p>When kube-controller-manager performs pod gc, prioritize deleting Evicted pods. (kube-controller-manager)</p>

TKE kubernetes 1.24.4 revisions

Note:

Kubernetes 1.24 has removed support for Docker through Dockershim. For the container runtime on new nodes, please use Containerd. Images built with Docker can continue to be used. For more information, please refer to

[Dockershim Removal FAQ](#).

Time	Version	Update Contents

2024-04-10	v1.24.4-tke.14	Merge PR119561 to fix the issue where kubelet using gpu might cause panic. (kubelet) Merge PR113145 to fix the issue where forced pod deletion might leave behind a sandbox, preventing pod reconstruction. (kubelet)
2024-02-27	v1.24.4-tke.13	Sandbox optimization for super nodes retained. (kube-scheduler) Super nodes allow scheduling Pods in TKE trunking ENI mode. (kube-scheduler) Optimization of computing resource policies for super nodes. (kube-scheduler)
2024-02-01	v1.24.4-tke.12	kube-proxy Optimization. (kube-proxy) job-controller support settings for multi-worker parameters --concurrent-job-syncs. (kube-controller-manager)
2023-12-04	v1.24.4-tke.11	Pods with added annotation support mounting cgroups. (kubelet) When deleting resources, add userAgent reporting and record detailed logs. (kube-apiserver)
2023-10-12	v1.24.4-tke.10	Port CVE-2022-3162 . (kube-apiserver) Read cache optimization. (kube-apiserver) Extend the self-signed certificate validity from 1 year to 5 years. (kube-apiserver) Merge PR116327 to optimize memory usage in GetList. (kube-apiserver) Enhance monitoring metrics collection. (kube-scheduler)
2023-07-21	v1.24.4-tke.9	Enable kube-proxy to detect changes in node PodCIDs and ensure consistency through restarts upon changes. (kube-proxy) Fix the issue where kubelet may panic when assigning devices to pods. (kubelet) Error in fixing the scheduler iterator, avoiding score exceeding the range of a single plugin due to adding duplicate scorePlugin. (kube-scheduler) Port CVE-2022-3294 (kube-apiserver)
2023-05-30	v1.24.4-tke.7	Support pod resource in-place updates. (kube-apiserver, kube-scheduler, kubelet) Optimize scheduling logs. (kube-scheduler) Merge PR117550 , to add --concurrent-cron-job-syncs parameter to set the worker count for the cronjob controller. (kube-controller-manager)
2023-03-17	v1.24.4-tke.6	Ignore TKEDefaultQuota conflict errors, and increase worker count to resolve resource quota evaluates timeout errors. (kube-apiserver) Support specified Pods during scale-down. (kube-controller-manager) Prevent Pods using CBS from being scheduled to external CHC nodes. (kube-scheduler) Merge PR110951 to avoid the issue where sporadic PV mount operations get stuck, causing Pods associated with that PV to remain in the Creating state. (kubelet) Merge PR111635 to fix the issue where ipvs real server weight is 0 after kube-proxy restarts. (kube-proxy) Support native node scaling feature. (kube-scheduler) Support binding LB addresses to IPVS NICs. (kube-proxy)

		<p>Enhancements to kube-apiserver priority and fairness features: (kube-apiserver)</p> <p>a. Support rate limiting based on userAgent, support Token Bucket for rate limiting, support rate limiting, circuit breaking for any type of component.</p> <p>b. When dumping priority levels, sort by pl name; add PL debug information: Total Requests Processed, Total Requests Rejected, Total Requests Timed Out, Total Requests Cancelled.</p> <p>Record metrics when list/watch is too old, too large, and Excessive Objects events (kube-apiserver): <code>watch_too_old_objects_events_total</code> , <code>list_too_large_objects_events_total</code> , <code>list_too_many_objects_events_total</code> , <code>watch_too_many_objects_events_total</code></p> <p>Support network metrics for Containerd runtime. (kubelet)</p> <p>Support Tencent Cloud CBS CSI migration. (kube-apiserver,kube-controller-manager,kubelet)</p> <p>When kube-controller-manager performs pod gc, prioritize deleting Evicted pods. (kube-controller-manager)</p>
2023-02-16	v1.24.4-tke.5	<p>Support exec/logs commands for super node DaemonSet Pod. (kube-apiserver)</p> <p>Super node supports fixed IP address. (kube-scheduler)</p> <p>Port 1.26 to support multiple HPAController workers. (kube-controller-manager)</p> <p>Port the community heuristic feature to set watch channel size and support setting certain resources' watch channel size via env. (kube-apiserver)</p> <p>Optimize self-adaptive node iptables mode Script, and fix misjudgments in some cases. (kube-proxy)</p>
2022-12-08	v1.24.4-tke.3	<p>Support Disk Usage Metrics upon Containerd runtime. (kubelet)</p> <p>Support extended scheduler Prebind and Unreserve operations. (kube-scheduler)</p> <p>Support TKE Serverless super nodes. (kube-controller-manager,kube-scheduler)</p>
2022-05-07	v1.24.4-tke.1	<p>Update the Startup mode for running kube-proxy in an Image manner, and automatically adapt to the iptables running mode of the node, to support the operating system that runs iptables in <code>nf_tables</code> mode by default. (kube-proxy)</p> <p>Allows special IP ranges used by TKE managed clusters. (kube-apiserver)</p> <p>revert pr63066, fixing the issue with LB health check and IPVS. (kube-proxy)</p> <p>Merge PR112299 to optimize the gzip compression level in apiserver.</p>

TKE kubernetes 1.22.5 revisions

Time	Version	Update Contents
2024-04-10	v1.22.5-tke.24	<p>Merge PR119561 to fix the issue where kubelet using gpu might cause panic. (kubelet)</p> <p>Merge PR106570 to fix the issue with kubelet setting cpushares. (kubelet)</p>

2024-02-27	v1.22.5-tke.23	Super nodes allow scheduling Pods in TKE trunking ENI mode. (kube-scheduler) Optimization of computing resource policies for super nodes. (kube-scheduler)
2024-02-01	v1.22.5-tke.22	job-controller supports settings for multi-worker parameters --concurrent-job-syncs. (kube-controller-manager)
2023-12-04	v1.22.5-tke.21	Port pr108753 to fix the issue where pod memory and other metrics couldn't be retrieved through kubelet /stats/summary when using systemd cgroup. (kubelet) Pods with added annotation support mounting cgroupfs. (kubelet) When deleting resources, add userAgent reporting and record detailed logs. (kube-apiserver)
2023-10-12	v1.22.5-tke.20	Extend the self-signed certificate validity from 1 year to 5 years. (kube-apiserver) Merge PR116327 to optimize memory usage in GetList. (kube-apiserver) Merge PR107900 to fix the issue where static pods may fail to start under certain scenarios. (kubelet) Revert the fix for the upgrade issue caused by ubuntu16.04lxcfs, as the fix is no longer needed and may lead to e2e test failures. (kubelet) Port CVE-2022-3162 . (kube-apiserver)
2023-08-02	v1.22.5-tke.19	Read cache optimization. (kube-apiserver)
2023-07-21	v1.22.5-tke.18	Fix the issue where kubelet may panic when assigning devices to pods. (kubelet) Fix the bug in the super node's definition policy. (kube-scheduler) Port CVE-2022-3294 (kube-apiserver)
2023-05-26	v1.22.5-tke.16	Port PR105606 & PR105872 to fix the goroutine leak issue caused by DeleteCollection. (kube-apiserver) Port PR117550 , to add --concurrent-cron-job-syncs parameter to set the worker count for the cronjob controller. (kube-controller-manager)
2023-04-20	v1.22.5-tke.15	Skip the creation of new container-related files when cadvisor collects metrics to avoid reporting null values and causing breakpoints. (kubelet) Merge PR112913 to fix the issue with cpu manager calculating allocatable cpusets. (kubelet) Improve the processing logic for in-place Pod configuration changes. (kube-apiserver,kube-scheduler,kubelet) Merge PR107670 to fix a potential timing bug upon Pod deletion that may lead to sandbox leakage. (kubelet)
2023-03-24	v1.22.5-tke.14	Merge PR108202 to fix the missing nodeName field in EndpointSlice after upgrading from 1.20 to 1.22. (kube-apiserver) Merge PR110245 to add a fallback logic to kube-proxy that uses <code>endpoint.DeprecatedTopology[v1.LabelHostname]</code> when <code>endpoint.NodeName</code> is empty, to avoid issues of service inaccessibility after upgrading from 1.20 to 1.22. (kube-proxy)

2023-03-17	v1.22.5-tke.13	<p>Merge PR110951 to avoid the issue where sporadic PV mount operations get stuck, causing Pods associated with that PV to remain in the Creating state. (kubelet)</p> <p>Fix the issue where some GET requests were not logged as WATCH. (kube-apiserver)</p> <p>When kube-controller-manager performs pod gc, prioritize deleting Evicted pods. (kube-controller-manager)</p> <p>Support binding LB addresses to IPVS NICs. (kube-proxy)</p>
2023-03-06	v1.22.5-tke.12	<p>Merge PR111635 to fix the issue where ipvs real server weight is 0 after kube-proxy restarts. (kube-proxy)</p>
2023-02-16	v1.22.5-tke.11	<p>Super node supports fixed IP address. (kube-scheduler)</p> <p>Port 1.26 to support multiple HPAController workers. (kube-controller-manager)</p> <p>Port the community heuristic feature to set watch channel size and support setting certain resources' watch channel size via env. (kube-apiserver)</p> <p>Port PR108366, and fix kubelet reporting pod status as finished too early, leading to insufficient resources when new pods are running. (kubelet)</p> <p>Fix inotify leaks causing kubelet memory leaks. (kubelet)</p> <p>Remove logs related to <code>watch_too_old_objects_events_total</code>, <code>list_too_many_objects_events_total</code>, <code>watch_too_many_objects_events_total</code>, avoid printing excessive information. (kube-apiserver)</p> <p>Optimize self-adaptive node script in iptables mode, and fix misjudgments in some cases. (kube-proxy)</p>
2023-01-12	v1.22.5-tke.9	<p>Support native node scaling feature. (kube-scheduler)</p> <p>Add a <code>delegated_to_storage</code> Tag to the <code>list_too_many_objects_events_total</code> metric, indicating whether the request is forwarded to etcd. (kube-apiserver)</p>
2023-01-05	v1.22.5-tke.8	<p>Enhancements to kube-apiserver priority and fairness features: (kube-apiserver)</p> <ol style="list-style-type: none"> When applying rate limiting, consider the number of objects of the list request's resource type. The more objects, the lower the concurrent QPS, and the less resources APIServer will occupy, significantly reducing resource usage and effectively enhancing cluster stability to avoid avalanches. If business components issue a large number of list requests, upgrades may result in increased latency, with APIServer returning rate limiting and retries, etc. Support rate limiting based on userAgent, token bucket rate limiting, and rate limiting or circuit breaking for any type of component. Support rate limiting for update requests related to objects involved in a watch to avoid generating too many update events at once. <p>Super nodes support running DaemonSet Pods when needed. (kube-controller-manager)</p> <p>Support <code>exec/logs</code> commands for super node DaemonSet Pod. (kube-apiserver)</p> <p>Record metrics when encountering list/watch too old, too large, or too many objects events (kube-apiserver): <code>watch_too_old_objects_events_total</code> , <code>list_too_large_objects_events_total</code> ,</p>

		<p><code>list_too_many_objects_events_total</code> , <code>watch_too_many_objects_events_total</code></p> <p>Merge PR108754 to fix the issue of reporting not registered error when occasionally mounting ConfigMap/Secrets volumes. (kubelet)</p> <p>Support in-place updates for Pod resources. (kube-apiserver, kubelet)</p>
2022-12-08	v1.22.5-tke.7	Fix the problem of container network interface default name being set to eth1. (kubelet)
2022-11-09	v1.22.5-tke.6	<p>Optimize monthly subscription super node scheduling. (kube-scheduler)</p> <p>Optimize cloud-native node stability. (kubelet)</p> <p>Optimize local replica count scheduling for super nodes.</p> <p>Add QoSIgnoreInitContainer switch, which ignores resources of init containers when setting pod QoS upon being turned on. (kube-apiserver, kubelet)</p>
2022-10-13	v1.22.5-tke.5	<p>Serverless clusters support daemonset. (kube-apiserver)</p> <p>Merge PR112299 to optimize the gzip compression level in apiserver. (kube-apiserver)</p> <p>Fixed CVE-2022-3172. (kube-apiserver)</p> <p>Add <code>calling_webhook_timeout_error</code> metric. (kube-apiserver)</p> <p>Ignore TKEDefaultQuota conflict errors, increase worker count to resolve resource quota evaluates timeout errors. (kube-apiserver)</p> <p>Merge PR110294 to fix the issue where the Job activeDeadlineSeconds setting was ineffective. (kube-controller-manager)</p> <p>Merge PR111773 to fix memory leak issues during scheduler preemption. (kube-scheduler)</p> <p>Support containerd disk metrics. (kubelet)</p> <p>Merge PR108831 to fix kubelet panic when multiple Pods are created simultaneously. (kubelet)</p> <p>Support specified Pods during scale-down. (kube-controller-manager)</p> <p>Optimize daemonset pod scheduling performance by only processing assigned nodes. (kube-scheduler)</p> <p>Set CBS CSI Migration to completed status. (kube-controller-manager)</p> <p>Support extended scheduler Prebind and Unreserve operations. (kube-scheduler)</p> <p>Optimize HPA for super nodes in Serverless Clusters. (kube-controller-manager)</p>
2022-05-07	v1.22.5-tke.1	<p>Allows special IP ranges used by TKE managed clusters. (kube-apiserver)</p> <p>Restore pr63066 to fix the issue with LB health check and IPVS. (kube-proxy)</p> <p>Merge pr90260 to fix the missing Cluster Network Monitoring issue in containerd clusters. (kubelet)</p> <p>Fix the Exit Issues of Pods caused by the upgrade of lxcfs under ubuntu16. (kubelet)</p> <p>Prevent Pods using CBS from being scheduled to external CHC nodes. (kube-scheduler)</p> <p>Support Tencent Cloud CBS CSI Migration. (kube-controller-manager,kubelet)</p> <p>Merge pr106906 to detect whether the subpath of the Network Storage Volume has been deleted, to avoid Pods remaining in terminating Status. (kubelet)</p>

Update the Startup mode for running kube-proxy as an image, automatically adapting to the iptables Running Mode of the node, to support Operating Systems that run iptables in nf_tables mode by default. (kube-proxy)

TKE kubernetes 1.20.6 revisions

Time	Version	Update Contents
2024-04-10	v1.20.6-tke.43	Customizable Eviction Delay by Definition, determining the order of availability zone eviction. (kube-scheduler)
2024-02-27	v1.20.6-tke.42	Sandbox optimization for super nodes retained. (kube-scheduler) Workload Mounted Disk Performance Optimization. (kube-controller-manager) Super nodes allow scheduling Pods in TKE trunking ENI mode. (kube-scheduler) Super Node Computing Resource Policy Optimization. (kube-scheduler)
2024-02-01	v1.20.6-tke.41	Native Node Support for CVM Mode. (kubelet) Merge PR102014 to fix the issue that CPUs bound to init containers cannot be used after release. (kubelet) job-controller support settings for multi-worker parameters --concurrent-job-syncs. (kube-controller-manager)
2023-12-04	v1.20.6-tke.40	Pods with added annotation support mounting cgroups. (kubelet) scheduler patch pod optimization. (kube-scheduler) When deleting resources, add userAgent reporting and record detailed logs. (kube-apiserver)
2023-10-12	v1.20.6-tk.39	Extend the self-signed certificate validity from 1 year to 5 years. (kube-apiserver) Revert the fix for the upgrade issue caused by ubuntu16.04lxcfs, as the fix is no longer needed and leads to e2e test failures. (kubelet)
2023-08-02	v1.20.6-tke.38	Merge PR103320 , Bump SMD to v4.1.2 to pick up #102749 fix. Read cache optimization. (kube-apiserver)
2023-07-20	v1.20.6-tke.37	Merge PR116327 , Optimize memory usage of List. (kube-apiserver) Make hpa controller's qps/burst configurable. (kube-controller-manager) Port PR107831 , Include pod UID in kubelet's secret/configmap cache. (kubelet) Reinforce cadvisor collection. (kubelet) Optimize the performance of the attach detach controller and accelerate processing speed during high-frequency node updates. (kube-controller-manager) Merge CVE-2022-3294 (kube-apiserver) Fix the bug in the super node's Definition policy. (kube-scheduler)
2023-05-26	v1.20.6-tke.34	Fix the issue where the watch_cache_capacity metrics for CRD instances are all Unstructured in their resource names, making them indistinguishable and affecting

		<p>each other. (kube-apiserver)</p> <p>Port PR116167, prioritize deleting Evicted pods during pod GC. (kube-controller-manager)</p> <p>Port PR97845, add apiserver_terminated_watchers_total metric, representing the number of watchers closed due to not being processed in time. (kube-apiserver)</p> <p>Port PR105606&PR105872 to fix the goroutine leak issue caused by DeleteCollection. (kube-apiserver)</p> <p>Port PR117550, adding --concurrent-cron-job-syncs parameter to set the worker count for the cronjob controller. (kube-controller-manager)</p> <p>Improve the processing logic for in-place Pod configuration changes. (kube-apiserver,kube-scheduler,kubelet)</p>
2023-03-10	v1.20.6-tke.33	<p>Merge PR110951, to avoid the issue where sporadic PV mount operations get stuck, causing Pods associated with that PV to remain in the Creating state. (kubelet)</p> <p>Merge PR106295, Device Manager reads the status file while adapting to both v1 and v2 formats. (kubelet)</p> <p>Fix the issue where some GET requests were not logged as WATCH. (kube-apiserver)</p>
2023-03-06	v1.20.6-tke.32	<p>Merge PR111635 to fix the issue where ipvs real server weight is 0 after kube-proxy restarts. (kube-proxy)</p>
2023-02-13	v1.20.6-tke.31	<p>Port 1.26 to support multiple HPAController workers. (kube-controller-manager)</p> <p>Port the community heuristic feature to set watch channel size and support setting certain resources' watch channel size via env. (kube-apiserver)</p> <p>Fix inotify leaks causing kubelet memory leaks. (kubelet)</p> <p>Remove logs related to watch_too_old_objects_events_total, list_too_many_objects_events_total, watch_too_many_objects_events_total, avoid printing excessive information. (kube-apiserver)</p> <p>Score super nodes based on the super node <code>eks.tke.cloud.tencent.com/eklet-priority</code> annotation. (kube-scheduler)</p> <p>Fix the bug with super node's fixed IP. (kube-scheduler)</p> <p>Optimize self-adaptive node iptables mode Script, fixing misjudgments in some cases. (kube-proxy)</p> <p>Fix the issue where the metric rest_client_rate_limiter_duration_seconds was unregistered. (kube-apiserver, kube-controller-manager, kube-scheduler, kubelet, kube-proxy)</p>
2023-01-05	v1.20.6-tke.30	<p>Merge PR104833 to optimize the controller's lock in API Priority and Fairness to a read-write lock. (kube-apiserver)</p> <p>Super nodes support running DaemonSet Pods when needed. (kube-controller-manager)</p> <p>Merge PR102576, and no longer use RemoveAll when deleting legacy Pod volumes to avoid deleting volume-related data. (kubelet)</p>

		<p>Allow DaemonSet pods with specific annotations to be scheduled on super nodes. (kube-controller-manager)</p> <p>Add a switch for the local replica count policy for super nodes. (kube-scheduler)</p> <p>Enhancements to kube-apiserver priority and fairness features: (kube-apiserver)</p> <p>a. When applying rate limiting, consider the number of objects of the list request's resource type. The more objects, the lower the concurrent QPS, and the less resources APIServer will occupy, significantly reducing resource usage and effectively enhancing cluster stability to avoid avalanches. If business components issue a large number of list requests, upgrades may result in increased latency, with APIServer returning rate limiting and retries, etc.</p> <p>b. Support rate limiting based on userAgent, token bucket rate limiting, and rate limiting or circuit breaking for any type of component.</p> <p>c. Support rate limiting for update requests related to objects involved in a watch to avoid generating too many update events at once.</p> <p>Super node supports fixed IP address. (kube-scheduler)</p> <p>Add a switch for features related to super nodes. (kube-scheduler)</p> <p>Record metrics when list/watch is too old, too large, and excessive objects (kube-apiserver): <code>watch_too_old_objects_events_total</code> , <code>list_too_large_objects_events_total</code> , <code>list_too_many_objects_events_total</code> , <code>watch_too_many_objects_events_total</code></p> <p>Add watch cache capacity metrics: <code>watch_cache_capacity</code>. (kube-apiserver)</p> <p>Use ConfigMap informer to optimize super node scheduling performance. (kube-scheduler)</p> <p>Support setting the minimum number of watch cache entries through the environment variable <code>DEFAULT_LOWER_BOUND_CAPACITY</code>. (kube-apiserver)</p> <p>Add AllowLocalConnection switch, which allows external requests to access local addresses when enabled. (kubelet)</p>
<p>2022-11-09</p>	<p>v1.20.6-tke.28</p>	<p>Optimize monthly subscription super node scheduling. (kube-scheduler)</p> <p>Optimize cloud-native node stability. (kubelet)</p> <p>Fix native node in-place upgrade and downgrade kubelet abnormally exits. (kubelet)</p> <p>Optimize super node local replica number scheduling. (kube-scheduler)</p> <p>Add QosIgnoreInitContainer switch, which ignores resources of init containers when setting pod QoS upon being turned on. (kube-apiserver, kubelet)</p>
<p>2022-10-13</p>	<p>v1.20.6-tke.27</p>	<p>Optimize scheduling policy for super node IP management. (kube-scheduler)</p> <p>Merge PR112299 to optimize the gzip compression level in apiserver. (kube-apiserver)</p> <p>Fix CVE-2022-3172. (kube-apiserver)</p> <p>Add <code>calling_webhook_timeout_error</code> metric. (kube-apiserver)</p> <p>Ignore TKEDefaultQuota conflict errors, and increase worker count to resolve resource quota evaluates timeout errors. (kube-apiserver)</p> <p>Optimize Serverless clusters local replica number scheduling. (kube-scheduler)</p> <p>Cancel fixed eip forced scheduling to the original super node. (kube-scheduler)</p>

		<p>Merge PR99324, and immediately reset after the kube-controller-manager token becomes invalid. (kube-controller-manager)</p> <p>Merge PR101155 to support setting multiple ServiceAccount Token Issuers. (kube-apiserver)</p> <p>Add StatusDetaching status for CBS Disks. (kube-controller-manager)</p> <p>Extend the timeout for updating resources when scheduling Static IP to super nodes, to prevent update failures during large-scale concurrency. (kube-scheduler)</p> <p>Serverless clusters support daemonset. (kube-apiserver)</p>
2022-09-07	v1.20.6-tke.24	<p>Optimize scheduler preemption to avoid crashes. (kube-scheduler)</p> <p>Optimize super node scheduling. (kube-scheduler)</p> <p>Support in-place updates for Pod resources. (kube-apiserver, kubelet)</p> <p>Support Pods under kube-system to be scheduled to prepaid super nodes. (kube-scheduler)</p> <p>Optimize super node HPA. (kube-controller-manager)</p> <p>Merge PR110294 to fix the issue where the Job activeDeadlineSeconds setting was ineffective. (kube-controller-manager)</p> <p>Merge PR111773 to fix the memory leak issues during scheduler preemption. (kube-scheduler)</p> <p>Merge PR97348 to fix the issue where HPA scaling numbers are incorrect when StabilizationWindowSeconds is set. (kube-controller-manager)</p> <p>Optimize monthly subscription super node scheduling.</p> <p>Merge PR108831 to fix the issue where creating multiple Pods simultaneously causes kubelet to panic. (kubelet)</p> <p>Repair cronjob Pod creation failure when Pod name/uid is empty. (kube-controller-manager)</p>
2022-07-27	v1.20.6-tke.21	<p>CBS supports native nodes. (kubelet)</p> <p>Optimize TKE Serverless super node HPA.</p>
2022-06-16	v1.20.6-tke.20	<p>Support monthly subscription super nodes. (kube-scheduler, kube-controller-manager)</p> <p>When using docker and overlay2, disk usage is retrieved through fs quota to enhance performance. (kubelet)</p> <p>Optimize daemonset pod scheduling performance by only processing assigned nodes. (kube-scheduler)</p> <p>Optimize TKE Serverless scheduling. (kube-scheduler)</p> <p>TKE Serverless: Support mounting PVC after pod creation. (kube-scheduler)</p> <p>When scheduling to TKE Serverless nodes, you can ignore hugepages resources through a feature toggle. (kube-scheduler)</p>
2022-04-22	v1.20.6-tke.17	<p>TKE Serverless: Support retaining sandbox features. (kube-scheduler)</p> <p>Merge pr101093 to avoid the issue where startupProbe does not run after container restart. (kubelet)</p>
2022-03-	v1.20.6-	<p>Resolve the issue where inline CSI and ephemeral generic temporary volumes</p>

24	tke.16	become unavailable after upgrading to 1.20. (kube-apiserver, kube-controller-manager, kube-scheduler, kubelet, kubectrl)
2022-03-18	v1.20.6-tke.15	Support specifying a pod during scaling down. (kube-controller-manager) Merge pr106906 to check if the subpath of network storage volumes has been deleted, avoiding pods being stuck in terminating state. (kubelet) When scheduling anti-affinity based on hostname, ignore TKE Serverless super nodes. (kube-scheduler) Support upgrading from tke1.18 to 1.20. (kube-apiserver, kube-controller-manager, kubelet) Port pr108325 to fix the issue where kubelet might panic during startup due to the deletion of the pause container. (kubelet) Support extended scheduler Prebind and Unreserve operations. (kube-scheduler)
2022-01-20	v1.20.6-tke.12	TKE Serverless Rescheduling Optimization: For the same availability zone's super nodes that have been evicted, reduce their score. (kube-scheduler) The apiserver supports integration of ExternalName type external services. (kube-apiserver) Support binding LB addresses to IPVS NICs. (kube-proxy)
2021-12-09	v1.20.6-tke.9	Optimize TKE Serverless super node scheduling and HPA. (kube-controller-manager, kube-scheduler) Fix the inconsistency in calculating CPU resources by TKE Serverless with the frontend. (kube-scheduler)
2021-12-02	v1.20.6-tke.8	Optimize grpc logs to avoid excessive logging when kubelet collects volume status. (kubelet) Avoid scheduling Pods using cbs to external CHC Nodes. (kube-scheduler)
2021-11-26	v1.20.6-tke.7	Add support for custom installation of other CNI in hybrid cloud external nodes. (kube-controller-manager) Avoid unnecessary processing of updates after Pod is Assumed. (kube-scheduler) Merge pr99336 to improve the synchronization mechanism of node information when kubelet starts. (kubelet)
2021-10-13	v1.20.6-tke.6	Merge 89465 to fix the issue of incorrect calculation of the number of instances by HPA based on pod metrics during rolling updates. (kube-controller-manager)
2021-09-27	v1.20.6-tke.5	Support for collecting disk usage metrics for containerd runtime. (kubelet)
2021-09-23	v1.20.6-tke.4	Fix the issue of no data for storage metrics when using cgroup v2. (kubelet) Fix CVE-2021-25741 to prevent illegal access to host files through symbolic links. (kubelet)
2021-07-19	v1.20.6-tke.3	From TKE Cluster scaling to TKE Serverless, when batch scheduling Pods, it can correctly perceive the remaining IPs in the subnet and schedule the appropriate

		<p>number of Pods to super nodes. (kube-scheduler)</p> <p>Port the upstream modifications for kubelet and cadvisor to fix the issue with metric collection statistics when using cgroupv2. (kubelet)</p>
2021-06-21	v1.20.6-tke.2	By default, CSIMigration and CSIMigrationQcloudCbs are enabled to mount CBS disks in CSI mode.
2021-05-25	v1.20.6-tke.1	<p>revert pr63066, fix the issue with LB health checks and IPVS. (kube-proxy)</p> <p>Merge pr90260 to fix the issue of missing network monitoring in containerd cluster. (kubelet)</p> <p>Fix the issue caused by the upgrade of lxcfs under ubuntu16 where Pods exit. (kubelet)</p> <p>Merge pr72914 to fix the issue where deleting a Pod and immediately creating and scheduling it to the same node might prevent it from being successfully mounted. (kube-controller-manager)</p> <p>Solve the issue where creating containers in CentOS leads to cgroup leaks. (kubelet)</p> <p>Merge pr98262 to support dynamic adjustment of log levels in kube-controller-manager. (kube-controller-manager)</p> <p>Merge pr97752 to fix the issue where NewReplicaSet is displayed when describing deployments. (kubectl)</p> <p>Merge pr94833 to fix the issue where the image tag in the status does not match when a Pod has multiple tags. (kubelet)</p> <p>Merge pr100060 and automatically delete directories of volumes left by orphaned Pods. (kubelet)</p> <p>kube-controller-manager supports super nodes. (kube-controller-manager)</p> <p>kube-scheduler supports retaining a fixed number of local replicas when scaling out to TKE Serverless in a hybrid cloud. (kube-scheduler)</p> <p>Support cbs csi migration. (kube-controller-manager, kubelet)</p> <p>Merge pr93260 to solve the problem of node startup slowing down caused by AWS Credential Provider. (kubelet)</p> <p>Add command-line parameter eks-config-namespace to the scheduler: Specify the namespace where the configuration related to scaling out TKE Serverless is located. (kube-scheduler)</p> <p>TKE supports hybrid cloud nodes. (kube-controller-manager)</p>

(Discontinued) TKE kubernetes 1.18.4 revisions

Time	Version	Update Contents
2024-04-10	v1.18.4-tke.44	Mounting PVC optimization. (kube-scheduler)
2024-02-	v1.18.4-	Super nodes allow scheduling Pods in TKE trunking ENI mode. (kube-scheduler)

27	tke.43	Optimization of computing resource policies for super nodes. (kube-scheduler)
2024-02-01	v1.18.4-tke.42	Native Node Support for CVM Mode. (kubelet) Merge PR102014 to fix the issue that CPUs bound to init containers cannot be used after release. (kubelet)
2023-12-04	v1.18.4-tke.41	CVE-2022-3294 (kube-apiserver) Read cache optimization. (kube-apiserver) Pods with added annotation support mounting cgroupfs. (kubelet) When deleting resources, add userAgent reporting and record detailed logs. (kube-apiserver)
2023-07-20	v1.18.4-tke.39	Support scheduling daemonset pods to monthly subscription super nodes. (kube-scheduler) Port PR107831 to include pod UID in kubelet's secret/configmap cache. (kubelet) Fix pod mounting disk anomaly issue. (kube-controller-manager/kubelet)
2023-05-26	v1.18.4-tke.37	Improve the processing logic for in-place Pod configuration changes. (kube-apiserver,kube-scheduler,kubelet)
2023-05-16	v1.18.4-tke.36	When kube-controller-manager performs pod gc, prioritize deleting Evicted pods. (kube-controller-manager) Fix the issue where some GET requests were not logged as WATCH. (kube-apiserver)
2023-03-03	v1.18.4-tke.35	Add <code>delegated_to_storage</code> tag to the <code>list_too_many_objects_events_total</code> metric, indicating whether the request is passed through to backend storage. (kube-apiserver) Port PR89160 to only delete the soft links related to a container after it has completely exited. (kubelet) Support exec/logs commands for super node DaemonSet Pod. (kube-apiserver) Optimize super node scheduling, add a unified switch <code>EnableSuperNodeFeatures</code> . (kube-scheduler) Fix inotify leaks causing kubelet memory leaks. (kubelet) Remove logs related to <code>watch_too_old_objects_events_total</code> , <code>list_too_many_objects_events_total</code> , <code>watch_too_many_objects_events_total</code> , to avoid printing excessive information. (kube-apiserver) Optimize self-adaptive node iptables mode Script, fixing misjudgments in some cases. (kube-proxy) Fix the issue where the metric <code>rest_client_rate_limiter_duration_seconds</code> was unregistered. (kube-apiserver, kube-controller-manager, kube-scheduler, kubelet, kube-proxy) Merge PR110951 to avoid the issue where sporadic PV mount operations get stuck, causing Pods associated with that PV to remain in the Creating state. (kubelet)

		Merge PR111635 to fix the issue where ipvs real server weight is 0 after kube-proxy restarts. (kube-proxy)
2023-01-05	v1.18.4-tke.34	<p>Super nodes support running DaemonSet Pods when needed. (kube-controller-manager)</p> <p>Merge PR102576 to no longer use RemoveAll when deleting legacy Pod volumes to avoid deleting volume-related data. (kubelet)</p> <p>Support lowering resource limits through in-place updates. (kube-apiserver)</p> <p>Fix the issue to not delete terminated containers during log rotation. (kubelet)</p> <p>Record metrics when list/watch is too old, too large, or there are too many objects (kube-apiserver): <code>watch_too_old_objects_events_total</code> , <code>list_too_large_objects_events_total</code> , <code>list_too_many_objects_events_total</code> , <code>watch_too_many_objects_events_total</code></p>
2022-11-09	v1.18.4-tke.33	<p>Optimize monthly subscription super node scheduling. (kube-scheduler)</p> <p>Optimize local replica count scheduling for super nodes.</p>
2022-10-13	v1.18.4-tke.30	<p>Merge PR112299 to optimize the gzip compression level in apiserver. (kube-apiserver)</p> <p>Fix CVE-2022-3172. (kube-apiserver)</p> <p>Add <code>calling_webhook_timeout_error</code> metric. (kube-apiserver)</p> <p>Ignore TKEDefaultQuota conflict errors and increase worker count to resolve resource quota evaluates timeout errors. (kube-apiserver)</p> <p>Add StatusDetaching status for CBS disks. (kube-controller-manager)</p> <p>Serverless clusters support daemonset. (kube-apiserver)</p>
2022-09-07	v1.18.4-tke.28	<p>Optimize List performance in large-scale clusters. (kube-apiserver)</p> <p>Optimize super node scheduling. (kube-scheduler)</p> <p>Support Pods under kube-system to be scheduled to prepaid super nodes. (kube-scheduler)</p> <p>Optimize super node HPA. (kube-controller-manager)</p> <p>Support in-place updates for Pod resources. (kube-apiserver, kubelet)</p> <p>Merge PR97348 to fix the issue where HPA scaling numbers are incorrect when StabilizationWindowSeconds is set. (kube-controller-manager)</p> <p>Optimize monthly subscription super node scheduling.</p> <p>Merge PR108831 to fix the issue where creating multiple Pods simultaneously causes kubelet to panic. (kubelet)</p>
2022-07-27	v1.18.4-tke.26	CBS supports native nodes. (kubelet)
2022-03-18	v1.18.4-tke.23	Merge pr92878 to only print a warning message when setting ConfigMap/Secret volume Ownership takes more than 30 seconds, to avoid generating excessive log information. (kubelet)

		<p>Merge pr106906 to check if the subpath of network storage volumes has been deleted, to prevent pods from being stuck in terminating status. (kubelet)</p> <p>When scheduling anti-affinity based on hostname, ignore TKE Serverless super nodes. (kube-scheduler)</p> <p>Merge pr93026 to fix the issue where DefaultPodTopologySpread cannot retrieve replicaset information. (kube-scheduler)</p>
2022-01-20	v1.18.4-tke.20	<p>TKE Serverless Rescheduling Optimization: For the same availability zone's super nodes that have been evicted, lower their score. (kube-scheduler)</p> <p>apiserver supports integrating ExternalName 556 type external services. (kube-apiserver)</p> <p>Support binding LB addresses to IPVS NICs. (kube-proxy)</p>
2021-12-09	v1.18.4-tke.17	<p>Solve the issue where kube-controller-manager's access to api-server is rate-limited when there are a large number of volume attachment objects in the cluster. (kube-controller-manager)</p> <p>Merge PR95650, and HPA ignores deleted Pods when calculating the number of replicas. (kube-controller-manager)</p> <p>Fix the inconsistency in calculating CPU resources by TKE Serverless with the frontend. (kube-scheduler)</p>
2021-12-02	v1.18.4-tke.16	<p>Fix the bug when scheduling to super nodes. (kube-scheduler)</p> <p>Optimize the super node scheduling algorithm. (kube-scheduler)</p>
2021-11-26	v1.18.4-tke.15	<p>Merge pr96444 to return retry if there are errors while syncing RBAC policy. (kube-apiserver)</p> <p>Add support for custom installation of other CNI in hybrid cloud external nodes. (kube-controller-manager)</p> <p>Support cloud gaming Android container group core binding requirement. (kubelet)</p> <p>Support extended scheduler Prebind and Unreserve operations. (kube-scheduler)</p> <p>Merge pr99336 to improve the synchronization mechanism of node information when kubelet starts. (kubelet)</p> <p>Repair CVE-2021-25741, prevent illegal access to host files through a symbolic link. (kubelet)</p> <p>Optimize the error information caused by cbs disk creation failure leading to scheduling timeout. (kube-scheduler)</p> <p>Optimize grpc logs to avoid excessive logging when kubelet collects volume status. (kubelet)</p> <p>Avoid scheduling Pods using cbs to external CHC Nodes. (kube-scheduler)</p>
2021-08-23	v1.18.4-tke.14	<p>Scale-out from TKE Cluster to TKE Serverless: Support fixed IP. (kube-scheduler)</p> <p>Scale-out from TKE Cluster to TKE Serverless: When matching TKE Serverless fixed IP, skip other pre-selection policies. (kube-scheduler)</p> <p>From TKE Cluster scaling to TKE Serverless: Scheduling optimization for TKE Serverless nodes includes Node Resource-Aware Re-scheduling for TKE Serverless nodes; Priority Model Scheduling for TKE Serverless nodes;</p>

		<p>Optimization of Preferred/Pre-selected Policies for TKE Serverless nodes. (kube-scheduler)</p> <p>Record loaded ipvs kernel modules to avoid kube-proxy crashing in ipvs mode. (kube-proxy)</p> <p>Prevent panic when writing to the CPU manager status file fails. (kubelet)</p>
2021-07-22	v1.18.4-tke.13	Merge PR91859 to fix the issue where kube-apiserver panics when CRD types have only one letter. (kube-apiserver)
2021-07-13	v1.18.4-tke.12	<p>From TKE Cluster scaling to EKS: When batch scheduling Pods, it correctly perceives remaining Subnet IPs and schedules the appropriate number of Pods to super nodes. (kube-scheduler)</p> <p>Support for collecting disk usage metrics for Containerd runtime. (kubelet)</p> <p>Support specified Pods during scale-down. (kube-controller-manager)</p>
2021-06-05	v1.18.4-tke.11	TKE supports hybrid cloud nodes. (kube-controller-manager)
2021-05-14	v1.18.4-tke.9	<p>Port pr93370 to support CronJobControllerV2. (kube-controller-manager)</p> <p>Merge pr100376 to enable HTTP/2 health checks to avoid issues where connections cannot recover after being lost. (kube-apiserver, kube-controller-manager, kube-scheduler, kubelet, kube-proxy, kubectl)</p> <p>Merge pr100317 to fix CVE-2021-25735 where node updates could bypass the Validating Admission Webhook. (kube-apiserver)</p> <p>From TKE cluster scaling out to TKE Serverless support for ComputeResource and TKE Serverless ClusterIP, and HPA. (kube-controller-manager, kube-scheduler)</p>
2021-04-02	v1.18.4-tke.8	<p>Merge pr97752 to fix the issue where NewReplicaSet appears as <none> when describing a deployment. (kubectl)</p> <p>Merge pr93808 to fix the issue of redundant information returned by executing kube-scheduler --version. (kube-scheduler)</p> <p>Merge pr91590 to fix the warning of ports already allocated when using NodePort type with multi-protocol Service. (kube-apiserver)</p> <p>Merge pr98262 to support dynamic adjustment of log levels in kube-controller-manager. (kube-controller-manager)</p> <p>Merge pr95154 to fix the issue of kube-scheduler snapshot including nodes that are in the process of being deleted. (kube-scheduler)</p> <p>Merge pr95711 to fix the issue of high CPU usage when executing the kubectl drain command. (kubectl)</p> <p>Merge pr96602 to fix the memory leaks in apiserver during timing jumps. (kube-apiserver)</p> <p>Merge pr97023 to delete related metadata directories when uninstalling an emptyDir type volume. (kubelet)</p> <p>Merge pr97527 to fix unsynchronized map access operations in cpumanager. (kubelet)</p>

		<p>Merge pr100190 to automatically delete orphaned Pod's leftover volume directories. (kubelet)</p> <p>Merge pr92614, and once all containers in a Pod with RestartPolicyOnFailure exit successfully, no new Sandbox will be created. (kubelet)</p> <p>Merge pr94833 to fix the issue where, if a Pod's image has multiple tags, the tags in the status do not match. (kubelet)</p>
2020-12-28	v1.18.4-tke.6 (Starting from this version, ARM Clusters are supported)	<p>Add metrics for QcloudCbs. (kube-controller-manager)</p> <p>Fix the issue of superfluous spaces when checking the serial value while mounting a cbs disk. (kubelet)</p>
2020-12-21	v1.18.4-tke.5	<p>Merge pr94712 to fix CVE-2020-8564 - Docker configuration leakage when the file format is incorrect and logLevel >= 4. (kubelet)</p> <p>Merge pr95316 to fix CVE-2020-8565 - Log token leakage due to an incomplete fix of CVE-2019-11250. (logLevel >= 9) (kube-apiserver, kubectl)</p> <p>Merge pr95245 to fix CVE-2020-8566 - Exposure of Ceph RBD adminSecrets in logs when logLevel >= 4. (kube-controller-manager)</p> <p>Fix the issue where restarting kubelet causes the Pod's readiness check to fail. (kubelet)</p> <p>Merge pr90825 to fix the issue where a race condition could cause the fifo queue Pop operation in client-go to get stuck, leading to pods remaining in a pending state indefinitely. (kubelet)</p> <p>The scheduler supports super nodes. (kube-scheduler)</p> <p>kube-controller-manager supports super nodes. (kube-controller-manager)</p> <p>Set the instance-type Tag based on the actual model of the node, no longer fixed to QCLOUD. (kubelet)</p> <p>Add the CBS section in OpenAPI. (kube-apiserver)</p> <p>Merge pr91126 to fix the issue where scheduler cache is inconsistent when Pods have the same name but different UIDs. (kube-scheduler)</p> <p>Merge pr93387 to fix the issue in the scheduler where disordered node cache information prevents daemonset pods from being scheduled on certain nodes. (kube-scheduler)</p> <p>Merge pr89465 to fix the issue of incorrect instance count calculation based on pod metrics during rolling updates. (kube-controller-manager)</p>
2020-10-13	v1.18.4-tke.3	<p>Merge pr89629 to fix the issue where containers with subpath mounted from configmap never restart successfully after configmap changes. (kubelet)</p> <p>QcloudCbs supports BulkVolumeVerification. (kube-controller-manager)</p> <p>Merge pr94430 to fix the issue where client-go reflector couldn't detect the "Too large resource version" error. (kubelet)</p>

2020-08-12	v1.18.4-tke.2	Merge pr93403 to remove the erroneous print information from kubelet regarding the update of Pod Conditions that do not belong to kubelet. (kubelet)
2020-08-04	v1.18.4-tke.1	<p>Revert pr63066 to fix the issue with LB health checks and IPVS. (kube-proxy)</p> <p>Merge pr72914 to fix the issue where immediately creating and scheduling to the same node after deleting a Pod may result in unsuccessful mounting. (kube-controller-manager)</p> <p>Fix the issue where creating containers in CentOS leads to cgroup leaks. (kubelet)</p> <p>Fix for Pod exit issue caused by the upgrade of lxcfs under Ubuntu16. (kubelet)</p> <p>metadata increment Cache and Timed out. cloud-provider supports adding the node name as hostname. (kubelet)</p> <p>metadata adds local cache. (kubelet)</p> <p>Incorporate CBS and related fix code. (kubelet)</p> <p>Merge pr90260 to fix the missing Cluster Network Monitoring issue in containerd clusters. (kubelet)</p> <p>TKE supports the awareness of the maximum number of qcloudcbs that can be mounted on a single node. For versions 1.12 and above, it is maxAttachCount-2; for version 1.10, it is now defaulted to 18. (kube-scheduler)</p> <p>CBS intree resolves the issue of continuing to Uninstall disks when a disk does not exist, causing a large number of invalid requests. (kubelet)</p> <p>Merge pr2359 to fix the monitoring gap caused by being unable to obtain the docker root. (kubelet)</p> <p>kube-scheduler supports dynamic configuration of log levels. (kube-scheduler)</p> <p>Bypass the issue of missing device path (/dev/disk/by-id/virtio-xxx/...) encountered with CBS, allowing users to use CBS normally. (kubelet)</p> <p>TKE perceives the maximum number of qcloudcbs that can be mounted on a single node; kubelet does not patch the node. (kubelet)</p> <p>Merge pr89296, and no longer record whether iptables random-fully parameter is enabled. (kube-proxy)</p> <p>Fix aws issue, pr92162. (kubelet)</p> <p>Merge pr91277 to avoid the issue of CLB health checks leading to a large number of TLS handshake error logs in kube-apiserver. (kube-apiserver)</p> <p>Merge pr91500 to fix the issue of missing KUBERNETES_SERVICE_HOST environment variable. (kubelet)</p> <p>Merge pr92537 to fix the issue where client-go reflector couldn't recover from the "Too large resource version" error. (kube-apiserver, kube-controller-manager, kube-scheduler, kubelet, and kube-proxy)</p> <p>Merge pr92969 to fix the CVE-2020-8559 issue of escalating privileges from a compromised node to compromise other nodes. (kube-apiserver)</p> <p>Merge pr92921 to fix the CVE-2020-8557, a DOS attack issue by writing to "/etc/hosts" and exhausting disk space. (kubelet)</p>

(Discontinued) TKE kubernetes 1.16.3 revisions

Time	Version	Update Contents
2024-04-10	v1.16.3-tke.40	Native nodes support the cvm mode.(kubelet) Volume Reconciliation Optimization.(kube-controller-manager) Super nodes allow the scheduling of pods in TKE trunking ENI mode.(kube-scheduler) Scheduler Cache Issue Fixed.(kube-scheduler) Optimize PVC mounting. (kube-scheduler)
2023-03-24	v1.16.3-tke.35	Optimize installation script. (kubelet)
2023-03-03	v1.16.3-tke.34	Super nodes support running DaemonSet Pods when needed. (kube-controller-manager) Merge PR102576 , and no longer use RemoveAll when deleting legacy Pod volumes to avoid deleting volume-related data. (kubelet) In-place update supports reducing resource limits. (kube-apiserver) Support exec/logs commands for super node DaemonSet Pod. (kube-apiserver) Fix inotify leaks causing kubelet memory leaks. (kubelet) Optimize self-adaptive node iptables mode Script, fixing misjudgments in some cases. (kube-proxy) Merge PR111635 to fix the issue where ipvs real server weight is 0 after kube-proxy restarts. (kube-proxy)
2022-10-13	v1.16.3-tke.32	Merge PR112299 to optimize the gzip compression level in apiserver. (kube-apiserver) Fix CVE-2022-3172. (kube-apiserver) Add calling_webhook_timeout_error metric. (kube-apiserver) Ignore TKEDefaultQuota conflict errors, and add worker count to resolve resource quota evaluates timeout errors. (kube-apiserver) CPU Manager cleans up terminated containers before processing Pods, to avoid CPU allocation failure. (kubelet) Add StatusDetaching state for CBS disks. (kube-controller-manager) Serverless clusters support daemonset. (kube-apiserver) Support in-place updates for Pod resources. (kube-apiserver, kubelet)
2022-07-27	v1.16.3-tke.28	When scheduling based on hostname anti-affinity, ignore TKE Serverless virtual nodes. (kube-scheduler) TKE Serverless supports retaining sandbox features. (kube-scheduler) CBS supports native nodes. (kubelet)
2022-03-18	v1.16.3-tke.27	Support specified Pods during scale-down. (kube-controller-manager) Optimize the super node scheduling algorithm. (kube-scheduler)
2022-01-20	v1.16.3-tke.25	Support binding LB addresses to IPVS NICs. (kube-proxy)

		<p>The apiserver supports integration with ExternalName type external services. (kube-apiserver)</p> <p>Optimize TKE Serverless scheduling. (kube-scheduler)</p>
2021-12-09	v1.16.3-tke.24	Fix the issue where the TKE Serverless local replica count policy does not work on statefulset type Pods. (kube-scheduler)
2021-12-02	v1.16.3-tke.23	<p>Support extended scheduler Prebind and Unreserve operations. (kube-scheduler)</p> <p>Avoid scheduling Pods using cbs to external CHC Nodes. (kube-scheduler)</p> <p>Fix the bug when scheduling to super nodes. (kube-scheduler)</p>
2021-09-03	v1.16.3-tke.22	Avoid panic when errors occur when writing the cpu manager status file. (kubelet)
2021-08-17	v1.16.3-tke.21	<p>Optimizations for scheduling TKE Serverless nodes: Node Resource-Aware Re-scheduling for TKE Serverless nodes; Priority Model Scheduling for TKE Serverless nodes; Optimized Optimal Selection / Pre-selection Policy for TKE Serverless nodes. (kube-scheduler)</p> <p>Port 87692 to fix the issue where the scheduler's pending_pods and schedule_attempts_total metrics have no data. (kube-scheduler)</p>
2021-07-19	v1.16.3-tke.20	<p>Port 87688 and 87693 to optimize Node Authorizer performance. (kube-apiserver)</p> <p>Scaling from TKE clusters to TKE Serverless: Accurately perceive the remaining IPs of the subnet when batch scheduling Pods, scheduling the correct number of Pods to super nodes. (kube-scheduler)</p> <p>Merge pr88507 to fix the issue where podIP and podIPs are inconsistent when updating Pod status. (kube-apiserver)</p>
2021-05-24	v1.16.3-tke.17	<p>Port pr93370 to support CronJobControllerV2. (kube-controller-manager)</p> <p>Scaling from TKE clusters to TKE Serverless supports retaining the local replica count. (kube-scheduler)</p>
2021-05-06	v1.16.3-tke.16	Update the startup mode when running kube-proxy in image method to automatically adapt to the iptables running mode of the node, to support operating systems that by default use the nf_tables mode for iptables.
2021-04-14	v1.16.3-tke.15	<p>Merge pr97752 to fix the issue where NewReplicaSet appears as <none> when describing a deployment. (kubectl)</p> <p>Merge pr92614, when the restart policy is RestartPolicyOnFailure and all containers in a Pod exit successfully, a new Sandbox is no longer created. (kubelet)</p> <p>Merge pr91590 to fix the issue of port already allocated warnings when using NodePort type services with multiple protocols. (kube-apiserver)</p> <p>Merge pr98262 to support for dynamic log level adjustment in kube-controller-manager. (kube-controller-manager)</p>

		Merge pr95301 to automatically delete the volume directories left by orphan Pods. (kubelet)
2020-12-28	v1.16.3-tke.14	Add metrics for QcloudCbs. (kube-controller-manager) Fix the issue of extraneous spaces when viewing the serial value during mount of cbs disks. (kubelet)
2020-12-21	v1.16.3-tke.13	Merge pr94712 to fix CVE-2020-8564 - Docker configuration leakage when the file format is incorrect and logLevel >= 4. (kubelet) Merge pr95316 to fix CVE-2020-8565 - Log token leakage due to incomplete fix of CVE-2019-11250 (logLevel >= 9). (kube-apiserver, kubectrl) Merge pr95245 to fix CVE-2020-8566 - Exposure of Ceph RBD adminSecrets in logs when logLevel >= 4. (kube-controller-manager) Merge pr86191 to fix the issue where pods may be in an Error Status after a node restart. (kubelet) Merge pr86140 to fix the issue where the controller manager fails to properly handle Timeout Errors, preventing the creation of Scale-out pods. (kube-controller-manager) Merge pr90825 to address the issue where FIFO Queue Pop Operations in client-go might get stuck due to race conditions, causing pods to remain in a Pending State indefinitely. (kubelet) The scheduler supports super nodes. (kube-scheduler) kube-controller-manager supports super nodes. (kube-controller-manager) Set the instance-type tag based on the actual model of the node, no longer fixed to QCLOUD. (kubelet) Add the CBS section in OpenAPI. (kube-apiserver) Merge pr81344 to fix the issue where CPU Manager does not support SourcesReady. (kubelet) Merge pr91126 to fix the inconsistency in scheduler cache when Pods have the same name but different UIDs. (kube-scheduler) Merge pr89224 to fix the issue where the absence of a check in NodeInfo caused kube-scheduler to crash unexpectedly. (kube-scheduler) Merge pr89465 to fix the issue where HPA calculates the wrong number of instances during rolling updates based on pod metrics. (kube-controller-manager)
2020-10-13	v1.16.3-tke.11	Merge pr92971 to fix CVE-2020-8559 issue of escalating privileges from a compromised node to invade other nodes. (kube-apiserver) Merge pr92924 to fix CVE-2020-8557 DOS attack issue by writing to /etc/hosts to deplete disk space. (kubelet) Merge pr93403 to remove erroneous print information in kubelet for updating Pod Conditions not owned by kubelet. (kubelet) Merge pr89629 to resolve the issue where the container restart always fails after configmap changes when mounting subpath. (kubelet) QcloudCbs supports BulkVolumeVerification. (kube-controller-manager)

		Merge pr84998 to fix the issue where the corresponding node lease objects might be recreated after a node is deleted, leading to garbage data. (kubelet)
2020-07-28	v1.16.3-tke.10	Merge pr91277 to fix the issue where CLB health checks cause kube-apiserver to generate a large number of TLS handshake error logs. (kube-apiserver) Merge pr91500 to fix the issue of missing KUBERNETES_SERVICE_HOST environment variable. (kubelet)
2020-06-17	v1.16.3-tke.9	Temporary fix for an AWS issue pr92162 . Stops registering AWS Credential Provider to avoid the node startup slowdown it causes.
2020-06-11	v1.16.3-tke.8	Merge pr85993 to support setting the gateway address for kubernetes using CNI results.
2020-06-10	v1.16.3-tke.7	Merge pr90260 to fix the issue of missing network monitoring for containerd cluster. Merge pr89515 to fix the issue where HPA calculates the wrong number of instances during rolling updates. Merge pr91252 to ignore Pod Condition updates from other components to avoid unnecessary scheduling. Merge pr89794 to clean up the error logs of kube-controller-manager to prevent Half-Blind SSRF attacks in CVE-2020-8555.
2020-05-18	v1.16.3-tke.6	TKE can sense the maximum number of qcloudcbs mountable on a single node, does not support dynamically acquiring the maximum value.
2020-04-20	v1.16.3-tke.5	Merge pr69047 to solve the backward compatibility issue with node.Spec.Unschedulable (this fix was overridden when merging in-tree CBS code).
2020-04-14	v1.16.3-tke.4	Merge pr87913 to fix CVE-2020-8551: Kubelet DoS attack issue. Merge pr87669 to fix CVE-2020-8552: apiserver DoS attack issue. TKE supports sensing the maximum quantity of qcloudcbs that can be mounted on a single node (For version 1.12 and above, it's maxAttachCount-2; for version 1.10, the default is currently 18). Merge pr87467 to fix the issue of authorized users sending malicious YAML causing kubectl to consume excessive CPU during YAML parsing.
2020-03-11	v1.16.3-tke.3	cbs intree fixes the issue of continuing to uninstall disks when the disk does not exist, leading to a large number of invalid requests. Metadata adds local caching.
2020-02-14	v1.16.3-tke.2	Merge pr2359 to solve the monitoring missing issue caused by the inability to obtain docker root. Merge pr86583 to improve the log output level when iptables does not support random-fully, avoiding excessive log generation. kube-scheduler supports dynamic setting of log levels.

		<p>Bypass the missing device path (/dev/disk/by-id/virtio-xxx/...) issue with CBS, allowing users to use CBS normally.</p> <p>Merge pr86230 to skip updating the assumed pod's scheduling during the pod scheduling process.</p>
2020-01-06	v1.16.3-tke.1	<p>Merge pr79036 to address the issue where, when the cpumanager is enabled, if the pod's QoS is Guaranteed, then the CPU quota is disabled.</p> <p>Merge pr84167 to fix the issue of the apiserver health check failing due to incorrect ETCD key prefix.</p> <p>Revert pr63066 to fix the issues with CLB health checks and IPVS.</p> <p>Merge pr72914 to fix the issue where immediately creating and scheduling a Pod to the same node after deletion may lead to unsuccessful mounting.</p> <p>Fix the issue of cgroup leakage caused by creating containers under CentOS.</p> <p>Fix the issue of pod exit caused by lxcfs upgrade under Ubuntu16.</p> <p>metadata increases cache and timeout, cloud-provider adds support for using node name as hostname.</p> <p>Revert pr79036 to address the issue where, when the cpumanager is enabled, if the pod's QoS is Guaranteed, the CPU quota is disabled.</p> <p>Bypass the missing device path (/dev/disk/by-id/virtio-xxx/...) issue with CBS, allowing users to use CBS normally.</p>

(Discontinued) TKE kubernetes 1.14.3 revisions

Time	Version	Update Contents
2024-04-10	v1.14.3-tke.31	<p>Volume Reconciliation Optimization.(kube-controller-manager)</p> <p>Super nodes allow the scheduling of pods in TKE trunking ENI mode.(kube-scheduler)</p>
2023-03-03	v1.14.3-tke.28	<p>Added StatusDetaching state for CBS disks. (kube-controller-manager)</p> <p>Fixed CVE-2022-3172. (kube-apiserver)</p> <p>Port 16+ webhook metrics to add <code>calling_webhook_timeout_error</code> metric. (kube-apiserver)</p> <p>Port <code>etcd_request_duration_seconds</code> metric. (kube-apiserver)</p> <p>Optimize self-adaptive node iptables mode Script to fix misjudgments in some cases. (kube-proxy)</p> <p>Merge PR111635 to fix the issue where ipvs real server weight is 0 after kube-proxy restarts. (kube-proxy)</p>
2022-04-13	v1.14.3-tke.27	<p>Merge pr78428, avoid writing warning messages when exporting iptables regulations, to prevent kube-proxy panic during recovery. (kube-proxy)</p>
2022-03-18	v1.14.3-tke.26	<p>Support specified Pods during scale-down. (kube-controller-manager)</p> <p>Optimize the super node scheduling algorithm. (kube-scheduler)</p>

		Merge pr80851 to fix CVE-2019-11247, prevent unauthorized access to CRD resources. (kube-apiserver)
2022-01-20	v1.14.3-tke.24	Support binding LB addresses to IPVS NICs. (kube-proxy) The apiserver supports integration with ExternalName type external services. (kube-apiserver) Optimize TKE Serverless scheduling. (kube-scheduler)
2021-12-02	v1.14.3-tke.23	Scaling from TKE clusters to TKE Serverless: Accurately perceive the remaining IPs of the subnet when batch scheduling Pods, scheduling the correct number of Pods to super nodes. (kube-scheduler) TKE Serverless Node Scheduling Optimization: Resource-aware re-scheduling for TKE Serverless nodes; Priority model scheduling for TKE Serverless nodes; Optimized preferred/pre-selected policies for TKE Serverless nodes. (kube-scheduler) Support extended scheduler Prebind and Unreserve operations. (kube-scheduler) Avoid scheduling Pods using cbs to external CHC Nodes. (kube-scheduler) Fix the bug when scheduling to super nodes. (kube-scheduler)
2021-05-06	v1.14.3-tke.22	Update the startup mode for running kube-proxy as an image, to automatically adapt to the iptables operating mode of the node, supporting operating systems that default to using the nf_tables mode for iptables.
2021-04-14	v1.14.3-tke.21	Merge pr97752 to fix the issue where NewReplicaSet is displayed as <none> when describing a deployment. (kubectl) Merge pr78999 to fix the case sensitivity issue in protocol determination during graceful shutdown. (kube-proxy) Merge pr91590 to fix the issue of port already allocated warnings when using NodePort type services with multiple protocols. (kube-apiserver) Merge pr98262 to support for dynamic log level adjustment in kube-controller-manager. (kube-controller-manager) Merge pr95301 to automatically delete the volume directories left by orphan Pods. (kubelet)
2020-12-28	v1.14.3-tke.19	Add metrics for QcloudCbs. (kube-controller-manager) Fix the issue of extraneous spaces when viewing the serial value during mount of cbs disks. (kubelet)
2020-12-21	v1.14.3-tke.18	Merge pr94712 to fix CVE-2020-8564 - Docker configuration leakage when the file format is incorrect and logLevel >= 4. (kubelet) Merge pr95316 to fix CVE-2020-8565 - Log token leakage due to an incomplete fix of CVE-2019-11250 (logLevel >= 9). (kube-apiserver, kubectl) Merge pr95245 to fix CVE-2020-8566 - Exposure of Ceph RBD adminSecrets in logs when logLevel >= 4. (kube-controller-manager) Merge pr86140 to fix the issue where the controller manager fails to properly handle timeout errors, preventing the creation of scale-out pods. (kube-controller-

		<p>manager)</p> <p>The scheduler supports super nodes. (kube-scheduler)</p> <p>kube-controller-manager supports super nodes. (kube-controller-manager)</p> <p>Set the instance-type Tag based on the actual model of the node, no longer fixed to QCLOUD. (kubelet)</p> <p>Merge pr79338 to not enable the pids cgroup subsystem when neither SupportPodPidsLimit nor SupportNodePidsLimit is enabled. (kubelet)</p> <p>Merge pr89224 to fix the issue where the absence of a check in NodeInfo caused kube-scheduler to crash unexpectedly. (kube-scheduler)</p> <p>Merge pr89465 to fix the issue where HPA calculates the wrong number of instances during rolling updates based on pod metrics. (kube-controller-manager)</p>
2020-10-13	v1.14.3-tke.17	<p>Merge pr74781 to change the default update policy of ConfigMap and Secret from Cache to Watch. (kubelet)</p> <p>Merge pr93403 to remove erroneous print information in kubelet for updating Pod Conditions not owned by kubelet. (kubelet)</p> <p>Merge pr89629 to resolve the issue where the container restart always fails after configmap changes when mounting subpath. (kubelet)</p> <p>Merge pr80942 to fix the issue where rules are not deleted after a service is deleted in ipvs mode. (kube-proxy)</p> <p>QcloudCbs supports BulkVolumeVerification. (kube-controller-manager)</p>
2020-08-04	v1.14.3-tke.16	<p>Merge pr78883 to fix the bug where a default value is automatically added to pod.spec.container.SecurityContext.ProcMount.</p>
2020-07-28	v1.14.3-tke.15	<p>Merge pr76518 and pr82514 to limit the return size of HTTP and exec probes to prevent the consumption of a significant amount of node memory. (kubelet)</p> <p>Merge pr91277 to avoid issues where CLB health checks cause kube-apiserver to generate a large number of TLS handshake error logs. (kube-apiserver)</p> <p>Merge pr91500 to fix the issue of missing KUBERNETES_SERVICE_HOST environment variable. (kubelet)</p> <p>Merge pr77475 to fix the issue where Cronjob cannot schedule when the number of Jobs exceeds 500. (kube-controller-manager)</p>
2020-06-10	v1.14.3-tke.14	<p>Merge pr85027 to fix the issue where HPA calculates the wrong number of instances during rolling updates.</p> <p>Merge pr79708 to use spec.replicas to calculate the current number of HPA replicas.</p> <p>Merge pr91252 to ignore Pod Condition updates from other components to avoid unnecessary scheduling.</p> <p>Merge pr89794 to clean up the error logs of kube-controller-manager to prevent Half-Blind SSRF attacks in CVE-2020-8555.</p>
2020-06-04	v1.14.3-tke.13	<p>Merge pr90260 to fix the issue of missing network monitoring for containerd cluster.</p>

		Merge pr79451 to fix the issue of missing cluster not recreate a Sandbox after a failure when restartPolicy is set to Never.
2020-05-18	v1.14.3-tke.12	TKE can perceive the maximum number of qcloudcbs that can be mounted on a single node, but does not support dynamically obtaining the maximum value.
2020-04-14	v1.14.3-tke.11	Merge pr75442 to correct the bandwidth unit from Kb to b. Merge pr87669 to fix CVE-2020-8552: apiserver DoS attack issue. TKE supports perceiving the maximum quantity of qcloudcbs that can be mounted on a single node (For version 1.12 and above, it's maxAttachCount-2; for version 1.10, the current default is 18).
2020-04-14	v1.14.3-tke.10	CBS intree fixes the issue of a large number of invalid requests caused by continuing to Uninstall a disk when the disk does not exist.
2020-01-13	v1.14.3-tke.9	Merge pr2359 to address the monitoring deficiency issues caused by the inability to obtain the docker root. Merge pr86583 to improve the log output level when iptables does not support random-fully, avoiding excessive log generation. kube-scheduler supports dynamic setting of log levels. Bypass the missing device path (/dev/disk/by-id/virtio-xxx/...) issue with CBS, allowing users to use CBS normally. Merge pr86230 to skip updating the assumed pod's scheduling during the pod scheduling process.
2019-12-23	v1.14.3-tke.8	Revert pr79036 to resolve the issue where, when cpumanager is enabled, the cpu quota is turned off if the Pod's QoS is Guaranteed.
2019-12-17	v1.14.3-tke.7	Metadata enhancements include cache and timeout. Fix the issue where upgrading lxcfs under Ubuntu16 causes Pods to exit. Restart kubelet to avoid pods not ready due to readiness.
2019-11-28	v1.14.3-tke.6	cloud-provider adds support for using the node name as hostname.
2019-11-18	v1.14.3-tke.5	Merge pr83435 to fix the issue where an attacker could send specially crafted malicious YAML or JSON attack payloads causing kube-apiserver CPU or memory exhaustion and inability to provide services. Merge pr84167 to fix the issue where an incorrect Etcd key prefix causes the apiserver health check to fail. Merge pr75622 to fix the issue of delayed synchronization of changes to sts to pods (about 20s) in clusters with a large number of sts (>2000) workloads.
2019-10-23	v1.14.3-tke.4	Merge pr79036 to fix the issue of turning off the cpu quota when CPU Manager is enabled and the Pod's QoS is Guaranteed.
2019-09-	v1.14.3-	Merge pr63066 to fix the issue of health check failure in IPVS mode for CLB.

10	tke.3	
2019-09-06	v1.14.3-tke.2	Fix cve-2019-9512&cve-2019-9514 HTTP/2 DDoS Security Vulnerabilities. Merge pr72914 to fix the issue that deleting a Pod and then immediately creating and scheduling it to the same node may cause a volume mount to fail. Fix the issue of cgroup leakage caused by creating containers under CentOS.

(Discontinued) TKE kubernetes 1.12.4 revisions

Time	Version	Update Contents
2023-03-03	v1.12.4-tke.32	Fix CVE-2022-3172. (kube-apiserver) Port 16+ webhook metrics to add <code>calling_webhook_timeout_error</code> metric. (kube-apiserver) Port <code>etcd_request_duration_seconds</code> metric. (kube-apiserver) Optimize self-adaptive node iptables mode Script to fix misjudgments in some cases. (kube-proxy) Merge PR111635 to fix the issue where ipvs real server weight is 0 after kube-proxy restarts. (kube-proxy)
2022-04-13	v1.12.4-tke.31	Merge pr78428 to avoid writing warning messages when exporting iptables rules, causing kube-proxy to panic during recovery. (kube-proxy)
2022-01-20	v1.12.4-tke.30	Support binding LB addresses to IPVS NICs. (kube-proxy)
2021-05-06	v1.12.4-tke.28	Update the startup mode of kube-proxy when running as a container, automatically adapting to the iptables running mode of the node, to support operating systems that default to running iptables in <code>nf_tables</code> mode.
2020-12-28	v1.12.4-tke.27	Add metrics for QcloudCbs. (kube-controller-manager) Fix the issue of superfluous spaces when checking the serial value while mounting a 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	Merge pr79495 to fix the issue of webhook call failures when CRD has multiple versions. (kube-apiserver)
2020-10-13	v1.12.4-tke.24	Merge pr93403 to remove the erroneous print information from kubelet regarding the update of Pod Conditions that do not belong to kubelet. (kubelet)
2020-08-04	v1.12.4-tke.23	Merge pr78881 to fix the bug that by default adds a default value to <code>pod.spec.container.SecurityContext.ProcMount</code> .

2020-07-28	v1.12.4-tke.22	Merge pr91277 to avoid the issue of CLB health checks leading to a large number of TLS handshake error logs in kube-apiserver. (kube-apiserver) Merge pr91500 to fix the issue of missing KUBERNETES_SERVICE_HOST environment variable. (kubelet)
2020-06-10	v1.12.4-tke.21	Merge pr73915 to avoid watcher receiving events before the start of watch. Merge pr91252 to ignore Pod Condition updates produced by other components to avoid unnecessary scheduling. Merge pr73915 to clean up the error logs of kube-controller-manager to avoid Half-Blind SSRF attacks in CVE-2020-8555.
2020-06-04	v1.12.4-tke.20	Merge pr90260 to fix the missing Cluster Network Monitoring issue in containerd. Merge pr79451 to fix the issue where kubelet fails to recreate SandBox after failure when restartPolicy is set to Never.
2020-05-18	v1.12.4-tke.19	Merge pr77802 to disable graceful termination for UDP traffic. Merge pr68741 to fix the issue with symbolic links to /var/lib/kubelet and the use of subpaths causing mount point leakage and pods remaining in terminating state after deletion. TKE can detect the maximum number of qcloudcbs that can be mounted on a single node, and does not support dynamically acquiring the maximum value.
2020-04-14	v1.12.4-tke.18	Merge pr73401 , pr73606 , and pr76060 to delete DaemonSet Pods assigned to non-existent nodes. Merge pr68619 to fix the issue of the cpumanager dirty data. Merge pr87669 to fix CVE-2020-8552: apiserver DoS attack issue. TKE supports detecting the maximum number of qcloudcbs that a single node can mount (for version 1.12 and above, it is maxAttachCount-2; for version 1.10, the default is currently 18).
2020-02-14	v1.12.4-tke.17	Upgrade the CBS V2 interface to V3. CBS intree resolves the issue of a large number of invalid requests caused by continuing to Uninstall a disk when the disk does not exist.
2020-01-13	v1.12.4-tke.16	Merge pr2359 to fix the issue of missing monitoring due to the inability to obtain the docker root. Merge pr86583 to improve the log output level when iptables does not support random-fully, avoiding excessive log generation. kube-scheduler supports dynamic setting of log levels. Bypass the issue where the device path (/dev/disk/by-id/virtio-xxx/...) is missing for cbs, allowing users to use cbs normally. Merge pr86230 to skip the update of the assumed pod's scheduling during the pod scheduling process.
2019-12-	v1.12.4-	Revert pr79036 to fix the issue where, when cpumanager is enabled, the cpu

23	tke.15	quota is turned off if the Pod's QoS is Guaranteed.
2019-12-17	v1.12.4-tke.14	Metadata enhancements include cache and timeout. Fix the issue where upgrading lxcfs under Ubuntu16 causes Pods to exit. Restart kubelet to avoid pods not ready due to readiness.
2019-11-28	v1.12.4-tke.13	cloud-provider adds support for using the node name as hostname.
2019-11-18	v1.12.4-tke.12	Merge pr75622 to fix the issue of delayed synchronization of changes to sts to pods (about 20s) in clusters with a large number of sts (>2000) workloads.
2019-10-23	v1.12.4-tke.11	Merge pr79036 to fix the issue where and when CPU Manager is enabled, the CPU quota is turned off if the Pod's QoS is Guaranteed. Merge pr72866 to add a --metrics-port command-line parameter to kube-proxy, and solve the bug where --metrics-bind-address could not include port.
2019-09-06	v1.12.4-tke.10	Address the cve-2019-9512&cve-2019-9514 HTTP/2 DDoS security vulnerability. Merge pr72914 to fix the issue where deleting a Pod and immediately creating and scheduling it to the same node might cause a failure to mount the volume. Merge pr71834 to fix the issue where, in IPVS mode, sessionAffinity set to ClientIP could cause access failure to the RS.
2019-08-09	v1.12.4-tke.9	Fix the issue where creating containers in CentOS leads to cgroup leaks.
2019-08-08	v1.12.4-tke.8	Merge pr72118 to fix the issue where StatefulSets based on CBS cannot be mounted when rescheduled to the same node.
2019-07-17	v1.12.4-tke.7	Merge pr75037 to address the security vulnerabilities in the kubectl cp command.
2019-07-16	v1.12.4-tke.6	Resolve compatibility issues between the tlinux kernel version and IPVS, and fix the issue where health checks fail in IPVS mode for CLB.
2019-07-09	v1.12.4-tke.5	Merge pr72361 to fix the possible Deadlock issue in kube-proxy.
2019-06-25	v1.12.4-tke.4	Resolve compatibility issues between the tlinux kernel version and IPVS.
2019-06-17	v1.12.4-tke.3	Merge pr71114 to fix throughput issues with IPVS.
2019-06-04	v1.12.4-tke.2	Merge pr74755 to fix the kubelet hang issue. Merge pr69047 to fix the backward compatibility issue with node.Spec.Unschedulable.

(Discontinued) TKE kubernetes 1.10.5 revisions

Time	Version	Update Contents
2021-05-06	v1.10.5-tke.20	Update the startup mode of kube-proxy when running as a container, automatically adapting to the iptables running mode of the node, to support operating systems that default to running iptables in nf_tables mode.
2020-06-10	v1.10.5-tke.19	Merge pr90260 to fix the missing Cluster Network Monitoring issue in containerd. Merge pr91252 to ignore Pod Condition updates produced by other components to avoid unnecessary scheduling. Merge pr89794 to clean up the error logs of kube-controller-manager to prevent Half-Blind SSRF attacks in CVE-2020-8555.
2020-05-18	v1.12.4-tke.19	Merge pr61549 , add volumeSpec data to the mountedPods cache to resolve issues with multiple pods using the same volume not being able to be deleted normally.
2020-04-29	v1.10.5-tke.17	Merge pr75622 to fix the issue of significant latency (~20s) in syncing sts to Pods in clusters with a large number of sts (>2000) workloads.
2020-04-14	v1.10.5-tke.16	Merge pr68619 to fix the cpumanager dirty data issues. Merge pr87669 to fix CVE-2020-8552: apiserver DoS attack issue. TKE supports detecting the maximum number of qcloudcbs that a single node can mount (for version 1.12 and above, it is maxAttachCount-2; for version 1.10, the default is currently 18).
2020-02-14	v1.10.5-tke.15	Upgrade the CBS V2 interface to V3. CBS intree resolves the issue of a large number of invalid requests caused by continuing to uninstall a disk when the disk does not exist.
2020-01-13	v1.10.5-tke.14	Merge pr2359 to fix the issue of missing monitoring due to the inability to obtain the docker root. Merge pr86583 to improve the log output level when iptables does not support random-fully, avoiding excessive log generation. kube-scheduler supports dynamic setting of log levels. Bypass the issue where the device path (/dev/disk/by-id/virtio-xxx/...) is missing for cbs, allowing users to use cbs normally. Merge pr86230 to skip the update of the assumed pod's scheduling during the pod scheduling process.
2019-12-23	v1.10.5-tke.13	Revert pr79036 to fix the issue where, when cpumanager is enabled, the cpu quota is turned off if the Pod's QoS is Guaranteed.

2019-12-13	v1.10.5-tke.12	kubelet does not delete node when checking externalID. Metadata enhancements include cache and timeout. Fix the issue where upgrading lxcfs under Ubuntu16 causes Pods to exit. Restart kubelet to avoid pods not ready due to readiness.
2019-11-18	v1.10.5-tke.11	Remove reverse detection from kube-controller-manager.
2019-10-23	v1.10.5-tke.10	Merge pr79036 to fix the issue where and when CPU Manager is enabled, the CPU quota is turned off if the Pod's QoS is Guaranteed. Merge pr72866 to add a <code>--metrics-port</code> command-line parameter to kube-proxy, and solve the bug where <code>--metrics-bind-address</code> could not include port.
2019-09-06	v1.10.5-tke.9	Address the cve-2019-9512&cve-2019-9514 HTTP/2 DDoS security vulnerability. Merge pr72914 to fix the issue where deleting a Pod and immediately creating and scheduling it to the same node might cause a failure to mount the volume. Merge 67430 to fix the data structure rollback in the event of a failure in updateContainerCPUSet.
2019-08-08	v1.10.5-tke.8	Merge pr72118 to fix the issue where kubelet reports "resource name may not be empty" when attempting to Mount the same device immediately after an Unmount.
2019-07-17	v1.10.5-tke.7	Merge pr75037 to fix the security vulnerabilities in the kubectl cp command.
2019-06-25	v1.10.5-tke.6	Resolve compatibility issues between the tlinux kernel version and IPVS.
2019-06-17	v1.10.5-tke.5	Merge pr71114 to fix throughput issues with IPVS.
2019-03-19	v1.10.5-tke.4	Merge pr65092 to fix the panic issue in apiserver when handling specific requests.
2019-02-19	v1.10.5-tke.3	Merge pr67288 to fix the connection leakage issue when apiserver acts as a proxy.
2018-09-28	v1.10.5-tke.2	Move the logic of creating CLB out of controller-manager (implemented through a separate service controller).
2018-09-27	v1.10.5-tke.1	backport pr63321 to fix the issue of excessively long Terminating time when there are multiple business containers in a pod.
2018-09-	v1.10.5-	When kubelet update status times out, controller-manager probes the kubelet

21	qcloud-rev1	port.
----	-------------	-------

(Discontinued) TKE kubernetes 1.8.13 revisions

Time	Version	Update Contents
2020-01-13	v1.8.13-tke.7	Merge pr2359 to fix the issue of missing monitoring due to the inability to obtain the docker root. Bypass the issue where the device path (/dev/disk/by-id/virtio-xxx/...) is missing for cbs, allowing users to use cbs normally.
2019-12-13	v1.8.13-tke.6	kubelet does not delete node when checking externalID. Metadata enhancements include cache and timeout. Fix the issue where upgrading Ixcfs under Ubuntu16 causes Pods to exit. Restart kubelet to avoid pods not ready due to readiness.
2019-11-18	v1.8.13-tke.5	Remove reverse detection from kube-controller-manager. Add metric to cbs pvc.
2018-09-28	v1.8.13-tke.2	Move the logic of creating CLB out of controller-manager (implemented through a separate service controller).
2018-09-27	v1.8.13-tke.1	Disable kmem statistics to prevent leakage of cgroup quantities. Reduce resourcequota conflicts when creating pods.
2018-09-21	v1.8.13-qcloud-rev1	When kubelet update status times out, controller-manager probes the kubelet port.

(Discontinued) TKE kubernetes 1.7.8 revisions

Time	Version	Update Contents
2019-12-17	v1.7.8-tke.4	kubelet does not delete node when checking externalID. Metadata enhancements include cache and timeout. Fix the issue where upgrading Ixcfs under Ubuntu16 causes Pods to exit. Restart kubelet to avoid pods not ready due to readiness.
2018-09-28	v1.7.8-tke.2	Resolve the conflict issue between controller-manager and external service controller.
2018-09-27	v1.7.8-tke.1	Move the logic of creating CLB out of controller-manager (implemented through a separate service controller).

2018-09-21	v1.7.8-qcloud-rev1	When kubelet update status times out, controller-manager probes the kubelet port.
------------	--------------------	---

Add-On Version Maintenance Description

Last updated : 2023-08-09 15:35:08

Change Log

TKE provides enhanced add-ons to extend cluster features in various scenarios, including network, storage, monitoring, image, scheduling, and GPU. You can view the current add-on versions and manually upgrade them on the **Add-On Management** page in the TKE cluster details.

Upgrade notice

1. The upgrading action is irreversible.
2. The add-ons can only be upgraded to a later version. By default, they are upgraded to the latest version compatible with the Kubernetes version.
3. **The TKE team no longer offers technical support for disused add-on versions.** We recommend you upgrade them in time.

Version Iteration Records

March 2023

Add-On	Release Date	Version	Update	Limits and Impact
CFS-CSI CFS	2023-03-07	v1.1.0	Supported inline mounting. Updated load configuration to improve add-on stability.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

February 2023

Add-On	Release Date	Version	Update	Limits and Impact
CBS-CSI (CBS)	2023-02-21	v1.1.0	Supported automatically perceiving the maximum number of CVMs that can be mounted.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

			Optimized the cloud disk unmounting logic to ensure the normal state of the cloud disk.	
Craned (Recommended Request add-on)	2023-02-24	v1.3.0	Improved the add-on stability.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
qGPU (GPU isolation add-on)	2023-02-08	v1.0.13	When processing a Pod with zero qGPU resource, TKE directly returns all nodes or scores all "0" to avoid processing a large number of nodes taking too long and blocking the scheduling of the qGPU pod.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

January 2023

Add-On	Release Date	Version	Update	Limits and Impact
QOSAgent (QoS assurance add-on)	2023-01-01	v1.0.1	Fixed the default value setting issues of some parameters.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
DeScheduler (Rescheduler add-on)	2023-01-12	v1.0.2	Fixed the issue where metrics were not available due to inconsistent names and IPs of cluster nodes.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
SecurityGroupPolicy (Security group policy plugin)	2023-01-15	v1.0.5	Fixed Pod creation failure caused by binding multiple duplicate security groups.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

December 2022

Add-On	Release Date	Version	Update	Limits and Impact
Craned (Recommended Request add-on)	2022-12-09	v1.2.0	Optimized the memory recommendation logic and increased the accuracy of recommendation data.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
qGPU (GPU isolation add-on)	2022-12-20	v1.0.12	Supported 16 cards. Optimized logs to avoid printing distracting information.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

November 2022

Add-On	Release Date	Version	Update	Limits and Impact
CBS-CSI (CBS)	2022-11-15	v1.0.8	Supported LVM multi-disk management.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
CBS-CSI (CBS)	2022-11-23	v1.0.9	Supported LVM striping. Increased the csi-resizer quota.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

October 2022

Add-On	Release Date	Version	Update	Limits and Impact
CBS-CSI (CBS)	2022-10-20	v1.0.7	Supported XFS file system snapshots. Supported creating 1-GB cloud disks.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

			Supported volumes of the block type.	
qGPU (GPU isolation add-on)	2022-10-24	v1.0.10	<p>Fixed the issue where, on some multi-card models, the GPU card serial number and the secondary device number were inconsistent, causing the qGPU card to be bound with a wrong graphics card and the business to not work properly.</p> <p>For private deployment, supported the installation of a run package of the base version that does not support hybrid deployment, without modifying kernel parameters or restarting nodes.</p> <p>Fixed the issue where the system reported qgpu-core missing when qgpu-memory was not set.</p>	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
qGPU (GPU isolation add-on)	2022-10-28	v1.0.11	<p>When deleting a Pod, the metrics associated with that Pod are removed directly in bulk, rather than one by one, to avoid the problem of metrics not being removed in some cases.</p>	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

September 2022

Add-On	Release Date	Version	Update	Limits and Impact
qGPU (GPU isolation add-on)	2022-09-06	v1.0.7	Fixed the issue where the qGPU scheduler did not clear deleted node information.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
qGPU (GPU isolation)	2022-09-21	v1.0.8	Added the <code>/metrics</code> API to	This upgrade doesn't affect the existing business. As

<p>add-on)</p>			<p>expose the GPU information (node_gpu_core_request, node_gpu_memory_request_bytes; type: Gauge; labels: "node", "gpu", "namespace", "pod", "container") allocated to Pod containers. When the Pod is deleted, the related container metrics are deleted as well.</p> <p>For qGPU devices left behind for various reasons, they are removed in the gc performed every minute to avoid affecting the creation of new qGPU devices.</p> <p>Fixed the issue where the creation of qGPU devices might fail after a Pod with the same name was created and scheduled to the same node immediately after Pod deletion.</p> <p>Supported native node installation using qGPUs (native node qGPUs do not support offline hybrid deployment).</p>	<p>the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.</p>
<p>qGPU (GPU isolation add-on)</p>	<p>2022-09-29</p>	<p>v1.0.9</p>	<p>Fixed the regression issue introduced by crd, which caused the qGPU scheduler to start scheduling new Pods after restarting without synchronizing Pod allocation information, resulting in inconsistency. Supported clearing related data before updating Pod container metrics, avoiding dirty data.</p>	<p>This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.</p>
<p>CBS-CSI (CBS)</p>	<p>2022-09-13</p>	<p>v1.0.6</p>	<p>Supported scale-out during recovery from a snapshot. Supported XFS file system scale-out. Supported adding super nodes to a TKE cluster.</p>	<p>This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.</p>
<p>Craned</p>	<p>2022-09-20</p>	<p>v1.1.0</p>	<p>Adjusted the interval of requesting Cloud Monitor metrics to five</p>	<p>This upgrade doesn't affect the existing business. As</p>

(Recommended Request add-on)			minutes and reduced the recommended frequency to once every 12 hours.	the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
CraneScheduler (Dedicated scheduler for native nodes)	2022-09-21	v1.1.3	Fixed the failure in updating node metrics to nodes due to the missing of Helm repository parameters. Integrated DeScheduler into CraneScheduler and added the DeScheduler draining watermark. The dynamic scheduler watermark and DeScheduler target watermark are consistent with each other.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
CFS-CSI CFS	2022-09-23	v1.0.8	Supported adding super nodes to a TKE cluster.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
Cerberus (Image signature verification add-on)	2022-09-26	v0.0.1	Image signature verification add-on.	This upgrade doesn't affect the existing business.

August 2022

Add-On	Release Date	Version	Update	Limits and Impact
Nginx Ingress (Open source layer-7 network access layer)	2022-08-17	v1.2.0	Migrated APIs to eliminate the dependency on ENIs.	This upgrade doesn't affect the existing business.
pod-identity-webhook (Pod identity injector)	2022-08-25	v1.0.0	Added chart to allow a Pod to be bound with an sa. An sa can play the role of cam to access other resources.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

SecurityGroupPolicy (Security group policy plugin)	2022-08-25	v1.0.0	Added chart to allow a batch of TKE Serverless Pods to be bound with the same security group to use the security group capability of allowing traffic.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
QOSAgent (QoS assurance add-on)	2022-08-30	v1.0.0	Added chart to initially support kernel-based QoS capabilities.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
qGPU (GPU isolation add-on)	2022-08-31	v1.0.6	Supported using the kubeconfig file of a node to visit the APIServer (in this case, the <code>crd</code> parameter cannot be enabled). Supported installing qGPUs with TencentOS 2.4 (TK4).	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

July 2022

Add-On	Release Date	Version	Update	Limits and Impact
CBS-CSI (CBS)	2022-07-06	v1.0.5	Supported cloud disk creation when cloud labels are duplicated. Fixed the online expansion time sequence issue. Revoked csi-node secret resource permissions.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
CFS-CSI CFS	2022-07-14	v1.0.7	Adapted the dynamic storage feature to Serverless clusters. Adapted the shared storage feature to Serverless clusters.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
qGPU	2022-	v1.0.5	Supported installing and	This upgrade doesn't affect the existing

(GPU isolation add-on)	07-20		using qGPUs with public images. Supported disabling the feature of labeling nodes (<code>-enable-label=false</code>).	business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
------------------------	-------	--	--	---

June 2022

Add-On	Release Date	Version	Update	Limits and Impact
DeScheduler (Rescheduler add-on)	2022-06-07	v1.0.1	Supported TMP authentication: Added <code>auth</code> authentication to <code>prom-probe</code> . Passed in environment variables such as <code>token</code> and <code>appid</code> to DeScheduler and int containers and decoded them. Added the Prometheus client authentication feature to DeScheduler.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
qGPU (GPU isolation add-on)	2022-06-08	v1.0.3	Updated the image of qGPU manager to <code>tkeimages/elastic-gpu-agent:v1.0.2</code> . Updated the image of qGPU scheduler to <code>tkeimages/elastic-gpu-scheduler:v1.0.2</code> . Supported using GPU CRD to manage GPU resources.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
DynamicScheduler (Dynamic scheduler)	2022-06-30	v1.0.2	Supported TMP authentication: Added <code>auth</code> authentication to <code>probe-prometheus</code> . Passed in environment variables such as <code>token</code> and <code>appid</code> to node-annotator and int containers and decoded them.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

			<p>Added the Prometheus client authentication feature to node-annotator, and updated the image to v3.2.1.</p> <p>Fixed the bug where data couldn't be queried for PromQL statements using the IP as the node exporter reporting label.</p>	
--	--	--	--	--

May 2022

Add-On	Release Date	Version	Update	Limits and Impact
CBS-CSI (CBS)	2022-05-06	v1.0.3	Supported taint and toleration configuration. Added the <code>type</code> launch parameter.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
COS-CSI (COS)	2022-05-06	v1.0.1	Supported taint and toleration configuration.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
CFS-CSI CFS	2022-05-06	v1.0.4	Supported taint and toleration configuration. Added priority and <code>hostAliases</code> to the <code>tcfs</code> configuration file.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
CFS-CSI CFS	2022-05-24	v1.0.5	Supported Serverless cluster <code>cfs</code> provisioner.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
CBS-CSI (CBS)	2022-05-31	v1.0.4	Optimized the add-on startup logic. Adjusted the default number of concurrent <code>csi-attacher</code> requests to 50.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

April 2022

Add-On	Release Date	Version	Update	Limits and Impact
CFS-CSI CFS	2022-04-12	v1.0.2	Supported the idempotency for <code>amount</code> .	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
Nginx Ingress (Open source layer-7 network access layer)	2022-04-18	v1.1.0	<p><code>eks.tke.cloud.tencent.com/norm</code> : "true"</p> <p><code>eks.tke.cloud.tencent.com/disable-metrics</code> : "true"</p> <p>Considering the case of using super nodes in TKE clusters, annotations that were originally effective under TKE Serverless were opened to be effective under TKE and TKE Serverless scenarios.</p>	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
CFS-CSI CFS	2022-04-19	v1.0.3	<p>Added the resource label field to <code>tcfs</code> crd.</p> <p>Installed no <code>tcfs</code> resources on Kubernetes 1.12 or earlier.</p> <p>Optimized the registration and startup of <code>cfs-csi</code> startServer.</p>	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
qGPU (GPU isolation add-on)	2022-04-21	v1.0.2	<p>Updated the image version of <code>qGPU</code> manager. Supported automatic settings of GPU driver version and other information on the current node.</p> <p>Updated the <code>qgpu-manager</code> ClusterRole to add operation permissions for nodes.</p>	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
CBS-CSI (CBS)	2022-04-24	v1.0.2	<p>Canceled the directory clearance logic in the <code>NodeUnpublishVolume</code> API.</p> <p>Supported getting the driver letter through the serial number.</p>	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend

			Retained the corresponding CRD resources while deleting the add-on.	you upgrade it during off-peak hours.
--	--	--	---	---------------------------------------

March 2022

Add-On	Release Date	Version	Update	Limits and Impact
CBS-CSI (CBS)	2022-03-16	v1.0.1	Supported in-place lossless migration of workloads using <code>intree cbs</code> to CSI while upgrading the cluster from v1.18 to v1.20.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.
CFS-CSI CFS	2022-03-24	v1.0.1	Supported using automatically generated subdirectories for data isolation in the shared storage instance during dynamic creation.	This upgrade doesn't affect the existing business. As the add-on may be unavailable during the upgrade, we recommend you upgrade it during off-peak hours.

CLS Add-on Version Description

Last updated : 2024-12-27 14:17:26

CLS Add-ons Overview

The CLS add-ons are deployed by CLS on each standard node in your cluster when you enable CLS in TKE. They are used to collect application logs generated by TKE and write them to consumers in Tencent Cloud, including CLS and Kafka.

The CLS add-ons are as follows:

Name	Resource Type	Description
tke-log-agent	DeamonSet	Each log-agent Pod contains a controller container and a loglistener sidecar container, which are responsible for collecting logs generated by all containers on the node.
cls-provisioner	Deployment	There is one instance per cluster, responsible for converting the CRD configuration into a collection configuration comprehensible by the loglistener to communicate with CLS.
logconfigs.cls.cloud.tencent.com	CRD	-

log-agent Release Notes

v1.1.17

Category	Content
Feature	cls-provisioner supports vmagent container. log-agent generates the loglistener configuration file in encryption mode by default.
Bugfix	-

v1.1.16

Category	Content

Feature	<p>cls-provisioner allows to enable Dynamic Index when you are creating a topic.</p> <p>cls-provisioner allows to add a finalizer to logconfig.</p> <p>The max_metadata_wait_time_s parameter added to loglistener can be configured by using environment variables.</p> <p>When the key-value index is text, the word segmentation symbol can be empty.</p>
Bugfix	<p>Fixed the log-agent problem of incorrect container judgment during metadata generation. For example, different containers of the same Pod exist in the /data path, such as A container volume mounted to /data and B having a normal directory /data.</p> <p>Fixed the cls-provisioner problem of potential residual dirty data during switching between host_file and container_file.</p> <p>Fixed the problem of collection configuration being deleted or overwritten in self-built mode.</p> <p>Fixed the problem where the kafka collector fails to retrieve the latest target metadata, and added a regular update mechanism for metadata.</p> <p>Fixed the problem where metadata is not generated as expected when different collection rules are used to collect different files under the same directory.</p>

v1.1.15

Category	Content
Feature	<p>Kafka collector supports custom metadata.</p> <p>Kafka collector allows to output metadata in string format.</p> <p>Added the cluster_id field to Kafka collector metadata.</p> <p>Supported advanced filters.</p>
Bugfix	<p>Fixed the resource compatibility problem for existing add-ons, and added non-CLS regions through the add-on configuration CLOUD_API_REGION. log-agent adds docker DataRoot to avoid triggering dynamic mounting for docker runtime with standard configuration.</p> <p>Adapted to standard output path mapping for clusters earlier than version 1.14. log-agent no longer automatically mounts docker dataroot.</p> <p>Fixed the problem where file events fail to be processed due to full retryChan.</p> <p>Fixed the problem where standard output logs fail to be collected due to empty containerOpt.</p> <p>Fixed the cls-provisioner OOM problem caused by frequent ns updates.</p>

v1.1.14

Category	Content

Feature	CLS add-ons expose metrics related to collection rule status. The default port is 8088.
Bugfix	-

v1.1.13

Category	Content
Feature	Elastic Pod cls-provisioner switches to norm authentication.
Bugfix	Fixed the problem of accidental path deletion introduced by multiple paths.

v1.1.12

Category	Content
Feature	TKE multi-path collection supports path wildcards.
Bugfix	Adapted to clusters earlier than version 1.19 to address log loss caused by kubelet mapping failure of the standard data log path.

v1.1.11

Category	Content
Feature	<p>Added source resources during collection configuration creation to distinguish the TKE collection configuration, CLS console, and direct access to TKE.</p> <p>Supported namespace and excluded namespace using regular expression.</p> <p>Added advanced collection configuration to collect CIsAgentParseFailMerge, and changed parameter types to fix the parameter loss problem during the JSON marshal/unmarshal process.</p> <p>Supported custom cls-provisioner and log-agent installation namespaces, which is kube-system by default.</p> <p>Supported multi-path collection for container files and host files.</p>
Bugfix	<p>Fixed the problem where the log-agent does not delete symbolic links when the job Pod ends, causing the number of files and directories monitored by loglistener to reach the upper limit.</p> <p>Fixed the problem where kafkaListener fails to collect logs due to excessively long logs in a single line, and updated the maximum length of a single line to 2 MB.</p>

	Fixed the problem where mapping the log directory fails when the PVC StorageClass is empty.
--	---

v1.1.10

Category	Content
Feature	-
Bugfix	Fixed the problem that occurs when log-agent handles the inverse selection logic of container notin.

v1.1.9

Category	Content
Feature	Added AdvancedConfig, including the MaxDepth and FileTimeout parameters. Added the HotPeriod parameter to support the settlement rules of transitioning standard topic data to a low frequency. Supported container selection and exclusion, container metadata collection, and tag-based collection.
Bugfix	Fixed the problem of abnormal prefix removal for paths in the blacklist. Fixed the log-agent problem where symbolic links are deleted and recreated while the Pod remains unchanged. Fixed the panic problem in cls-provisioner caused by failure to query topics. Added validation for logconfig. Invalid logconfigs will not be processed to prevent program panic or processing errors.

v1.1.8

Category	Content
Feature	-
Bugfix	Fixed the problem of no retry after workload is not supported.

v1.1.7

Category	Content
Feature	<p>Added the logConfig-max-threshold parameter. When the logconfig reaches the threshold, the informer is used to query the workload information after you upgrade or restart the add-on.</p> <p>Optimized the method to obtain docker root. You can obtain it via the API preferentially.</p> <p>Custom metadata is available to kafka collector, and SASL authentication is supported.</p>
Bugfix	<p>Fixed the problem where a kubelet parameter error occurred when starting the log-agent.</p> <p>Fixed the problem where a logset API calling error occurred when replacing the topicID.</p> <p>CLS is not case-sensitive. Fixed the case-sensitive issue of cls-provisioner.</p> <p>The log-agent supports CSI by default. The Provisioner of StorageClasses is no longer judged and compared separately.</p>

v1.1.6

Category	Content
Feature	<p>Added multi-line Regex matching with container standard output for the kafka collector.</p> <p>You can configure tag API address through environment variables.</p>
Bugfix	<p>Fixed the problem that the agent automatically created a topic but failed to sync it to the cluster tag in regions where CLS was not available, such as Shenzhen and Tianjin.</p> <p>Fixed the problem where the kafka collector's metadata prefix captured the format issue of containerd.</p> <p>Fixed the problem where the incremental and full options were missing when cls-provisioner was synced to CLS.</p> <p>You can query the workload information through informer during logconfig loading to reduce the load on apiserver.</p>

v1.1.5

Category	Content
Feature	Added São Paulo and Shanghai Auto-Driving Cloud Zone that can be shipped to by CLS.
Bugfix	-

v1.1.4

Category	Content
Feature	Supported collecting standard output logs of the init container. Supported the GBK encoding format for CLS log parsing. Supported namespace label selectors for CLS log collection rules.
Bugfix	Fixed the issue of frequent retries of the event queue. Optimized the queue processing algorithm for the LogConfig and Pod in the log-agent to prevent the queue from being blocked by a large number of repeated events. Fixed the issue where the standard output logs of the container were not collected when there were only namespaces but no labels.

v1.1.3

Category	Content
Feature	Supported index settings for CLS STANDARD_IA.
Bugfix	-

v1.1.2

Category	Content
Feature	Supported configuring a blocklist in CLS collection rules for container files and host files.
Bugfix	-

v1.1.1

Category	Content
Feature	Supported the multi-core CPU for the LogListener. Adapted the LogListener memory to the upper limit of 100 MB memory.
Bugfix	Fixed the issue where the log-agent failed to update the volume link and did not retry after the returned error was not processed during the update.

v1.1.0

Category	Content
Feature	Supported inheriting TKE clusters' tags when a topic was automatically created by CLS. Supported cross-region log shipping for CLS.
Bugfix	-

v1.0.9

Category	Content
Feature	-
Bugfix	Fixed the issue where an empty topic ID caused the logconfig to be deleted and recreated during topic replacement. Fixed the issue where the logconfigpro sync crashed as the logconfigpro informer was not started.

v1.0.8

Category	Content
Feature	Blocked the collection of loglistener logs under kube-system by default. Modified the index creation policy, so that the default index was created only when a topic was automatically created, and the topic index was not modified in other scenarios. Supported the Kafka collector to add metadata information to messages. Supported the following parsing modes in the Kafka collector: full text in a single line, JSON, and full text in multiple lines.
Bugfix	Fixed the issue where the container couldn't be specified when collecting the standard output in the workload scenario. Added the docker client to get Storage Driver, so that if there was no configuration file, the info information could be obtained through the client to get Storage Driver. Fixed the error of the specified metadatalabel when collecting container files. Fixed the scheme to get the kubelet root directory. Fixed the collection configuration match error caused by the incorrectly set prefix of the old collection configuration to be deleted.

	Fixed the issue where the timestampKey timestamp set for messages in the current Kafka collector was invalid.
--	---

v1.0.7

Category	Content
Feature	Supported creating key-value indexes during topic creation by cls-provisioner, including index name, type, word segmentation, and statistics collection status switch; if not supported, the pod_name,namespace,container_name index would be enabled by default. Supported specifying metadata labels to write the specified Pod labels into metadata collection; if not supported, all Pod labels would be collected as metadata. Supported customizing the CLS TencentCloud API service backend address.
Bugfix	-

V1.0.6

Category	Content
Feature	Supported modifying the kubelet root directory and docker root directory in log-agent.
Bugfix	-

v1.0.5

Category	Content
Feature	Supported the <code>label !=</code> operation (exclude labels) in the log collection configuration. Supported collecting only incremental logs by CLS. Supported selecting multiple namespaces and excluding namespaces in the log collection configuration. Supported configuring Pod labels with the same key but different values in log-agent. Supported configuring loglistener parameters.
Bugfix	Fixed a known issue with log-agent using configmap as source. Fixed an issue where the collector configuration was empty and caused verification failures under some conditions.

	Fixed the issue where the collector failed to delete the configuration when deleting the log rule. Fixed compatibility issues with logConfig configuration.
--	--

v1.0.1

Category	Content
Feature	Switched the API for cls-provisioner to access CLS to TencentCloud API. Supported TKE log collection and delivery to CKafka (for more information, see Configure Log Collection via the Console).
Bugfix	-

v0.2.28

Category	Content
Feature	-
Bugfix	Fixed the issue where one Pod corresponded to multiple logconfig files.

v0.2.27

Category	Content
Feature	-
Bugfix	Fixed the issue where the extraction mode configured on a topic was overwritten in some scenarios.

v0.2.26

Category	Content
Feature	-

Bugfix	Fixed the issue where the metadata couldn't be created in some cases when the collection configuration of the stdout type was deleted.
--------	--

v0.2.25

Category	Content
Feature	-
Bugfix	Fixed the log-agent panic issue in some cases. Fixed the issue of soft connection deletion caused by workload cache. Fixed the metadata file creation failure.

v0.2.24

Category	Content
Feature	-
Bugfix	Fixed the issue where metadata was accidentally deleted during container restart in a Pod. Supported automatically cleaning up LogAgentRootDir before log-agent start so as to avoid dirty data. Fixed the panic of log-agent caused by extreme scenarios. Fixed the startup failure caused by log-agent's repeated mounting of the <code>/data</code> directory.

Change Records of VPC-CNI Component

Last updated : 2024-11-14 15:43:10

The VPC-CNI component contains three Kubernetes cluster components: `tke-eni-agent` , `tke-eni-ipamd` , and `tke-eni-ip-scheduler` . Generally, their versions are the same. However, `tke-eni-ip-scheduler` is less modified, so its version may be a little earlier.

Checking the Component Version

The component version is the image tag. You can check it via the kubernetes API.

```
# Checking the version of tke-eni-agent
kubectl -nkube-system get ds tke-eni-agent -o jsonpath=
{.spec.template.spec.containers[0].image}
# Checking the version of tke-eni-ipamd
kubectl -nkube-system get deploy tke-eni-ipamd -o jsonpath=
{.spec.template.spec.containers[0].image}
# Checking the version of tke-eni-ip-scheduler
kubectl -nkube-system get deploy tke-eni-ip-scheduler -o jsonpath=
{.spec.template.spec.containers[0].image}
```

Change Records

Version Number	Release Date	Updates	Impacts
v3.4.7	2022-09-07	<p>Supports the preferential scheduling policy of ip-scheduler, where Pods with static IP addresses are preferentially scheduled to the ENIs matching the subnet.</p> <p>eni-ipamd supports the dry run to sync existing custom resources (CRs) and promptly discover change exceptions.</p> <p>Optimizes the polling logic for ENI-IP address binding to reduce the errors caused by ENIs/IP addresses that are being bound.</p> <p>Fixed the occasional issue where internally allocated IP addresses are leaked when shared ENIs are released in non-static IP address mode.</p>	No impact on services

v3.4.6	2022-07-26	Supports the native node pool.	No impact on services
v3.4.5	2022-06-28	<p>The non-static IP address mode of shared ENIs supports IPv4/IPv6 dual-stack. In dual-stack mode, each Pod will be allocated an IPv6 IP address and an IPv4 IP address.</p> <p>Fixed the issue where the EIP becomes invalid due to <code>nodeLost</code> on super nodes. After the fix, the EIP will be bound again.</p>	No impact on services
v3.4.4	2022-06-06	<p>By default, the EIP is tagged with <code>tke-clusterId</code> and <code>tke-created-eip</code> and inherits the TKE cluster's tag.</p> <p>Supports unbinding ENIs in instances that have been shut down.</p> <p>Optimizes ip-scheduler and fixed the issue of slow start due to too many subnets.</p>	No impact on services
v3.4.3	2022-04-13	<p>eni-ipamd and ip-scheduler support disabling subnets. Disabled subnets can be allocated only to specified objects by setting the <code>--only-nominated-eni-subnets</code> startup parameter.</p> <p>The static IP address mode supports specifying subnets for Pods through the <code>tke.cloud.tencent.com/nominated-eni-subnets</code> annotation. Multiple subnets need to be separated by comma.</p> <p>eni-agent supports protecting key kernel parameters of the system and adopting new TLinux features to prevent kernel parameters (<code>rp_filter</code> and <code>ip_forward</code>) from being modified.</p> <p>Fixed the occasional issue where the eni-ip resource of the node fails to be registered due to kubelet restart during node initialization in shared ENI mode.</p> <p>Fixed the issue where the IP garbage collection mechanism fails due to dockershim or containerd restart during container running.</p>	No impact on services
v3.4.2	2022-03-04	<p>The non-static IP address mode supports specifying the ENI and subnet of the node.</p> <p>eni-agent supports automatically setting <code>ip_forward</code> and <code>rp_filter</code> kernel parameters on schedule to avoid network failures due to their changes.</p>	No impact on services

		<p>Optimizes the scheduling performance. In shared ENI mode, if an ENI is being bound, the polling wait occurs to reduce scheduling failures.</p> <p>Fixed the occasional issue where eni-ip extension resources are lost due to high node loads.</p> <p>Attempts to delete and recreate the ENI and IP address that are pending for a long time; fixed the issue where the ENI and IP address become unavailable for a long time due to underlying failures.</p>	
v3.4.1	2022-01-21	<p>Supports scheduling Pods to TKE Serverless nodes and maintaining the IP address in static IP address mode.</p> <p>Supports specifying the EIP through the <code>`tke.cloud.tencent.com/eip-id-list`</code> annotation.</p> <p>Supports binding dedicated ENIs to security groups in non-static IP address mode.</p> <p>Upgrades the CRD API to v1 and supports Kubernetes 1.22.</p> <p>Fixed the occasional issue where the IP status is not synced in static IP address mode.</p>	No impact on services
v3.4.0	2021-12-08	<p>Supports static IP addresses with multiple ENIs.</p> <p>Supports underlay connection in and off the hybrid cloud and elastic Pod deployment.</p> <p>Fixed the issue of incorrect CNI data plane settings due to occasional CNI concurrency in the same Pod.</p>	No impact on services
v3.3.9	2021-11-09	<p>Fixed repeat creation of an EIP caused by the network.</p> <p>Pods with independent ENIs in non-static IP address mode can be bound to an EIP.</p> <p>Optimizes the mechanism of expansion resources for eni-agent to make the management of expansion resources more stable and robust.</p> <p>Fixed the issues caused by inconsistency between quota set for the node and the actual quota.</p> <p>Optimizes IP garbage collection mechanism for eni-agent. If there is a dirty container in the Pod that is being created, the reclaimed IPs will be allocated to a new container in the Pod.</p> <p>Optimizes the calculating method for resources of the used IPs and ENIs in non-static IP address mode.</p> <p>Fixed the issue of inaccurate calculation of resources caused by the Pod status of <code>Error</code> , <code>Evicted</code> and <code>Completed</code> etc.</p>	No impact on services

v3.3.8	2021-08-17	<p><code>--master</code> can configure the backend kube-apiserver address without relying on kube-proxy.</p> <p>eni-agent supports <code>--kube-client-qps</code> and <code>--kube-client-burst</code> to configure QPS and Burst of kube client, and the default values increase to 10 and 20 respectively.</p> <p>If eni-agent finds that the updated expansion resources are less than original ones, it will update the latest expansion resources information in the node status to prevent issues caused by async updating of kubelet.</p>	No impact on services
v3.3.7	2021-08-13	<p>eni-ipamd supports <code>--enable-node-condition</code> and <code>--enable-node-taint</code>. If <code>eni-ip</code> or <code>direct-eni</code> is missed on the node after enabling, the condition or taints of the node will be set.</p> <p>EIP supports parsing new API parameters in json format.</p> <p>Fixed the issue where the allocated IPs may be reclaimed improperly by garbage collection of eni-agent in containerd runtime.</p> <p>Fixed ipamd panic that may be caused by the EIP API.</p> <p>Fixed the issue where an ENI is unbound because <code>disable-node-eni</code> annotation is set improperly when the non-static IP mode is upgrading.</p>	No impact on services
v3.3.6	2021-07-26	<p>Fixed the issue where the allocated IPs and routes may be reclaimed improperly because of the garbage collection mechanism of eni-agent.</p> <p>Fixed the issue where IPs may be released before the Pod when deleting deployment and other upper-layer resources after <code>--enable-ownerref</code> is enabled for eni-ipamd.</p>	No impact on services
v3.3.5	2021-07-20	<p>Fixed the issue where locally stored data of the Pod cannot be deleted because of improper deletion of the IPs or ENIs of the Pod with a shared ENI/exclusive ENI in non-static IP address mode.</p> <p>Fixed the issue where CNI information of a shared ENI/exclusive ENI does not store and verify ENI information of the Pod in non-static IP address mode.</p>	No impact on services

v3.3.4	2021-07-07	<p>Fixed the issue where the component continues trying to unbind the ENI in the condition that the CVM has shut down.</p> <p>Fixed the panic caused by the concurrent writes of asynchronous logs.</p> <p>Optimizes the ENI synchronization logic in non-static IP address mode to ensure internal data consistency and prevent the ENIs in use from being unbound.</p> <p>Fixed the issue where the existing nodes cannot allocate IPs caused by insufficient IPs in the subnet of the cluster upgrading from v3.2 in non-static IP address mode.</p> <p>Fixed the issue where the ENI may be incorrectly released when the primary IP of the existing ENI is being used by the Pod.</p>	No impact on services
v3.3.3	2021-06-07	Supports hybrid cloud ipam, and it can work in collaboration with the cilium overlay/underlay mode.	No impact on services
v3.3.2	2021-06-01	<p>ip-scheduler supports occupancy caused by insufficiency of default resources, and does not support occupancy caused by insufficiency of IP resources.</p> <p>The security group feature logic of the shared ENI is reconstructed. It supports strong synchronization with the security group set on the node to ensure that the binding sequence and priority of security groups is consistent with that in user's settings.</p> <p>Supports the cilium cni-chain mode.</p> <p>For eni-agent, <code>hostPort</code> field can be configured for the Pod after <code>--port-mapping</code> is enabled.</p> <p>The annotation <code>tke.cloud.tencent.com/claim-expired-duration</code> can be added to the Pods to reclaim static IPs in specific time. The annotation only affects the added Pods.</p>	No impact on services
v3.3.1	2021-05-11	<p>Multiple ENIs can be used in shared ENI non-static IP address mode.</p> <p>Tencent Cloud API can call API QPS limits, and the limit for a single cluster is 50 QPS by default (limit by the type of CVM, VPC and TKE).</p> <p>Changes of IP quota can be perceived after upgrading of non-static IP address mode.</p> <p>The annotation <code>tke.cloud.tencent.com/desired-route-</code></p>	No impact on services

		<p><code>eni-pod-num</code> can be added for <code>node</code> . The desired number of route-eni ip can be written and the node quota will be adjusted automatically by the component after the writing.</p> <p>Fixed the issue of VPC task polling timeout caused by the fact that the VPC task does not exist.</p> <p>Fixed the issue of eni-ipamd panic caused by failure of task creation for the ENI.</p> <p>Optimizes routing reconciliation logic and only clears the IP routes managed by eni-agent.</p> <p>Fixed the issue of exceptional panic occurred at the time of ENI releasing in the independent ENI non-static IP address mode caused by the fact that the ENI has already been released.</p>	
v3.3.0	2021-04-13	Supports customized GR mode. Multiple CIDR blocks can be set in a node and a cluster.	No impact on services
v3.2.6	2021-03-31	<p>Reduces the time of retrying for binding an ENI in exclusive ENI mode and improves binding efficiency.</p> <p>Reduces failures of concurrent binding and unbinding of ENIs, and improves the efficiency of binding and unbinding through concurrency control.</p> <p>Optimizes subnet allocation logic for an ENI in non-static IP address mode. Fixed the issue where some nodes cannot obtain IPs in the condition that IPs are sufficient when the nodes are added concurrently.</p> <p>The garbage collection mechanism of eni-agent supports self-awareness of the underlying runtime and supports containerd.</p>	No impact on services
v3.2.5	2021-02-22	<p><code>dnsConfig</code> is added when <code>eni-ipamd</code> and <code>ip-scheduler</code> are deployed to avoid the issues caused by the DNS that are created by users.</p> <p>In the shared ENI static IP address mode, the information of <code>subnetID</code> of the ENI that is bound to each node will be synced to the label of the node, and the key is <code>tke.cloud.tencent.com/route-eni-subnet-ids</code> .</p> <p><code>eni-agent</code> will try to obtain the reasons for failures of IP allocation and return them to the CNI plugin to make them reflect in the Pod event.</p> <p>A bare Pod can specify an IP through the annotation <code>tke.cloud.tencent.com/nominated-vpc-ip</code> .</p>	No impact on services

		<p>eni-agent supports periodic test for the connection with APIServer. It will restart automatically if a timeout occurs.</p> <p>Fixed the waste of IPs caused by internal data inconsistency.</p>	
--	--	--	--

Runtime Version Maintenance Description

TKE Runtime Revision Version History

Last updated : 2024-12-26 17:29:16

Containerd

TKE containerd 1.6.9 patch releases

Time	Version	Update
2023-01-09	containerd-1.6.9-tke.2	Fixed a community bug #7848 that restarting containerd and then kubelet causes Pods to restart.
2022-12-01	containerd-1.6.9-tke.1	Same as the community edition, added installation and upgrade scripts.

TKE containerd 1.4.3 patch releases

Time	Version	Update
2022-10-27	containerd-1.4.3-tke.2	Upgraded ttrpc in containerd to avoid ttrpc bugs from causing Pod Terminating.
2021-06-22	containerd-1.4.3-tke.1	Supported customizing sandbox ENI names by modifying the containerd configuration as [plugins.cri.cni] interface_name=eth1.

(Discontinued) TKE containerd 1.3.4 patch releases

Time	Version	Update
2022-01-19	containerd-1.3.4-tke.6	Fixed a bug that container status errors cause Pod Terminating after containerd is restarted.
2021-11-18	containerd-1.3.4-tke.5	Supported enabling writable cgroup and starting the systemd container without privileges enabled in containerd.
2021-06-25	containerd-1.3.4-tke.4	Supported using the xfs quota to restrict the capacity of the rwyler container in containerd.
2020-12-08	containerd-1.3.4-tke.3	Supported using the signal configured in the image when a container is killed in containerd.

2020-09-10	containerd-1.3.4-tke.2	Supported customizing sandbox ENI names by modifying the containerd configuration as [plugins.cri.cni] interface_name=eth1.
2020-04-27	containerd-1.3.4-tke.1	Fixed a bug that not handling the runc return value (delete container not found) leads to container deletion failure and Pod Terminating.

Docker

TKE docker 19.03.9 patch releases

Time	Version	Update
2022-11-01	docker v19.3.9-tke.1	Fixed a bug that Pods remain in Terminating after a containerd event is lost in docker, to ensure that containers can be properly killed when there is data inconsistency between docker and containerd.
2021-01-04	docker 19.03.9	Same as the community edition.

(Discontinued) TKE Docker 18.06.3 patch releases

Time	Version	Update
2022-03-10	docker v18.6.3-ce-tke.4	Upgraded runc to 1.0.0-rc95 to avoid CVE-2019-16884 . Supported correctly identifying the runc version in docker, reducing the output of error logs. You can confirm the commit ID of runc through docker info. In earlier versions, runc 1.0.2 was used and the commit was shown as N/A.
2021-06-22	docker v18.6.3-ce-tke.3	Ensured that containers can still be deleted after events are lost. You can verify this with the command <pre>id=\$(docker run -d busybox top); kill -9 \$(pidof docker-containerd) && pkill top; docker stop \$id</pre> . An undeletable container will be obtained in earlier versions of docker, while containers can be stopped in the later version of docker.
2021-01-14	docker 18.06.3-tke.2	Fixed a gRPC community bug #2818 that deleting a container could cause docker to get stuck for a long time and stop handling subsequent events. Added a timeout in docker18 to avoid not handling events.
2020-08-13	docker 18.06.3-tke.1	Fixed a bug that restarting containerd could cause gRPC to receive a cancel request. Fixed event blockage caused by unclosed exec connections.

