

# Tencent Cloud Observability Platform Prometheus Monitoring Product Documentation



©2013-2022 Tencent Cloud. All rights reserved.



#### **Copyright Notice**

©2013-2024 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

#### 🔗 Tencent Cloud

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

**Prometheus Monitoring** 

Product Introduction

Overview

Use Cases

Strengths

Concepts

Use Limits

Features

Service Regions

Access Guide

Scrape Configuration Description

**Custom Monitoring** 

**EMR** Integration

Flink Integration

Java Application Integration

Spring Boot Integration

JVM Integration

Go Application Integration

Exporter Integration

Elasticsearch Exporter Integration

Kafka Exporter Integration

MongoDB Exporter Integration

PostgreSQL Exporter Integration

NGINX Exporter Integration

Redis Exporter Integration

MySQL Exporter Integration

Consul Exporter Integration

Memcached Exporter Integration

Integration with Other Exporters

CVM Node Exporter

Health Check

**Cloud Monitoring** 

Read Cloud-Hosted Prometheus Instance Data via Remote Read

Agent Self-Service Access

Instructions for Installing Components in the TKE Cluster

Security Group Open Description **Operation Guide** Instance **Creating Instance** Searching for Instance **Renaming Instance Terminating Instance Rebooting Instance** Modifying Storage Period Viewing Instance's Basic Information TKE Integration with TKE Integration Center Data Multi-Write **Recording Rule** Overview **Rule Management** List of Default Recording Rules **Instance Diagnosis Alerting Rule** Overview Alerting Rule Description **Creating Alerting Rule Disabling Alerting Rule** Rule Type Description(old) Notification Template Tag Examples Using Tag Editing Tag **Use Limits** Access Control Overview Setting Policy Granting Policy Description of Role Permissions Related to Service Authorization Grafana **API** Guide



Overview
Writing Data
Querying Monitoring Data
TKE Metrics
Free Metrics in Pay-as-You-Go Mode
Recommended Common Metrics for TKE
Recommended Common Metrics for TKE
Data Collection Configuration
Configuring Necessary Monitoring Metrics
Adjusting Collection Interval
Resource Usage and Billing Overview
Practical Tutorial
Migration from Self-Built Prometheus
Custom Integration with CVM
TKE Monitoring
Enabling Public Network Access for TKE Serverless Cluster
Connecting TMP to Local Grafana
Terraform
Terraform Overview
Managing Prometheus Instances Using Terraform
Managing the Integration Center of Prometheus Instances Using Terraform
Collecting Container Monitoring Data Using Terraform
Configuring Alarm Policies Using Terraform
FAQs
Basic Questions
Integration with TKE Cluster
Product Consulting

Use and Technology

# Prometheus Monitoring Product Introduction Overview

Last updated : 2024-08-07 21:50:20

TencentCloud Managed Service for Prometheus (TMP) provides the highly available Prometheus service as well as the open-source visualization tool Grafana and Cloud Monitor alarms while inheriting the monitoring capabilities of the open-source Prometheus, which reduce your development and OPS costs.

### **Prometheus Overview**

Prometheus is an open-source monitoring system. Similar to Kubernetes inspired by Google's Borgman monitoring system, it was inspired by Google's Borgman monitoring system. It was created and developed in 2012 by SoundCloud's internal engineers and released in January 2015. In May 2016, it became the second project after Kubernetes to officially join Cloud Native Computing Foundation (CNCF). Nowadays, it is usually used for monitoring in the most common Kubernetes container management systems.

Prometheus has the following features:

Custom multidimensional data models (a time series data entry is composed of a metric and a key-value pair called label).

Flexible and powerful query language PromQL, which can use multidimensional data to complete complex monitoring queries.

Independence from distributed storage and support for operations based on one single master node.

Time series data collection through HTTP pull.

Data push through Pushgateway.

Acquisition of collection target servers through dynamic scrape configuration or static configuration.

Integration with Grafana to easily support various visual charts and dashboards.

### Features

According to the layering of monitoring, TMP covers business monitoring, application layer monitoring, middleware monitoring, and system layer monitoring. Together with Cloud Monitor's alarming capabilities and open-source Grafana, it can provide a one-stop all-round monitoring system to help your quickly identify and locate business problems and reduce the impact of various faults on your business.

System layer monitoring: CPU, memory, disk, network, etc.

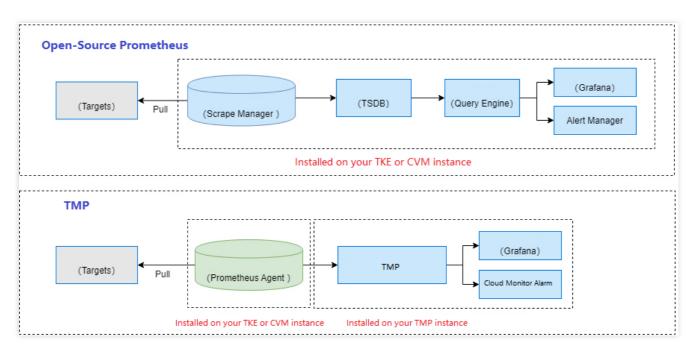
Intermediate component layer monitoring: Kafka, MySQL, Redis, etc.

Application layer monitoring: application services such as JVM, HTTP, and RPC.

Business monitoring: golden business metrics such as the number of logins and the number of orders.

# Strengths

TMP has the following strengths over the open-source Prometheus:



Lighter, more stable, and more available.

Fully compatible with the open-source Prometheus ecosystem.

Free of manual setup, saving the development costs.

Highly integrated with TKE, saving the development costs of integrating with Kubernetes.

Integrated with Cloud Monitor's alarming system, saving the costs of developing alarm notifications.

Integrated with commonly used Grafana dashboards and alerting rule templates.

# Use Cases

Last updated : 2024-08-07 21:50:26

#### **Integrated Monitoring**

TMP provides the one-stop out-of-the-box Prometheus monitoring service, which is natively integrated with Grafana dashboards and Cloud Monitor alarms for various monitoring scenarios such as basic services, application layer, and container services.

#### **Application Service Monitoring**

#### Scenario 1

An application provides external API services but their quality information cannot be secured. TMP can be integrated according to the development language to monitor the access request volume, delay, and success rate of APIs in real time.

#### Scenario 2

TMP also detects service exceptions to help you understand which APIs an exception responds to, which servers an exception occurs on, and whether an exception occurs on individual servers or in the entire cluster.

#### Scenario 3

For Java applications, TMP can monitor the GC, memory, and thread status of individual servers to help you fully understand the internal status of JVM.

#### **CVM Monitoring**

If your service is deployed in CVM, you must modify the Prometheus scrape configuration almost every time the service is scaled. For this kind of scenario, with the aid of the tagging capability of the Tencent Cloud platform and the tag discovery capability of the Prometheus agent, you only need to properly associate tags with CVM instances to easily manage and monitor target objects, which helps you get rid of continuously updating the configuration manually; for example:

1. Service A is deployed on 2 CVM instances at the same time, and the instances are associated with the tag "service name: A";

2. The original number of CVM instances cannot meet the requirements of a business campaign, and 3 more instances should be added. At this time, you only need to associate the tag "service name: A" with these new instances, and then the agent will automatically discover them and actively scrape their monitoring metrics;

3. After the campaign is over, 3 CVM instances are removed, and the agent can automatically discover the instance deactivation and stop scraping their monitoring metrics.

#### **Custom Monitoring**



You can use TMP to customize the reported metric monitoring data so as to monitor internal status of applications or services, such as the number of processed requests and the number of orders. You can also monitor the processing duration of some core logic, such as requesting external services.

# Strengths

Last updated : 2024-08-07 21:50:34

#### Lightweight Service

Compared with the open-source Prometheus monitoring service, TMP features a more lightweight overall structure. You can use a TMP instance simply after creating it in Cloud Monitor, where an agent can complete data scrape by using less than 1 GB of memory.

#### **High Stability and Reliability**

TMP only occupies megabytes of resources, much fewer than the open-source Prometheus does. It also integrates Tencent Cloud storage services and its own replica capabilities to reduce the number of system interruptions and provide services with a higher availability.

#### **Great Openness**

TMP offers the out-of-the-box Grafana service. It also integrates a wealth of Kubernetes basic monitoring services and dashboards for common service monitoring, which can be quickly used after activation.

#### Low Costs

TMP provides the native Prometheus monitoring service. After you purchase a TMP instance, you can quickly integrate it with TKE to monitor services running on Kubernetes, which eliminates the costs of setup, OPS, and development.

#### **High Scalability**

TMP features an unlimited data storage capacity not restricted to local disks. It can be dynamically scaled based on Tencent Cloud's proprietary sharding and scheduling technologies to meet your elastic needs. It also supports CLB for better load balancing. This helps solve the pain points that the open-source Prometheus cannot be scaled horizontally.

#### **High Compatibility**

TMP is fully compatible with Prometheus protocols, support core APIs, custom multidimensional data models, flexible query language PromQL, and collection target discovery through dynamic scrape configuration or static configuration, so you can easily migrate to it for smooth access.

# Concepts

Last updated : 2024-08-07 21:50:40

This document describes the basic concepts involved in using the TMP service, so that you can check and understand related concepts.

Concept	Description	
Exporter	Exporter is a component that collects monitoring data and provides data externally based on Prometheus monitoring specifications. There are currently hundreds of official or third-party tools available, see Exporters and Integrations.	
Job	A collection of configurations for a set of targets. After the scrape interval is defined, access restrictions will will be applied to the scraping job for a given set of targets.	
TMP instance	The logical unit for managing monitoring data collection and data storage analysis provided by TMP.	
TMP probe	Kubernetes cluster deployed on either the user or Tencent Cloud side. It is responsible for automatic discovery of collection targets, collection metrics and remote writing to other libraries.	
PromQL	The query language for the TMP service, which supports instant query and timespan query. It has a variety of built-in functions and operators, so raw data can be aggregated, sliced, predicted, and federated.	
Target	The collection target to be scraped by the Prometheus Agent. The collection target either exposes its own operation and business metrics or serves as a proxy for exposing the operation and business metrics of a monitored object.	
Alerting rule	The alert configuration of alerting rule formats in TMP, Which can be described by PromQL.	
Label	An pair of key-value used to describe the metric.	
Scrape configuration	One of the features in TMP, which can automatically discover collection targets without static configuration. It supports Kubernetes SD, Consul, Eureka, and other ways of service discovery. It also allows collection targets to be exposed via Service Monitor and Pod Monitor.	
Recording rule	The recording rule capacity in TMP. With recording rule, raw data can be processed into new metrics through PromQL to improve query efficiency.	
Integration Center	It integrates all the services supported by TMP. You can install the corresponding services as instructed on the page. After successful installation, you can view the monitoring data on the monitoring panel.	
Alerting rule	It is used to define how to trigger and send an alert.	

Cloud product monitor	TMP integrates the monitoring data of Tencent Cloud products. You can quickly install the Agent to view the monitoring data.
Metrics	Metrics are used to collect a series of labeled data exposed by the target, which can fully reflect the operation or business status of the monitored object.
TPS	The total number of reported data points per second. It is an important metric to measure the processing power of the system.
Series limit	The maximum number of metrics. The upper limit of series = (single metric $\times$ the dimension combination of the metric) $\times$ the number of metrics.

# Use Limits

Last updated : 2024-08-07 21:50:46

# **Instance** Limits

Each instance can have up to 4.5 million series. Free trial instance is limited to 2 million series. If you need to adjust the limit for a paid instance, contact us. In the case of adequate resources, we will adjust the related limit for you properly.

#### Note:

A series consists of a metric name and label. The same metric name and labels form a unique series.

# **Custom Reporting Limits**

If you use TMP's custom monitoring feature to monitor data, there will be the following limits on metrics (series with a unique \_\_\_\_\_name\_\_\_\_).

Data reporting must carry a metric name, i.e., the \_\_\_\_\_name\_\_\_ label, which can contain only ASCII letters,

characters, digits, underscores, and colons and must start with a letter and match the regex [a-zA-Z\_:][a-zA-

Z0-9\_: ] \* . For more information, see Metric names and labels.

Each metric can have up to 32 labels.

The label name can contain only ASCII letters, digits, and underscores. It must match the regex [a-Za-Z\_] [a-Za-

 Z0-9\_] \*
 . Labels beginning with
 \_\_\_\_\_
 are reserved for internal use.

The label name and label value can contain up to 1,024 and 2,048 characters respectively.

Under the same metric, the dimension combinations of labels cannot exceed 100,000. When the histogram has many buckets, the histogram-type metrics cannot be adjusted.

Total number of data points reported per second: A paid instance cannot exceed 300,000, and a free trial instance cannot exceed 100,000.

#### Note

The role of labels: In Prometheus, data is stored as time series, which are uniquely identified by the metric name and a series of labels (key-value pairs). Different labels represent different time series, so you can query the specified data by label. The more labels you add, the finer the query dimension.

# **Prometheus Query Limits**

To ensure the query efficiency and better user experience, Prometheus query has the following limits (which don't apply to metadata such as queries about labels and don't affect the Grafana metrics browser feature).

The number of time series involved in a single query cannot exceed 100,000.

The amount of data involved in a single query cannot exceed 100 MB.

There is no limit on the query frequency, but if the concurrency exceeds 15, there may be a certain queuing delay in slower large queries (the probability is low though). Large queries with a time span of more than two weeks will have a higher delay.

The above limits also apply to alarm rules and recording rules. We recommend that you limit the query scope based on your business scenario or appropriately split queries in other ways. You can also use the method of splitting first and then aggregating, such as aggregating recorded results again.

# **Other Configurations Limits**

Configuration limits:

The maximum of alarm rules can be configured for each instance: 150.

The maximum of recording rules can be configured for each instance: 150.

# Features

Last updated : 2024-08-07 21:50:52

# **Monitoring Object Access**

Feature	Description	
Instance creation	You can create a TMP instance in multiple region.	
Integration center	The integration center supports quick installation and custom access of various components. After successful installation, you can view monitoring data in Grafana.	
Health check	Health check detects the service connectivity on a regular basis to monitor the service health, helping you stay up to date with the service health in real time and promptly discover exceptions to improve the SLA.	
Custom monitoring	You can customize monitoring data reporting and monitor key business metrics, such as the number of requests, orders placed, and time consumed requesting external services.	

# Monitoring Metrics for Collection

Feature	Description	
Scrape configuration	ServiceMonitor: A built-in scrape configuration feature in TMP, which is automatically enabled when connected. Currently, the default objects for scrape metrics collection are pods contained in all namespaces under the Kubernetes cluster. ServiceMonitor: It collects the monitoring data in the corresponding endpoints of services based on Prometheus Operator in the K8s ecosystem. PodMonitor: It collects the corresponding monitoring data in pods based on Prometheus Operator in the K8s ecosystem.	
Targets	You can visually view the target being scraped through Targets and check whether the scraping status is normal. You can also view the metrics exposed in the target.	

# Monitoring Data

Feature	Description



Getting Remote Write address	The Remote Write feature can store TMP data as a remote database. You can use the Remote Write address to store the monitoring data of self-built Prometheus in the TMP instance to achieve remote storage, and visualize it in the same Grafana system.
Recordir rule	A recording rule allows you to calculate some commonly used or complex metrics in advance and then store the calculated data in new data metrics. In this way, querying the calculated data will be faster than querying the original data. This can solve the problems of complicated user configuration and slow query.

# Monitoring Data Display

Feature	Description
Grafana	There are numerous embedded Grafana dashboards available. You can also customize one if necessary. Plugins that are often used on the Grafana official website are also preset, and you can install them quickly in the console.
Instance monitoring	It supports the monitoring of TMP instance status and usage, including TMP instance storage, alert sending, Grafana requests and the number of dashboards, helping you check the usage of TMP instances in real time.
HTTP API	This API is used to get the TMP data address. You can use this address to connect the monitoring data of the TMP instance to the self-built Grafana dashboard for data display, or perform secondary development of the TMP monitoring data.

# Alerts

Feature	Description
Creating alerting rules	There are numerous preset alerting rules available. You can also customize one for a specific monitoring object. TMP integrates the alarm notification template of Tencent Cloud Observability Platform (TCOP), which notifies you to take measures in time when certain metrics are abnormal.
Managing alerting rules	You can perform operations such as enabling, disabling, editing, and deleting alerting rules. You can also quickly import other instance alerts.

# Service Regions

Last updated : 2024-08-08 09:55:01

#### Note:

Service regions refer to locations of physical data centers or servers. You can regard them as the regions or countries where a service or resource is available. Once a resource is successfully created in a data center in a specific region, the region of the resource usually cannot be changed.

Regions and AZs supported by TMP are as follows:

Region	AZ
South Chipa (Cuanazhou) an guangzhou	Guangzhou Zone 3 ap-guangzhou-3
South China (Guangzhou) ap-guangzhou	Guangzhou Zone 4 ap-guangzhou-4
	Shanghai Zone 2 ap-shanghai-2
East China (Shanghai) an changhai	Shanghai Zone 3 ap-shanghai-3
East China (Shanghai) ap-shanghai	Shanghai Zone 4 ap-shanghai-4
	Shanghai Zone 5 ap-shanghai-5
	Hong Kong (China) Zone 1 ap-hongkong-1
Hong Kong (China), Macao (China), and Taiwan (China) (Hong Kong (China)) ap-hongkong	Hong Kong (China) Zone 2 ap-hongkong-2
	Hong Kong (China) Zone 3 ap-hongkong-3
	Beijing Zone 3 ap-beijing-3
	Beijing Zone 4 ap-beijing-4
North China (Beijing) ap-beijing	Beijing Zone 5 ap-beijing-5
	Beijing Zone 6 ap-beijing-6
	Beijing Zone 7 ap-beijing-7
Southwast China (Chanadu) an abanadu	Chengdu Zone 1 ap-chengdu-1
Southwest China (Chengdu) ap-chengdu	Chengdu Zone 2 ap-chengdu-2
Southwest China (Chongqing) ap-chongqing	Chongqing Zone 1 ap-chongqing-1
East China (Nanjing) ap-nanjing	Nanjing Zone 1 ap-nanjing-1



	Nanjing Zone 2 ap-nanjing-2
North America (Toronto) na-Toronto	Toronto Zone 1 na-Toronto-1
	Silicon Valley Zone 1 na-siliconvalley-1
West US (Silicon Valley) na-siliconvalley	Silicon Valley Zone 2 na-siliconvalley-2
Europe (Frankfurt) eu-frankfurt	Frankfurt Zone 1 eu-frankfurt-1
Fact US (Virginia) no achburn	Virginia Zone 1 na-ashburn-1
East US (Virginia) na-ashburn	Virginia Zone 2 na-ashburn-2
South America (Sao Paulo) sa-saopaulo	Sao Paulo Zone 1 sa-saopaulo-1
	Singapore Zone 1 ap-singapore-1
	Singapore Zone 2 ap-singapore-2
Southeast Asia (Singapore) ap-singapore	Singapore Zone 3 ap-singapore-3
	Singapore Zone 4 ap-singapore-4
Asia Pacific (Seoul) ap-seoul	Seoul Zone 1 ap-seoul-1
Asia Dasifa (Mumbai) an hambay	Mumbai Zone 1 ap-bombay-1
Asia Pacific (Mumbai) ap-bombay	Mumbai Zone 2 ap-bombay-2
Asia Pacific (Bangkok) ap-bangkok	Bangkok Zone 2 ap-Bangkok-2
Asia Pacific (Tokyo) ap-tokyo	Tokyo Zone 1 ap-tokyo-1
	Tokyo Zone 2 ap-tokyo-2
Southoast Asia Dasifia (Jakarta) and jakarta	Jakarta Zone 1 spa-Jakarta-1
Southeast Asia Pacific (Jakarta) spa-jakarta	Jakarta Zone 2 spa-Jakarta-2

# Access Guide Scrape Configuration Description

Last updated : 2024-07-23 18:10:38

### Overview

Prometheus mainly uses PULL to scrape the monitoring APIs exposed by the target service; therefore, you need to configure the corresponding scrape task to request the monitoring data and write it into the storage provided by Prometheus. Currently, Prometheus provides the configurations of the following tasks:

Native job configuration: the native scrape job configuration of Prometheus is provided.

PodMonitor: It collects the corresponding monitoring data in Pods based on Prometheus Operator in the K8s ecosystem.

ServiceMonitor: It collects the monitoring data in the corresponding Endpoints of Services based on Prometheus Operator in the K8s ecosystem.

#### Note:

Configuration items in [] are optional.

# Native job configuration

The relevant configuration items are as detailed below:



```
# Scrape task name. `label(job=job_name)` will be added to the corresponding metric
job_name: <job_name>
# Scrape task interval
[ scrape_interval: <duration> | default = <global_config.scrape_interval> ]
# Scrape request timeout period
[ scrape_timeout: <duration> | default = <global_config.scrape_timeout> ]
# Scrape task request URI path
```

```
🕗 Tencent Cloud
```

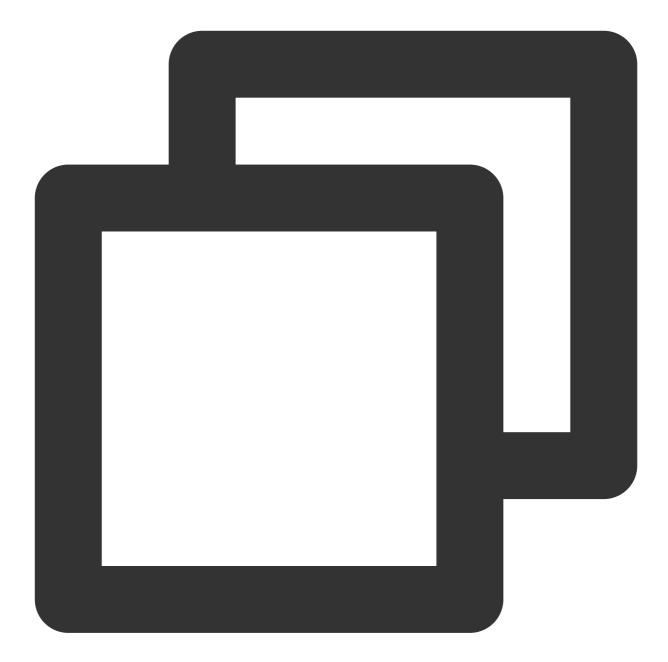
```
[ metrics_path: <path> | default = /metrics ]
# Solve the conflict between the scraped label and the label added to Prometheus on
# true: Retain the scraped label and ignore the label conflicting with Prometheus o
# false: Add `exported_<original-label>` before the scraped label to add the label
[ honor_labels: <boolean> | default = false ]
# Whether to use the time generated on the scrape target
# true: Use the time on the target
# false: Directly ignore the time on the target
[ honor_timestamps: <boolean> | default = true ]
# Scrape protocol: HTTP or HTTPS
[ scheme: <scheme> | default = http ]
# URL parameter of the scrape request
params:
  [ <string>: [<string>, ...] ]
# Use `basic_auth` to set `Authorization` in the scrape request header. `password`
basic_auth:
  [ username: <string> ]
  [ password: <secret> ]
  [ password_file: <string> ]
# Use `bearer_token` to set `Authorization` in the scrape request header. `bearer_t
[ bearer_token: <secret> ]
# Use `bearer_token` to set `Authorization` in the scrape request header. `bearer_t
[ bearer_token_file: <filename> ]
# Specify whether the scrape connection passes through a TLS secure channel and con
tls_config:
  [ <tls_config> ]
# Use a proxy service to scrape metrics on the target and enter the corresponding p
[ proxy_url: <string> ]
# Use static configuration to specify the target. For more information, see the des
static_configs:
  [ - <static_config> ... ]
# Set the CVM scrape configuration. For more information, see the description below
cvm_sd_configs:
  [ - <cvm_sd_config> ... ]
# After scraping the data, change the label on the target through the relabeling me
```



```
# For more information on `relabel_config`, see the description below
relabel_configs:
    [ - <relabel_config> ... ]
# After the data is scraped and before it is written, use the relabeling mechanism
# For more information on `relabel_config`, see the description below
metric_relabel_configs:
    [ - <relabel_configs ... ]
# Limit of data points in one scrape. 0: no limit. Default value: 0
[ sample_limit: <int> | default = 0 ]
# Limit of targets in one scrape. 0: no limit. Default value: 0
[ target_limit: <int> | default = 0 ]
```

#### static\_config configuration

The relevant configuration items are as detailed below:



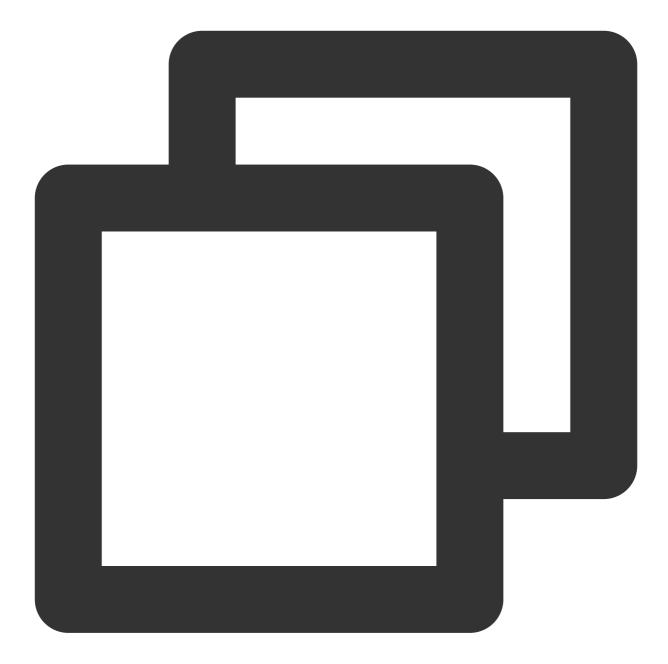
```
# Specify the corresponding target host value, such as `ip:port`
targets:
  [ - '<host>' ]
# Add the corresponding label to all targets, which is similar to a global label
labels:
  [ <labelname>: <labelvalue> ... ]
```

#### cvm\_sd\_config configuration

CVM scrape configuration uses TencentCloud API to automatically get the CVM instance list, and the CVM instance's private IP is used by default. Scrape configuration will generate the following meta labels, which can be used in relabeling configuration.

Label	Description
meta_cvm_instance_id	Instance ID
meta_cvm_instance_name	Instance name
meta_cvm_instance_state	Instance status
meta_cvm_instance_type	Instance model
meta_cvm_OS	Instance OS
meta_cvm_private_ip	Private IP
meta_cvm_public_ip	Public IP
meta_cvm_vpc_id	VPC ID
meta_cvm_subnet_id	Subnet ID
meta_cvm_tag_ <tagkey></tagkey>	Instance tag value
meta_cvm_region	Instance region
meta_cvm_zone	Instance AZ

CVM scrape configuration description:



# Tencent Cloud region. For the region list, visit https://cloud.tencent.com/docume
region: <string>

```
# Custom endpoint.
[ endpoint: <string> ]
```

```
# Credential information for accessing TencentCloud API. If it is not set, the valu
# Leave it empty if you use a CVM scrape task in **Integration Center** for configu
[ secret_id: <string> ]
[ secret_key: <secret> ]
```

```
# CVM list refresh interval
[ refresh_interval: <duration> | default = 60s ]
# Port for scraping metrics
ports:
    - [ <int> | default = 80 ]
# CVM list filtering rule. For more information on the supported filtering rules, v
filters:
    [ - name: <string>
        values: <string>, [...] ]
```

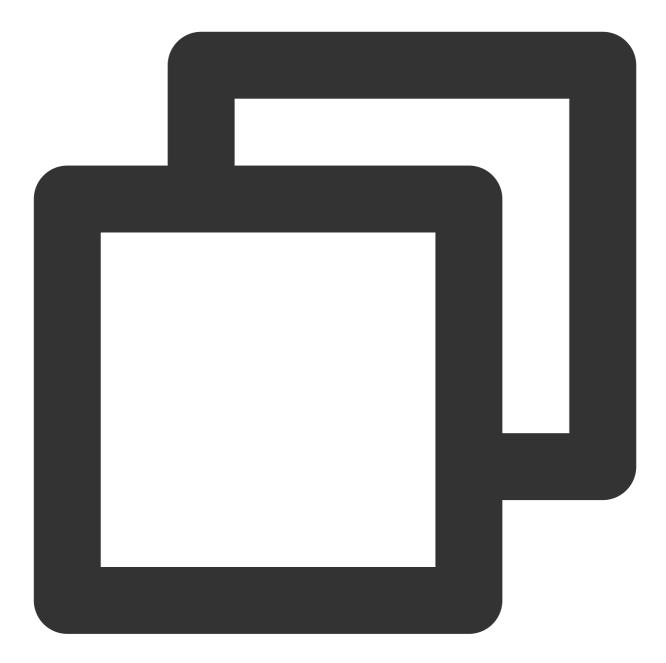
#### Note:

If a CVM scrape task in **Integration Center** is used to configure cvm\_sd\_configs , the integration automatically uses the preset role authorization of the service for security considerations. You don't need to manually enter the secret\_id , secret\_key , and endpoint parameters.

#### Sample

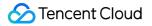
Static configuration

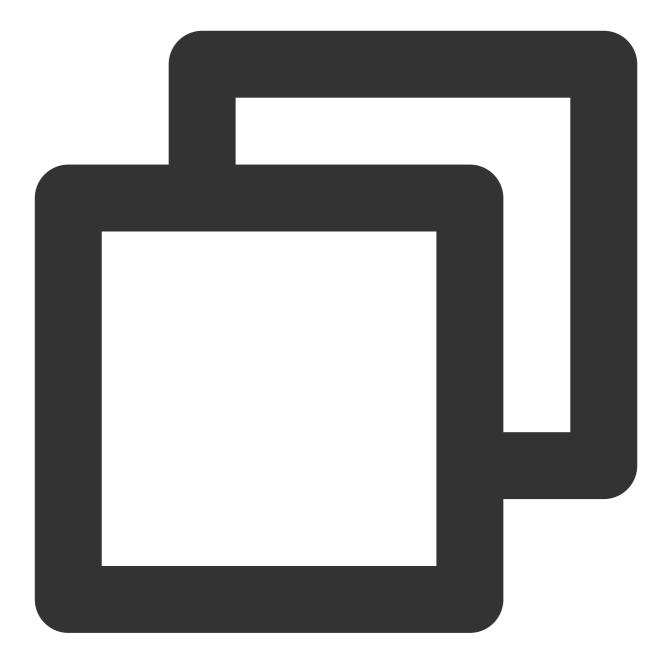




```
job_name: prometheus
scrape_interval: 30s
static_configs:
- targets:
- 127.0.0.1:9090
```

**CVM scrape configuration** 





```
job_name: demo-monitor
cvm_sd_configs:
- region: ap-guangzhou
  ports:
  - 8080
  filters:
  - name: tag:service
    values:
    - demo
relabel_configs:
- source_labels: [__meta_cvm_instance_state]
```

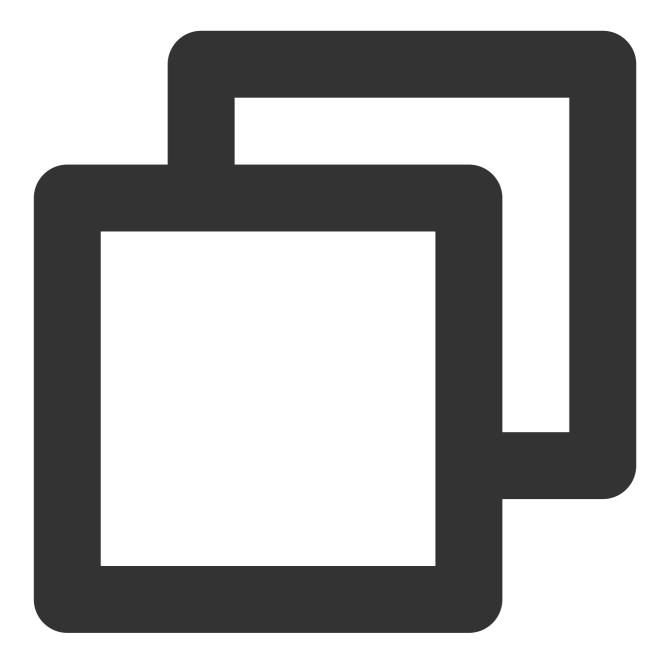


```
regex: RUNNING
action: keep
- regex: __meta_cvm_tag_(.*)
replacement: $1
action: labelmap
- source_labels: [__meta_cvm_region]
target_label: region
action: replace
```

# PodMonitor

The relevant configuration items are as detailed below:





# Prometheus Operator CRD version
apiVersion: monitoring.coreos.com/v1
# Corresponding K8s resource type, which is PodMonitor here
kind: PodMonitor
# Corresponding K8s metadata. Here, only the `name` is concerned. If `jobLabel` is
metadata:
name: redis-exporter # Enter a unique name
namespace: cm-prometheus # The namespace is fixed. Do not change it
# Describe the selection of the scrape target Pod and the configuration of the scra
spec:
# Enter the target Pod label. PodMonitor will use the corresponding value as the



```
# If Pod YAML configuration is to be viewed, use the value in `pod.metadata.label
# If `Deployment/Daemonset/Statefulset` is to be viewed, use `spec.template.metad
[ jobLabel: string ]
# Add the label on the corresponding Pod to the target label
[ podTargetLabels: []string ]
# Limit of data points in one scrape. 0: no limit. Default value: 0
[ sampleLimit: uint64 ]
# Limit of targets in one scrape. 0: no limit. Default value: 0
[ targetLimit: uint64 ]
# Configure the Prometheus HTTP port to be exposed and scraped. You can configure
podMetricsEndpoints:
[ - <endpoint_config> ... ] # For more information, see the endpoint description
# Select the namespace where the Pod to be monitored resides. If it is not specif
[ namespaceSelector: ]
  # Whether to select all namespaces
  [ any: bool ]
  # List of namespace to be selected
  [ matchNames: []string ]
# Enter the label of the Pod to be monitored to locate the target Pod. For more i
selector:
  [ matchExpressions: array ]
    [ example: - {key: tier, operator: In, values: [cache]} ]
  [ matchLabels: object ]
    [ example: k8s-app: redis-exporter ]
```

#### Sample



```
apiVersion: monitoring.coreos.com/v1
kind: PodMonitor
metadata:
   name: redis-exporter # Enter a unique name
   namespace: cm-prometheus # The namespace is fixed. Do not change it
spec:
   podMetricsEndpoints:
    - interval: 30s
    port: metric-port # Enter the name of the corresponding port of the Promethe
    path: /metrics # Enter the value of the corresponding path of the Prometheus
    relabelings:
```

- action: replace
sourceLabels:
- instance
regex: (.*)
targetLabel: instance
replacement: 'crs-xxxxxx' # Change it to the corresponding Redis instance I
- action: replace
sourceLabels:
- instance
regex: (.*)
targetLabel: ip
replacement: '1.x.x.x' # Change it to the corresponding Redis instance IP
namespaceSelector: # Select the namespace where the Pod to be monitored resid
matchNames:
- redis-test
selector: # Enter the label value of the Pod to be monitored to locate the t
matchLabels:
k8s-app: redis-exporter

# ServiceMonitor

The relevant configuration items are as detailed below:



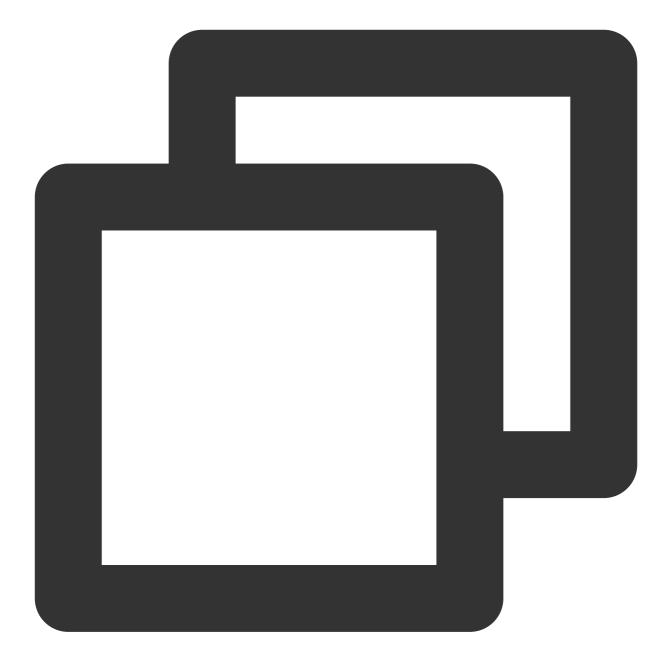


# Prometheus Operator CRD version
apiVersion: monitoring.coreos.com/v1
# Corresponding K8s resource type, which is ServiceMonitor here
kind: ServiceMonitor
# Corresponding K8s metadata. Here, only the `name` is concerned. If `jobLabel` is
metadata:
name: redis-exporter # Enter a unique name
namespace: cm-prometheus # The namespace is fixed. Do not change it
$\ensuremath{\sharp}$ Describe the selection of the scrape target Pod and the configuration of the scra
spec:



```
# Enter the target Pod label (metadata/labels). ServiceMonitor will use the corre
[ jobLabel: string ]
# Add the label on the corresponding Service to the target label
[ targetLabels: []string ]
# Add the label on the corresponding Pod to the target label
[ podTargetLabels: []string ]
# Limit of data points in one scrape. 0: no limit. Default value: 0
[ sampleLimit: uint64 ]
# Limit of targets in one scrape. 0: no limit. Default value: 0
[ targetLimit: uint64 ]
# Configure the Prometheus HTTP port to be exposed and scraped. You can configure
endpoints:
[ - <endpoint_config> ... ] # For more information, see the endpoint description
# Select the namespace where the Pod to be monitored resides. If it is not specif
[ namespaceSelector: ]
  # Whether to select all namespaces
  [ any: bool ]
  # List of namespace to be selected
  [ matchNames: []string ]
# Enter the label of the Pod to be monitored to locate the target Pod. For more i
selector:
  [ matchExpressions: array ]
    [ example: - {key: tier, operator: In, values: [cache]} ]
  [ matchLabels: object ]
    [ example: k8s-app: redis-exporter ]
```

Sample

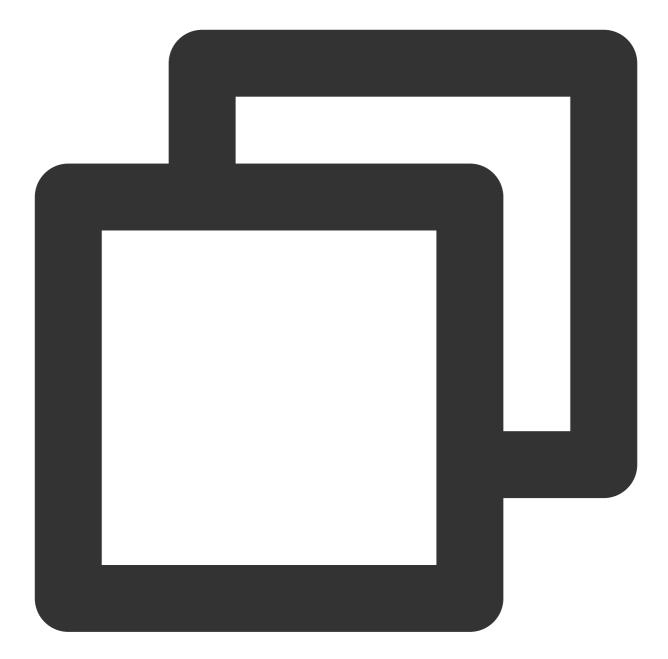




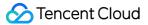
```
path: /metrics
relabelings:
# ** There must be a label named `application`. Here, suppose that K8s has a
# Use the `replace` action of `relabel` to replace it with `application`
- action: replace
sourceLabels: [__meta_kubernetes_pod_label_app]
targetLabel: application
# Select the namespace where the Service to be monitored resides
namespaceSelector:
matchNames:
- golang-demo
# Enter the label value of the Service to be monitored to locate the target Ser
selector:
matchLabels:
app: golang-app-demo
```

### endpoint\_config configuration

The relevant configuration items are as detailed below:



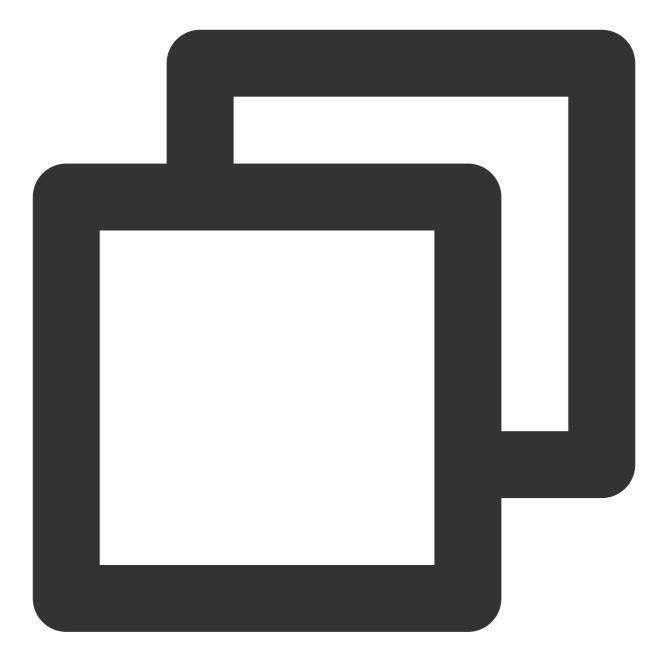
```
# Corresponding port name. Note that it is not the port number here. Default value:
# ServiceMonitor: `Service>spec/ports/name`
# PodMonitor description:
# If Pod YAML configuration is to be viewed, use the value in `pod.spec.container
# If `Deployment/Daemonset/Statefulset` is to be viewed, use `spec.template.spec.co
[ port: string | default = 80]
# Scrape task request URI path
[ path: string | default = /metrics ]
# Scrape protocol: HTTP or HTTPS
[ scheme: string | default = http]
# URL parameter of the scrape request
```



```
[ params: map[string][]string]
# Scrape task interval
[ interval: string | default = 30s ]
# Scrape task timeout period
[ scrapeTimeout: string | default = 30s]
# Specify whether the scrape connection passes through a TLS secure channel and con
[ tlsConfig: TLSConfig ]
# Read the value of the bearer token through the corresponding file and add it to t
[ bearerTokenFile: string ]
# You can use the corresponding K8s secret key to read the bearer token. Note that
[ bearerTokenSecret: string ]
# Solve the conflict between the scraped label and the label added to Prometheus on
# true: Retain the scraped label and ignore the label conflicting with Prometheus o
# false: Add `exported_<original-label>` before the scraped label to add the label
[ honorLabels: bool | default = false ]
# Whether to use the time generated on the scrape target
# true: Use the time on the target
# false: Directly ignore the time on the target
[ honorTimestamps: bool | default = true ]
# `basic auth` authentication information. Enter the corresponding K8s secret key v
[ basicAuth: BasicAuth ]
# Use a proxy service to scrape metrics on the target and enter the corresponding p
[ proxyUrl: string ]
# After scraping the data, change the label on the target through the relabeling me
# For more information on `relabel_config`, see the description below
relabelings:
[ - <relabel_config> ...]
# After the data is scraped and before it is written, use the relabeling mechanism
# For more information on `relabel_config`, see the description below
metricRelabelings:
[ - <relabel_config> ...]
```

### relabel\_config configuration

The relevant configuration items are as detailed below:



```
# Specify which labels are to be taken from the original labels for relabeling. The
# The corresponding configuration item for PodMonitor/ServiceMonitor is `sourceLabe
[ source_labels: '[' <labelname> [, ...] ']' ]
# Define the separator symbol for concatenating the labels to be relabeled. Default
[ separator: <string> | default = ; ]
# If `action` is ` replace` or `hashmod`, you need to use the `target_label` to spe
# The corresponding configuration item for PodMonitor/ServiceMonitor is `targetLabe
```

```
[ target_label: <labelname> ]
```



```
# Regex for regular match of the values of source labels
[ regex: <regex> | default = (.*) ]
```

# Calculate the modulus of the MD5 value of the source label. The modulo operation [ modulus: <int> ]

# If `action` is `replace`, use `replacement` to define the expression to be replac
[ replacement: <string> | default = \$1 ]

# Perform an action based on the value matched by the regex. Valid values of `actio
# replace: Replace the matched value with that defined in `replacement` if the rege
# keep: Drop the value if the regex has no matches

# drop: Drop the value if the regex has any match

# hashmod: Calculate the modulus of the MD5 value of the source label based on the

 $\ensuremath{\texttt{\#}}$  labelmap: Use `replacement` to replace the corresponding label name if the regex

 $\ensuremath{\texttt{\#}}$  labeldrop: Delete the corresponding label name if the regex has any match

# labelkeep: Delete the corresponding label name if the regex has no matches

[ action: <relabel\_action> | default = replace ]

# **Custom Monitoring**

Last updated : 2024-08-07 21:53:55

# Overview

You can use TMP to customize the reported metric monitoring data so as to monitor internal status of applications or services, such as the number of processed requests and the number of orders. You can also monitor the processing duration of some core logic, such as requesting external services.

This document uses Go as an example to describe how to use TMP to customize reported metrics, visualization, and alerting.

# Supported Programming Languages

Official SDKs from the native Prometheus community: Go Java or Scala Python Ruby Third-Party SDKs for other programming languages: **Bash** С C++ **Common Lisp** Dart Elixir Erlang Haskell Lua for NGINX Lua for Tarantool .NET/C# Node.js Perl PHP R Rust

For more information, please see CLIENT LIBRARIES.

# Data Model

Prometheus has multidimensional analysis capabilities. A data model consists of the following parts:

Metric Name + Labels + Timestamp + Value/Sample

Metric Name: monitoring object (for example, <a href="http\_request\_total">http\_request\_total</a> indicates the current total number of HTTP requests received by the system).

Labels: characteristics dimensions of the current sample, which are in K/V structure. Through such dimensions,

Prometheus can filter, aggregate, and perform other operations on the sample data.

Timestamp: a timestamp accurate down to the millisecond

Value: a float64 value, which indicates the current sample value.

Metric Name/Labels can contain only ASCII characters, digits, underscores, and colons and must comply with the regular expression [a-zA-Z\_:][a-zA-Z0-9\_:]\*.

For more information on a data model, please see DATA MODEL.

For the best practice of metric and label naming, please see METRIC AND LABEL NAMING.

## Metric Tracking Method

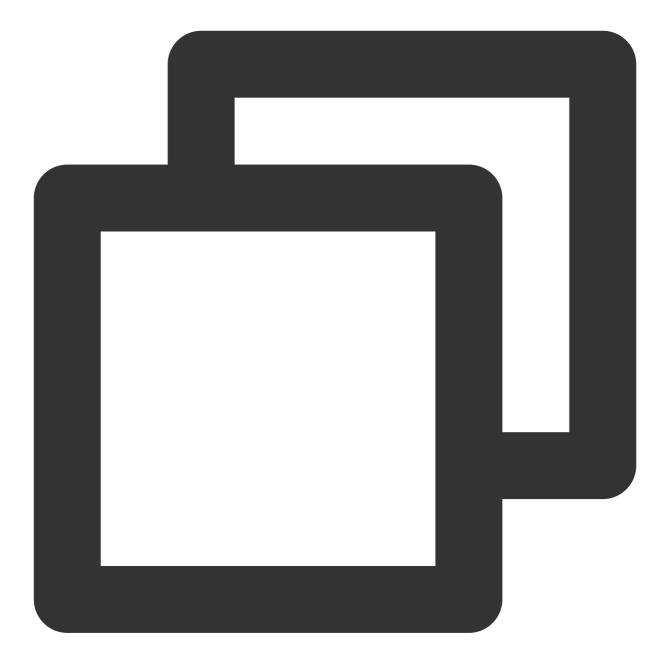
Prometheus provides four metric types for different monitoring scenarios: Counter, Gauge, Histogram, and Summary, as described below. For more information, please see METRIC TYPES.

The Prometheus community provides SDKs for multiple programing languages, all of which are basically similar in usage but differ mostly in syntax. This document uses Go as an example to describe how to report custom monitoring metrics.

### Counter

A metric in Counter type increases monotonically and will be reset after service restart. You can use counters to monitor the numbers of requests, exceptions, user logins, orders, etc. You can use a counter to monitor the number of orders as follows:



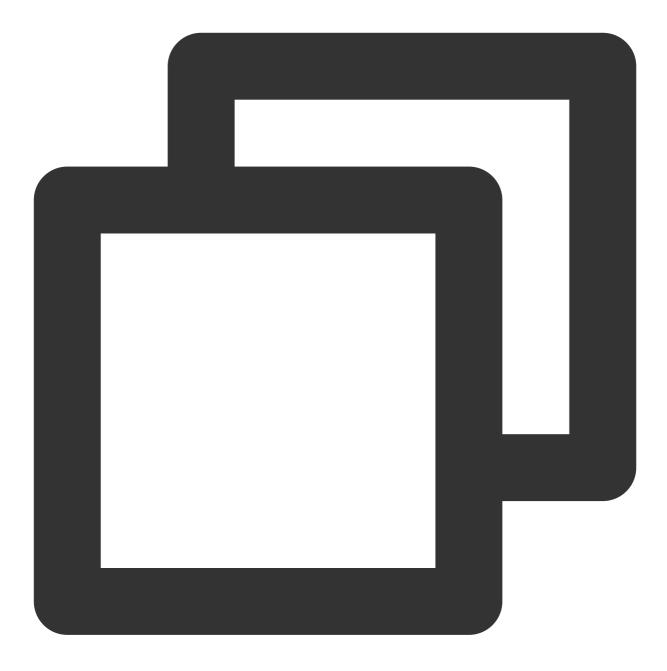


```
package order
import (
    "github.com/prometheus/client_golang/prometheus"
    "github.com/prometheus/client_golang/prometheus/promauto"
)
// Define the counter object to be monitored
var (
    opsProcessed = promauto.NewCounterVec(prometheus.CounterOpts{
        Name: "order_service_processed_orders_total",
```

```
Help: "The total number of processed orders",
}, []string{"status"}) // Processing status
)
// Process the order
func makeOrder() {
    opsProcessed.WithLabelValues("success").Inc() // Success
    // opsProcessed.WithLabelValues("fail").Inc() // Failure
    // Order placement business logic
}
```

For example, you can use the <code>rate()</code> function to get the order increase rate:



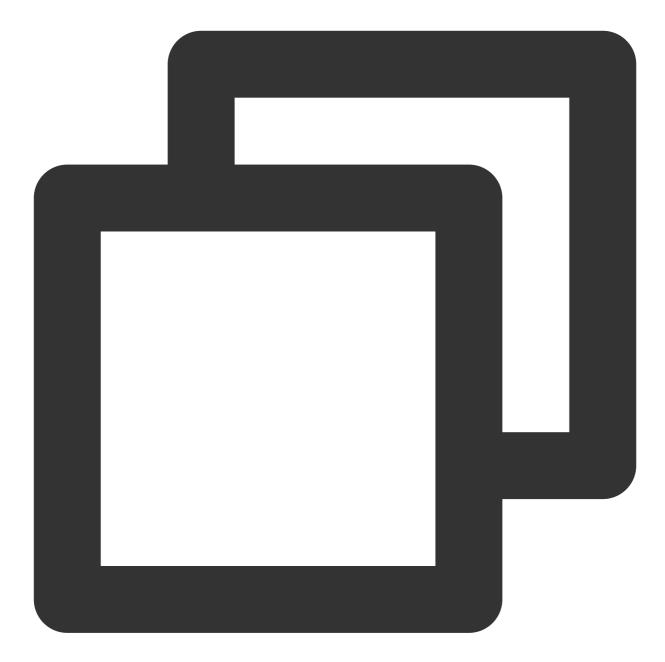


rate(order\_service\_processed\_orders\_total[5m])

### Gauge

A gauge is a current value, which can be increased or reduced during metric timestamping. You can use gauges to monitor the current memory utilization, CPU utilization, current number of threads, queue size, etc. You can use a gauge to monitor the size of an order queue as follows:



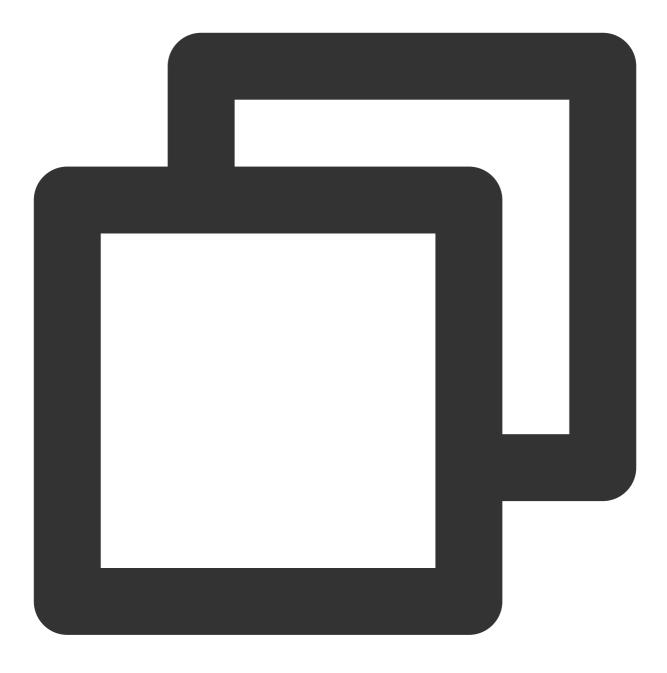


```
package order
import (
    "github.com/prometheus/client_golang/prometheus"
    "github.com/prometheus/client_golang/prometheus/promauto"
)
// Define the gauge object to be monitored
var (
    queueSize = promauto.NewGaugeVec(prometheus.GaugeOpts{
        Name: "order_service_order_queue_size",
```

```
Help: "The size of order queue",
   }, []string{"type"})
)
type OrderQueue struct {
   queue chan string
}
func newOrderQueue() *OrderQueue {
   return &OrderQueue{
        queue: make(chan string, 100),
    }
}
// Produce an order message
func (q *OrderQueue)produceOrder() {
   // Produce an order message
    // Increase the queue size by 1
    queueSize.WithLabelValues("make_order").Inc() // Order placement queue
    // queueSize.WithLabelValues("cancel_order").Inc() // Order cancellation queue
}
// Consume an order message
func (q *OrderQueue)consumeOrder() {
    // Consume an order message
   // Reduce the queue size by 1
    queueSize.WithLabelValues("make_order").Dec()
}
```

You can use the gauge metric to directly view the current size of each type of queue of an order:





order\_service\_order\_queue\_size

#### Histogram

Prometheus calculates the sample distribution based on the configured Bucket to generate a histogram, which can be processed subsequently and is generally used for duration monitoring. For example, you can use a histogram to calculate the latencies of P99, P95, and P50 and monitor the numbers of processed items. With histograms, you don't need to use counters to count items. In addition, you can use histograms to monitor metrics such as API response time and database access time.

A histogram can be used in a similar way to a summary, so you can directly refer to the summary usage.

### Summary

A summary is similar to a histogram, as it also calculates the sample distribution, but their differences lie in that a summary calculates the distribution (P99/P95/Sum/Count) on the client and therefore uses more client resources, and the data cannot be calculated and processed in an aggregated manner subsequently. You can use summaries to monitor metrics such as API response time and database access duration.

You can use a summary to monitor the order processing duration as follows:



package order



```
import (
    "net/http"
    "time"
    "github.com/prometheus/client_golang/prometheus"
    "github.com/prometheus/client_golang/prometheus/promauto"
    "github.com/prometheus/client_golang/prometheus/promhttp"
)
// Define the summary object to be monitored
var (
    opsProcessCost = promauto.NewSummaryVec(prometheus.SummaryOpts{
        Name: "order_service_process_order_duration",
        Help: "The order process duration",
    }, []string{"status"})
)
func makeOrder() {
    start := time.Now().UnixNano()
    // The order placement logic processing is completed, and the processing durati
    defer opsProcessCost.WithLabelValues("success").Observe((float64)(time.Now().Un
    // Order placement business logic
    time.Sleep(time.Second) // Simulate the processing duration
}
```

You can use a summary metric to directly view the average order placement processing duration:



order\_service\_processed\_order\_duration\_sum / order\_service\_processed\_order\_duration

### **Exposing Prometheus metrics**

Use promhttp.Handler() to expose the metric tracking data to the HTTP service.





```
package main
import (
    "net/http"
    "github.com/prometheus/client_golang/prometheus/promhttp"
)
func main() {
    // Business code
```

}

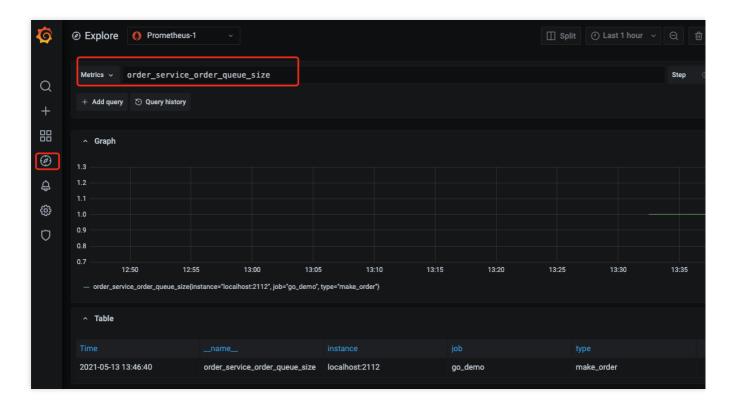
```
// Expose Prometheus metrics in the HTTP service
http.Handle("/metrics", promhttp.Handler())
// Business code
```

# **Collecting Data**

After the tracking of custom metrics for your business is completed and the application is released, you can use Prometheus to collect the monitoring metric data. For more information, please see Go Integration.

# Viewing Monitoring Data and Alerts

Open the Grafana service that comes with TMP and use Explore to view the monitoring metric data as shown below. You can also customize Grafana monitoring dashboards.



You can use Prometheus together with the alarming capabilities of Cloud Monitor to trigger alerts for custom monitoring metrics in real time. For more information, please see Alert Overview and Usage.

# EMR Integration Flink Integration

Last updated : 2024-08-07 21:59:47

# Overview

When using Flink, you need to monitor its task running status to know whether the tasks run normally and troubleshoot faults. TMP integrates Pushgateway to allow Flink to write metrics and provides an out-of-the-box Grafana monitoring dashboard for it.

# Prerequisites

- 1. The EMR product you purchased includes the Flink component, and a Flink task is running in your instance.
- 2. You have created a TKE cluster in the region and VPC of your TMP instance.

# Directions

### Integrating product

### Getting Pushgateway access configuration

1. Log in to the **EMR console**, select the corresponding **instance**, and select **Basic Info** > **Instance Info** to get the Pushgateway address and token.

Service address			
Token	***** 🗖		
Remote Write address	htt		api/v1/prom/write
HTTP API	httr		api/v1 🗖
Pushgateway address		Б	

2. Get the APPID on the Account Info page.

### **Modifying Flink configuration**

1. Log in to the EMR console, select the corresponding instance, and select Cluster Service.

2. Find the **Flink** configuration item and select **Configuration Management** in the **Operation** column on the right to enter the configuration management page.

3. On the right of the page, click **Add Configuration Item** and add the following configuration items one by one:

Configuration Item	Default Value	Data Type	Description	Suggestion
metrics.reporter.promgateway.class	None	String	Name of the Java class for exporting metrics to Pushgateway	-
metrics.reporter.promgateway.jobName	None	String	Push task name	Specify an eas understandab string
metrics.reporter.promgateway.randomJobNameSuffix	true	Boolean	Whether to add a random string after the task name	Set it to `true`. no random stri is added, metr of different Flir tasks will overwrite each other
metrics.reporter.promgateway.groupingKey	None	String	Global label	Add the EMR



			added to each metric in the format of `k1=v1;k2=v2`	instance ID to distinguish between the d of different instances, suc as `instance_id=€ xxx`
metrics.reporter.promgateway.interval	None	Time	Time interval for pushing metrics, such as 30s	We recommer you set the val to about 1 min
metrics.reporter.promgateway.host	None	String	Pushgateway service address	It is the service address of the TMP instance the console
metrics.reporter.promgateway.port	-1	Integer	Pushgateway service port	It is the port of TMP instance the console
metrics.reporter.promgateway.needBasicAuth	false	Boolean	Whether the Pushgateway service requires authentication	Set it to `true`, the Pushgatev of TMP require authentication
metrics.reporter.promgateway.user	None	String	Username for authentication	It is your `APP
metrics.reporter.promgateway.password	None	String	Password for authentication	It is the access token of the TI instance in the console
metrics.reporter.promgateway.deleteOnShutdown	true	Boolean	Whether to delete the corresponding metrics on the Pushgateway after the Flink task is completed	Set it to `true`



Below is a sample configuration:

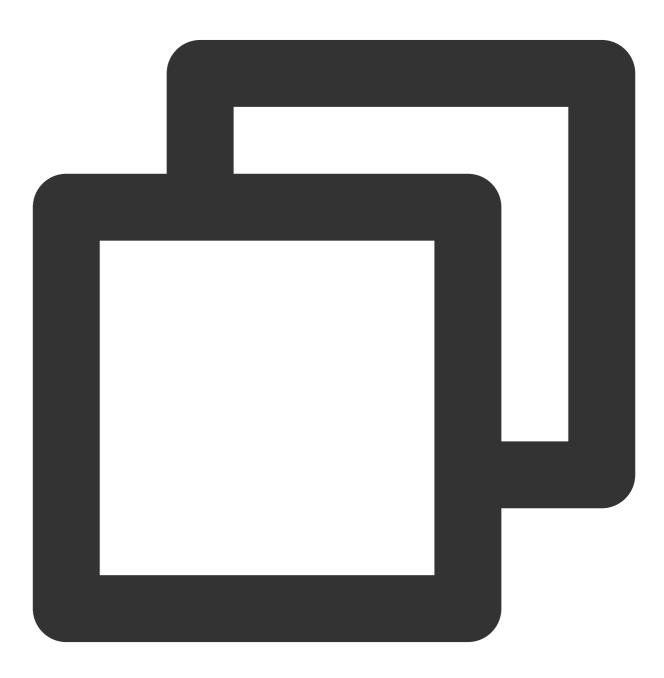


metrics.reporter.promgateway.class: org.apache.flink.metrics.prometheus.PrometheusP metrics.reporter.promgateway.jobName: climatePredict metrics.reporter.promgateway.randomJobNameSuffix:true metrics.reporter.promgateway.interval: 60 SECONDS metrics.reporter.promgateway.groupingKey:instance\_id=emr-xxxx metrics.reporter.promgateway.host: 172.xx.xx.xx metrics.reporter.promgateway.port: 9090 metrics.reporter.promgateway.needBasicAuth: true metrics.reporter.promgateway.user: appid metrics.reporter.promgateway.password: token

#### Installing Flink Pushgateway plugin

The Pushgateway plugin in the official package currently does not support configuring the authentication information, but TMP requires authentication before data can be written. Therefore, we recommend you use the JAR package we provide. We have also submitted a pull request for supporting authentication to the Flink team.

1. To prevent class conflicts, if you have already used the official Flink plugin, run the following command to delete it first:

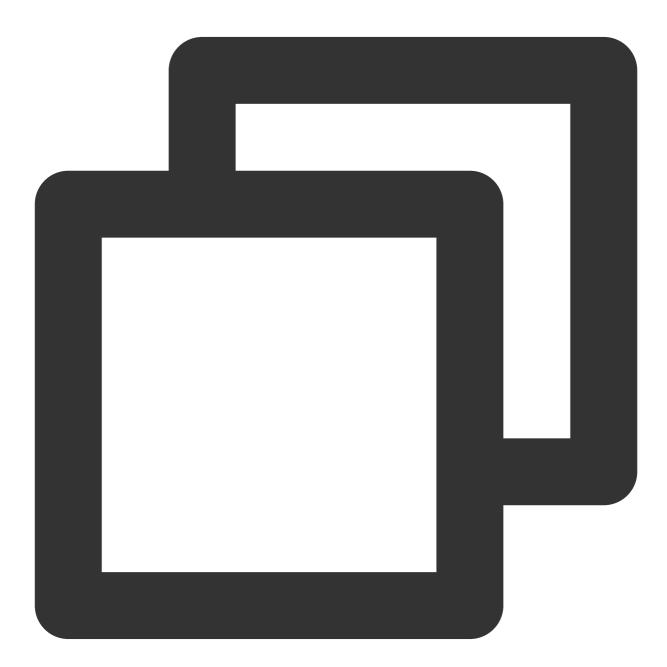


```
cd /usr/local/service/flink/lib
rm flink-metrics-prometheus*jar
```

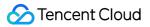
2. In the EMR console, select the corresponding instance, and select Cluster Resource > Resource

Management > Master to view the master node.

3. Click the instance ID to go to the CVM console, log in to the CVM instance, and run the following command to install the plugin:

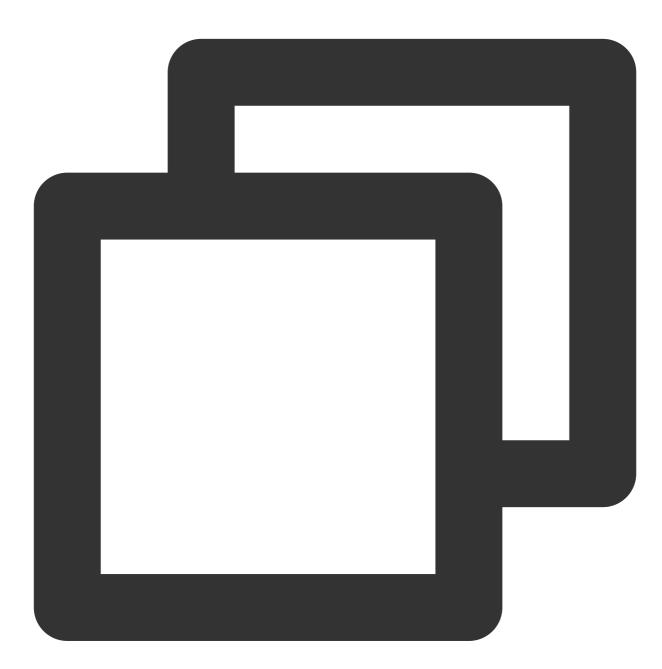


```
cd /usr/local/service/flink/lib
wget https://rig-1258344699.cos.ap-guangzhou.myqcloud.com/flink/flink-metrics-prome
```



### Verifying

1. Run the flink run command on the master node to submit a new task and view the task log:



grep metrics /usr/local/service/flink/log/flink-hadoop-client-\*.log

2. If the log contains the following content, the configuration is successfully loaded:

2020-12-11 16:09:04,114 INFO org.apache.flink.configuration.GlobalConfiguration	- Loading configuratio
trics.reporter.promgateway.class, org.apache.flink.metrics.prometheus.PrometheusPushGatewa	ayReporter
2020-12-11 16:09:04,114 INFO org.apache.flink.configuration.GlobalConfiguration	– Loading configuratio
trics.reporter.promgateway.groupingKey, instance_id=emr-	
2020-12-11 16:09:04,114 INFO org.apache.f <u>lin</u> k.configuration.GlobalConfiguration	– Loading configuratio
trics.reporter.promgateway.host, 1	
2020–12–11 16:09:04,114 INFO org.apache.flink.configuration.GlobalConfiguration	- Loading configuratio
trics.reporter.promgateway.interval, 60 SECONDS	
2020–12–11 16:09:04,115 INFO org.apache.flink.configuration.GlobalConfiguration	– Loading configuratio
trics.reporter.promgateway.jobName,	
2020–12–11 16:09:04,115 INFO org.apache.flink.configuration.GlobalConfiguration	– Loading configuratio
trics.reporter.promgateway.needBasicAuth, true	
2020–12-11 16:09:04,115 INFO org.apache.flink.configuration.GlobalConfiguration	– Loading configuratio
trics.reporter.promgateway.password, ******	
2020–12–11 16:09:04,115 INFO org.apache.flink.configuration.GlobalConfiguration	– Loading configuratio
trics.reporter.promgateway.port, 9090	
2020–12–11 16:09:04,115 INFO org.apache.flink.configuration.GlobalConfiguration	– Loading configuratio
<pre>trics.reporter.promgateway.randomJobNameSuffix, true</pre>	
2020–12–11 16:09:04,116 INFO org.apache.flink.configuration.GlobalConfiguration	- Loading configuratio
trics.reporter.promgateway.user,	

#### Note:

As tasks previously submitted in the cluster use the old configuration file, their metrics are not reported.

#### Viewing monitoring information

1. In **Integration Center** in the target TMP instance, find Flink monitoring, install the corresponding Grafana dashboard, and then you can enable the Flink monitoring dashboard.

2. Enter Grafana and click



to expand the Flink monitoring panel.

Ъ	Flink
	Flink Cluster Flink
	Flink Job Flink
	Flink Job List Flink
	Flink Task Flink

3. Click Flink Job List to view the monitoring information.



DetaBource	Prometheus -	DAR REF D	env 👘	÷				
$V_{-}$								
					Job ID 💎	Job 🛱 🖓		
emr					d0b80600b85c29x2c03d64de83ce44e	PrometheusAnother.lob	1 day	1

4. Click a **job name** or **job ID** in the table to view the job monitoring details.

Completed	etereteriteta (m. 1.78 day	0		3
<b>307.67</b> к	<b>307.67</b> к	0		0
<b>0</b>	2	2		1
ElekkhetisilinnoslopMacEunction 1.226 GB Beik, Diseatinglink 1.226 GB Bencer, Randembourdefunction 0 8	41581872 144315359 0	357 MB 05 1,230 GB	63167823 0 144312011	

5. Click **Flink Cluster** in the top-right corner to view the Flink cluster monitoring information.



6. Click a task name in the table to view the task monitoring details.

1/2018/05/0     1/2018/05/0     1/2/11/05/0     1/2/11/05/0     1/2/11/05/0	Outsdource Prometheus - emrs, - PrometheusAnotherJob - Fil	inkMetricsExposingMapFunction - All -	All - TaskManager All -	
00       FieldAnnicsEgnestingMagFunction       7219942       421004       1         1       FieldAnnicsEgnestingMagFunction       7215977       4191972       1         1       FieldAnnicsEgnestingMagFunction       FieldAnnicsEgnestingMagFunction <td< td=""><td>~ Operator</td><td></td><td></td><td></td></td<>	~ Operator			
00       PiekkenickSpestelykkylunction       7219942       42004       1         1       FiekkenickSpestelykkylunction       7215977       4191927       1         4       FiekkenickSpestelykkylunction       7215977       4191927       1         5       FiekkenickSpestelykkylunction       7215977       4191927       1         6       FieldenickSpestelykkylunction       7215977       4191927       1         6       FieldenickSpestelykkylunction       7015977       4191927       4191927       4191927         6       FieldenickSpestelykkylunction       FieldenickSpestelykkylunction       7015977       4191927       4191927       4191927         6       FieldenickSpestelykkylunction       FieldenickSpestelykkylunction       10191927       7019777       4191927       4191927<				
1     Pailaberichpossightigfunction     7215977     1181822     1       1     Pailaberichpossightigfunction     1     1     1     1       1     Pailaberichpossightigfunction     1     1     1     1       1     Pailaberichpossightigfunction     1     1     1     1     1       1     Pailaberichpossightigfunction     1     1     1     1     1       1     Pailaberichpossightigfunction     Pailaberichpossightigfunction     1     1     1     1       1     Pailaberichpossightigfunction     Pailaberichpossightigfunction     Pailaberichpossightigfunction     1     1     1     1       1     Pailaberichpossightigfunction     Pailaberichpossightigfunction     Pailaberichpossightigfunction     1     1     1     1       1     Pailaberichpossightigfunction     Pailaberichpossightigfunction     1     1     1     1     1       1     Pailaberichpossightigfunction     1				17
Bit scends//         Git scends//         Git scends//         Git scends//           6.8 scends//         No data         443 scends//         443 scends// <td< td=""><td></td><td></td><td></td><td></td></td<>				
1 Stausdah       OC Stausdah	1 FinkMetricsExposingMapFunction	72155977	41581872	1
1 Structuly       OC Structuly				
1 Structuly       OC Structuly				
1 Sueuduh       602 suuduh       602 suuduh       602 suuduh       603 suuduh <td></td> <td></td> <td></td> <td></td>				
1 Structuly       OC Structuly				
1 Stausdah       OC Stausdah				
6.1 munchish       No data       46.3 munchish				
No. data         Alt 3 month/h	1.0 vecords/s	470.0 records/s	المعالمة الم	470.0 records/s
No dra         469 munduh         469 munduh<	6.5 excerta/a	469.5 mcondu/s	AN AN ILLAND A MARK	AND S monthly
fewerich     483 month/n     483 month/n     483 month/n     483 month/n       -1.8 month/n     -1.8 month/n     483 month/n     483 month/n     483 month/n       -1.8 month/n     -1.9 month/n     -1.9 month/n     -1.9 month/n     483 month/n       -1.8 month/n     -1.9 month/n     -1.9 month/n     -1.9 month/n     483 month/n       -1.8 month/n     -1.9 month/n     -1.9 month/n     -1.9 month/n     483 month/n       -1.8 month/n     -1.9 month/n     -1.9 month/n     -1.9 month/n     483 month/n       -1.8 month/n     -1.9 month/n     -1.9 month/n     -1.9 month/n     -1.9 month/n       -1.8 month/n     -1.9 month/n     -1.9 month/n     -1.9 month/n     -1.9 month/n       -1.8 month/n     -1.9 month/n     -1.9 month/n     -1.9 month/n     -1.9 month/n				
-13 months' -13 months' -13 months' -13 months' - 13 m		469.0 recordu/s	A THE REAL PROPERTY OF THE REA	449.0 recorde/s
-19 monds/t         1200 0000         1201 0000		468.5 recordu/s	· · · · · · · · · · · · · · · · · · ·	448.5 records/s
1.8 models         1.2 mod	-0.3 HORDAY			
- TaskManager 25 115 ma 115 ma 16 m	-1.0 records/s	12/08 00:00 12/09 00:00 12/10 00:00 12/		12/04 00:00 12/09 00:00 12/10 00:00 12/11 00:00 12/12 00:00 12/13 00:00 12/14 00
		- 0-FinkMetricsDoosingMapFunction - 1-FinkMetricsDoosingMa	apfunction .	<ul> <li>OFInidMetricsDeposingMapFunction — 1+FinidMetricsDeposingMapFunction</li> </ul>
25 16m	~ TaskManager			
12.5 mg				
				1.0 ms
Q5	2.0 12.5 ma			65m
14 Mönu				
14				• • • • • • • • • • • • • • • • • • •
10m 41 41	5.0ma	(1. 27) mmm	45	-05 ma
	0.5 25ms			
9 1200 1200 1201 1011 1011 1011 1011 101	0 12/08 12/09 12/10 12/11 12/12 12/13 12/14 0ms	99 12/10 12/11 12/12 12/13 12/14	-1.0	12/13 12/14 -1.0 ms 12/09 12/19 12/11 12/12 12/13 1

### Integrating with alert feature

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click **Alerting Rule** and add the corresponding alerting rules. For more information, please see Creating Alerting Rule.

# Java Application Integration Spring Boot Integration

Last updated : 2024-08-07 22:00:28

## Overview

When using Spring Boot as the development framework, you need to monitor the status of applications such as JVM and Spring MVC. TMP collects data such as JVM data based on the Spring Boot Actuator mechanism. With the Grafana dashboard that comes with TMP, you can conveniently monitor the status of Spring Boot applications. This document uses deploying a Spring Boot application in TKE as an example to describe how to use TMP to monitor the application status.

## Prerequisites

Create a TKE cluster. Use a private image repository to manage application images. The image is developed based on the Spring Boot framework.

## Directions

#### Note:

Spring Boot provides the Actuator component to monitor applications, which reduces the development costs. Therefore, Actuator is directly used in this document to track Spring Boot metrics. You should use Spring Boot v2.0 or above in the following steps, as lower versions may have different configurations.

# If you use Spring Boot v1.5 for integration, the integration process will differ from that for v2.0, and you should note the following:

- 1. The address for accessing prometheus metrics is different from that for v2.0. On v1.5, the default address is /prometheus , i.e., http://localhost:8080/prometheus .
- 2. If error 401 is reported, it indicates no permissions (Whitelabel Error Page). On v1.5, security control is enabled for the management API by default, so you need to set management.security.enabled=false .
- 3. If bootstrap.yml is used to configure parameters in the project, modifying management in it will not work, which should be modified in application.yml due to the Spring Boot start and load sequence.

4. You cannot add metric common tag through YML; instead, you can add it only by adding a bean to the code.

### Modifying application dependencies and configuration

#### Step 1. Modify POM dependencies

If spring-boot-starter-web is already imported in this project, add the actuator/prometheus Maven dependency to the pom.xml file.



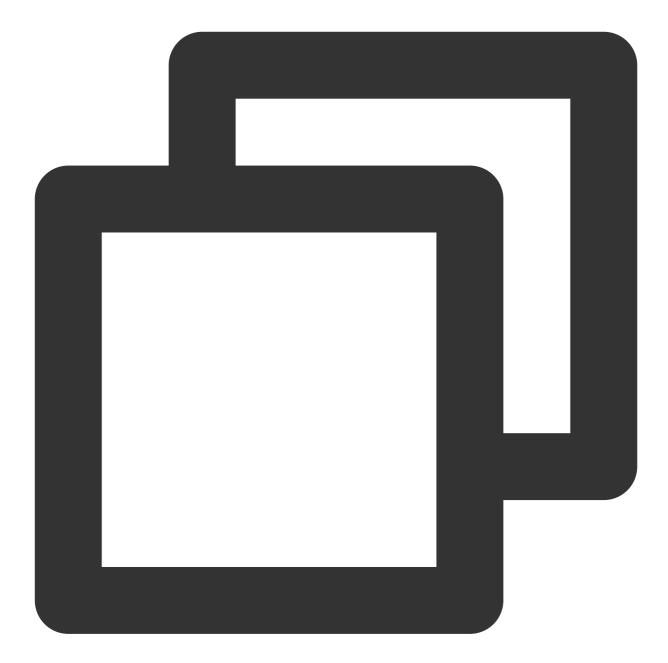
<dependency>

```
<groupId>org.springframework.boot</groupId>
        <artifactId>spring-boot-starter-actuator</artifactId>
        </dependency>
        <groupId>io.micrometer</groupId>
        <artifactId>micrometer-registry-prometheus</artifactId>
        </dependency>
        </depndency>
```

#### Step 2. Modify the configuration

Edit the application.yml file in the resources directory and modify the actuator configuration to expose the metric data in the Prometheus protocol.





```
management:
endpoints:
web:
exposure:
include: prometheus # Web access path for opening Prometheus
metrics:
# We recommend you enable the following options to monitor P99 and P95 latencie
distribution:
sla:
http:
server:
```

```
requests: 1ms,5ms,10ms,50ms,100ms,200ms,500ms,1s,5s
# Add special labels to Prometheus
tags:
    # You must add the corresponding application name, as the corresponding monit
    application: spring-boot-mvc-demo
```

### Step 3. Perform local verification

In the current directory of the project, run mvn spring-boot:run . If you can access the metric data of the
Prometheus protocol through http://localhost:8080/actuator/prometheus , the relevant dependency
configuration is correct.

#### Note:

The default configurations of the port and path are used in the same, which should be replaced with those in your actual project.

#### **Releasing application to TKE**

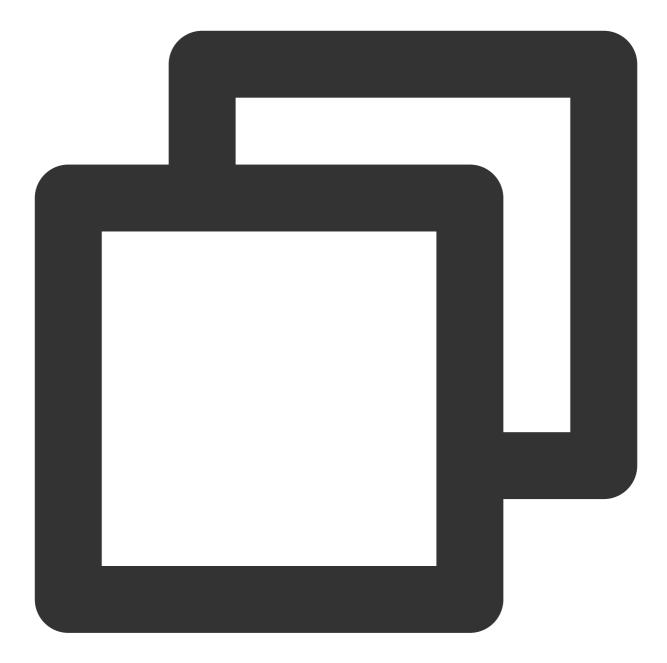
#### Step 1. Configure a Docker image environment locally

If you have already configured a Docker image environment locally, proceed to the next step; otherwise, configure one as instructed in Getting Started.

#### Step 2. Package and upload the image

1. Add Dockerfile in the root directory of the project. You can add it by referring to the following sample code and modify Dockerfile based on your actual project:





```
FROM openjdk:8-jdk
WORKDIR /spring-boot-demo
ADD target/spring-boot-demo-*.jar /spring-boot-demo/spring-boot-demo.jar
CMD ["java","-jar","spring-boot-demo.jar"]
```

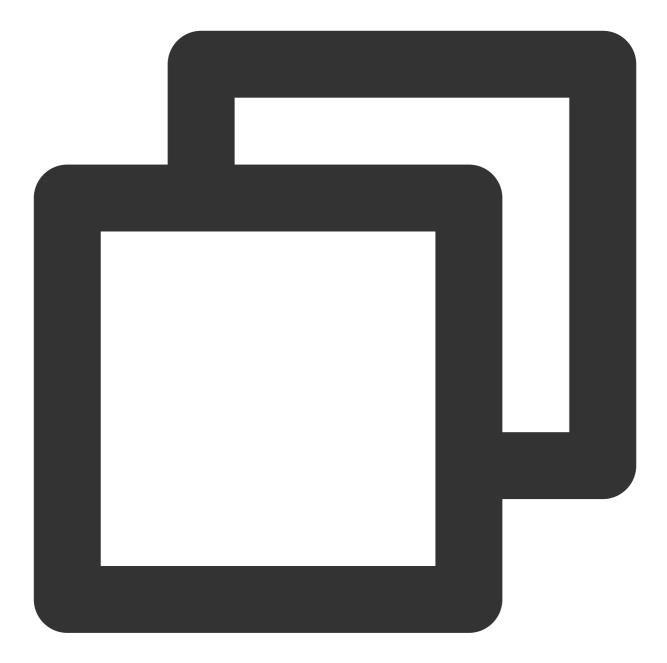
2. Package the image by running the following command in the project root directory. You need to replace namespace, ImageName, and image tag as needed in your actual project.





```
mvn clean package
docker build . -t ccr.ccs.tencentyun.com/[namespace]/[ImageName]:[image tag]
docker push ccr.ccs.tencentyun.com/[namespace]/[ImageName]:[image tag]
```

For example:



mvn clean package
docker build . -t ccr.ccs.tencentyun.com/prom\_spring\_demo/spring-boot-demo:latest
docker push ccr.ccs.tencentyun.com/prom\_spring\_demo/spring-boot-demo:latest

#### Step 3. Deploy the application

1. Log in to the TKE console and select the container cluster for deployment.

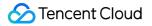
2. Click **Workload** > **Deployment** to enter the Deployment management page and select the corresponding namespace to deploy the service. Here, a workload is created in the console, and Service access is also enabled. You can also create one on the command line.

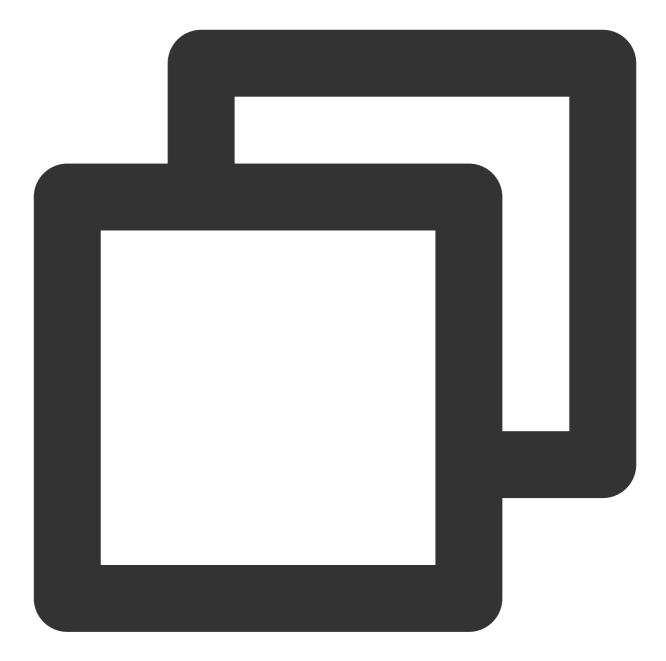
Workload name	spring-mvc-demo								
	The maximum length of 40 ch	aracters, can only contain lowercase letters, numbers and separators ("-"), and must start with a lowercase letter, and end with a number or a lowercase letter							
describe	Please enter the description information, no more than 1000 characters								
Label	k8s-app = s	pring-mvc-demo X							
	New variable								
	Can only contain letters, numb	sers and separators ("-", ", ", ", ",", "/"), and must start and end with letters and numbers							
Namespaces	default	v							
type	Deployment (Scalable D     DaemonSet (Run Pod o     StatefulSet (operating P)     CronJob (run regularly a     Job (single task)	n each host) od with stateful set)							
Data volume (optional)	Add data volume								
	Provide storage for the container. Currently, it supports temporary paths, host paths, cloud hard disk data volumes, file storage NFS, configuration files, and PVCs. It also needs to be more of the container. Guidelines for use 🖄								
Instance content container		✓ ×							
	name	spring-mvc-demo							
		Up to 63 characters, can only contain lowercase letters, numbers and separators ("-"), and cannot start or end with separators							
	Mirror image	ccr.ccrd.tencent.com/ Select mirror							
	Mirror version (Tag)	If not filled, the default is latest							
	Image pull strategy	Always IfNotPresent Never							

Service       Enable         Service access method       Access only within the cluster       Host port access       Public network LB access       Intranet LB accesshow to         That is, the ClusteriP type will provide an entry that can be accessed by other services or containers in the cluster. It support cluster to ensure service network isolation.
That is, the ClusterIP type will provide an entry that can be accessed by other services or containers in the cluster. It support
Headless Service ⑦ (Headless Service only supports selection during creation, and does not support changing the acc
Port Mapping protocol () Container port () Service port ()
TCP  The port that the application in the It is recommended to be consister
Add port mapping
show advanced settings

3. Add K8s labels to the corresponding Service. If the workload is created on the command line, you can directly add labels. Here, the configuration is adjusted in the TKE console. Select the TKE cluster that needs to be adjusted. Click Services and Routes > Service to enter the Service management page. Select the corresponding namespace to adjust the Service YAML configuration as shown below:

			Namespace	default 🔻
Туре 🗡	Selector	IP Address(j)	Time Created	Operation
	-			Update access m
		- 1/6	100	Update access m
				Type T     Selector     IP Address()     Time Created       III     IIII     IIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIII





```
apiVersion: v1
kind: Service
metadata:
   labels: # Add the corresponding labels based on the actual conditions
   k8sapp: spring-mvc-demo
   name: spring-mvc-demo
   namespace: spring-demo
spec:
   ports:
    - name: 8080-8080-tcp # Corresponding `port` value in the ServiceMonitor scrape
   port: 8080
```

```
protocol: TCP
targetPort: 8080
selector:
   k8s-app: spring-mvc-demo
   qcloud-app: spring-mvc-demo
sessionAffinity: None
type: ClusterIP
```

#### Step 4. Add a scrape task

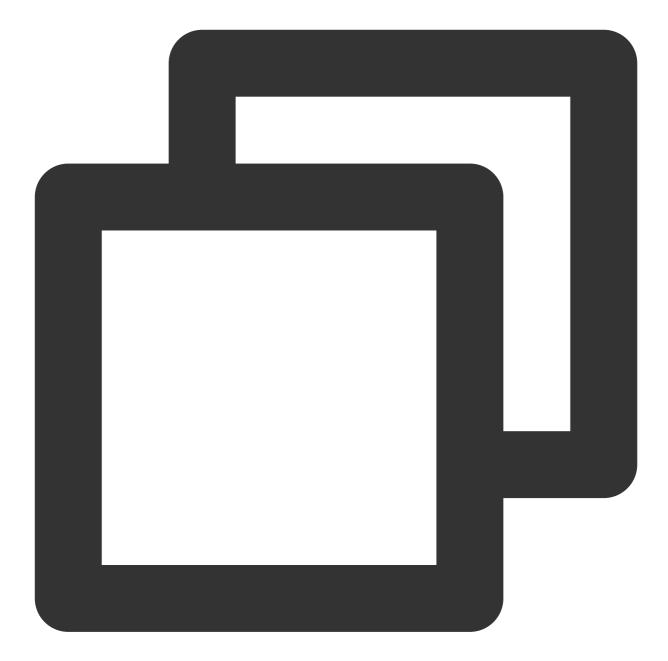
1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click a cluster ID in the TKE cluster list to enter the Integrate with TKE page.

3. In **Scrape Configuration**, add a ServiceMonitor. Currently, TMP supports discovering the corresponding target instance address through labels; therefore, you can add some specific K8s labels to some services, which will be automatically identified by TMP after configuration, eliminating your need to add scrape tasks for all services one by one. The configuration information for the above sample is as follows:

#### Note:

Here, note that the port value is the spec/ports/name value in the Service YAML configuration file.



```
apiVersion: monitoring.coreos.com/v1
kind: ServiceMonitor
metadata:
   name: spring-mvc-demo # Enter a unique name
   namespace: cm-prometheus # The namespace is fixed. Do not change it
spec:
   endpoints:
    - interval: 30s
    port: 8080-8080-tcp # Enter the name of the corresponding port of the Prometh
    path: /actuator/prometheus # Enter the value of the corresponding path of th
    namespaceSelector: # Select the namespace where the Service to be monitored re
```



```
matchNames:
    - spring-demo
selector: # Enter the label value of the Service to be monitored to locate the
    matchLabels:
        k8sapp: spring-mvc-demo
```

### Step 5. View the monitoring information

Access the Grafana address of your TMP instance to view the application monitoring dashboard in Dashboards >

#### Manage > Application.

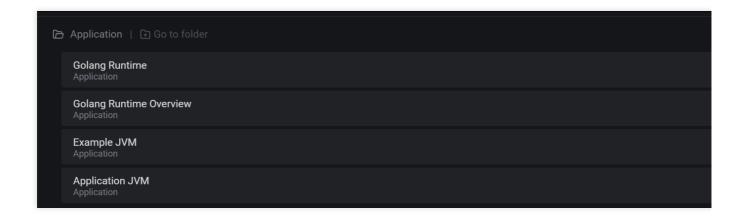
Spring MVC application: monitoring data of MVC status, such as the request latency, number of requests, success rate, and exception distribution.

Spring MVC API: API-level monitoring data, which supports multiple APIs to help you locate faulty APIs.

Tomcat: monitoring dashboard of internal Tomcat status, such as thread usage.

Application JVM: monitoring data of the status of all instances under an application. If you find a faulty instance, you can view its monitoring information at any time.

Instance JVM: detailed monitoring data of a single instance JVM.



# **JVM** Integration

Last updated : 2024-08-07 22:00:35

# Overview

When using the Java programming language, you need to monitor JVM performance. TMP collects the JVM monitoring data exposed by applications and provides an out-of-the-box Grafana dashboard for it. This document uses deploying a Java application in TKE as an example to describe how to use TMP to monitor the application status.

### Note:

If you have already used Spring Boot as the development framework, please see Spring Boot Integration.

# Prerequisites

Create a TKE cluster. Use a private image repository to manage application images.

# Directions

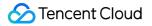
#### Note:

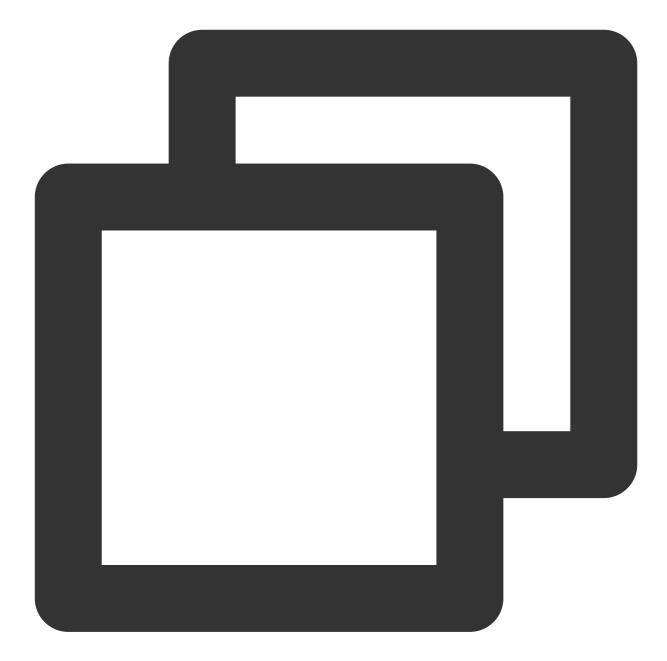
As a major programming language, Java has a comprehensive ecosystem, where Micrometer has been widely used as a metric timestamping SDK. This document uses Micrometer as an example to describe how to monitor JVM.

### Modifying application dependencies and configuration

#### Step 1. Modify POM dependencies

Add Maven dependencies to the pom.xml file and adjust the version as needed as follows:



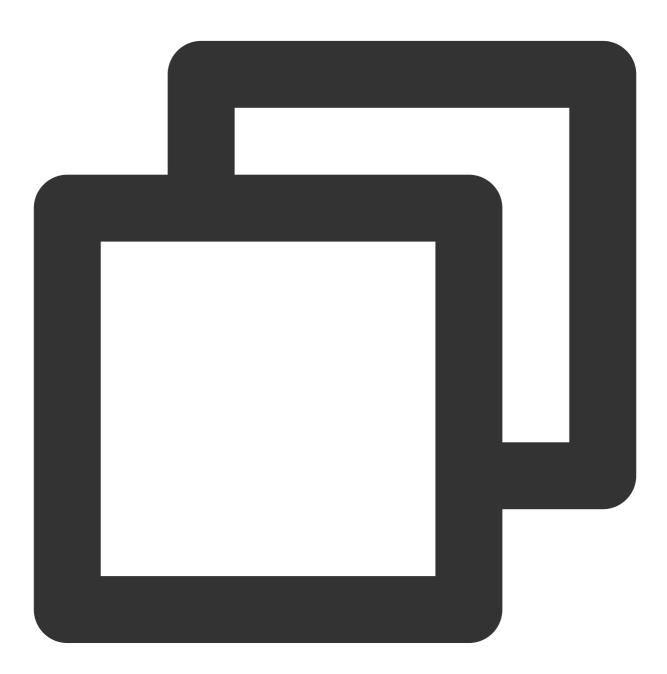


```
<dependency>
<groupId>io.prometheus</groupId>
<artifactId>simpleclient</artifactId>
<version>0.9.0</version>
</dependency>
<dependency>
<groupId>io.micrometer</groupId>
<artifactId>micrometer-registry-prometheus</artifactId>
<version>1.1.7</version>
</dependency>
```



#### Step 2. Modify the code

When the project is started, add the corresponding monitoring configuration. In addition, Micrometer also provides the collection of some common metrics, which are in the io.micrometer.core.instrument.binder package and can be added as needed as follows:



public class Application {
 // It can be used in custom monitoring as a global variable
 public static final PrometheusMeterRegistry registry = new PrometheusMeterRegis
 static {

```
// Add a global Prometheus label. We recommend you add the corresponding ap
    registry.config().commonTags("application", "java-demo");
}
public static void main(String[] args) throws Exception {
    // Add JVM monitoring
    new ClassLoaderMetrics().bindTo(registry);
    new JvmMemoryMetrics().bindTo(registry);
    new JvmGcMetrics().bindTo(registry);
    new ProcessorMetrics().bindTo(registry);
    new JvmThreadMetrics().bindTo(registry);
    new UptimeMetrics().bindTo(registry);
    new FileDescriptorMetrics().bindTo(registry);
    System.gc(); // Test GC
    try {
        // Expose the Prometheus HTTP service. If it already exists, you can us
        HttpServer server = HttpServer.create(new InetSocketAddress(8080), 0);
        server.createContext("/metrics", httpExchange -> {
            String response = registry.scrape();
            httpExchange.sendResponseHeaders(200, response.getBytes().length);
            try (OutputStream os = httpExchange.getResponseBody()) {
                os.write(response.getBytes());
            }
        });
        new Thread(server::start).start();
    } catch (IOException e) {
        throw new RuntimeException(e);
    }
}
```

#### Note:

}

As monitoring of JVM GC pauses is implemented through the GarbageCollector Notification mechanism, the monitoring data will be generated only after a GC occurs. The above sample actively calls <code>System.gc()</code> to make the test more straightforward.

#### Step 3. Perform local verification

After the application is started locally, you can access the metric data of the Prometheus protocol through http://localhost:8080/metrics .

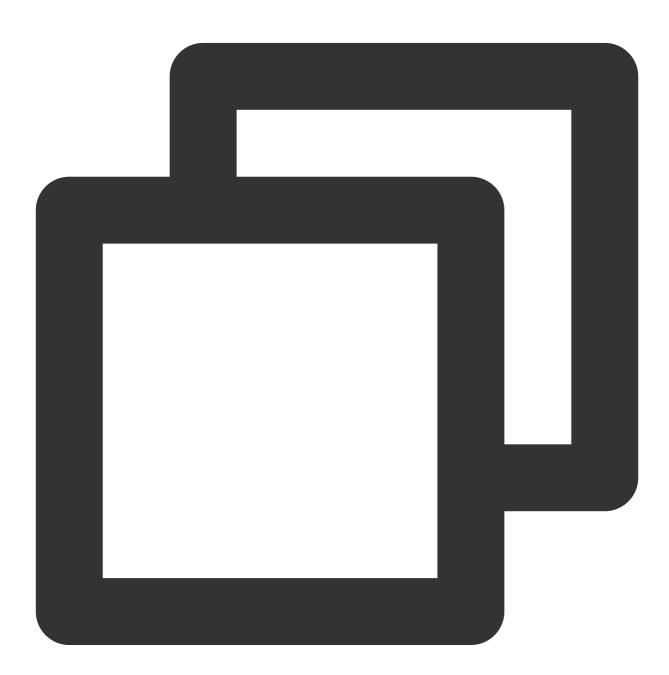
### **Releasing application to TKE**

### Step 1. Configure a Docker image environment locally

If you have already configured a Docker image environment locally, proceed to the next step; otherwise, configure one as instructed in Getting Started.

#### Step 2. Package and upload the image

1. Add Dockerfile in the root directory of the project. Please modify it based on your actual project conditions as follows:

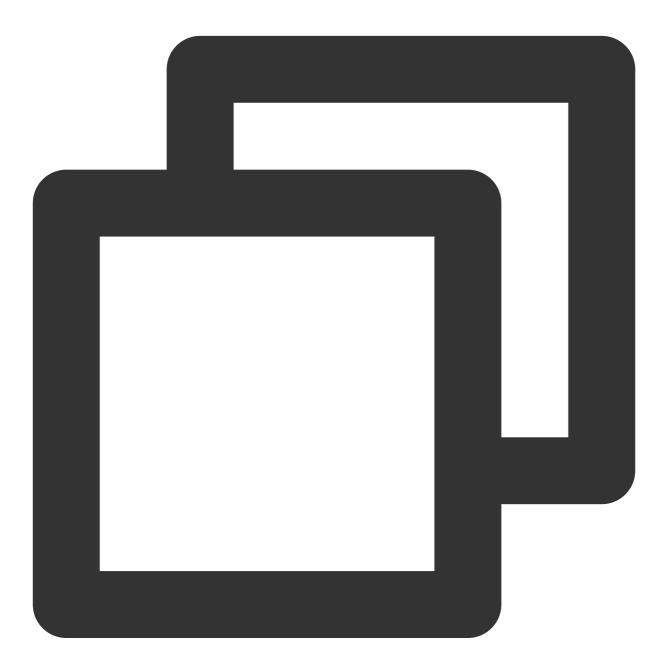


FROM openjdk:8-jdk WORKDIR /java-demo ADD target/java-demo-\*.jar /java-demo/java-demo.jar



```
CMD ["java","-jar","java-demo.jar"]
```

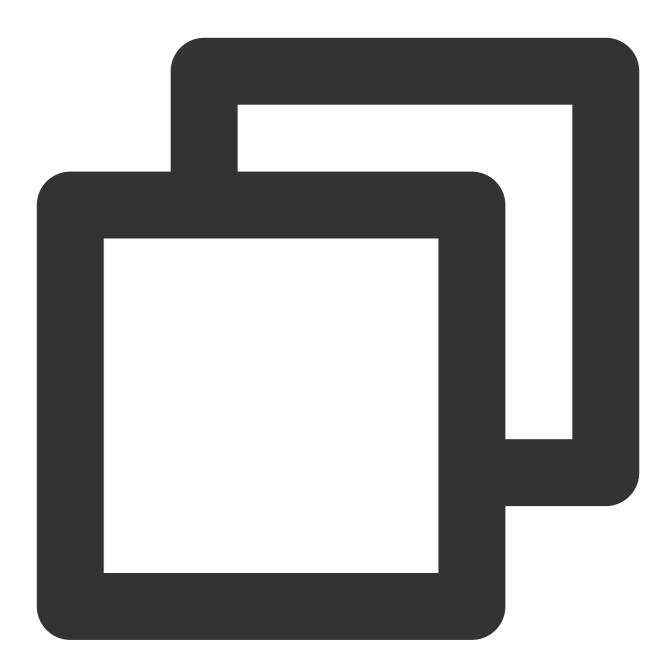
2. Package the image by running the following command in the project root directory. You need to replace namespace, ImageName, and image tag as needed.



mvn clean package
docker build . -t ccr.ccs.tencentyun.com/[namespace]/[ImageName]:[image tag]
docker push ccr.ccs.tencentyun.com/[namespace]/[ImageName]:[image tag]



<b>Below is a sample:</b>



mvn clean package
docker build . -t ccr.ccs.tencentyun.com/prom\_spring\_demo/java-demo:latest
docker push ccr.ccs.tencentyun.com/prom\_spring\_demo/-demo:latest

### Step 3. Deploy the application

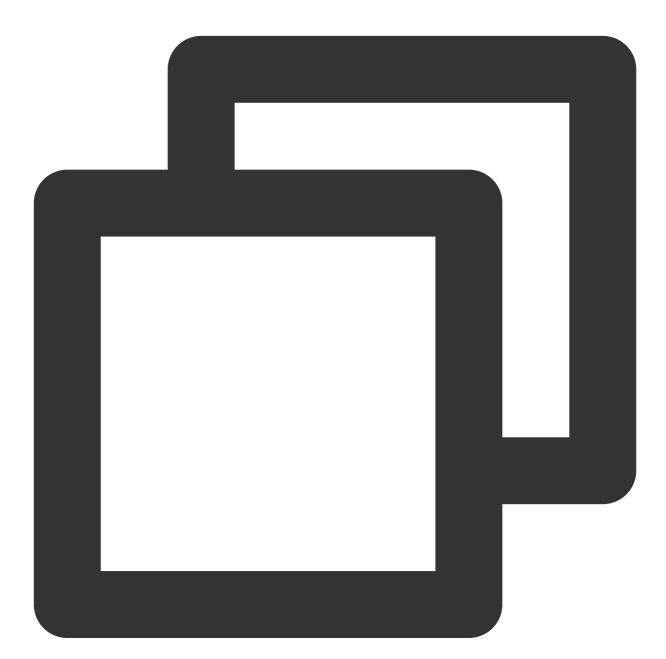
1. Log in to the TKE console and select the container cluster for deployment.

2. Select *Workload*\* > **Deployment** to enter the Deployment management page and select the corresponding

namespace to deploy the service. Use the following YAML configuration to create the corresponding Deployment:

#### Note:

If you want to create in the console, please see Spring Boot Integration.



```
apiVersion: apps/v1
kind: Deployment
metadata:
labels:
k8s-app: java-demo
name: java-demo
```

```
namespace: spring-demo
spec:
   replicas: 1
    selector:
     matchLabels:
       k8s-app: java-demo
    template:
      metadata:
        labels:
          k8s-app: java-demo
    spec:
      containers:
      - image: ccr.ccs.tencentyun.com/prom_spring_demo/java-demo
        imagePullPolicy: Always
        name: java-demo
        ports:
        - containerPort: 8080
          name: metric-port
        terminationMessagePath: /dev/termination-log
        terminationMessagePolicy: File
      dnsPolicy: ClusterFirst
      imagePullSecrets:
      - name: qcloudregistrykey
      restartPolicy: Always
      schedulerName: default-scheduler
      terminationGracePeriodSeconds: 30
```

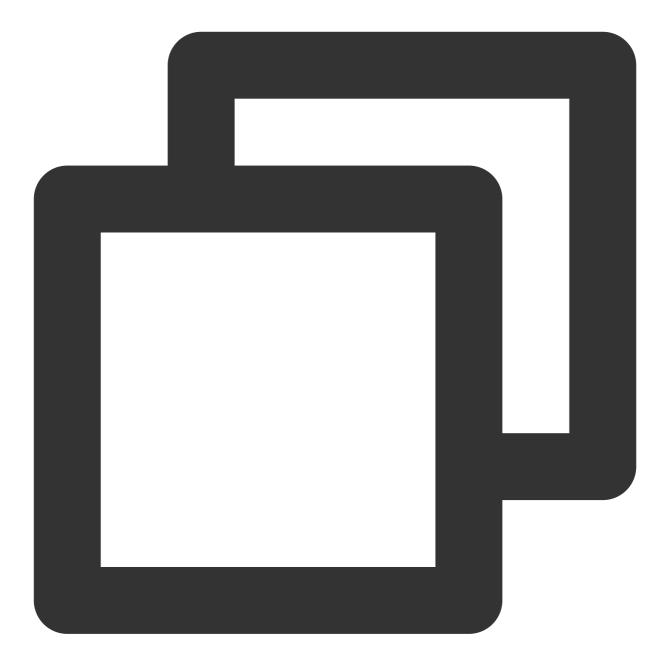
#### Step 4. Add a scrape task

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click a cluster ID in the TKE cluster list to enter the Integrate with TKE page.

3. In **Scrape Configuration**, add Pod Monitor to define a Prometheus scrape task. Below is a sample YAML configuration:





```
apiVersion: monitoring.coreos.com/v1
kind: PodMonitor
metadata:
   name: java-demo
   namespace: cm-prometheus
spec:
   namespaceSelector:
    matchNames:
        - java-demo
   podMetricsEndpoints:
        - interval: 30s
```

```
path: /metrics
port: metric-port
selector:
matchLabels:
    k8s-app: java-demo
```

#### Step 5. View the monitoring information

1. In **Integration Center** in the target TMP instance, find JVM monitoring, install the corresponding Grafana dashboard, and then you can enable the JVM monitoring dashboard.

2. Access the Grafana address of your TMP instance to view the application monitoring dashboard in **Dashboards** > **Manage** > **Application**.

**Application JVM**: monitoring data of the status of all instances under an application. If you find a faulty instance, you can view its monitoring information at any time.

**Instance JVM**: detailed monitoring data of a single instance JVM.

Datasource default ~ App java-demo ~							
	应用所在实例 JVM 监控						
实例	Uptime	CPU 最大使用率% ~	GC 总次数	GC 总耗时	Heap 使用		
15 00 1 42 Adag	13.04 hour	0.02%	0	0 s	1.03%		
<u>59.1 m wa</u> a	26.89 s	0.00%			0.20%		





# Go Application Integration

Last updated : 2024-08-07 22:01:00

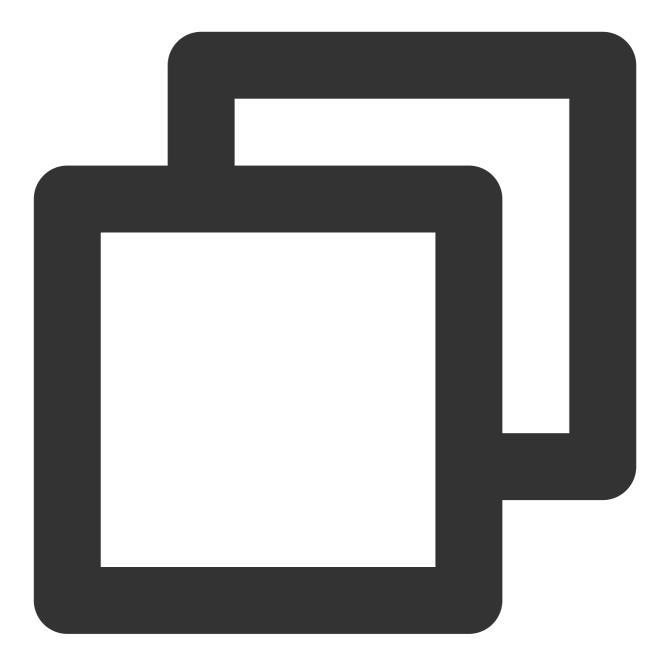
Prometheus provides an official Go library to collect and expose the monitoring data. This document describes how to use it to expose the Go runtime data and use TMP to collect metrics and display data with some basic samples. **Note:** 

For Go client API documentation, please see Prometheus Go client library.

# Installation

You can run the following go get commands to install the relevant dependencies:





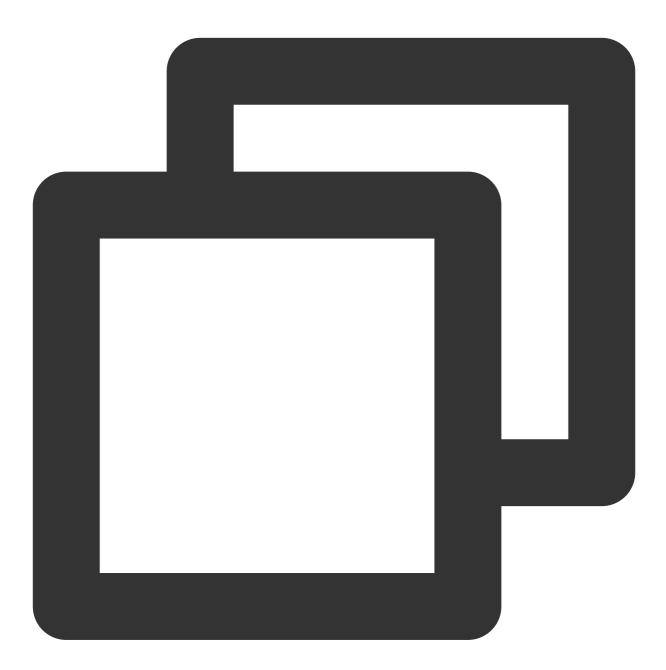
go get github.com/prometheus/client\_golang/prometheus
go get github.com/prometheus/client\_golang/prometheus/promauto

go get github.com/prometheus/client\_golang/prometheus/promhttp

# Start (Runtime Metrics)

1. Prepare an HTTP service with the commonly used path /metrics . You can directly use the Handler function provided in prometheus/promhttp .

The following is a sample Go application, which exposes some default metrics (including runtime, process, and build metrics) through http://localhost:2112/metrics :



```
package main
import (
    "net/http"
    "github.com/prometheus/client_golang/prometheus/promhttp"
)
```

```
func main() {
    http.Handle("/metrics", promhttp.Handler())
    http.ListenAndServe(":2112", nil)
}
```

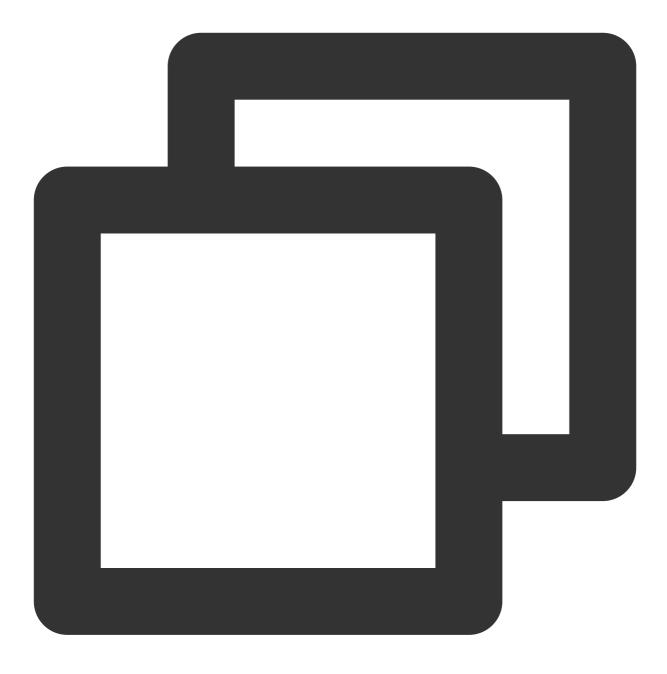
2. Run the following command to start the application:



go run main.go

3. Run the following command to access the basic built-in metric data:





curl http://localhost:2112/metrics

# **Application Layer Metrics**

1. The above sample only exposes some basic built-in metrics. For metrics at the application layer, you need to add them additionally (we will provide some SDKs in the future for easier integration). The following sample exposes a



Counter metric named myapp\_processed\_ops\_total to count the currently completed operations. The operation is performed once every 2 seconds, and the count increases by 1 each time:

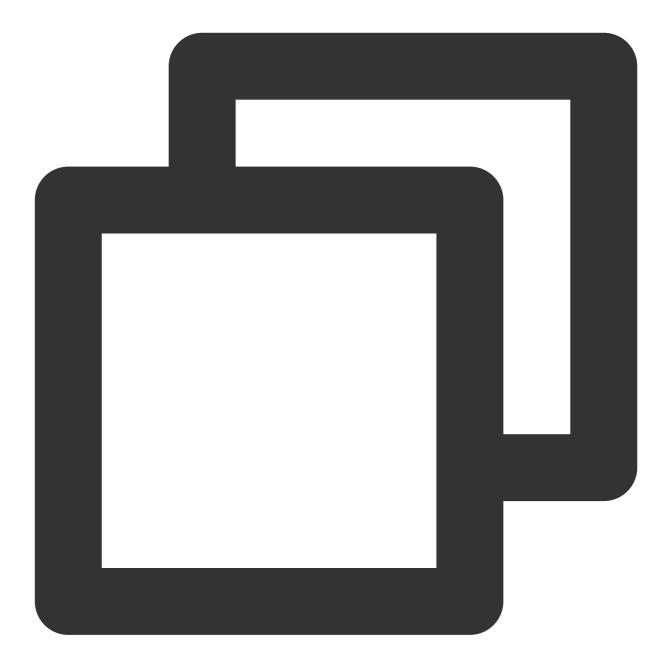


```
package main
import (
    "net/http"
    "time"
    "github.com/prometheus/client_golang/prometheus"
    "github.com/prometheus/client_golang/prometheus/promauto"
```

```
"github.com/prometheus/client_golang/prometheus/promhttp"
)
func recordMetrics() {
        go func() {
                for {
                        opsProcessed.Inc()
                         time.Sleep(2 * time.Second)
                }
        }()
}
var (
        opsProcessed = promauto.NewCounter(prometheus.CounterOpts{
                Name: "myapp_processed_ops_total",
                Help: "The total number of processed events",
        })
)
func main() {
        recordMetrics()
        http.Handle("/metrics", promhttp.Handler())
        http.ListenAndServe(":2112", nil)
}
```

2. Run the following command to start the application:

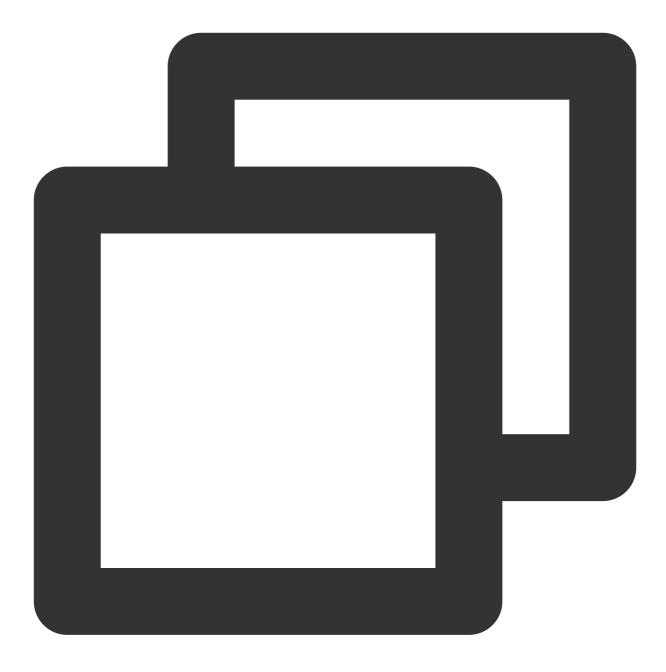




go run main.go

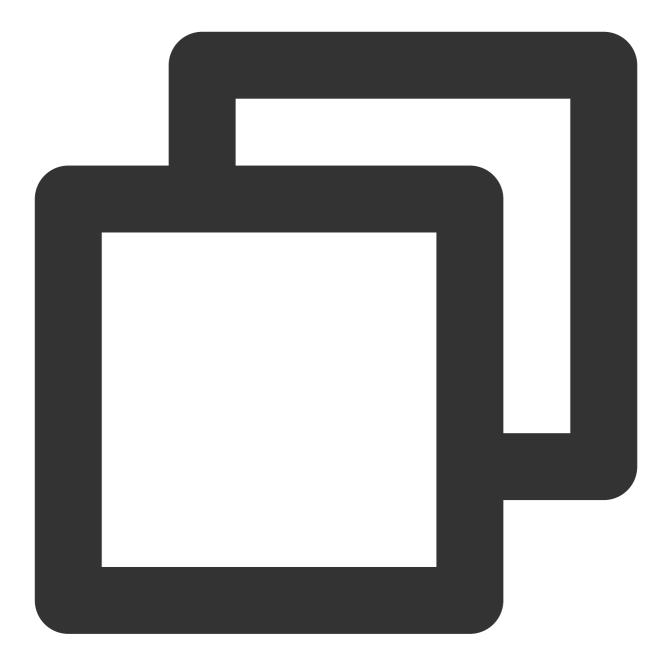
3. Run the following command to access the exposed metrics:





curl http://localhost:2112/metrics

From the output result, you can see the information related to the myapp\_processed\_ops\_total counter, including the help documentation, type information, metric name, and current value, as shown below:



# HELP myapp\_processed\_ops\_total The total number of processed events # TYPE myapp\_processed\_ops\_total counter myapp\_processed\_ops\_total 666

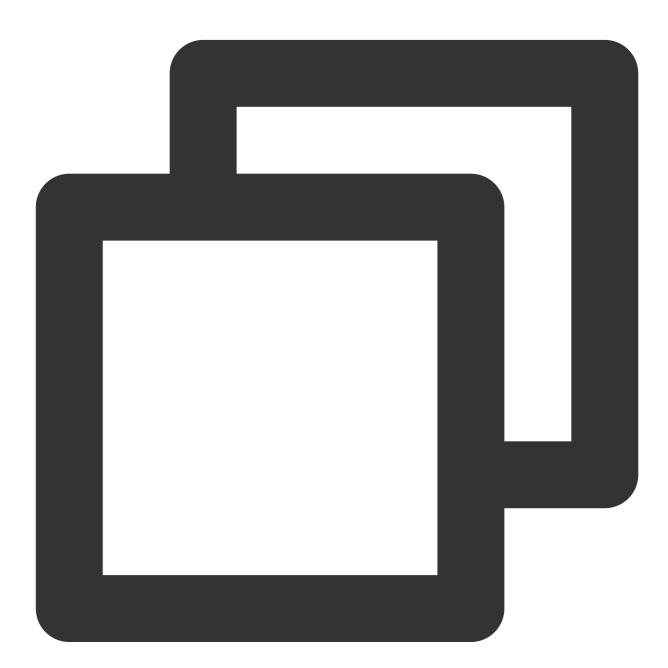
# Using TMP

Two samples are used above to show how to use the Prometheus Go library to expose application metric data. However, because the exposed data is in text format, you'll need to set up and maintain an additional Prometheus service to collect metrics, which may require additional Grafana dashboards for visual display.

In contrast, if you use TMP, you can directly skip the above steps and achieve the same purpose with just a few clicks. For more information, please see Getting Started.

## Packaging and deploying application

1. A Go application generally can use a Dockerfile in the following format (it should be modified as needed):



```
FROM golang:alpine AS builder
RUN apk add --no-cache ca-certificates \\
    make \\
```

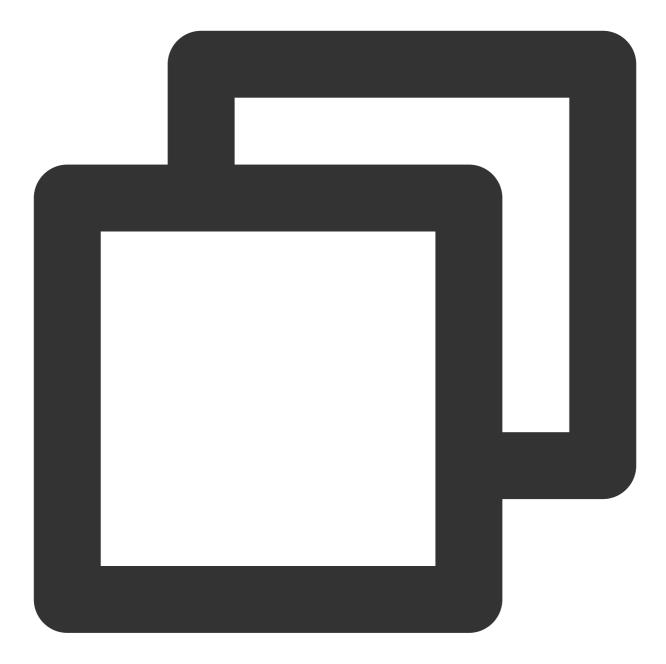
git

```
COPY . /go-build
RUN cd /go-build && \\
    export GO111MODULE=on && \\
    export GOPROXY=https://goproxy.io && \\
    go build -o 'golang-exe' path/to/main/
FROM alpine
RUN apk add --no-cache tzdata
COPY --from=builder /etc/ssl/certs/ca-certificates.crt /etc/ssl/certs
COPY --from=builder /go-build/golang-exe /usr/bin/golang-exe
ENV TZ Asia/Shanghai
CMD ["golang-exe"]
```

2. You can use an image from Tencent Cloud Image Registry or another public or self-built image registry.

3. You need to define a Kubernetes resource based on your application type. Here, a Deployment is used as shown below:



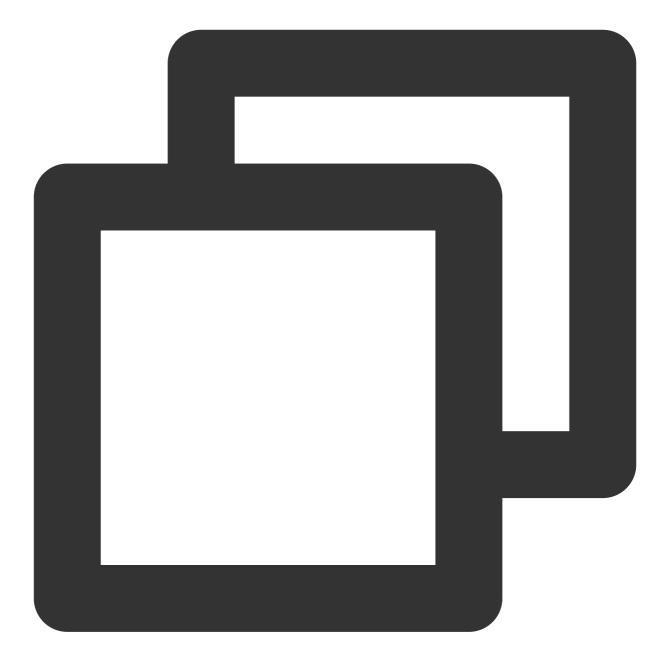


```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: golang-app-demo
  labels:
  app: golang-app-demo
spec:
  replicas: 3
  selector:
  matchLabels:
    app: golang-app-demo
```

```
template:
metadata:
   labels:
      app: golang-app-demo
spec:
   containers:
   - name: golang-exe-demo:v1
   image: nginx:1.14.2
   ports:
   - containerPort: 80
```

4. You also need a Kubernetes Service for scrape configuration and load balancing.





```
apiVersion: v1
kind: Service
metadata:
  name: golang-app-demo
spec:
  selector:
  app: golang-app-demo
  ports:
  - protocol: TCP
    port: 80
    targetPort: 80
```

#### Note:

You must add a label to identify the current application. The label name doesn't necessarily need to be app, but there must be a label with the similar meaning. You can add other extended labels by relabeling when adding a data collection task subsequently.

5. You can use the TKE console or directly use kubectl to submit the resource definitions to Kubernetes and wait for successful creation.

### Adding data collection task

After the service runs, you need to configure TMP to discover and collect the monitoring metrics in the following steps:

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click a **cluster ID** in the TKE cluster list to enter the **Integrate with TKE** page.

3. In **Scrape Configuration**, add a ServiceMonitor. Currently, TMP supports discovering the corresponding target instance address through labels; therefore, you can add some specific K8s labels to some services, which will be automatically identified by TMP after configuration, eliminating your need to add scrape tasks for all services one by one. The configuration information for the above sample is as follows:

#### Note:

The port value is the spec/ports/name value in the Service YAML configuration file.





```
path: /metrics
relabelings:
# ** There must be a label named `application`. Here, suppose that K8s has a
# Use the `replace` action of `relabel` to replace it with `application`
- action: replace
    sourceLabels: [__meta_kubernetes_pod_label_app]
    targetLabel: application
# Select the namespace where the Service to be monitored resides
namespaceSelector:
    matchNames:
    - golang-demo
    # Enter the label value of the Service to be monitored to locate the target S
selector:
    matchLabels:
    app: golang-app-demo
```

#### Note:

You must configure the label named application in the sample; otherwise, you cannot use some other out-ofthe-box integration features of TMP. For more advanced usage, please see <u>ServiceMonitor</u> or <u>PodMonitor</u>.

#### Viewing monitoring information

1. In the TMP instance list, find the corresponding TMP instance, click

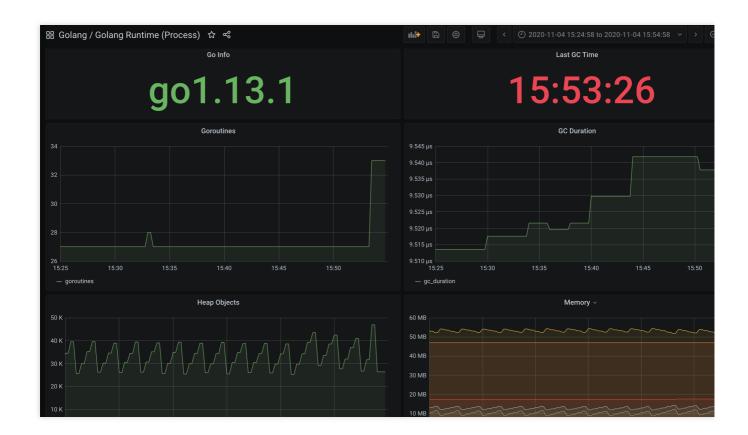
on the right of the instance ID to open your Grafana page, and enter your account and password to access the Grafana visual dashboard operation section.

2. Enter Grafana, click the



icon to expand the monitoring dashboard, and click the name of the corresponding monitoring chart to view the monitoring data.

器  Golang / Golang Runtime Ov	erview ☆ ペ				•••	@ 🖵 🕘 Last 30	minutes   ~
Datasource default ~ Cluster cl	ls-6 Application	1	×				
		CPU Usage	Memory(RSS) ~	Threads		GC Duration	
10902	5.20 day	0.00	62.07 MiB	18	33	20.52 µs	73.3
9. 0	5.20 day	0.38	2.02 GiB	39	57	188.75 µs	14.56
<u>9.</u> 02	5.20 day	0.00	59.36 MiB	16	33	21.10 µs	46.92
90	5.20 day	0.34	2.25 GiB	39	56	492.40 µs	15.7



## Summary

This document uses two samples to describe how to expose Go metrics to TMP and how to use the built-in visual charts to view monitoring data. This document only uses the Counter metrics. In other scenarios, you many need to use Gauge, Histogram, and Summary metrics. For more information, please see Metric Types.

For other use cases, TMP will integrate more frameworks to provide more out-of-the-box monitoring metrics, visual dashboards, and alerting templates.

# Exporter Integration Elasticsearch Exporter Integration

Last updated : 2024-08-07 22:01:17

## Overview

When using Elasticsearch, you need to monitor its running status, such as cluster and index status. TMP provides an exporter to monitor Elasticsearch and offers an out-of-the-box Grafana monitoring dashboard for it. This document describes how to deploy the Elasticsearch exporter and integrate it with the alert feature.

#### Note:

For easier export installation and management, we recommend you use TKE for unified management.

### Prerequisites

You have created a TKE cluster in the region and VPC of your TMP instance and created a namespace for the cluster. You have located and integrated the target TKE cluster in the **Integrate with TKE** section of the **target TMP instance** in the **TMP console**. For more information, please see Agent Management.

### Directions

### **Deploying exporter**

1. Log in to the TKE console.

2. Click the ID/name of the cluster whose access credential you want to get to enter the cluster management page.

3. Perform the following steps to deploy an exporter: Using Secret to manage Elasticsearch connection string > Deploying Elasticsearch exporter > Verifying.

### Using Secret to manage Elasticsearch connection string

1. On the left sidebar, select Workload > Deployment to enter the Deployment page.

In the top-right corner of the page, click Create via YAML to create a YAML configuration as detailed below:
 You can use Kubernetes Secrets to manage and encrypt passwords. When starting the Elasticsearch exporter, you can directly use the Secret key but need to adjust the corresponding URI. Below is a sample YAML configuration:

### Overview

When using Elasticsearch, you need to monitor its running status, such as cluster and index status. TMP provides an exporter to monitor Elasticsearch and offers an out-of-the-box Grafana monitoring dashboard for it. This document describes how to deploy the Elasticsearch exporter and integrate it with the alert feature.

### Note:

For easier export installation and management, we recommend you use TKE for unified management.

## Prerequisites

You have created a TKE cluster in the region and VPC of your TMP instance and created a namespace for the cluster. You have located and integrated the target TKE cluster in the **Integrate with TKE** section of the **target TMP instance** in the **TMP console**. For more information, please see Agent Management.

## Directions

### **Deploying exporter**

1. Log in to the TKE console.

2. Click the ID/name of the cluster whose access credential you want to get to enter the cluster management page.

3. Perform the following steps to deploy an exporter: Using Secret to manage Elasticsearch connection string >

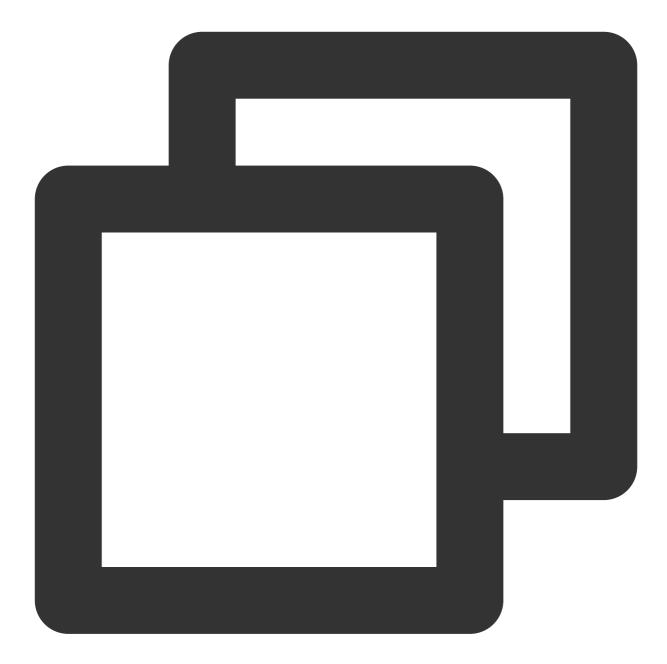
Deploying Elasticsearch exporter > Verifying.

### Using Secret to manage Elasticsearch connection string

1. On the left sidebar, select **Workload** > **Deployment** to enter the **Deployment** page.

In the top-right corner of the page, click Create via YAML to create a YAML configuration as detailed below:
 You can use Kubernetes Secrets to manage and encrypt passwords. When starting the Elasticsearch exporter, you can directly use the Secret key but need to adjust the corresponding URI. Below is a sample YAML configuration:



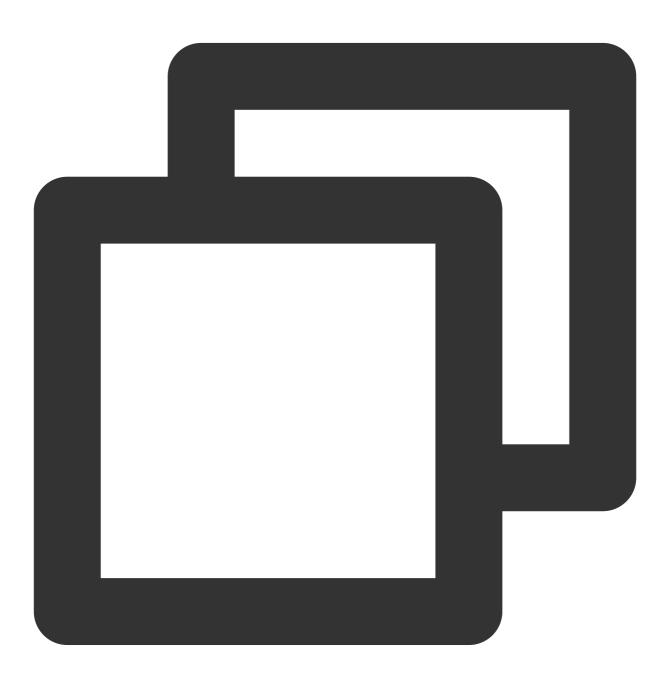


```
apiVersion: v1
kind: Secret
metadata:
   name: es-secret-test
   namespace: es-demo
type: Opaque
stringData:
   esURI: you-guess  # Corresponding Elasticsearch URI
```

The Elasticsearch connection string is in the format of
such as http://admin:pass@localhost:9200.

### **Deploying Elasticsearch exporter**

On the Deployment management page, click **Create** and select the target **namespace** to deploy the service. You can create in the console. Here, YAML is used to deploy the exporter. Below is a sample YAML configuration:



apiVersion: apps/v1 kind: Deployment metadata:

```
labels:
   k8s-app: es-exporter
 name: es-exporter
 namespace: es-demo
spec:
 replicas: 1
 selector:
    matchLabels:
     k8s-app: es-exporter
 template:
   metadata:
     labels:
       k8s-app: es-exporter
    spec:
      containers:
      - env:
          - name: ES_URI
            valueFrom:
              secretKeyRef:
                name: es-secret-test
                key: esURI
          - name: ES_ALL
            value: "true"
        image: bitnami/elasticsearch-exporter:latest
        imagePullPolicy: IfNotPresent
        name: es-exporter
        ports:
        - containerPort: 9114
         name: metric-port
        securityContext:
          privileged: false
        terminationMessagePath: /dev/termination-log
        terminationMessagePolicy: File
      dnsPolicy: ClusterFirst
      imagePullSecrets:
      - name: qcloudregistrykey
      restartPolicy: Always
      schedulerName: default-scheduler
      securityContext: {}
      terminationGracePeriodSeconds: 30
```

The above sample uses ES\_ALL to collect all monitoring metrics of Elasticsearch, which can be adjusted through the corresponding parameters. For detailed exporter parameters, please see elasticsearch\_exporter.

### Verifying

1. Click the newly created Deployment on the **Deployment** page to enter the Deployment management page.

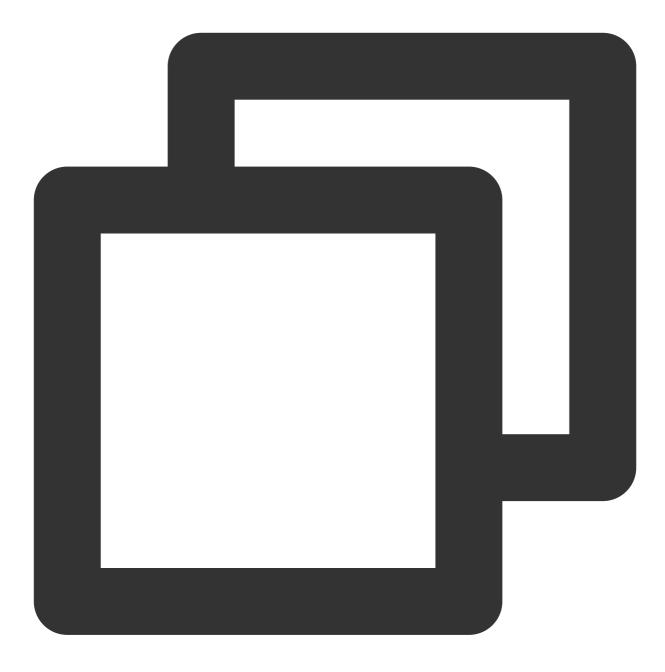
2. Click the **Log** tab, and you can see that the exporter is successfully started and its address is exposed as shown below:



3. Click the **Pod Management** tab to enter the Pod page.

4. In the **Operations** column on the right, click **Remote Login** to log in to the Pod. Run the following curl command with the address exposed by the exporter in the command line window, and you can get the corresponding Elasticsearch metrics normally. If no corresponding data is returned, please check whether the **connection string** is correct as shown below:





curl localhost:9114/metrics

The execution result is as shown below:

# HELP elasticsearch breakers estimated size bytes Estimated size in # TYPE elasticsearch breakers estimated size bytes gauge elasticsearch breakers estimated size bytes{breaker="accounting",cl 2.0102643e+07 elasticsearch breakers estimated\_size\_bytes{breaker="accounting",clu 1.9926654e+07 elasticsearch breakers estimated size bytes{breaker="accounting",clu 1.9685163e+07 elasticsearch breakers estimated size bytes{breaker="fielddata",clu: elasticsearch breakers estimated size bytes{breaker="fielddata",clu: elasticsearch breakers estimated size bytes{breaker="fielddata",clu: elasticsearch breakers estimated size bytes{breaker="in flight reque 0 elasticsearch breakers estimated size bytes{breaker="in flight requi 1167 elasticsearch breakers estimated size bytes{breaker="in flight requi 1167 elasticsearch breakers estimated size bytes{breaker="parent",cluste: 2.0102643e+07 alastissoarch broakars astimated size but as (broakar-"naront" alusta

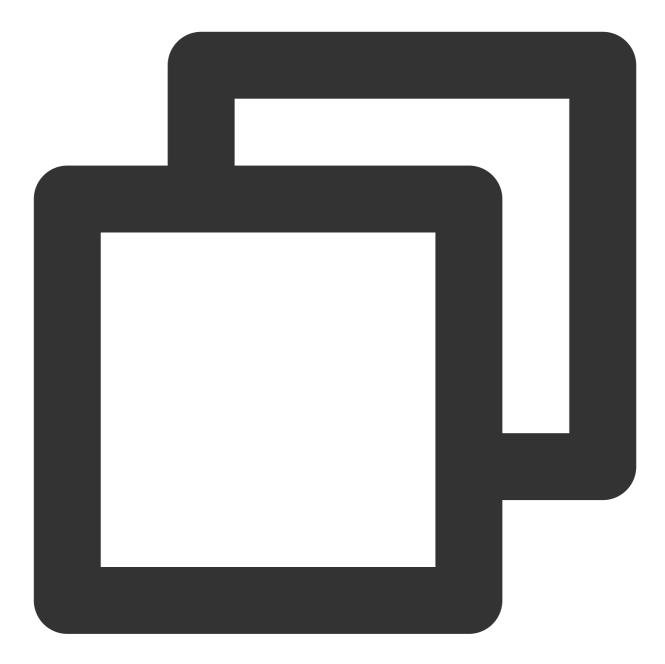
### Adding scrape task

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click a **cluster ID** in the TKE cluster list to enter the **Integrate with TKE** page.

3. In **Scrape Configuration**, add Pod Monitor to define a Prometheus scrape task. Below is a sample YAML configuration:





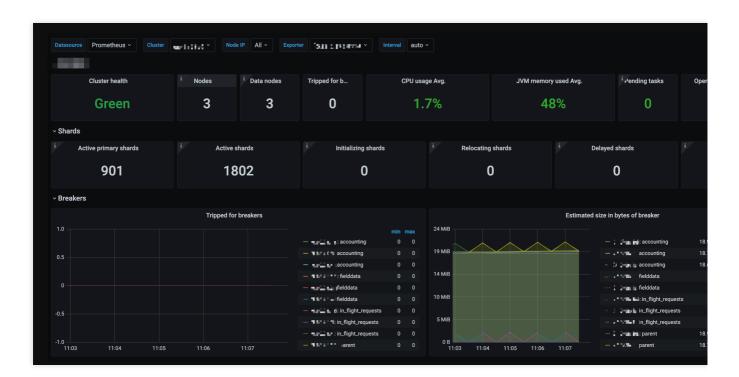
```
apiVersion: monitoring.coreos.com/v1
kind: PodMonitor
metadata:
   name: es-exporter
   namespace: cm-prometheus
spec:
   namespaceSelector:
    matchNames:
        - es-demo
   podMetricsEndpoints:
        - interval: 30s
```

```
path: /metrics
port: metric-port
selector:
matchLabels:
k8s-app: es-exporter
```

### Viewing monitoring information

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click **Integration Center** to enter the **Integration Center** page. Find Elasticsearch monitoring, install the corresponding Grafana dashboard, and then you can enable the Elasticsearch monitoring dashboard to view instance monitoring data as shown below:



### Integrating with alert feature

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click **Alerting Rule** and add the corresponding alerting rules. For more information, please see Creating Alerting Rule.

# Kafka Exporter Integration

Last updated : 2024-08-07 22:01:27

# Overview

When using Kafka, you need to monitor its running status, such as cluster status and message heap. TMP provides an exporter to monitor Kafka and offers an out-of-the-box Grafana monitoring dashboard for it. This document describes how to deploy the Kafka exporter and integrate it with the alert feature.

### Note:

For easier export installation and management, we recommend you use TKE for unified management.

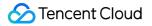
## Prerequisites

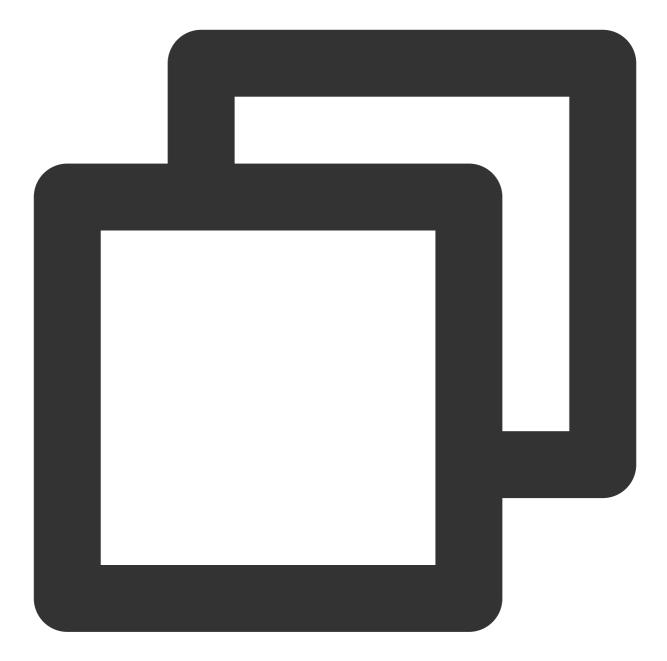
You have created a TKE cluster in the region and VPC of your TMP instance and created a namespace for the cluster. You have located and integrated the target TKE cluster in the **Integrate with TKE** section of the **target TMP instance** in the **TMP console**. For more information, please see Agent Management.

### Directions

### **Deploying exporter**

- 1. Log in to the TKE console.
- 2. Click the ID/name of the cluster whose access credential you want to get to enter the cluster management page.
- 3. On the left sidebar, select **Workload** > **Deployment** to enter the **Deployment** page.
- 4. On the Deployment management page, click **Create** and select the target **namespace** to deploy the service. You can create in the console. Here, YAML is used to deploy the exporter. Below is a sample YAML configuration:





```
apiVersion: apps/v1
kind: Deployment
metadata:
   labels:
      k8s-app: kafka-exporter # Rename the exporter based on the business needs. We r
   name: kafak-exporter # Rename the exporter based on the business needs. We recomm
   namespace: kafka-demo
spec:
   replicas: 1
   selector:
      matchLabels:
```

```
k8s-app: kafka-exporter # Rename the exporter based on the business needs. We
template:
 metadata:
    labels:
      k8s-app: kafka-exporter # Rename the exporter based on the business needs.
  spec:
    containers:
    - args:
      - --kafka.server=x.x.x.x:9092 # Corresponding Kafka instance address inform
      image: danielgsj/kafka-exporter:latest
      imagePullPolicy: IfNotPresent
      name: kafka-exporter
      ports:
      - containerPort: 9121
        name: metric-port # This name is required during scrape task configurati
      securityContext:
        privileged: false
      terminationMessagePath: /dev/termination-log
      terminationMessagePolicy: File
    dnsPolicy: ClusterFirst
    imagePullSecrets:
    - name: qcloudregistrykey
    restartPolicy: Always
    schedulerName: default-scheduler
    securityContext: {}
    terminationGracePeriodSeconds: 30
```

For detailed exporter parameters, please see kafka\_exporter.

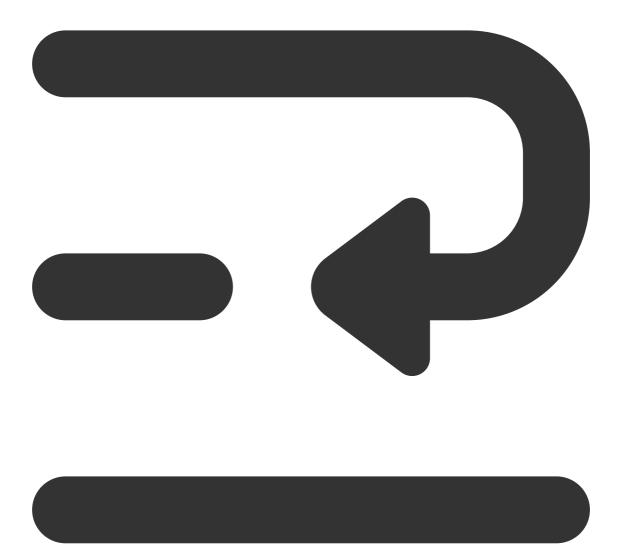
### Adding scrape task

1. Log in to the TMP console and select the target TMP instance to enter the management page.

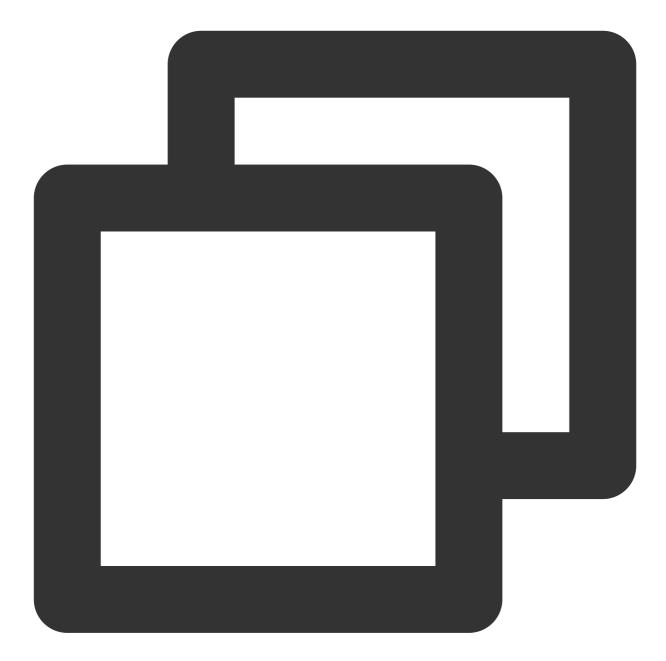
2. Click a cluster ID in the TKE cluster list to enter the Integrate with TKE page.

3. In Scrape Configuration, add Pod Monitor to define a Prometheus scrape task. Below is a sample YAML configuration:









```
apiVersion: monitoring.coreos.com/v1
kind: PodMonitor
metadata:
   name: kafka-exporter # Enter a unique name
   namespace: cm-prometheus # The namespace is fixed. Do not change it
spec:
   podMetricsEndpoints:
        - interval: 30s
        port: metric-port # Enter the name of the corresponding port of the Prometheu
        path: /metrics # Enter the value of the corresponding path of the Prometheus
        relabelings:
```

- action: replace
sourceLabels:
- instance
regex: (.*)
targetLabel: instance
replacement: 'ckafka-xxxxxx' # Change it to the corresponding Kafka instanc
- action: replace
sourceLabels:
- instance
regex: (.*)
targetLabel: ip
replacement: '1.x.x.x' # Change it to the corresponding Kafka instance IP
namespaceSelector:
matchNames:
- kafka-demo
selector: # Enter the label value of the Pod to be monitored to locate the tar
matchLabels:
k8s-app: kafka-exporter

As the exporter and Kafka are deployed on different servers, we recommend you use the Prometheus relabeling mechanism to add the Kafka instance information to the monitoring metrics so as to locate problems more easily.

### Viewing monitoring information

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click **Integration Center** to enter the **Integration Center** page. Find Kafka monitoring, install the corresponding Grafana dashboard, and then you can enable the Kafka monitoring dashboard to view instance monitoring data as shown below:

		<ul> <li>Author? In the elder deca</li> </ul>	12.00 12
	10.0	- venez cons	4.00 4
		<ul> <li>Ja vedlavá s</li> </ul>	3.00 3
		T - Mariana maria	3.00 3
		<ul> <li>Reflect in the otherway</li> </ul>	2.00 2
		<ul> <li>In region opening in the second</li> </ul>	
		<ul> <li>Native''''e dia state avita</li> </ul>	1.00 1
		<ul> <li>Citring to most time</li> </ul>	
		فيستعدون وبالإم والعام المتعاد	1.00 1
18:40 18:45 18:50 18:55 19:00 19:05 19:10 19:15 19:20 19:25 19:30 19:35	0 18:40 18:50 19:00 19:10 19:20 19:30	<ul> <li>Firmer of several costs as</li> </ul>	
ps-	100.00%		
cps		— fiervi čen	100.0
	95.00%	— An M(Anna A	100.0
ps	90.00%	— Anvi fors was	100.0
ccps	90.00%	- An of the se	100.0
**************************************	85.00%	<ul> <li>Relieving to of General</li> </ul>	100.0
		<ul> <li>Unit half has placed given it</li> </ul>	100.0
	80.00%	<ul> <li>Reflection involutions in</li> </ul>	100.0
18:40 18:45 18:50 18:55 19:00 19:05 19:10 19:15 19:20 19:25 19:30 19:35		<ul> <li>Het het het jierteljier det 24</li> <li>Red inginge ingest groute gena.</li> </ul>	100.0
hand r = -hand r = -hand r = -hand r = hand r = hand r = -hand r	75.00%	<ul> <li>All fait is the initial strip.</li> </ul>	100.0
(F) Webs Land Tang, — File? Webs Hand Tang J & — Reference "And	18:40 18:50 19:00 19:10 19:20 19:30		
(Consumer Group)		14.000	
		<ul> <li>Anthref in Jacober (ma)</li> </ul>	24.00 24
		<ul> <li>At receive at 12</li> </ul>	9.00 9
		- Andreas Trans	o nn 🛛 🕯

### Integrating with alert feature

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click **Alerting Rule** and add the corresponding alerting rules. For more information, please see Creating Alerting Rule.

# MongoDB Exporter Integration

Last updated : 2024-08-07 22:01:35

# Overview

When using MongoDB, you need to monitor its running status to know whether it runs normally and troubleshoot its faults. TMP provides an exporter to monitor MongoDB and offers an out-of-the-box Grafana monitoring dashboard for it. This document describes how to deploy the MongoDB exporter and integrate it with the alert feature.

For easier export installation and management, we recommend you use TKE for unified management.

# Prerequisites

You have created a TKE cluster in the region and VPC of your TMP instance.

You have located and integrated the target TKE cluster in the **Integrate with TKE** section of the **target TMP instance** in the **TMP console**. For more information, please see Agent Management.

### Directions

### **Deploying exporter**

1. Log in to the TKE console.

2. Click the ID/name of the cluster whose access credential you want to get to enter the cluster management page.

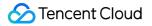
3. Perform the following steps to deploy an exporter: Using Secret to manage MongoDB connection string > Deploying MongoDB exporter > Verifying.

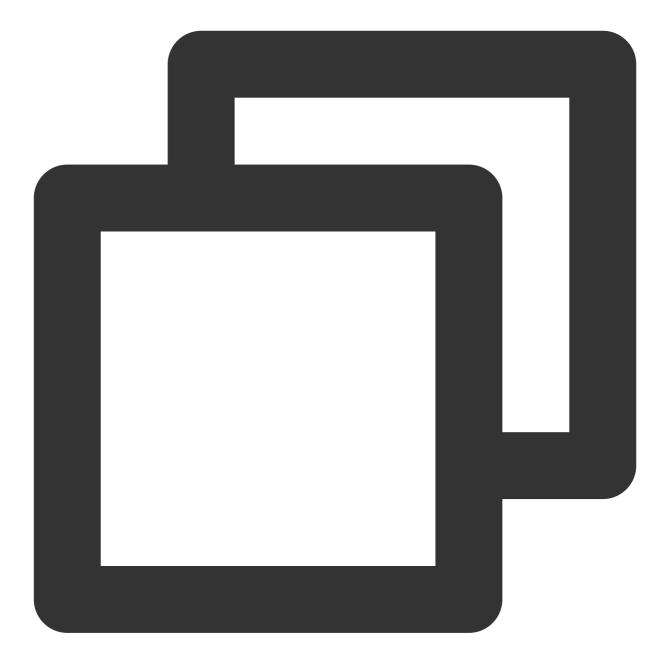
### Using Secret to manage MongoDB connection string

1. On the left sidebar, select Workload > Deployment to enter the Deployment page.

2. In the top-right corner of the page, click Create via YAML to create a YAML configuration as detailed below:

You can use Kubernetes Secrets to manage and encrypt passwords. When starting the MongoDB exporter, you can directly use the Secret key but need to adjust the corresponding URI. Below is a sample YAML configuration:



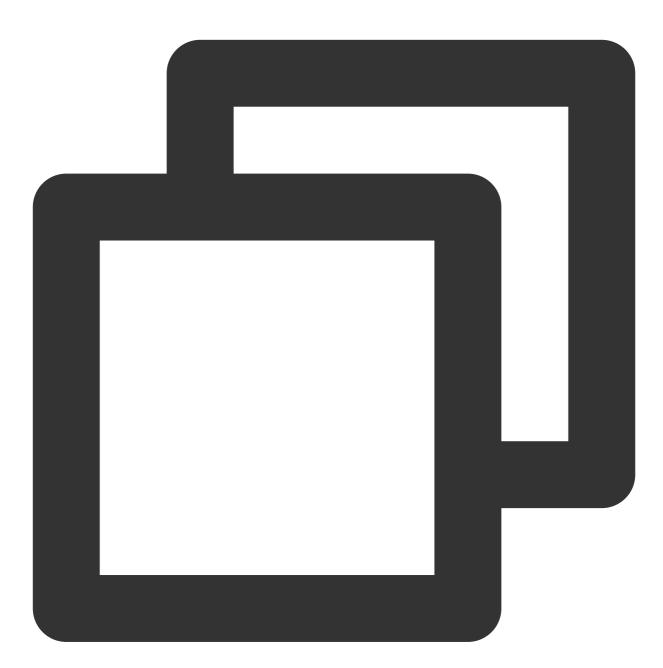


```
apiVersion: v1
kind: Secret
metadata:
  name: mongodb-secret-test
  namespace: mongodb-test
type: Opaque
stringData:
  datasource: "mongodb://{user}:{passwd}@{host1}:{port1},{host2}:{port2},{host3}:
```

### Deploying MongoDB exporter

### 🔗 Tencent Cloud

On the Deployment management page, click **Create** and select the target **namespace** to deploy the service. You can create in the console. Here, YAML is used to deploy the exporter. Below is a sample YAML configuration:



```
apiVersion: apps/v1
kind: Deployment
metadata:
   labels:
        k8s-app: mongodb-exporter # Rename the exporter based on the business needs. We
   name: mongodb-exporter # Rename the exporter based on the business needs. We reco
   namespace: mongodb-test
spec:
```

```
replicas: 1
selector:
  matchLabels:
    k8s-app: mongodb-exporter # Rename the exporter based on the business needs.
template:
 metadata:
    labels:
      k8s-app: mongodb-exporter # Rename the exporter based on the business needs
  spec:
    containers:
      - args:
                                     # Enable the collection of `Database` metric
          - --collect.database
          - --collect.collection  # Enable the collection of `Collection` metr
          - --collect.topmetrics
                                    # Enable the collection of `table top` metri
          - --collect.indexusage
                                    # Enable the collection of `per index usage
          - --collect.connpoolstats # Enable the collection of `MongoDB connpool
        env:
          - name: MONGODB_URI
            valueFrom:
              secretKeyRef:
                name: mongodb-secret-test
                key: datasource
        image: ssheehy/mongodb-exporter
        imagePullPolicy: IfNotPresent
        name: mongodb-exporter
        ports:
          - containerPort: 9216
            name: metric-port # This name is required during scrape task configu
        securityContext:
          privileged: false
        terminationMessagePath: /dev/termination-log
        terminationMessagePolicy: File
    dnsPolicy: ClusterFirst
    imagePullSecrets:
      - name: qcloudregistrykey
    restartPolicy: Always
    schedulerName: default-scheduler
    securityContext: { }
    terminationGracePeriodSeconds: 30
```

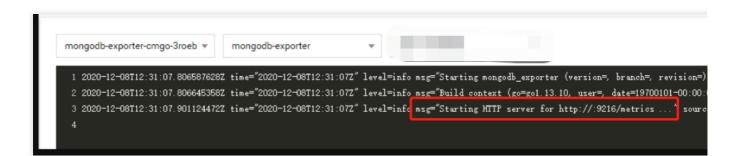
For detailed exporter parameters, please see mongodb\_exporter.

### Verifying

1. Click the newly created Deployment on the **Deployment** page to enter the Deployment management page.



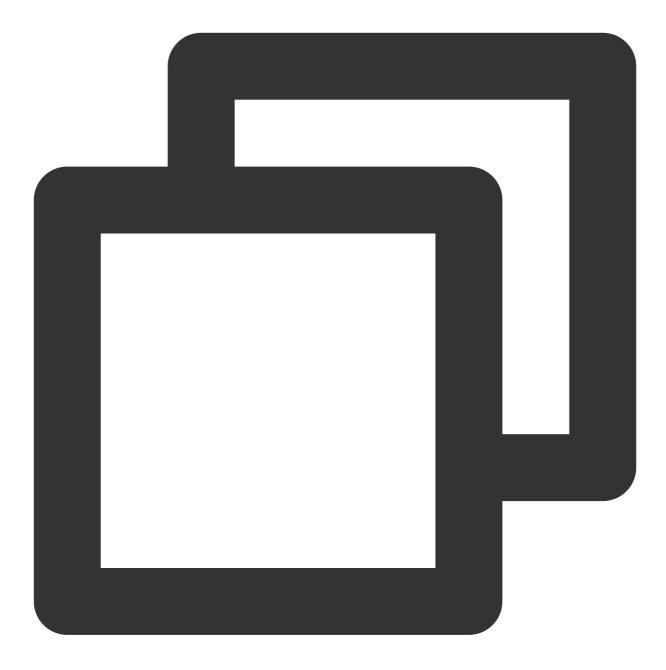
2. Click the **Log** tab, and you can see that the exporter is successfully started and its address is exposed as shown below:



3. Click the Pod Management tab to enter the Pod page.

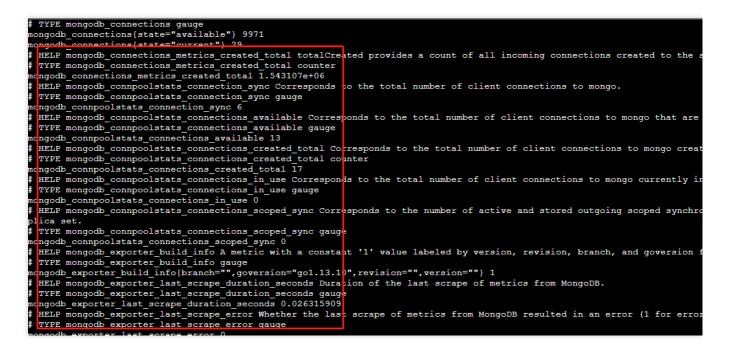
4. In the **Operations** column on the right, click **Remote Login** to log in to the Pod. Run the following wget command with the address exposed by the exporter on the command line, and you can get the corresponding MongoDB metrics normally. If no corresponding data is returned, please check whether the connection URI is correct as shown below:





wget 127.0.0.1:9216/metrics
cat metrics

The command execution result is as shown below:

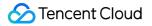


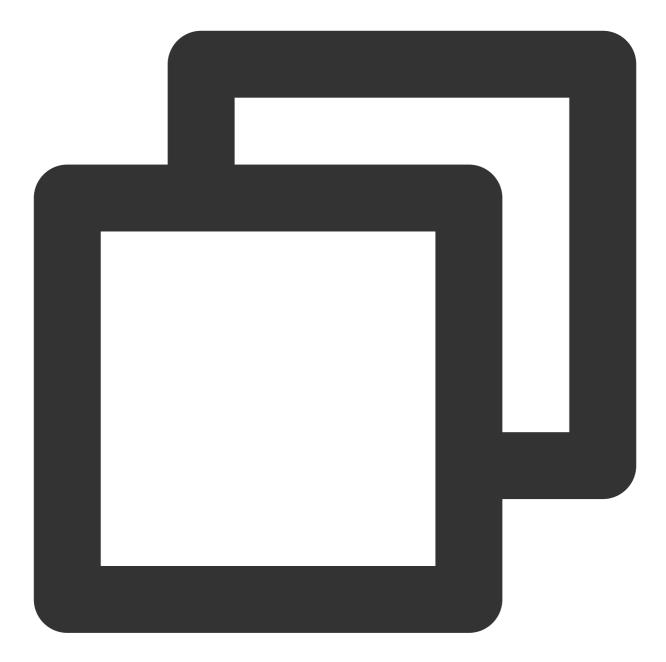
### Adding scrape task

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click a cluster ID in the TKE cluster list to enter the Integrate with TKE page.

3. In **Scrape Configuration**, add Pod Monitor to define a Prometheus scrape task. Below is a sample YAML configuration:





```
apiVersion: monitoring.coreos.com/v1
kind: PodMonitor
metadata:
   name: mongodb-exporter # Enter a unique name
   namespace: cm-prometheus # The namespace is fixed. Do not change it
spec:
   podMetricsEndpoints:
        - interval: 30s
        port: metric-port # Enter the name of the corresponding port of the Prometh
        path: /metrics # Enter the value of the corresponding path of the Prometheus
        relabelings:
```



As the exporter and MongoDB are deployed on different servers, we recommend you use the Prometheus relabeling mechanism to add the MongoDB instance information to the monitoring metrics so as to locate problems more easily.

### Viewing monitoring information

monitoring data as shown below:

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click Integration Center to enter the Integration Center page. Find MongoDB monitoring, install the corresponding Grafana dashboard, and then you can enable the MongoDB monitoring dashboard to view instance

**MongoDB Overview**: you can view the status of each instance, such as number of documents, connection utilization, and read/write time. You can click an instance to view its details.

**MongoDB Details**: you can view the detailed status of an instance, such as metadata overview, core metrics, command operations, request traffic, and top reads/writes.

88 Mongodb 2 44	witations	al*	D D Last 12 hours - Q C -
<b>4.6</b> week	32.5 мів	16	<b>280</b> к
1%	93%	0%	0 s
1 128.5 128.6 127.5 127.5 127.5 120.5 10.00 12.00 14.00 14.00 - read - with	1600 1800 2500	3.00 2.00 1.00 1.00 - cohi magi	1.60 11.00 20.00
15. 15. 10. 00 mi 0 mi 1000 1200 1600	M_MM_M/V/M/M/	1.00 opa 0.75 opa 0.05 opa 0.05 opa 0.05 opa 0.05 opa 10:00 12:00 14:00	1600 1800 2000
- Witk Wat Tree O ops 0.088 ops 0 ops	ий нис еду 25 m 15/4, 1004. О оря 0.0035 оря 0 оря	- Timeods O ops 0.011 ops 0 ops	nak mag mag baga Baga Baga Baga O ops 0.0035 ops 0 ops

You can click ! on the left of each chart to view the description.

### Integrating with alert feature

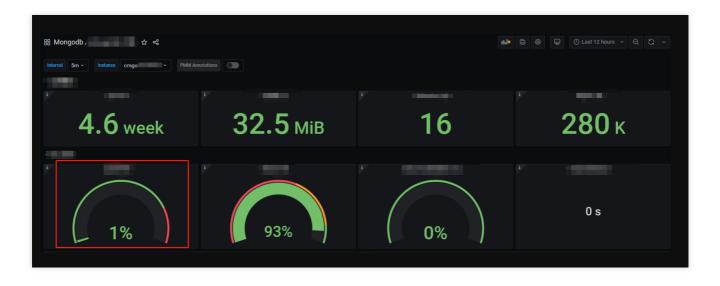
1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click **Alerting Rule** and add the corresponding alerting rules. For more information, please see Creating Alerting Rule.

## FAQs

### The client reported an error "client checkout connect timeout". What should I do?

This is probably because that the connection pool utilization has reached 100%, resulting in a connection creation failure. You can check the **Connection Utilization** metric in **MongoDB Details > Core Metrics** on the Grafana dashboard for troubleshooting.



#### Write keeps timing out. What should I do?

Check whether the cache utilization is excessive and whether the number of available transactions is 0. You can check the **Available WiredTiger Transactions**, **WiredTiger Cache Utilization**, and **GetLastError Write Time** metrics in **MongoDB Details** > **Core Metrics** on the Grafana dashboard for troubleshooting.



# PostgreSQL Exporter Integration

Last updated : 2024-08-07 22:01:43

# Overview

When using PostgreSQL, you need to monitor its running status to know whether it runs normally and troubleshoot its faults. TMP provides an exporter to monitor PostgreSQL and offers an out-of-the-box Grafana monitoring dashboard for it. This document describes how to deploy the PostgreSQL exporter and integrate it with the alert feature. **Note:** 

For easier export installation and management, we recommend you use TKE for unified management.

# Prerequisites

You have created a TKE cluster in the region and VPC of your TMP instance.

You have located and integrated the target TKE cluster in the **Integrate with TKE** section of the **target TMP instance** in the **TMP console**. For more information, please see Agent Management.

### Directions

### **Deploying exporter**

1. Log in to the TKE console.

2. Click the ID/name of the cluster whose access credential you want to get to enter the cluster management page.

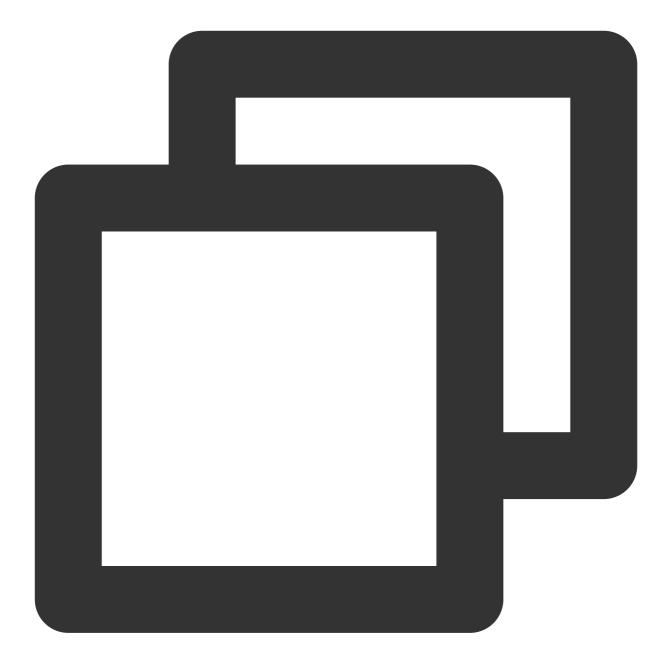
3. Perform the following steps to deploy an exporter: Using Secret to manage PostgreSQL password > Deploying PostgreSQL exporter > Getting metric.

### Using Secret to manage PostgreSQL password

1. On the left sidebar, select Workload > Deployment to enter the Deployment page.

2. In the top-right corner of the page, click **Create via YAML** to create a YAML configuration as detailed below: You can use Kubernetes Secrets to manage and encrypt passwords. When starting the PostgreSQL exporter, you can directly use the Secret key but need to adjust the corresponding password. Below is a sample YAML configuration:



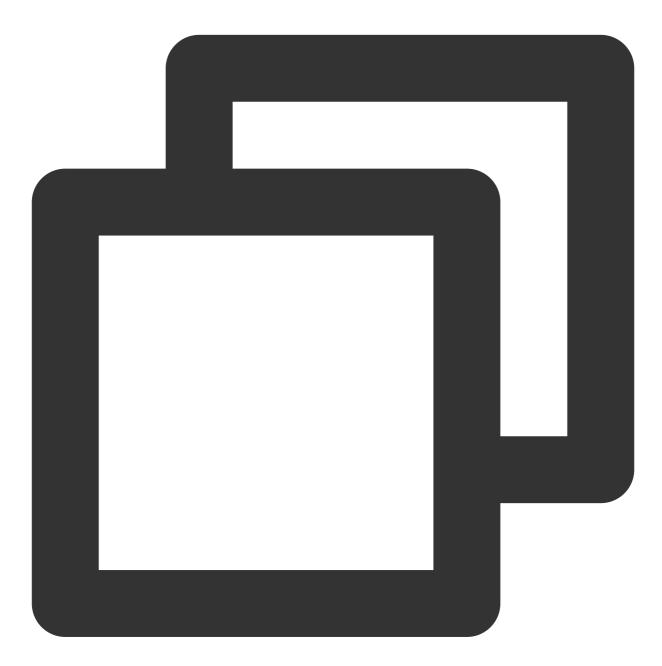


```
apiVersion: v1
kind: Secret
metadata:
name: postgres-test
type: Opaque
stringData:
username: postgres
password: you-guess # Corresponding PostgreSQL password
```

### Deploying PostgreSQL exporter

### 🔗 Tencent Cloud

On the Deployment management page, click **Create** and select the target **namespace** to deploy the service. You can create in the console. Here, YAML is used to deploy the exporter. Below is a sample YAML configuration (please directly copy the following content and adjust the corresponding parameters based on your actual business needs):



apiVersion: apps/v1
kind: Deployment
metadata:
 name: postgres-test
 namespace: postgres-test
 labels:
 app: postgres

```
app.kubernetes.io/name: postgresql
spec:
 replicas: 1
 selector:
   matchLabels:
      app: postgres
      app.kubernetes.io/name: postgresql
 template:
   metadata:
      labels:
        app: postgres
        app.kubernetes.io/name: postgresql
    spec:
      containers:
      - name: postgres-exporter
        image: wrouesnel/postgres_exporter:latest
        args:
          - "--web.listen-address=:9187"
          - "--log.level=debug"
        env:
          - name: DATA_SOURCE_USER
            valueFrom:
              secretKeyRef:
                name: postgres-test
                key: username
          - name: DATA_SOURCE_PASS
            valueFrom:
              secretKeyRef:
                name: postgres-test
                key: password
          - name: DATA_SOURCE_URI
            value: "x.x.x.x:5432/postgres?sslmode=disable"
        ports:
        - name: http-metrics
          containerPort: 9187
```

In the above sample, the username and password in Secret are passed in to the environment variables DATA\_SOURCE\_USER and DATA\_SOURCE\_PASS, so the username and password cannot be viewed in plaintext. You can also use DATA\_SOURCE\_USER\_FILE / DATA\_SOURCE\_PASS\_FILE to read the username and password from the file, or use DATA\_SOURCE\_NAME to put them in the connection string, such as postgresql://login:password@hostname:port/dbname .

#### Parameter description



The query part (after ?) in the DATA\_SOURCE\_URI / DATA\_SOURCE\_NAME connection string supports the following parameters (the latest supported parameters listed in Connection String Parameters shall prevail):

Parameter	Description
sslmode	Whether to use SSL. Valid values:
- disable	Do not use SSL
- require	Always use (skip verification)
- verify-ca	Always use (check whether the certificate provided by the server is issued by a trusted CA)
- verify-full	Always use (check whether the certificate provided by the server is issued by a trusted CA and whether the hostname matches the certificate)
fallback_application_name	Alternative application_name
connect_timeout	Maximum connection wait time in seconds. `0` indicates to wait infinitely
sslcert	Certificate file path. The file data must be in PEM format
SSHKey	Private key file path. The file data must be in PEM format
sslrootcert	Root certificate file path. The file data must be in PEM format

Other supported exporter parameters are as detailed below (for more information, please see PostgreSQL Server

### Exporter):

Parameter	Description	Environment Variable
web.listen- address	Listening address. Default value: :9487	PG_EXPORTER_WEB_LISTEN_ADDRESS
web.telemetry- path	Path under which to expose metrics. Default value: /metrics	PG_EXPORTER_WEB_TELEMETRY_PATH
extend.query- path	Path of a YAML file containing custom queries to run. For more information, please see queries.yaml	PG_EXPORTER_EXTEND_QUERY_PATH
disable-default- metrics	Uses only metrics supplied from queries.yaml	PG_EXPORTER_DISABLE_DEFAULT_METRICS
disable-settings- metrics	Skips scraping pg_settings metrics	PG_EXPORTER_DISABLE_SETTINGS_METRICS



auto-discover- databases	Whether to discover the databases in the PostgreSQL instance dynamically	PG_EXPORTER_AUTO_DISCOVER_DATABASES
dumpmaps	Prints the internal metric information to help troubleshoot custom queries (do not use it unless for debugging)	-
constantLabels	Custom label provided in the format of key=value. Multiple labels are separated with ,	PG_EXPORTER_CONSTANT_LABELS
exclude- databases	Database to be excluded. It takes effect only ifauto-discover-databases is enabled	PG_EXPORTER_EXCLUDE_DATABASES
log.level	Log level. Valid values: debug, info, warn, error, fatal	PG_EXPORTER_LOG_LEVEL

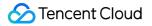
### Getting metric

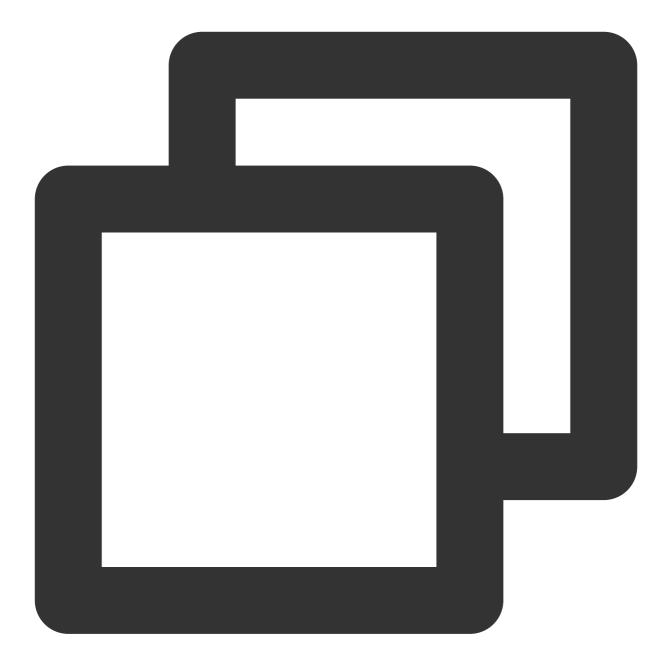
You cannot get the PostgreSQL instance operation time through curl http://exporter:9187/metrics . You can define a queries.yaml file to get this metric:

1. Create a ConfigMap containing queries.yaml .

2. Mount the ConfigMap to a directory in the exporter as a volume.

3. Use the ConfigMap through <u>--extend.query-path</u> to aggregate the information of the aforementioned Secret and Deployment. The YAML file after aggregation is as shown below:





```
# Note: the following document sample code creates a namespace named `postgres-test
apiVersion: v1
kind: Namespace
metadata:
    name: postgres-test
# The following document sample code creates a Secret containing a username and pas
```

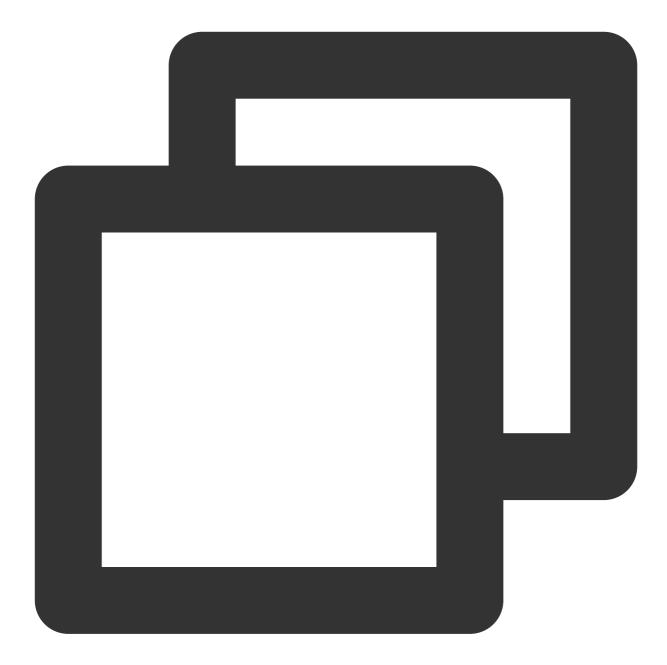
```
# The following document sample code creates a secret containing a username and pas
---
apiVersion: v1
kind: Secret
metadata:
```



```
name: postgres-test-secret
 namespace: postgres-test
type: Opaque
stringData:
 username: postgres
 password: you-guess
# The following document sample code creates a `queries.yaml` file containing custo
apiVersion: v1
kind: ConfigMap
metadata:
 name: postgres-test-configmap
 namespace: postgres-test
data:
 queries.yaml: |
   pg_postmaster:
      query: "SELECT pg_postmaster_start_time as start_time_seconds from pg_postmas
      master: true
      metrics:
        - start_time_seconds:
            usage: "GAUGE"
            description: "Time at which postmaster started"
# The following document sample code mounts the Secret and ConfigMap and defines ex
apiVersion: apps/v1
kind: Deployment
metadata:
 name: postgres-test
 namespace: postgres-test
 labels:
    app: postgres
   app.kubernetes.io/name: postgresql
spec:
  replicas: 1
  selector:
   matchLabels:
      app: postgres
      app.kubernetes.io/name: postgresql
 template:
    metadata:
      labels:
        app: postgres
        app.kubernetes.io/name: postgresql
    spec:
      containers:
```

```
- name: postgres-exporter
    image: wrouesnel/postgres_exporter:latest
    args:
      - "--web.listen-address=:9187"
      - "--extend.query-path=/etc/config/queries.yaml"
      - "--log.level=debug"
    env:
      - name: DATA SOURCE USER
        valueFrom:
          secretKeyRef:
            name: postgres-test-secret
            key: username
      - name: DATA_SOURCE_PASS
        valueFrom:
          secretKeyRef:
            name: postgres-test-secret
            key: password
      - name: DATA_SOURCE_URI
        value: "x.x.x.:5432/postgres?sslmode=disable"
    ports:
      - name: http-metrics
        containerPort: 9187
    volumeMounts:
      - name: config-volume
        mountPath: /etc/config
volumes:
  - name: config-volume
    configMap:
      name: postgres-test-configmap
```

4. Run curl http://exporter:9187/metrics , and you can use the custom queries.yaml to query the PostgreSQL instance start time as follows:



# HELP pg\_postmaster\_start\_time\_seconds Time at which postmaster started # TYPE pg\_postmaster\_start\_time\_seconds gauge pg\_postmaster\_start\_time\_seconds{server="x.x.x.x:5432"} 1.605061592e+09

#### Adding scrape task

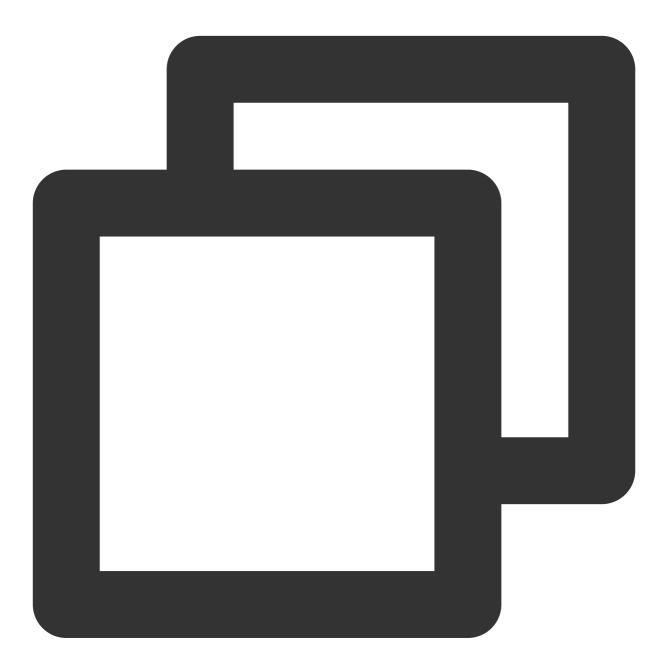
After the exporter runs, you need to configure TMP to discover and collect the monitoring metrics in the following steps:

1. Log in to the TMP console and select the target TMP instance to enter the management page.



2. Click a cluster ID in the TKE cluster list to enter the Integrate with TKE page.

3. In Scrape Configuration, add Pod Monitor to define a Prometheus scrape task. Below is a sample YAML configuration:



```
apiVersion: monitoring.coreos.com/v1
kind: PodMonitor
metadata:
   name: postgres-exporter
   namespace: cm-prometheus
spec:
   namespaceSelector:
```



```
matchNames:
  - postgres-test
podMetricsEndpoints:
- interval: 30s
  path: /metrics
  port: http-metrics # Port name of the aforementioned exporter container
  relabelings:
  - action: labeldrop
    regex: __meta_kubernetes_pod_label_(pod_|statefulset_|deployment_|controlle
  - action: replace
    regex: (.*)
    replacement: postgres-xxxxxx
    sourceLabels:
    - instance
    targetLabel: instance
selector:
  matchLabels:
    app: postgres
```

#### Note:

For more advanced usage, please see ServiceMonitor and PodMonitor.

#### Visualizing Grafana dashboard

#### Note:

You need to use the configuration in Getting metric to get the PostgreSQL instance start time.

1. In the TMP instance list, find the corresponding TMP instance, click



on the right of the instance ID to open your Grafana page, and enter your account and password to access the Grafana visual dashboard operation section.

2. Enter Grafana, click the



icon to expand the monitoring dashboard, and click the name of the corresponding monitoring chart to view the monitoring data.

<ul> <li>General Counters, CPU, N</li> </ul>	lemory and File Descriptor S	tats 💿 🖻							
Version			Current fetch data		Current insert data			Current update data	
9.5.4			4.404	MB	61 B			7 B	
i /	Average CPU Usage			Average Memory Us	age				Open File Desc
400 ms			60 kB					11.0	
300 ms			40 kB					10.0	
200 ms	<u> </u>	V	20 kB					9.0	
00 ms			0 В					8.0	
0 ns			14:40	14:50 15:00		15:20	15:30		
14:40 14:	50 15:00 15:10 min max	15:20 15:30	— Resident Mem		in max B 51.0 kB	<b>avg</b> 10.6 kB	current 819 B	14:4	0 14:50 15:00
— CPU Time		avg current 275 ms 259 ms	<ul> <li>— Kesident Mem</li> <li>— Virtual Mem</li> </ul>		B 971 B	213 B	0 B	— Open FD	
Settings Shared Buffers	Effective Cache	Maintenance Wo			Max WAL Size		Rande	om Page Cost	Seq Page C
512 MiB	4.000 GiB	64.0 M	iB 4.00	DO MiB	N/A			4	1
P Database Stats									
	Active sessions			Transactions					Update da
	max	avg current ~	10.0		avg	current	total	9.0 B	
	<ul> <li>postgres, s: active 1.0</li> </ul>	1.0 1.0	7.5	— template1 com				8.0 B	— postgres
				template0 com					
				postgres comm     template1 rollb		8.10 0	949.03 0	7.0 В	
			2.5	template1 rolls     template0 rolls		0	0	6.0 B	
				<ul> <li>postgres rollba</li> </ul>			51.60	0.0 0	

#### Integrating with alert feature

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click Alerting Rule and add the corresponding alerting rules. For more information, please see Creating Alerting

#### Rule.

#### Note:

TMP will provide more PostgreSQL alerting templates in the near future.

# **NGINX Exporter Integration**

Last updated : 2024-08-07 22:01:48

# Overview

NGINX exposes certain monitoring metrics through the stub\_status page. NGINX Prometheus Exporter collects the metrics of a single NGINX instance, converts them into monitoring data that can be used by Prometheus, and exposes the data to the Prometheus service for collection over the HTTP protocol. You can use the exporter to report the key monitoring metrics, which can be used for exception alerting and displayed on the dashboard.

# Directions

#### Using Docker container to run exporter

Method 1. Use nginx-prometheus-exporter to quickly deploy the exporter in a Docker container. Run the following Docker command:



\$ docker run -p 9113:9113 nginx/nginx-prometheus-exporter:0.8.0 -nginx.scrape-uri h

Method 2. Use the nginx-prometheus-exporter image to deploy the service in TKE and collect the monitoring data through TMP's self-discovery CRD PodMonitor or ServiceMonitor.

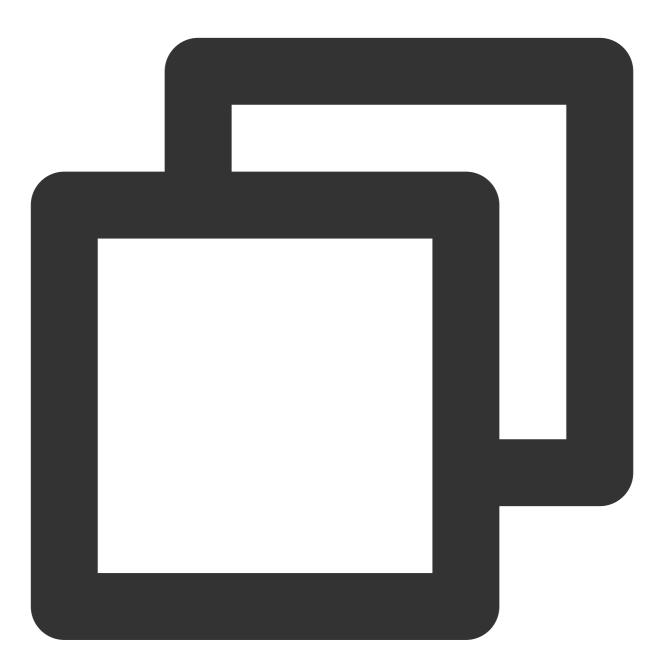
#### Using binary program to run exporter

#### Downloading and installing explorer

- 1. Download NGINX Prometheus Exporter from the community for your runtime environment.
- 2. Install NGINX Prometheus Exporter.

#### Enabling NGINX stub\_status feature

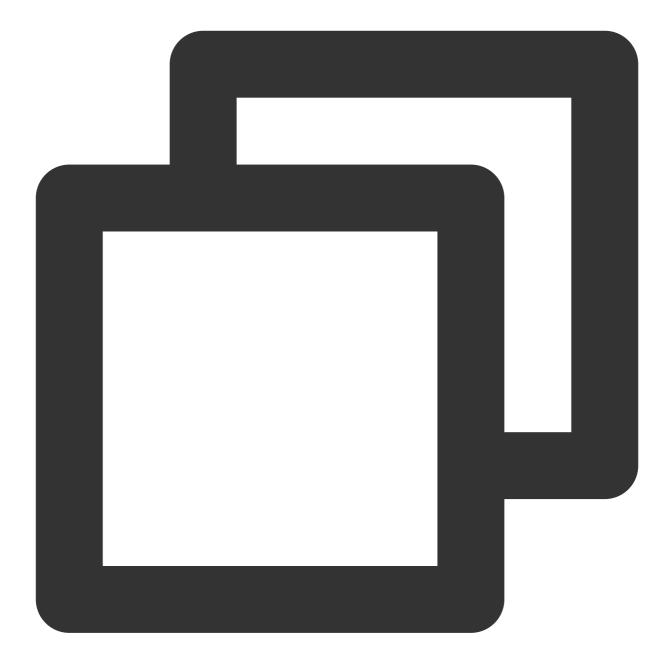
1. The stub\_status module of open-source NGINX provides a simple page to display the status data. Run the following command to check whether this module is enabled in NGINX:



nginx -V 2>&1 | grep -o with-http\_stub\_status\_module

If with-http\_stub\_status\_module is output on the terminal, the stub\_status module in NGINX is enabled.

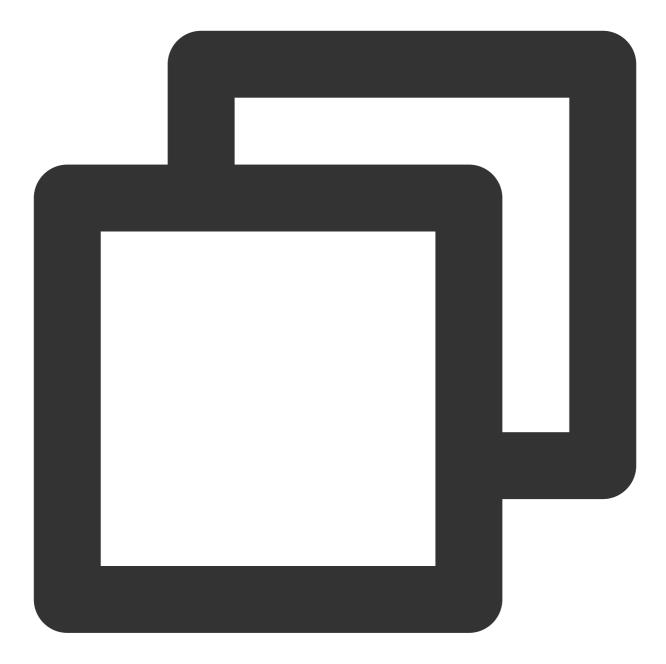
If no result is output, you can use the \_-with-http\_stub\_status\_module parameter to configure and compile NGINX again from the source code. Below is the sample code:



```
./configure \\
... \\--with-http_stub_status_module
make
sudo make install
```

2. After confirming that the stub\_status module is enabled, modify the NGINX configuration file to specify the URL of the stub\_status page as follows:

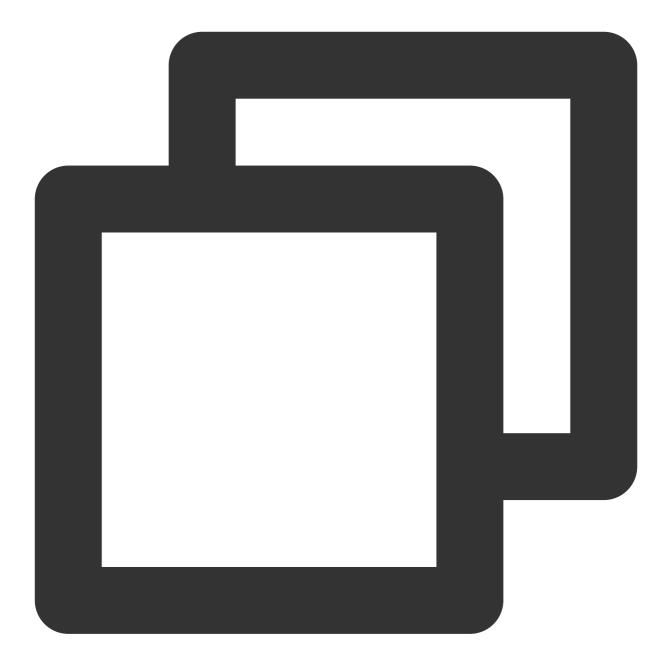




```
server {
  location /nginx_status {
    stub_status;
    access_log off;
    allow 127.0.0.1;
    deny all;
  }
}
```

3. Check NGINX and load it again to make the configuration take effect.

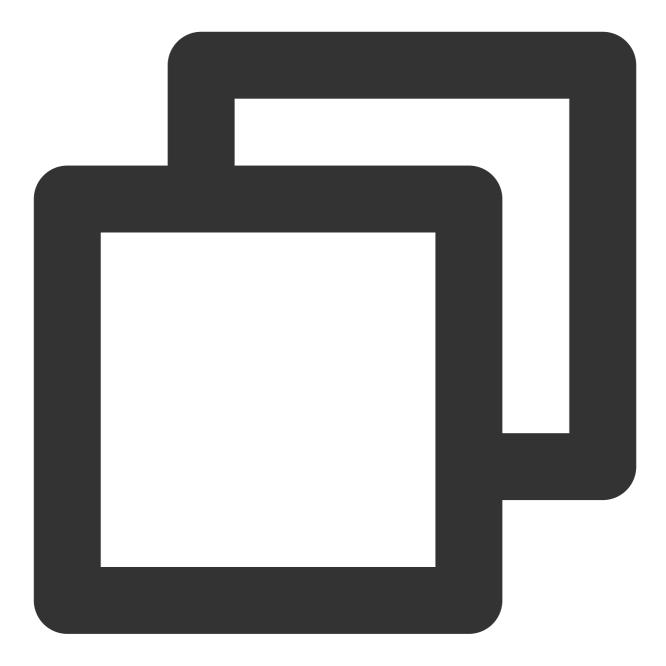




nginx -t nginx -s reload

4. After completing the above steps, you can view the NGINX metrics through the configured URL.

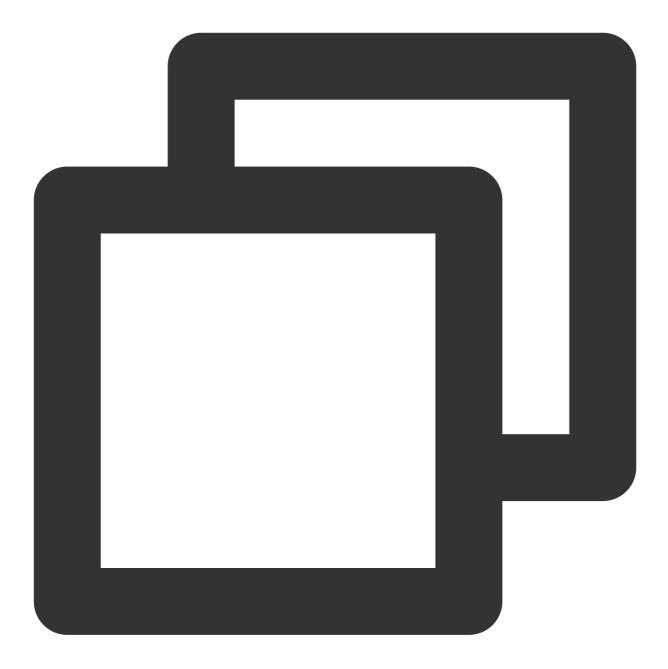




```
Active connections: 45
server accepts handled requests
1056958 1156958 4491319
Reading: 0 Writing: 25 Waiting : 7
```

#### **Running NGINX Prometheus Exporter**

Run the following command to start NGINX Prometheus Exporter:



\$ nginx-prometheus-exporter -nginx.scrape-uri http://<nginx>:8080/nginx\_status

#### **Reported metrics**

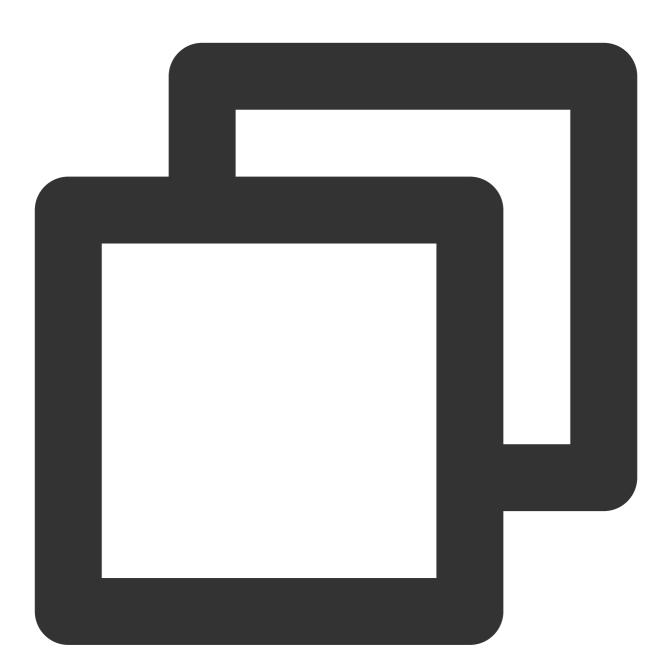
nginxexporter\_build\_info , which is the exporter compilation information.

#### All stub\_status metrics.

nginx\_up , which displays the last scrape status. 1 indicates success, while 0 indicates failure.

#### Configuring Prometheus scrape job

1. After NGINX Prometheus Exporter runs normally, run the following command to add a job to the Prometheus scrape task.

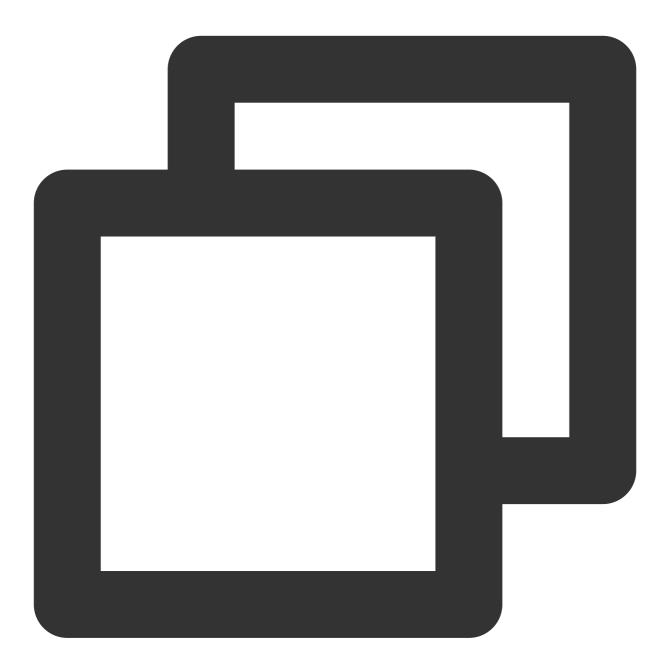


```
...
    job_name: 'nginx_exporter'
    static_configs:
        - targets: ['your_exporter:port']
```

2. Generally, the exporter and NGINX do not run together, so the instance of the reported data cannot describe the real instance. To facilitate data search and observation, you can modify the instance label and replace it with



the real IP to make the label more intuitive as follows:



```
...
- job_name: 'mysqld_exporter'
static_configs:
- targets: ['your_exporter:port']
relabel_configs:
- source_labels: [__address__]
regex: '.*'
target_label: instance
replacement: '10.0.0.1:80'
```

#### Enabling database monitoring dashboard

TMP provides a preconfigured NGINX exporter dashboard in Grafana. You can view the NGINX monitoring data in the following steps.

- 1. Log in to the TMP console.
- 2. Click



on the right of the corresponding instance ID to view the data.



# **Redis Exporter Integration**

Last updated : 2024-08-07 22:01:56

# Overview

When using Redis, you need to monitor its running status to know whether it runs normally and troubleshoot its faults. TMP provides an exporter to monitor Redis and offers an out-of-the-box Grafana monitoring dashboard for it. This document describes how to use TMP to monitor Redis.

#### Note:

For easier export installation and management, we recommend you use TKE for unified management.

# Prerequisites

You have created a TKE cluster in the region and VPC of your TMP instance and created a namespace for the cluster. You have located and integrated the target TKE cluster in the **Integrate with TKE** section of the **target TMP instance** in the **TMP console**. For more information, please see Agent Management.

### Directions

#### **Deploying exporter**

1. Log in to the TKE console.

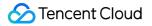
2. Click the ID/name of the cluster whose access credential you want to get to enter the cluster management page.

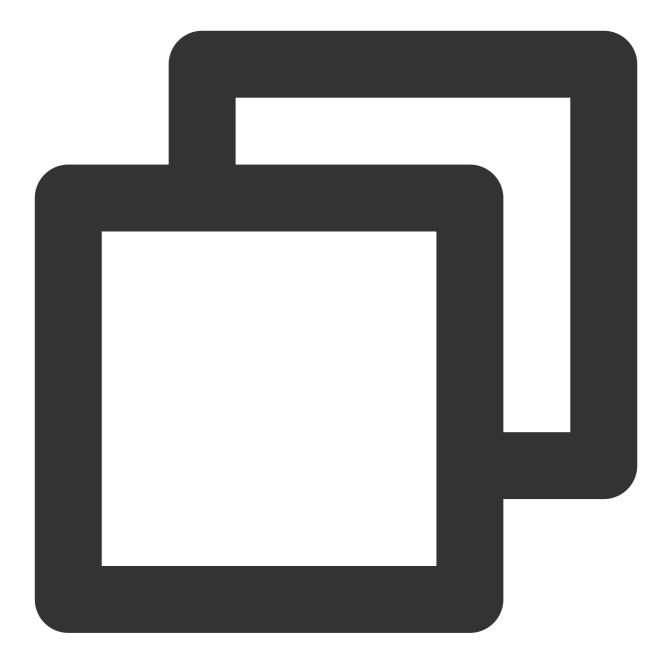
3. Perform the following steps to deploy an exporter: Using Secret to manage Redis password > Deploying Redis exporter > Verifying.

#### Using Secret to manage Redis password

1. On the left sidebar, select Workload > Deployment to enter the Deployment page.

2. In the top-right corner of the page, click **Create via YAML** to create a YAML configuration as detailed below: You can use Kubernetes Secrets to manage and encrypt passwords. When starting the Redis exporter, you can directly use the Secret key but need to adjust the corresponding password. Below is a sample YAML configuration:





```
apiVersion: v1
kind: Secret
metadata:
    name: redis-secret-test
    namespace: redis-test
type: Opaque
stringData:
    password: you-guess # Corresponding Redis password
```

#### **Deploying Redis exporter**

On the Deployment management page, click **Create** and select the target **namespace** to deploy the service. You can create in the console. Here, YAML is used to deploy the exporter. Below is a sample YAML configuration:

#### Note:

For more information on the detailed exporter parameters, please see redis\_exporter.



```
apiVersion: apps/v1
kind: Deployment
metadata:
   labels:
        k8s-app: redis-exporter # Rename the exporter based on the business needs. We r
        name: redis-exporter # Rename the exporter based on the business needs. We recomm
```

### 🔗 Tencent Cloud

```
namespace: redis-test
spec:
 replicas: 1
 selector:
   matchLabels:
     k8s-app: redis-exporter # Rename the exporter based on the business needs. We
 template:
   metadata:
     labels:
       k8s-app: redis-exporter # Rename the exporter based on the business needs.
    spec:
     containers:
      - env:
        - name: REDIS_ADDR
         value: ip:port # `ip:port` of the corresponding Redis instance
        - name: REDIS_PASSWORD
         valueFrom:
            secretKeyRef:
              name: redis-secret-test
             key: password
        image: ccr.ccs.tencentyun.com/redis-operator/redis-exporter:1.12.0
        imagePullPolicy: IfNotPresent
        name: redis-exporter
        ports:
        - containerPort: 9121
         name: metric-port # This name is required during scrape task configurati
        securityContext:
          privileged: false
        terminationMessagePath: /dev/termination-log
        terminationMessagePolicy: File
      dnsPolicy: ClusterFirst
      imagePullSecrets:
      - name: gcloudregistrykey
      restartPolicy: Always
      schedulerName: default-scheduler
      securityContext: {}
      terminationGracePeriodSeconds: 30
```

#### Verifying

1. Click the newly created Deployment on the **Deployment** page to enter the Deployment management page.

2. Click the **Log** tab, and you can see that the exporter is successfully started and its address is exposed as shown below:

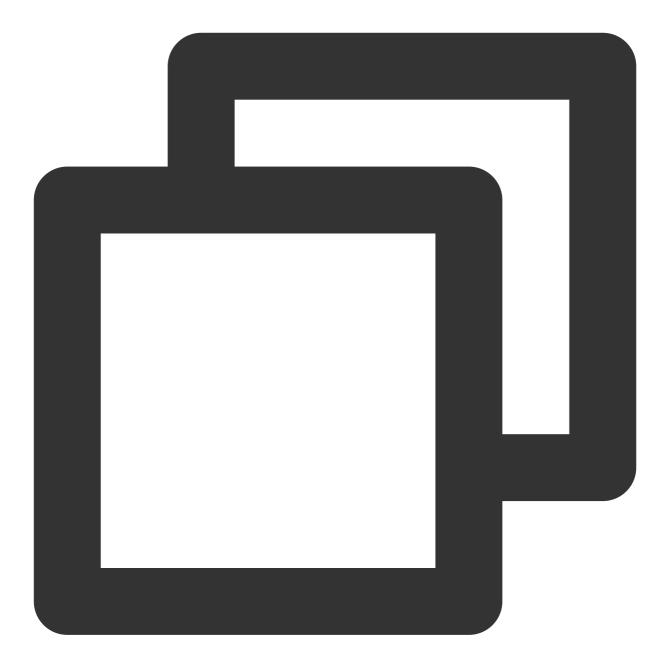


redis-exporter-		edis-exporter	•		<b>.</b>		
1 2020-12-07T1	3:28:24.5722823	93Z time="2020-12	2-07113:28:2	242" level=info	o msg="Redis Metrics	Exporter v1.12.0	bui
		93Z time="2020-12 GOARCH: amd64"	2-07113:28:2	24Z" level=info	o msg="Redis Metrics	Exporter v1.12.0	bui

3. Click the **Pod Management** tab to enter the Pod page.

4. In the **Operations** column on the right, click **Remote Login** to log in to the Pod. Run the following curl command with the address exposed by the exporter in the command line window, and you can get the corresponding Redis metrics normally. If no corresponding data is returned, please check whether <code>REDIS\_ADDR</code> and <code>REDIS\_PASSWORD</code> are correct as shown below:





curl localhost:9121/metrics

The command execution result is as shown below:

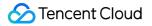
# TYPE redis keyspace hits total counter redis keyspace hits total 29916 HELP redis keyspace misses total counter redis late slow execution duration\_seconds the amount of time needed for last slow execution, in 7 TYPE redis last slow execution duration\_seconds gauge redis last slow execution duration\_seconds used to the last latency spike in seconds 7 TYPE redis latency spike duration seconds Length of the last latency spike in seconds 8 HELP redis latency spike duration seconds (event\_name="command") 0.011 redis latency spike duration seconds (event\_name="command") 0.022 HELP redis latency spike last when the latency spike last occurred F TYPE redis latency spike last when the latency spike last occurred F TYPE redis latency spike last quere "command") 1.604732448e+09 redis latency spike last (event name="command") 1.604736466e+09 HELP redis latest fork seconds gauge redis latency spike last fork seconds gauge redis latest fork seconds latest fork seconds metric F TYPE redis latest fork seconds latest fork seconds metric F TYPE redis latest fork seconds latest seconds hatest fork seconds metric F TYPE redis latest fork seconds latest latest fork seconds latest fork seconds latest F TYPE redis latest fork seconds latest seconds metric F TYPE redis latest fork seconds latest latest seconds latest fork seconds latest F TYPE redis latest fork seconds latest seconds latest seconds latest F TYPE redis latest repl offset seconds latest F TYPE redis latest repl offset seconds latest F TYPE redis master repl offset master repl offset metric # TYPE redis memory repl forset seconds replex metric # TYPE redis memory replex latest semeory max bytes metric # TYPE redis memory max bytes gauge redis memory max bytes gauge redis memory max bytes semeory max bytes metric # TYPE redis memory max bytes semeory max bytes metric # TYPE redis memory used bytes memory used bytes metric # TYPE redis memory used bytes memory

#### Adding scrape task

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click a cluster ID in the TKE cluster list to enter the Integrate with TKE page.

3. In **Scrape Configuration**, add Pod Monitor to define a Prometheus scrape task. Below is a sample YAML configuration:





```
apiVersion: monitoring.coreos.com/v1
kind: PodMonitor
metadata:
   name: redis-exporter # Enter a unique name
   namespace: cm-prometheus # The namespace is fixed. Do not change it
spec:
   podMetricsEndpoints:
        - interval: 30s
        port: metric-port # Enter the name of the corresponding port of the Promethe
        path: /metrics # Enter the value of the corresponding path of the Prometheus
        relabelings:
```

- action: replace
sourceLabels:
- instance
regex: (.*)
targetLabel: instance
replacement: 'crs-xxxxxx' # Change it to the corresponding Redis instance I
- action: replace
sourceLabels:
- instance
regex: (.*)
targetLabel: ip
replacement: '1.x.x.x' # Change it to the corresponding Redis instance IP
namespaceSelector: # Select the namespace where the Pod to be monitored resid
matchNames:
- redis-test
selector: # Enter the label value of the Pod to be monitored to locate the t
matchLabels:
k8s-app: redis-exporter

#### Note:

As the exporter and Redis are deployed on different servers, we recommend you use the Prometheus relabeling mechanism to add the Redis instance information to the monitoring metrics so as to locate problems more easily.

#### Viewing monitoring information

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click **Integration Center** to enter the **Integration Center** page. Find Redis monitoring, install the corresponding Grafana dashboard, and then you can enable the Redis monitoring dashboard to view instance monitoring data as shown below:



#### Integrating with alert feature

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click **Alerting Rule** and add the corresponding alerting rules. For more information, please see Creating Alerting Rule.

# MySQL Exporter Integration

Last updated : 2024-08-07 22:02:03

# Overview

The MySQL exporter is specially designed and developed by the Prometheus community to collect MySQL/MariaDB database monitoring metrics. The exporter reports core database metrics, which can be used for exception alerting and displayed on the monitoring dashboard. TMP supports integration with the MySQL exporter and provides an out-of-the-box Grafana monitoring dashboard.

Currently, the exporter supports MySQL 5.6 or above and MariaDB 10.1 or above. If MySQL or MariaDB is below 5.6 or 10.1 respectively, some monitoring metrics may fail to be collected.

#### Note:

For easier export installation and management, we recommend you use TKE for unified management.

# Prerequisites

You have created a TKE cluster in the region and VPC of your TMP instance and created a namespace for the cluster. You have located and integrated the target TKE cluster in the **Integrate with TKE** section of the **target TMP instance** in the **TMP console**. For more information, please see Agent Management.

## Directions

#### Authorizing in database

As the MySQL exporter monitors a database by querying its status data, you need to grant the exporter access to the corresponding database instance. The account and password should be set based on the actual conditions. The authorization steps are as follows:

1. Log in to the TencentDB for MySQL console.

2. On the instance list page, click the name of the database for which to authorize the exporter to enter the database details page.

3. Select **Database Management** > **Account Management** to enter the account management page and create an account for monitoring based on the actual business needs.

4. Click **Modify Permissions** in the **Operation** column on the right of the account to modify the corresponding permissions as shown below:



			×
	ALTER		ALTER ROUTINE
	CREATE		CREATE ROUTINE
	CREATE TEMPORARY TABLES		CREATE USER
	CREATE VIEW		DELETE
	DROP		EVENT
	EXECUTE		INDEX
	INSERT		LOCK TABLES
~	PROCESS	~	REFERENCES •
	RELOAD	~	REPLICATION CLIENT
	全部		
	100		

You can run the following command for authorization:



CREATE USER 'exporter'@'ip' IDENTIFIED BY 'XXXXXXX' WITH MAX\_USER\_CONNECTIONS 3; GRANT PROCESS, REPLICATION CLIENT, SELECT ON \*.\* TO 'exporter'@'ip';

#### Note:

We recommend you set the allowed maximum number of connections for the account to avoid any impact on the database due to monitoring data collection. However, not all database versions support this configuration, for example, MariaDB 10.1. For more information, please see Resource Limit Options.

#### **Deploying exporter**

1. Log in to the TKE console.

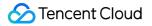
2. Click the ID/name of the cluster whose access credential you want to get to enter the cluster management page.

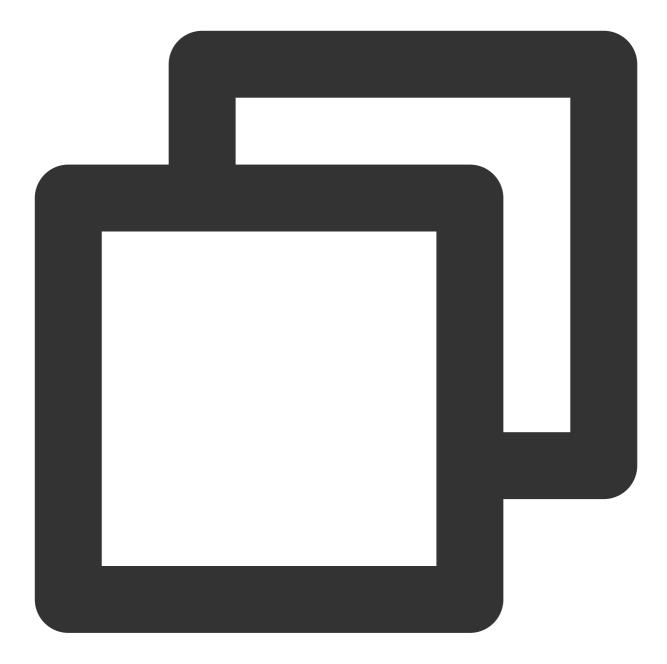
3. Perform the following steps to deploy an exporter: Using Secret to manage MySQL connection string > Deploying MySQL exporter > Verifying.

#### Using Secret to manage MySQL connection string

1. On the left sidebar, select **Workload** > **Deployment** to enter the **Deployment** page.

2. In the top-right corner of the page, click **Create via YAML** to create a YAML configuration as detailed below: You can use Kubernetes Secrets to manage and encrypt connection strings. When starting the MySQL exporter, you can directly use the Secret key but need to adjust the corresponding **connection string**. Below is a sample YAML configuration:



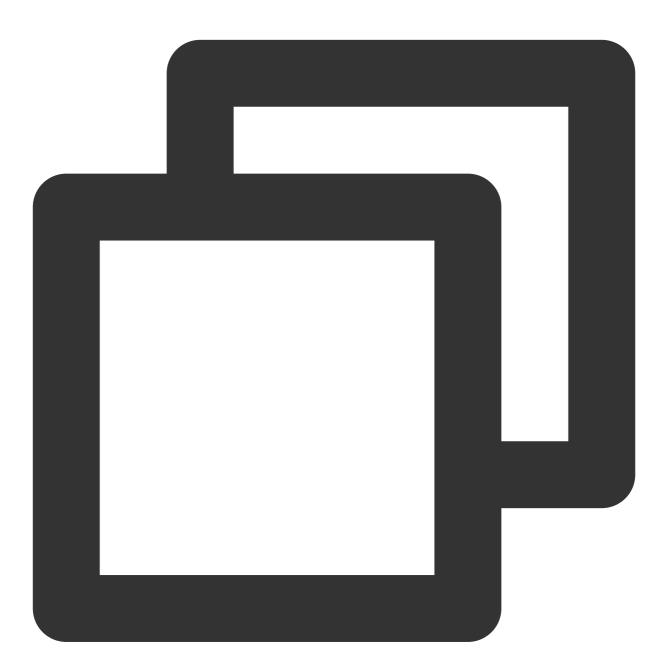


```
apiVersion: v1
kind: Secret
metadata:
    name: mysql-secret-test
    namespace: mysql-demo
type: Opaque
stringData:
    datasource: "user:password@tcp(ip:port)/" # Corresponding MySQL connection strin
```

#### Deploying MySQL exporter

### 🔗 Tencent Cloud

On the Deployment management page, select the target namespace to deploy the service. You can create in the console. Here, YAML is used to deploy the exporter. Below is a sample configuration:



```
apiVersion: apps/v1
kind: Deployment
metadata:
   labels:
        k8s-app: mysql-exporter  # Rename the exporter based on the business needs. We
        name: mysql-exporter  # Rename the exporter based on the business needs. We recom
        namespace: mysql-demo
spec:
```

```
replicas: 1
selector:
  matchLabels:
    k8s-app: mysql-exporter # Rename the exporter based on the business needs. W
template:
 metadata:
    labels:
      k8s-app: mysql-exporter # Rename the exporter based on the business needs.
  spec:
    containers:
    - env:
      - name: DATA_SOURCE_NAME
        valueFrom:
          secretKeyRef:
            name: mysql-secret-test
            key: datasource
      image: ccr.ccs.tencentyun.com/k8s-comm/mysqld-exporter:0.12.1
      imagePullPolicy: IfNotPresent
      name: mysql-exporter
      ports:
      - containerPort: 9104
        name: metric-port
      terminationMessagePath: /dev/termination-log
      terminationMessagePolicy: File
    dnsPolicy: ClusterFirst
    imagePullSecrets:
    - name: qcloudregistrykey
    restartPolicy: Always
    schedulerName: default-scheduler
    securityContext: {}
    terminationGracePeriodSeconds: 30
```

#### Verifying

1. Click the newly created Deployment on the **Deployment** page to enter the Deployment management page.

2. Click the **Log** tab, and you can see that the exporter is successfully started and its address is exposed as shown below:



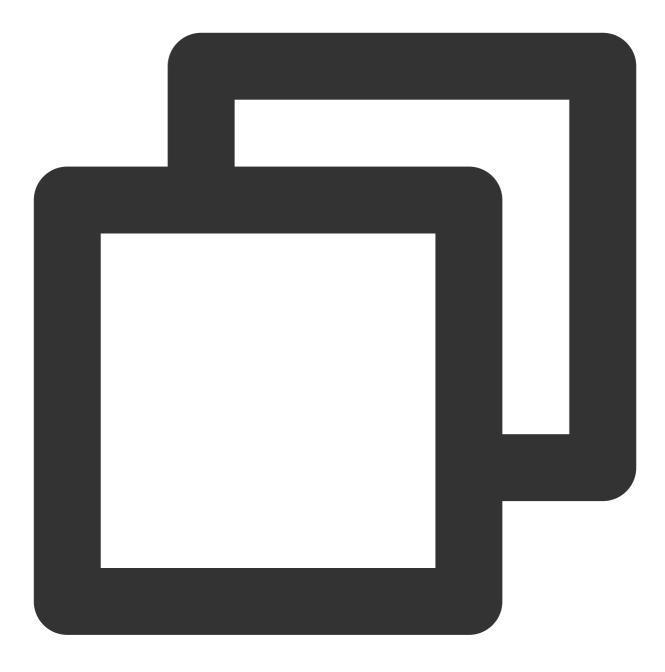
mysql-exporter-54dd5dc589-lz v mysql-exporter v	The second se
1 2020-12-08T09:55:18.315462103Z time="2020-12-08T09:55:18Z"	<pre>level=info msg="Starting mysgld_exporter (version=0.12.1, branch=HEAD, revision=48667bf7c3b438b5e93b2</pre>
source="mysqld_exporter.go:257"	
2 2020-12-08T09:55:18.315532352Z time="2020-12-08T09:55:18Z"	level=info msg="Build context (go=go1.12.7, date=20190729-12:35:58)" source="
3 2020-12-08T09:55:18.315537718Z time="2020-12-08T09:55:18Z"	level=info msg="Enabled scrapers:" source="mysqld_exporter.go:269"
4 2020-12-08T09:55:18.315541954Z time="2020-12-08T09:55:18Z"	level=info msg="collect.global_status" source="mysqld_exporter.go:273"
5 2020-12-08T09:55:18.315546174z time="2020-12-08T09:55:18z"	level=info msg="collect.global_variables" source="mysqld_exporter.go:273"
6 2020-12-08T09:55:18.315549924z time="2020-12-08T09:55:18z"	level=info msg="collect.slave_status" source="mysqld_exporter.go:273"
7 2020-12-08T09:55:18.315748537Z time="2020-12-08T09:55:18Z"	<pre>level=info msg="collect.info_schema.innodb_cmp" source="mysqld_exporter.go:273"</pre>
8 2020-12-08T09:55:18.315765268Z time="2020-12-08T09:55:18Z"	level=info msg="collect.info_schema.innodb_cmpmem" source="mysqld_exporter.go:273"
9 2020-12-08T09:55:18.315770376Z time="2020-12-08T09:55:18Z"	_level=info_msg="collect.info_schema.query_response_time"_source="mysqld_exporter.go:273"
10 2020-12-08T09:55:18.315774561Z time="2020-12-08T09:55:18Z"	level=info msg="Listening on :9104" source="mysqld exporter.go:283"
11	

3. Click the **Pod Management** tab to enter the Pod page.

4. In the **Operations** column on the right, click **Remote Login** to log in to the Pod. Run the following curl

command with the address exposed by the exporter in the command line window, and you can get the corresponding MySQL metrics normally. If no corresponding data is returned, please check whether the **connection string** is correct as shown below:





curl localhost:9104/metrics

The execution result is as shown below:

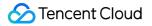
<pre>mysql_info_schema_innodb_cmpmem_pages_used_total {buffer_pool="0",page_size="4096"} 0</pre>
<pre>mysql_info_schema_innodb_cmpmem_pages_used_total {buffer_pool="0",page_size="8192"} 0</pre>
# HELP mysql_info_schema_innodb_cmpmem_relocation_ops_total Number of times a block of the size PAG
<pre># TYPE mysql_info_schema_innodb_cmpmem_relocation_ops_total counter</pre>
<pre>mysql_info_schema_innodb_cmpmem_relocation_ops_total{buffer_pool="0",page_size="1024"} 0</pre>
<pre>mysql_info_schema_innodb_cmpmem_relocation_ops_total{buffer_pool="0",page_size="16384"} 0</pre>
<pre>mysql_info_schema_innodb_cmpmem_relocation_ops_total{buffer_pool="0",page_size="2048"} 0</pre>
<pre>mysql_info_schema_innodb_cmpmem_relocation_ops_total{buffer_pool="0",page_size="4096"} 0</pre>
<pre>mysql_info_schema_innodb_cmpmem_relocation_ops_total{buffer_pool="0",page_size="8192"} 0</pre>
# HELP mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total Total time in seconds spent in
# TYPE mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total counter
<pre>mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total{buffer_pool="0",page_size="1024"} 0</pre>
<pre>mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total{buffer_pool="0",page_size="16384"} 0</pre>
<pre>mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total{buffer_pool="0",page_size="2048"} 0</pre>
<pre>mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total{buffer_pool="0",page_size="4096"} 0</pre>
<pre>mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total{buffer_pool="0",page_size="8192"} 0</pre>
# HELP mysql_up Whether the MySQL server is up.
# TYPE mysql_up gauge
mysql_up 1
# HELP mysql_version_info MySQL version and distribution.
# TYPE mysql_version_info gauge

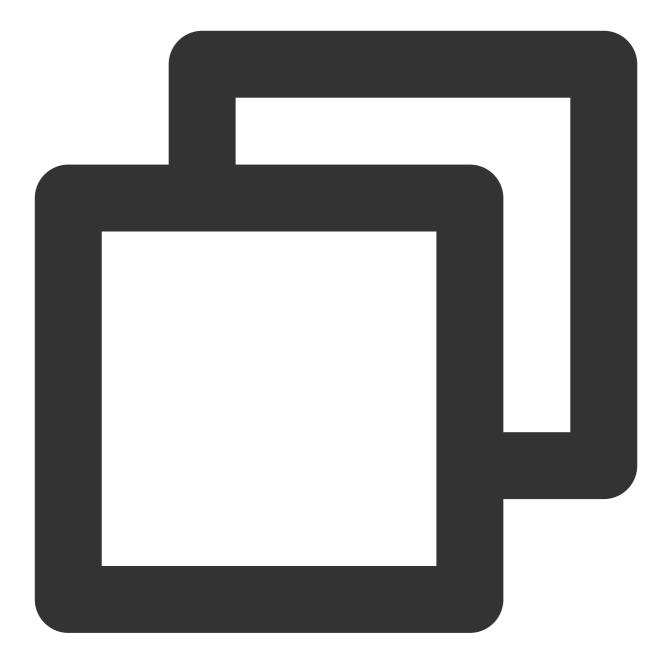
#### Adding scrape task

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click a cluster ID in the TKE cluster list to enter the Integrate with TKE page.

3. In **Scrape Configuration**, add Pod Monitor to define a Prometheus scrape task. Below is a sample YAML configuration:





```
apiVersion: monitoring.coreos.com/v1
kind: PodMonitor
metadata:
   name: mysql-exporter # Enter a unique name
   namespace: cm-prometheus # The namespace is fixed. Do not change it
spec:
   podMetricsEndpoints:
    - interval: 30s
    port: metric-port # Enter the name of the corresponding port of the Promet
   path: /metrics # Enter the value of the corresponding path of the Prometheus
   relabelings:
```

- action: replace
sourceLabels:
- instance
regex: (.*)
targetLabel: instance
replacement: 'crs-xxxxxx' # Change it to the corresponding MySQL instance I
- action: replace
sourceLabels:
- instance
regex: (.*)
targetLabel: ip
replacement: '1.x.x.x' # Change it to the corresponding MySQL instance IP
namespaceSelector: # Select the namespace where the Pod to be monitored resid
matchNames:
- mysql-demo
selector: # Enter the label value of the Pod to be monitored to locate the tar
matchLabels:
k8s-app: mysql-exporter

#### Viewing monitoring information

1. Log in to the TMP console and select the target TMP instance to enter the management page.

2. Click **Integration Center** to enter the **Integration Center** page. Find MySQL monitoring, install the corresponding Grafana dashboard, and then you can enable the MySQL monitoring dashboard to view instance monitoring data as shown below:





#### Integrating with alert feature

TMP has some built-in MySQL alerting rule templates. You can adjust the corresponding thresholds to add alerting rules based on your actual business conditions. For more information, please see Creating Alerting Rule.

Alert strategy / New					
Strategy template	MySQL/MySQL outage				
Oherhamister	Marcol obut dama				
Strategy name *	MySQL Shut down				
Rules PromQL *	mysql_up != 1				
	Click to preview rules 🗹				
duration	1 minu 🔻				
diation					
Alarm notification period 🛈	please choose 🔹				
Alarm Object (Summary) *	MySQL Not running				
Alarm message (Description) *	MySQL Not running, Instance: {{\$labels.ii	o otomo o l	n		
Alarm message (Description) *	MySQL Not running, instance. {{\$labels.li	istance	)) •		
Labels	severity:critical 🛞				
	Key : please enter	Value	please enter	save	
Annotations	Key : please enter	Value	please enter	save	
Alert notification *	Choose a template New 🛂				
	0 notification templates have been selected	I, 3 more	e can be selected		
					0
	Notification template name				Contains operations
			The current n	otification template list is empty, you can select th	e corresponding notification temp
save Cancel					

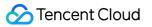
### MySQL Exporter Collection Parameter Description

The MySQL exporter uses various collectors to enable/disable data collection. The specific parameters are as listed below:

Parameter	MySQL Version	Description
collect.auto_increment.columns	5.1	Collects auto_increment columns

### 🔗 Tencent Cloud

collect.binlog_size	5.1	Collects the current size of all registered
collect.engine_innodb_status	5.1	Collects the status data from SHOW EN
collect.engine_tokudb_status	5.6	Collects the status data from SHOW EN
collect.global_status	5.1	Collects the status data from SHOW GL
collect.global_variables	5.1	Collects the status data from SHOW GL
collect.info_schema.clientstats	5.5	If userstat=1 is set, this parameter collection.
collect.info_schema.innodb_metrics	5.6	Collects the monitoring data from info
collect.info_schema.innodb_tablespaces	5.7	Collects the monitoring data from information_schema.innodb_sy
collect.info_schema.innodb_cmp	5.5	Collects the monitoring data of compress information_schema.innodb_cm
collect.info_schema.innodb_cmpmem	5.5	Collects the monitoring data of InnoDB b information_schema.innodb_cm
collect.info_schema.processlist	5.1	Collects the monitoring data of the thread information_schema.processli
collect.info_schema.processlist.min_time	5.1	Minimum time a thread must be in each
collect.info_schema.query_response_time	5.5	Collects query response time distribution to ON.
collect.info_schema.replica_host	5.6	Collects the status data from informa
collect.info_schema.tables	5.1	Collects the status data from informa
collect.info_schema.tables.databases	5.1	Sets the list of databases to collect table
collect.info_schema.tablestats	5.1	If userstat=1 is set, this parameter statistics.
collect.info_schema.schemastats	5.1	If userstat=1 is set, this parameter statistics.
collect.info_schema.userstats	5.1	If userstat=1 is set, this parameter statistics.
collect.perf_schema.eventsstatements	5.6	Collects the monitoring data from



		<pre>performance_schema.events_st</pre>
collect.perf_schema.eventsstatements.digest_text_limit	5.6	Sets the maximum length of the normaliz
collect.perf_schema.eventsstatements.limit	5.6	Limits the number of event statements. C
collect.perf_schema.eventsstatements.timelimit	5.6	Limits how old the 'last_seen' events stat 86400.
collect.perf_schema.eventsstatementssum	5.7	Collects the monitoring data from performance_schema.events_st summed .
collect.perf_schema.eventswaits	5.5	Collects the monitoring data from performance_schema.events_wa
collect.perf_schema.file_events	5.6	Collects the monitoring data from performance_schema.file_summ
collect.perf_schema.file_instances	5.5	Collects the monitoring data from performance_schema.file_summ
collect.perf_schema.indexiowaits	5.6	Collects the monitoring data from performance_schema.table_io_
collect.perf_schema.tableiowaits	5.6	Collects the monitoring data from performance_schema.table_io_
collect.perf_schema.tablelocks	5.6	Collects the monitoring data from performance_schema.table_loc
collect.perf_schema.replication_group_members	5.7	Collects the monitoring data from performance_schema.replicati
collect.perf_schema.replication_group_member_stats	5.7	Collects the monitoring data from performance_schema.replicati
collect.perf_schema.replication_applier_status_by_worker	5.7	Collects the monitoring data from performance_schema.replicati
collect.slave_status	5.1	Collects the monitoring data from SHOW
collect.slave_hosts	5.1	Collects the monitoring data from SHOW
collect.heartbeat	5.1	Collects the monitoring data from heartbe
collect.heartbeat.database	5.1	Database from where to collect heartbea

### 🕗 Tencent Cloud

collect.heartbeat.table	5.1	Table from where to collect heartbeat da
collect.heartbeat.utc	5.1	Uses UTC for timestamps of the current utc ). Default value: false.

#### **Global configuration parameters**

Item	Description
config.my-cnf	Path of .my.cnf file to read MySQL credentials from. Default value: ~/.my.cnf .
log.level	Log level. Default value: info.
exporter.lock_wait_timeout	Sets a lock_wait_timeout (in seconds) on the connection to avoid long metadata locking. Default value: 2.
exporter.log_slow_filter	Adds a log_slow_filter to avoid slow query logging of scrapes. Note: not supported by Oracle MySQL.
web.listen-address	Web port listening address.
web.telemetry-path	Metric API path.
version	Prints the version information.

#### Heartbeat detection

If collect.heartbeat is enabled, mysqld\_exporter will scrape replication delay measured by heartbeat mechanisms.

# **Consul Exporter Integration**

Last updated : 2024-08-07 22:02:11

### Overview

When using Consul, you need to monitor its running status to know whether it runs normally and troubleshoot its faults. TMP provides an exporter to monitor Consul and offers an out-of-the-box Grafana monitoring dashboard for it. This document describes how to use TMP to monitor Consul.

### Directions

- 1. Log in to the TMP console.
- 2. In the instance list, select the corresponding TMP instance.
- 3. Enter the instance details page and click Integration Center.
- 4. Select Consul in the Integration Center and click Install for integration.

#### **Configuration description**

name *	example
Consul ir	istance
address *	192.1.1.1
Label 🛈	+ Add to

Item	Description
Name	Unique integration name



Address	Address and port of the Consul instance to be collected
Label	Label with business meaning, which will be automatically added to Prometheus labels

#### Viewing monitoring information

You can clearly view the following monitoring metrics on the monitoring dashboard:

- 1. Status of Consul cluster nodes.
- 2. Status of services registered in Consul.



# Memcached Exporter Integration

Last updated : 2024-08-07 22:02:18

### Overview

When using Memcached, you need to monitor its running status to know whether it runs normally and troubleshoot its faults. TMP provides an exporter to monitor Memcached and offers an out-of-the-box Grafana monitoring dashboard for it. This document describes how to use TMP to monitor Memcached.

### Directions

- 1. Log in to the TMP console.
- 2. In the instance list, select the corresponding TMP instance.
- 3. Enter the instance details page and click Integration Center.
- 4. Select Memcached in the Integration Center and click Install for integration.

#### **Configuration description**

name *	example
Memcach	ed instance
address *	192.168.1.1:3600
Label 访	+ Add to

Item	Description
Name	Unique integration name

Address	Address and port of the Memcached instance to be collected
Label	Label with business meaning, which will be automatically added to Prometheus labels

#### Viewing monitoring information

You can clearly view the following monitoring metrics on the monitoring dashboard:

1. Memory utilization. The used memory and total memory are also displayed.

2. Current hit rate of Get commands. The hit and miss rates of Get commands during the service operation are also displayed.

3. Old data eviction rate and expired data reclaim rate of Memcached. The total numbers of evictions and reclaims during the service operation are also displayed.

4. Total amount of data stored in Memcached.

- 5. Number of bytes read from and written by the network.
- 6. Current number of open connections.
- 7. Ratio of Get and Set commands during the service operation.
- 8. Current generation rate of each command.





# Integration with Other Exporters

Last updated : 2024-08-07 22:02:28

### Overview

TMP currently provides integration methods for common basic components and corresponding out-of-the-box monitoring dashboards. As TMP is compatible with the native Prometheus, you can also install other exporters available in the community.

### Directions

If there is no integration method available for the basic component you want to use, you can integrate it as follows and customize a monitoring dashboard to meet your monitoring requirements:

1. Find your component in EXPORTERS AND INTEGRATIONS and integrate it as instructed.

2. Refer to the integration method for MySQL.

# CVM Node Exporter

Last updated : 2024-08-07 22:02:34

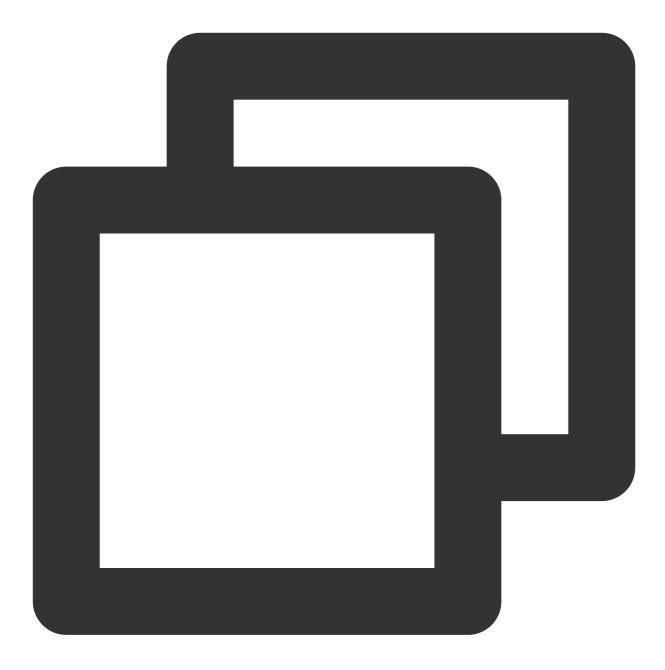
This document describes how to install Node Exporter to expose CVM basic metrics to TMP.

### Directions

#### Step 1. Download and install Node Exporter

Download and install Node Exporter (used to collect basic metric data) in the target CVM instance. Click here or run the following command for download:





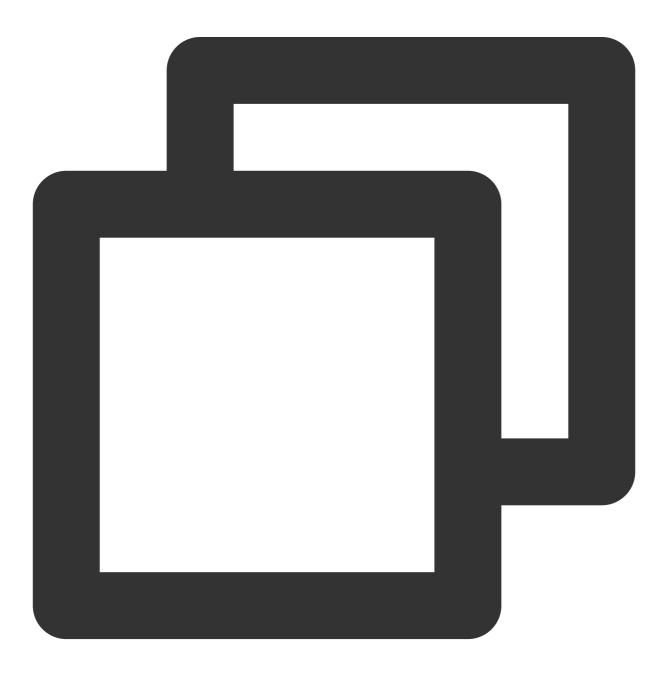
wget https://github.com/prometheus/node\_exporter/releases/download/v1.3.1/node\_expo

The file directory is as follows:

rw-rr	1	3434	3434	11357	Aug	6	2021	LICENSE
rwxr-xr-x	1	3434	3434	18494215	Aug	6	2021	node_exporter
- rw- r r	1	3434	3434	463	Aug	6	2021	NOTICE

#### Step 2. Run Node Exporter to collect basic monitoring data

1. Go to the target folder and run Node Exporter.





```
cd node_exporter-1.3.1.linux-amd64
./node_exporter
```

If the following result is displayed, basic monitoring data has been collected successfully.

rw-rr1 3434 3434 463 Aug 6 2021 NUIICE
root@VM-0-7-centos node_exporter-1.2.2.linux-amd64]# ./node_exporter
evel=info ts=2022-02-11T07:15:26.555Z caller=node exporter.go:182 msg="Starting node_exporter" version="(version=1.2.2, br
pn=26645363b486e12be40af7ce4fc91e731a33104e)"
.evel=info ts=2022-02-11T07:15:26.555Z caller=node_exporter.go:183 msg="Build context" build_context="(go=go1.16.7, user=rc
late=20210806-13:44:18)"
evel=warn ts=2022-02-11T07:15:26.555Z caller=node_exporter.go:185 msg="Node Exporter is running as root user. This exporte.
run as unpriviledged user, root is not required."
evel=info ts=2022-02-11T07:15:26.555Z caller=filesystem_common.go:110 collector=filesystem msg="Parsed flagcollector.fi
.nts-exclude" flag=^/(dev proc sys var/lib/docker/.+)(\$ /)
.evel=info ts=2022-02-11T07:15:26.555Z caller=filesystem_common.go:112 collector=filesystem msg="Parsed flagcollector.fi
exclude" flag=^(autofs binfmt_misc bpf cgroup2? configfs debugfs devpts devtmpfs fusectl hugetlbfs iso9660 mqueue nsfs ove
<pre>store rpc_pipefs securityfs selinuxfs squashfs sysfs tracefs)\$</pre>
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:108 msg="Enabled collectors"
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=arp
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=bcache
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=bonding
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=btrfs
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=conntrack
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=cpu
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=cpufreq
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=diskstats
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=edac
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=entropy
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=fibrechannel
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=filefd
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=filesystem
evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=hwmon
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=infiniband

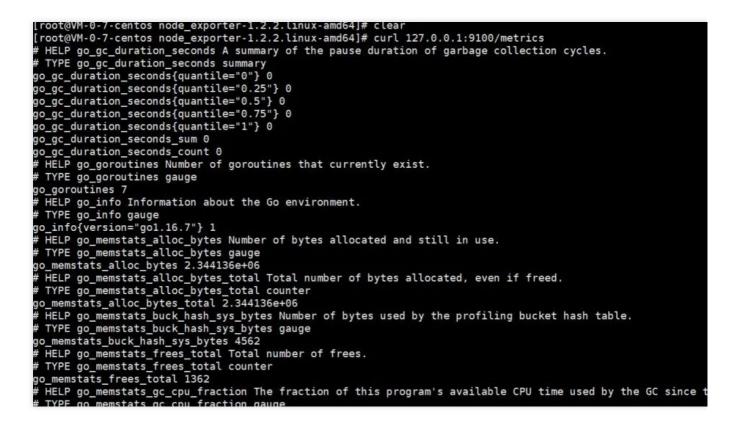
2. Run the following command to expose the basic monitoring data to port 9100:





curl 127.0.0.1:9100/metrics

You can see the following metric monitoring data that is exposed after the command is executed.



#### Step 3. Configure the collection

Log in to the TMP console, select Integration Center > CVM, and configure the information in Task Configuration as prompted.

Below is a sample configuration of a scrape task:





```
job_name: example-job-name
metrics_path: /metrics
cvm_sd_configs:
- region: ap-guangzhou
ports:
    - 9100
filters:
    - name: tag: Sample tag key
    values:
        - Sample tag value
relabel_configs:
```

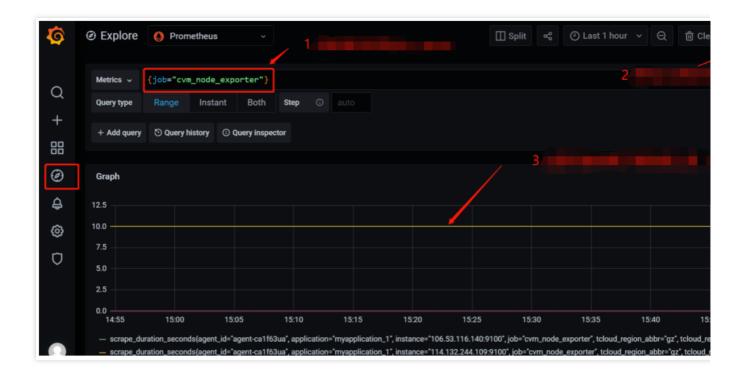


```
source_labels: [__meta_cvm_instance_state]
regex: RUNNING
action: keep
regex: __meta_cvm_tag_(.*)
replacement: $1
action: labelmap
source_labels: [__meta_cvm_region]
target_label: region
action: replace
```

#### Step 4. Check whether data is reported successfully

Log in to the TMP console and click the Grafana icon to enter Grafana.

Search for {job="cvm\_node\_exporter"} in **Explore** to see whether there is data, and if so, data is reported successfully.



#### Step 5. Configure the dashboard

Every product has some existing JSON files that can be directly imported into the dashboard.

1. **Download a dashboard file**: Go to the **Dashboard** page, search for <u>node\_exporter</u>, and select the latest dashboard for download.

🧔 Grafana Labs	Products 🗸	Open source 🗸	Learn 🗸	Company 🗸	Downloads	Contact us
All dashboards »	Node Exporter Full					
	ode Exporter Fu	// by rfraile				Down
	DASHBOARD t updated: 3 days ago	_				
Sta	rt with Grafana Cloud and tl	he new FREE tier. Incl	udes 10K series	Prometheus or Graphite M	etrics and 50gb Loki Logs	Ad
Overview	Revisions Review	S	-		Get this o	dashboard
					1860	
·····					Copied! C	lick to copy a
Nearly all default	values exported by Promet	heus node exporter	graphed.		Download	

2. Import a JSON file into the dashboard: Log in to the TMP console, select Basic Info > Grafana Address to enter Grafana. In the Grafana console, select Create > Import and upload the dashboard file in Upload JSON file.

Ø	Import dashboard from file or Grafana.com
Q	Options
	Name
+	Node Exporter Full
	Folder
Ø	General ~
¢	<b>Unique identifier (uid)</b> The unique identifier (uid) of a dashboard can be used for uniquely identify a dashboard
Ą	between multiple Grafana installs. The uid allows having consistent URL's for accessing dashboards so changing the title of a dashboard will not break any bookmarked links to
ŝ	that dashboard.
D	rYdddlPWk1
$\bigcirc$	Prometheus
	Prometheus
	Import Cancel

Ô	器 ′ Node Ex	porter(cvm) ☆ ペ		the	• B © <del>Q</del>	(2) Last 24	hours
Q	datasource default - Job node-exporter - Quick CPU / Mem / Disk	- Host:					
+ 8000000000000000000000000000000000000			RAM Used	SWAP Used	<sup>1</sup> Root FS Used	<sup>i</sup> .:PU C 1 <sup>i</sup> .kootF 49 GiB	i 
₽ ©	<ul> <li>Basic CPU / Mem / Net / Disk</li> <li>i CPU Basic</li> </ul>	c			Memory	Basic	
	100% 75% 50% 25%			954 MiB 715 MiB 477 MiB 238 MiB	nn a lanna lannar		

# Health Check

Last updated : 2024-08-07 22:02:40

### Overview

Health check detects the service connectivity on a regular basis to monitor the service health, helping you stay up to date with the service health in real time and promptly discover exceptions to improve the SLA.

### Directions

- 1. Log in to the TMP console.
- 2. In the instance list, select the corresponding TMP instance.
- 3. Enter the instance details page and click Integration Center.
- 4. Select Health Check in Integration Center to configure the detection of the corresponding service.

#### **Detection description**



ntegration list / Nev	,	
• The number	of remaining IPs in the current subnet	[ 2221 ]· 238
		[222]]. 200
Detect		
name *	ping-pp	
Probe configurat	ion	
Detection method *	http_get 💌	
Detection target *	https://console.cloud.tencent.com	×
	+ Add to	
Label 🚯	+ Add to	
save	Will incur additional costs ,	billing overview 🔼

Parameter	Description
Name	Unique detection task name, which corresponds to the detection group on the Grafana monitoring dashboard
Detection Method	Currently, the following detection methods are supported: http_get http_post tcp ssh ping
Detection Target	Address of the service to be detected
Label	Label with business meaning, which will be automatically added to Prometheus labels

#### Viewing monitoring information

You can clearly view the following status on the monitoring dashboard:

- 1. Service access latency and health status.
- 2. Latency in each processing phase of service access.
- 3. Expiration time of certificate in case of HTTPS
- 4. Status of various detection types.



# **Cloud Monitoring**

Last updated : 2024-08-07 22:02:48

### Overview

TMP collects, stores, and visualizes the basic monitoring data of Tencent Cloud products.

### Directions

- 1. Log in to the TMP console.
- 2. In the instance list, select the corresponding TMP instance.
- 3. Enter the instance details page and click Integration Center.

4. Select **Cloud Monitoring** in the integration center. Define the integration name, configure the Exporter, and select the corresponding cloud product.

#### Note

Time offset: Select the time range. End time = current time - time offset.

The integrated monitoring data contains the tag data of the selected cloud products. Any Chinese or special

characters in the tag will be filtered out.

Data collection frequency in this module: 1 minute.

# Read Cloud-Hosted Prometheus Instance Data via Remote Read

Last updated : 2024-08-07 22:02:54

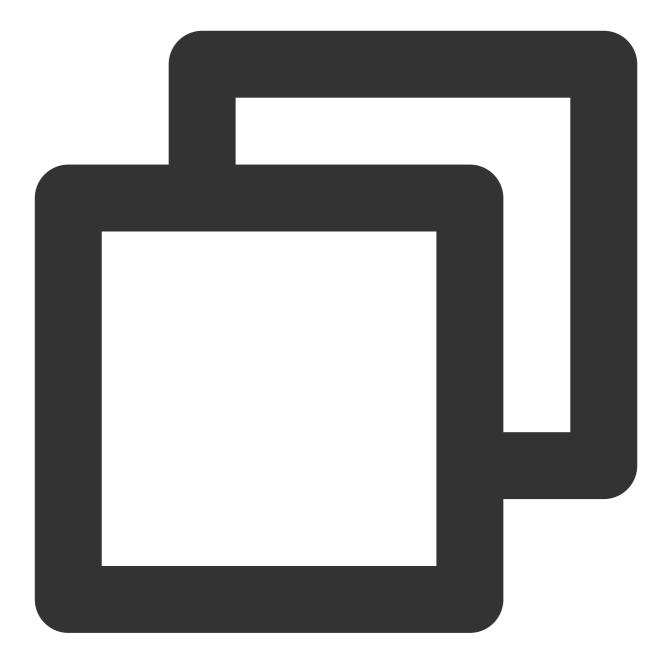
### Overview

TMP provides the remote read API, which supports organizing a series of data sources of the Prometheus protocol into a single data source for query. This document describes how to use self-built Prometheus to read data from a cloud-managed TMP instance through the remote read API.

### **Remote Read Configuration**

The recommended configuration for prometheus.yml is as follows:





```
remote_read:
- url: 'http://prom_ip:prom_port/api/v1/read'
    read_recent: true
    basic_auth:
        username: app_id
        password: token
```

It is recommended to use the Basic Auth method to access the cloud-managed TMP instance. The username is the account AppID and the password is the token obtained on **Basic Info** > **Service Address** in the Prometheus console.



Service Address				
Token	***** 🗗			
Remote Write Address	http://1			: 6
Remote Read Address	http:/		ſ	<b>D</b>
HTTP API	http		Б	
Pushgateway Address		Б		

### Note

**Configure** global:external\_labels carefully for TMP instances with remote read enabled:

As external\_labels will be appended to the query condition of remote read, an inaccurate label may prevent you from querying the necessary data.

The filter\_external\_labels: false configuration item can avoid adding external\_labels to the query condition (supported in v2.34 and later).

#### Avoid identical series:

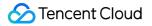
For two identical series, TMP will randomly select a series value at each time point to form a new series as the query result during query merging, which will lead to inaccurate query results.

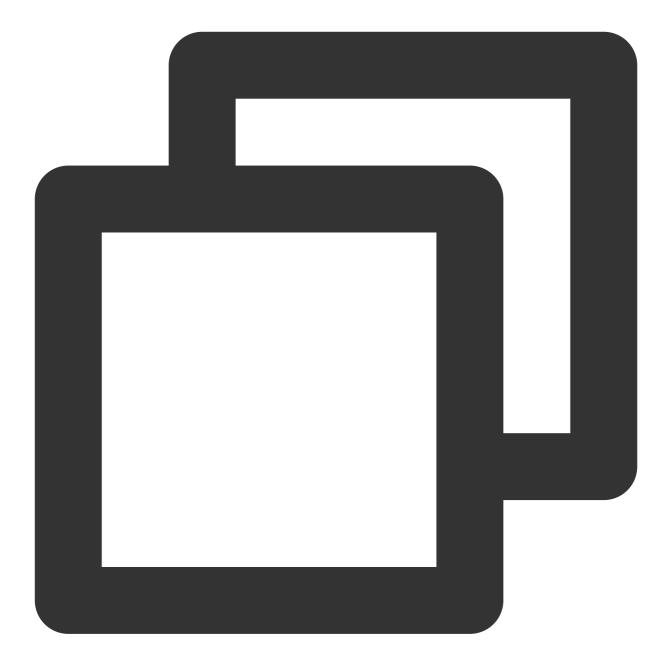
Since there is no multi-copy redundant storage in the design concept of TMP, identical series will not be supported.

### **Remote Read Configuration Items**

#### Note

The configuration items in [] are optional. This document shows Prometheus v2.40 configuration, and some configuration items may be missing in lower versions. For more information, see Prometheus official documentation.





```
# The API address of the target TMP instance for remote read
url: <string>
# Identify a unique remote read configuration name
[ name: <string> ]
# The PromQL must contain the following label filter conditions to perform remote r
required_matchers:
    [ <labelname>: <labelvalue> ... ]
# The timeout for remote read query
```

```
[ remote_timeout: <duration> | default = 1m ]
# Customize the headers attached to the remote read request. You can't overwrite th
headers:
  [ <string>: <string> ... ]
# Whether to perform remote read query in the time range with complete local data s
[ read recent: <boolean> | default = false ]
# Add Authorization header to each remote read request, and choose password or pass
basic auth:
  [ username: <string> ]
  [ password: <secret> ]
  [ password_file: <string> ]
# Customize authorization header configuration
authorization:
  # Authentication type
  [ type: <string> | default: Bearer ]
  # Authentication key. You can choose credentials or credentials_file.
  [ credentials: <secret> ]
  # Get the key from the file
  [ credentials_file: <filename> ]
# OAuth2.0 authentication, which cannot be used with basic_auth authorization at th
oauth2:
  [ <oauth2> ]
# TLS configuration
tls_config:
  [ <tls_config> ]
# Proxy URL
[ proxy_url: <string> ]
# Query whether the request accepts 3XX redirection
[ follow_redirects: <boolean> | default = true ]
# Whether to enable HTTP2
[ enable_http2: <bool> | default: true ]
# Whether to append `external_labels` for remote read
[ filter_external_labels: <boolean> | default = true ]
```

# Agent Self-Service Access

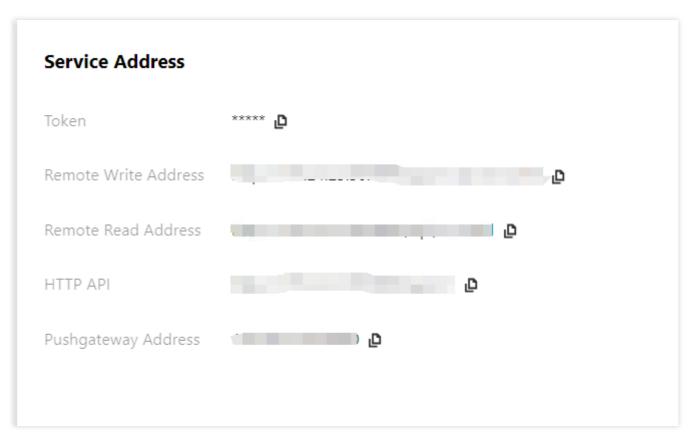
Last updated : 2024-08-15 16:10:55

### Application Scenario

To collect services on self-built IDC, deploy Agent and manage collection configurations, and report monitoring data to the cloud TMP. For cloud services, we recommend using Integration Center, which will manage Agent, offering automated integration for multiple middlewares and scraping tasks.

### **Obtaining Prometheus Instance Access Configuration**

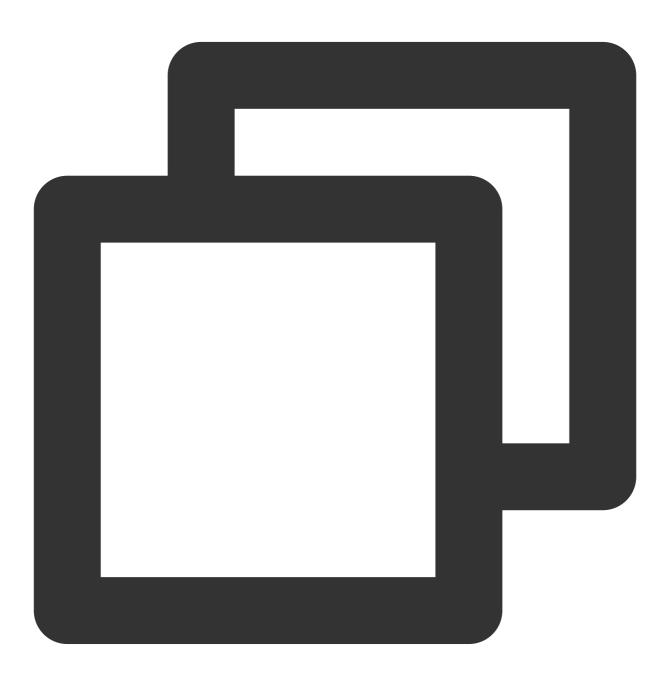
1. Go to Prometheus Monitoring Console, select the corresponding instance ID/Name, and on the **Basic Info** > **Service Address** page, obtain the Remote Write address and Token.



2. Obtain APPID on the Account Information page.

# Confirming the Network Environment and Connectivity with Cloud Instances

Based on the acquired RemoteWrite address, execute the following command. If the network is connected, the returned information will include 401 Unauthorized .



curl -v -X POST \${RemoteWriteURL}

### Installing and Starting vmagent

vmagent uses fewer resources and is widely used due to its compatibility with Prometheus collection configuration and Remote Write protocol. This document only describes common startup options for vmagent, managed through Systemd or Docker. For more detailed information, please see the official documentation.

#### **Common Startup Options**

-promscrape.noStaleMarkers: If the collection target disappears, a stale marker for all associated metrics is generated and written to remote storage by default. Setting this option disables this behavior and can reduce memory usage. -loggerTimezone: The time zone for the time in logs, for example, Asia/Shanghai, Europe/Berlin or Local (UTC by default).

-remoteWrite.tmpDataPath: The file path for temporary data storage to be written after collection.

-remoteWrite.url: The URL where data is written to remote storage.

-remoteWrite.basicAuth.username: Remote storage -remoteWrite.url corresponding basic auth username.
-remoteWrite.basicAuth.password: Remote storage -remoteWrite.url corresponding basic auth password.
-promscrape.config: Path of the collection configuration, which can be a file path or HTTP URL. For more details, please see Reference Documentation.

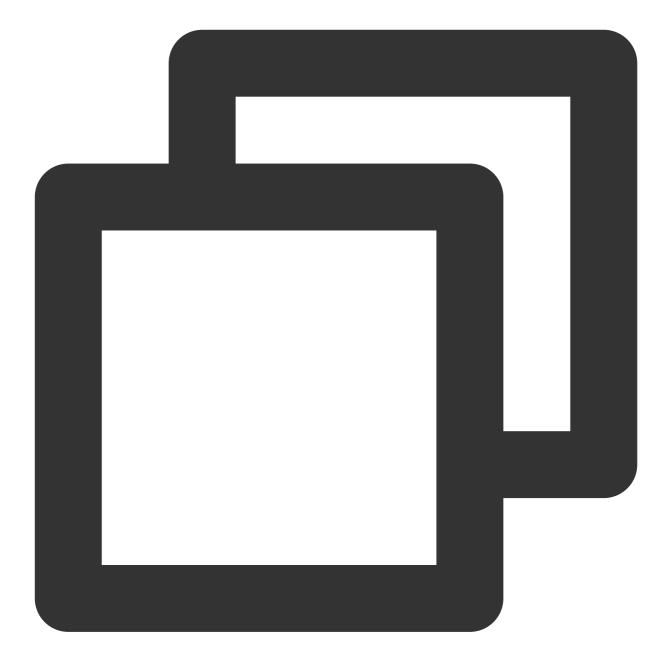
-promscrape.configCheckInterval: Interval for checking the -promscrape.config configuration changes. For configuration updates, please see Reference Documentation.

#### Managing via Docker

1. On the vmagent Release Page, select the image version. It is recommended to use latest.

2. Replace the Prometheus instance information in the script and start vmagent.

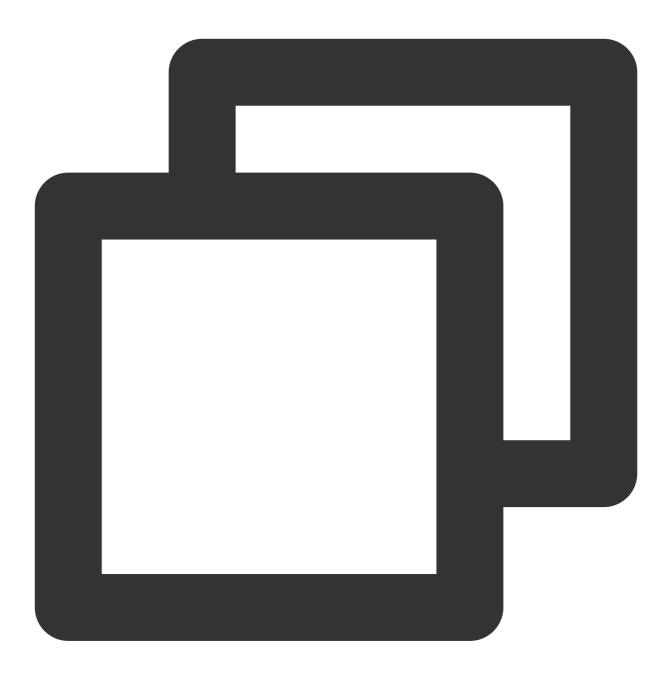




```
mkdir /etc/prometheus
touch /etc/prometheus/scrape-config.yaml
docker run -d --name vmagent --restart always --net host -v /etc/prometheus:/etc/pr
-promscrape.noStaleMarkers \\
-loggerTimezone=Local \\
-remoteWrite.url="${RemoteWriteURL}" \\
-remoteWrite.basicAuth.username="${APPID}" \\
-remoteWrite.basicAuth.password='${Token}' \\
-remoteWrite.tmpDataPath=/var/lib/vmagent \\
-promscrape.config=/etc/prometheus/scrape-config.yaml \\
-promscrape.configCheckInterval=5s
```



#### 3. View vmagent logs



docker ps docker logs vmagent

If it starts normally, executing the following command will return  $\hfill OK$  .





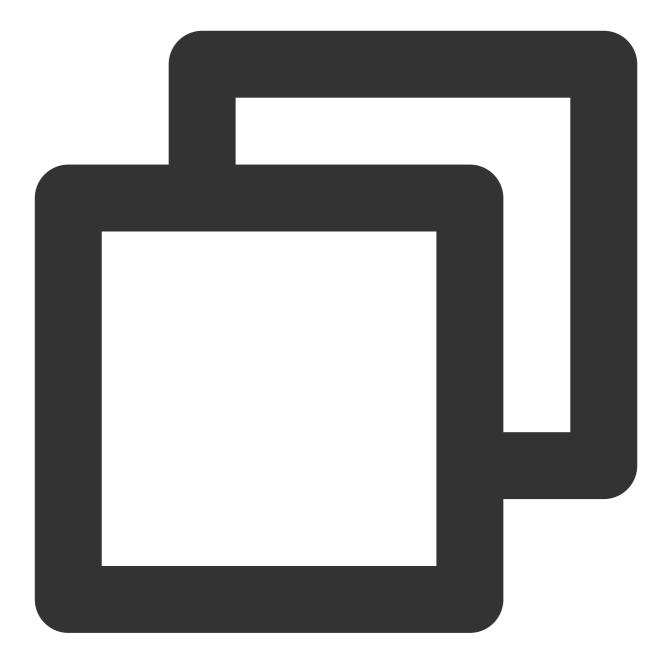
curl localhost:8429/health

#### Managing via Systemd

1. On the vmagent Release page, download the corresponding vmutils-\* compressed package according to your operating system and CPU architecture, and decompress it.

2. Replace the access information of the Prometheus instance in the script and start vmagent.





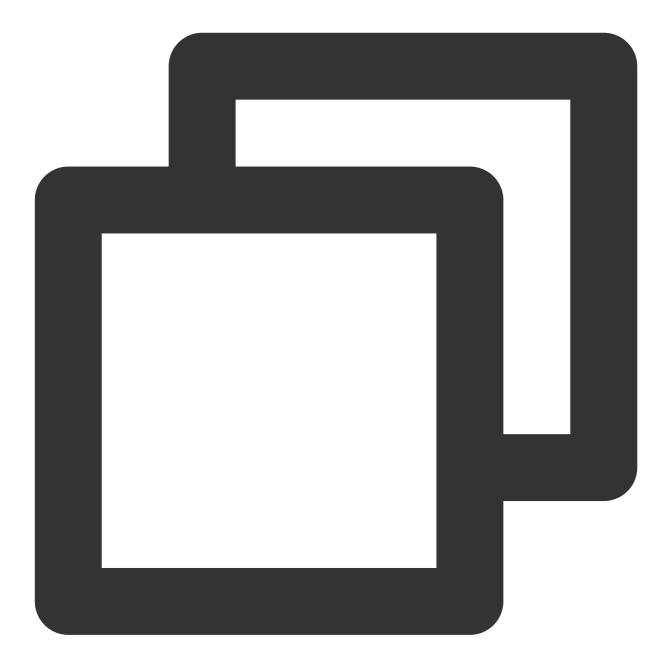
```
mkdir /etc/prometheus
touch /etc/prometheus/scrape-config.yaml
cat >/usr/lib/systemd/system/vmagent.service <<EOF
[Unit]
Description=VictoriaMetrics Agent
After=network.target
[Service]
```

```
LimitNOFILE=10240
ExecStart=/usr/bin/vmagent \\
-promscrape.noStaleMarkers \\
```

```
-loggerTimezone=Local \\
-remoteWrite.url="${RemoteWriteURL}" \\
-remoteWrite.basicAuth.username="${APPID}" \\
-remoteWrite.basicAuth.password="${Token}" \\
-remoteWrite.tmpDataPath=/var/lib/vmagent \\
-promscrape.config=/etc/prometheus/scrape-config.yaml \\
-promscrape.configCheckInterval=5s
Restart=always
RestartSec=10s
[Install]
WantedBy=multi-user.target
EOF
systemctl daemon-reload
systemctl enable vmagent
systemctl start vmagent
sleep 3
systemctl status vmagent
```

3. View logs





journalctl -u vmagent

If it starts normally, executing the following command will return  $\hfill \ensuremath{\mathsf{OK}}\xspace$  .





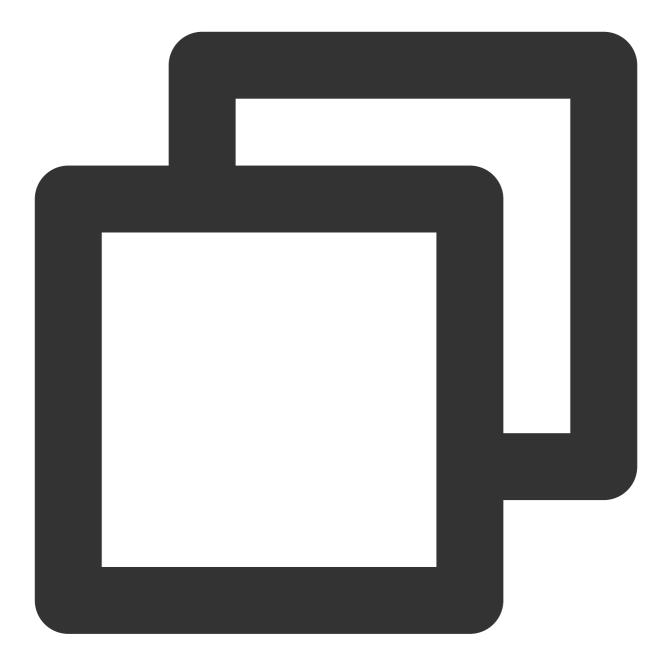
curl localhost:8429/health

### Managing the Configuration

### Modifying the Configuration File

Edit the collection configuration file /etc/prometheus/scrape-config.yaml to add/update/delete collection tasks. For Prometheus collection task configuration, see Official Documentation.





After the configuration is modified, it will only take effect after the time set by the option

```
promscrape.configCheckInterval .
```



### **Viewing Monitoring Target Information**

Execute the following command to view the collection target and check whether the configuration is effective and meets expectations.



curl localhost:8429/api/v1/targets

# Instructions for Installing Components in the TKE Cluster

Last updated : 2024-07-23 18:11:09

### Overview

This document describes the features, use permissions, and resource consumption of various components installed in the user's TKE cluster during the TKE Integration process of TMP.

### proxy-agent

### **Component Overview**

The TKE cluster has independent network environment. Therefore, the proxy-agent is deployed within the cluster to provide access proxies for collection components outside the cluster. On one hand, external collection components discover resources within the cluster through the proxy-agent service; on the other hand, they scrape metrics through the proxy-agent and write them to the time series storage of the Prometheus instance.

### **Resource Objects Deployed in the Cluster**

Namespace	Kubernetes Object Name	Туре	Resource Amount	Description
<prometheus instance ID&gt;</prometheus 	proxy-agent	Deployment	0.25C256Mi*2	Collection proxy
<prometheus instance ID&gt;</prometheus 	<prometheus instance ID&gt;</prometheus 	ServiceAccount	-	Permission carrier
-	<prometheus instance ID&gt;</prometheus 	ClusterRole	-	Collection permissions related
-	<prometheus instance ID&gt;-crb</prometheus 	ClusterRoleBinding	-	Collection permissions related

### **Component Permission Description**

#### **Permission Scenarios**

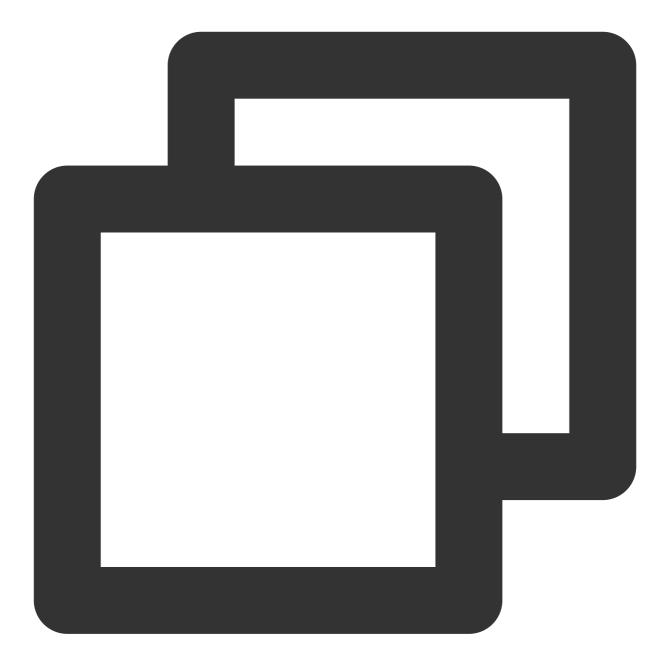
Feature Involved Objects Involved
-----------------------------------



		Operati Permis:
Collection configuration management	scrapeconfigs,servicemonitors,podmonitors,probes,configmaps,secrets,namespaces	get/list/
Service discovery	services,endpoints,nodes,pods,ingresses	get/list/v
Scraping some system component metrics	nodes/metrics,nodes/proxy,pods/proxy	get/list/
Scraping metrics with RBAC authentication	/metrics,/metrics/cadvisor	get

#### **Permission Definition**





```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
   name: prom-instance
rules:
        - apiGroups:
            - monitoring.coreos.com
        resources:
            - scrapeconfigs
            - servicemonitors
            - podmonitors
```



- probes
- prometheuses
- prometheusrules

verbs:

- get
- list
- watch

- apiGroups:

\_ ""

```
resources:
```

- namespaces
- configmaps
- secrets
- nodes
- services
- endpoints
- pods

verbs:

- get
- list
- watch
- apiGroups:

```
- networking.k8s.io
```

- resources:
  - ingresses

verbs:

- get
- list

```
- watch
```

```
- apiGroups: [ "" ]
```

```
resources:
```

- nodes/metrics
- nodes/proxy

```
- pods/proxy
```

verbs:

```
- get
```

- list
- watch
- nonResourceURLs: [ "/metrics", "/metrics/cadvisor" ]
  - verbs:
    - get

### tke-kube-state-metrics



#### **Component Overview**

tke-kube-state-metrics uses the open-source component kube-state-metrics, listens to the cluster's API server, and generates status metrics for various objects within the cluster.

### **Resource Objects Deployed in the Cluster**

Namespace	Kubernetes Object Name	Туре	Resource Amount	Description
kube- system	tke-kube-state- metrics	Statefulset	0.5C512Mi	Collection program
kube- system	tke-kube-state- metrics	ServiceAccount	-	Permission carrier
-	tke-kube-state- metrics	ClusterRole	-	Collection permissions related
-	tke-kube-state- metrics	ClusterRoleBinding	-	Collection permissions related
kube- system	tke-kube-state- metrics	Service	-	Collection agent corresponding service, for service discovery use
kube- system	tke-kube-state- metrics	ServiceMonitor	-	Collection configuration
kube- system	tke-kube-state- metrics	Role	-	Shard collection permission related
kube- system	tke-kube-state- metrics	RoleBinding	-	Shard collection permission related

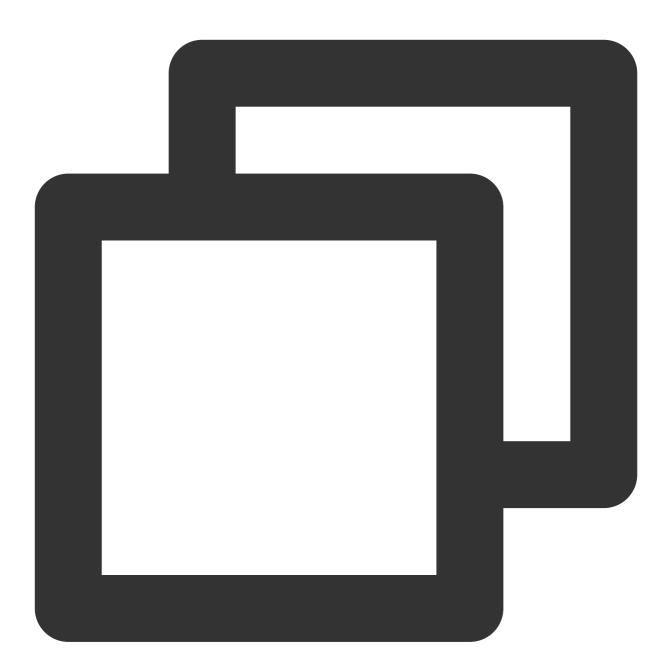
### **Component Permission Description**

#### Permission Scenarios

Feature	Involved Objects	Involved Operation Permissions
Listening to the status of various resources in the cluster	Most Kubernetes resources	list/watch
Get the shard number of the collection pod	statefulsets, pods	get



#### **Permission Definition**



🕗 Tencent Cloud

- secrets
- nodes
- pods
- services
- serviceaccounts
- resourcequotas
- replicationcontrollers
- limitranges
- persistentvolumeclaims
- persistentvolumes
- namespaces
- endpoints
- verbs:
  - list
  - watch
- apiGroups:
  - apps

```
resources:
```

- statefulsets
- daemonsets
- deployments
- replicasets
- verbs:
  - list
  - watch
- apiGroups:
  - batch

```
resources:
```

- cronjobs
- jobs
- verbs:
  - list
    - watch
- apiGroups:

```
- autoscaling
```

resources:

```
- horizontalpodautoscalers
```

- verbs:
  - list
    - watch
- apiGroups:
  - authentication.k8s.io
  - resources:

```
- tokenreviews
```

```
verbs:
```

```
- create
- apiGroups:
```

```
- authorization.k8s.io
 resources:
   - subjectaccessreviews
 verbs:
   - create
- apiGroups:
   - policy
  resources:
   - poddisruptionbudgets
 verbs:
   - list
    - watch
- apiGroups:
   - certificates.k8s.io
 resources:
   - certificatesigningrequests
 verbs:
   - list
    - watch
- apiGroups:
   - storage.k8s.io
  resources:
   - storageclasses
    - volumeattachments
 verbs:
   - list
    - watch
- apiGroups:
   - admissionregistration.k8s.io
  resources:
   - mutatingwebhookconfigurations
    - validatingwebhookconfigurations
 verbs:
   - list
   - watch
- apiGroups:
   - networking.k8s.io
 resources:
   - networkpolicies
    - ingresses
 verbs:
   - list
    - watch
- apiGroups:
   - coordination.k8s.io
 resources:
   - leases
```

```
verbs:
     - list
      - watch
  - apiGroups:
     - rbac.authorization.k8s.io
    resources:
      - clusterrolebindings
      - clusterroles
     - rolebindings
      - roles
   verbs:
     - list
     - watch
___
kind: Role
metadata:
 name: tke-kube-state-metrics
 namespace: kube-system
rules:
 - apiGroups:
      _ ""
    resources:
     - pods
   verbs:
     - get
  - apiGroups:
      - apps
   resourceNames:
     - tke-kube-state-metrics
    resources:
      - statefulsets
   verbs:
      - get
```

### tke-node-exporter

### **Component Overview**

tke-node-exporter uses the open-source project node\_exporter, deployed on each node in the cluster to collect hardware and Unix-like operating system metrics.

### **Resources Deployed in the Cluster**

Namespace	Kubernetes	Туре	Resource Amount	Description

©2013-2022 Tencent Cloud. All rights reserved.

	Object Name			
kube- system	tke-node- exporter	DaemonSet	0.1C180Mi*node amount	Collection program
kube- system	tke-node- exporter	Service	-	Collection program corresponding service, for service discovery use
kube- system	tke-node- exporter	ServiceMonitor	-	Collection configuration

### **Component Permission Description**

This component does not use any cluster permissions.

# Security Group Open Description

Last updated : 2024-08-15 16:32:27

### Overview

This document describes the port that needs to be opened for security groups of managed clusters and user clusters during the process of integrating TKE for TMP. It also describes solutions for security group related issues that arise when managed clusters and user clusters are bound.

### Managed Cluster

Managed cluster Security Groups are created by TMP and generally do not need modifications.

### Security Group

Rule	Protocol Port	Policy
Inbound rule	TCP:9093, 9090, 10901, 10902, 9990, 3000, 8080, and 8008	Allow
Inbound rule	TCP:8100-8200	Allow
Outbound rule	ALL	Allow

#### **Port Description**

Port	Function	Remarks
TCP:8008	proxy-server listens for the proxy- agent connection port	-
TCP:8080	Cluster internal API calls port	-
TCP:3000	grafana proxy port	-
TCP:9990	cm-notify synchronization port	About to be decommissioned
TCP:10901,10902	thanos sidecar listening address	-
TCP:9090	Configure reload port, and collect data query API	-



TCP:9093	Alarm port	-
TCP:8100-8200	proxy-server listening collection port	Since the collection port range is 100, the maximum number of associated clusters cannot exceed 100.

#### **Viewing Method**

log in to Prometheus Monitoring, select the instance's ID/Name > instance diagnostics, choose Integration Center for diagnostics, in the data collection architecture diagram you can see the Managed Cluster Security Group, click it to jump to the security group interface via hyperlink to view the Managed Cluster Security Group.

diagnostic collec	tion Integration Center	
resource utilization	Number of Pods 4/4 (1.75core 3.5 G)	data collection architecture diagram $\red{gram}$
collection configuration		Managed cluster for data collection components
Target allocation status	Allocated4items Not allocated0items	Available IP: 235 Security Group :
Target status	1 Up 1 Down	Number of Pods <b>7/7</b> (4.25 core 6.75 G)
Agent status	1 items	
Version	tmp-agent(v1.1.2) tmp-operator(v1.1.9) proxy-server(v1.0.8->v1.0.7)	

### User Cluster

The user cluster security group is specified when the user creates a node. If not specified, the default security group will be used.

### Security Group

Rule	Cluster Type	Protocol Port	Policy	Description
Outbound rule	-	TCP:8008	Allow	Ensure that the proxy-agent and proxy- server can establish a connection
Inbound	Standard	-		The standard cluster does not need



rule	cluster			opening ports.
Inbound rule	Independent cluster	TCP: 9092, 8180, 443, 10249, 9100, 60002, 10252, 10257, 10259, and 10251	Allow	The independent cluster needs to open additional master node-related ports to ensure proxy-agent can pull master node- related monitoring data

#### **Viewing Method**

log in to Prometheus Monitoring, select the instance ID/Name > **Data Collection**, and click the cluster ID/Name to jump to the cluster's TKE interface.

#### **Native Nodes**

Click **Node Management** > Worker Node > Node pool, and click Node Pool ID. In the Details page, you can see the security group. In the **Security group**, search by security group ID to view specific rules.

Node list Det	ails Operation logs	
Node pool inforr	nation	
Node pool name	np-emkhjrii(test)	Number of node
Node pool status	Running	Time created
Maintenance level	Medium	Deletion Protecti
K8s version	1.26.1-tke.7	Security reinforce
		Tag
Node launch con	figuration info	
Operating system	TencentOS Server 3.1	Model
Billing mode	Pay-as-you-go 🎤	System disk
Supported subnets	3 🖉	Data disk
Security group		Node name()
Bind an SSH key	· · · ·	
Operation config	uration	
Node self-heal	Enabled	Auto scaling 🛈
Calf has line and a	wudi	Scale-out policy
Self-healing rule		

#### **Common Nodes**



Click **Node Management** > **Worker Node** > **Node Pool**, and click Node Pool ID. In the Details page, hover over the Node ID and click **Details**:

Node pool information				
Node pool name	(wudi2)		Scaling group name	2
Node pool status	Running		Launch configuratio	on name(j)
Labels/Taints/Annotations	View		Number of nodes ir	n the scaling g
Number of manually-added no	odes(i) 0		Retry policy	
Auto scaling	On(Min nodes:0,Max nodes:1)		Tag	
Scaling mode 🛈	Release mode		Deletion Protection	
Instance creation policy 🛈	Preferred availability zone (subnet) first			
Removal Policy	Remove the latest instance			
-	<b>ils</b> TencentOS Server 3.2 (Final) ┏* Public image -Basic image		Runtime componen Subnet	ts <b>contain</b>
Node configuration deta	TencentOS Server 3.2 (Final) 🎤 Public image -Basic image		Subnet	its <b>contain</b>
Operating system 🕄 Model 🚯	TencentOS Server 3.2 (Final) Public image -Basic image SA2.MEDIUM2(Primary)		Subnet Custom data	View
Operating system 🛈	TencentOS Server 3.2 (Final) 🎤 Public image -Basic image		Subnet	
Operating system <b>③</b> Model <b>④</b> Data disk Custom Kubelet parameters	TencentOS Server 3.2 (Final) Public image -Basic image SA2.MEDIUM2(Primary) -		Subnet Custom data	View
Operating system <b>③</b> Model <b>④</b> Data disk Custom Kubelet parameters	TencentOS Server 3.2 (Final) Public image -Basic image SA2.MEDIUM2(Primary) - View existing node Remove More •	Removal	Subnet Custom data	View

After navigating to the Instance Details page, click **Security groups** to view specific security group information:

s-tke-np			Kunning		
-				hasing the instance, c	heck the password in <sup>Mes</sup>
Log in S	hutdown	Restart	Reset password	Terminate/Retu	rn More actions 🔻
Basic informat	on ENI	Public IF	P Monitoring	Security gro	
			5	Jecunty gro	oups Operation
Pound or					
Bound se	curity groups			Sort Configuration	Rule previe
Bound se		Security gro		Sort Configuration	
			oup ID/na Opera	Sort Configuration	Rule previe

#### Super Nodes

Click **Node Management** > **Worker Node** > **Node Pool**, and click Node Pool ID. In **Node pool information**, you can view the security group:

lode pool name			
lode pool status	Running		
Security group	sg-	2	
abels	View		
Taints	View		
Deletion Protection	Enabled		
Node type	Linux		

### **Related Issues**

### **Issue Description**

Abnormal binding status, "Install tmp-agent CR" step shows "context deadline exceeded":

202020000	84	2024-06-12 11:04:55	2024-06-12 11:05:00	N/A
tmp-agent CR		2024-06-12 11:04:57		{Reason:get re

### Troubleshooting

Is the VPC the Same or Interconnected?



1. Click the user cluster link, open the associated cluster, and view the cluster node network (i.e., vpcid):

Cluster information		Node and Network Inform
Cluster name	il 🖍	Number of nodes
Cluster ID	10.000	
Deployment type	General cluster	Default OS
Status	Running(j)	System image source
Region	South China(Guangzhou)	Node hostname naming pattern
Addition of Resource Allocated Project 🛈	DEFAULT PROJECT 🎤	Node network
Cluster specification	L5 🎤	Container network add-on
	The application size does not exceed the recommended management size. Up to 5 nodes, <b>150 Pods</b> , <b>128</b> <b>ConfigMap and 150 CRDs</b> are allowed under the current cluster specification. Please read Choosing Cluster Specification C carefully before you make the choice.	Container network
	Auto Cluster Upgrade 🚯	
	Check specification adjustment history	Network mode
Kubernetes version	Master 1.26.1-tke.3(Updates available) Upgrade	VPC-CNI mode
	Node 1.26.1-tke.7、1.26.1-	Service CIDR block
	tke.3(Updates available)Upgrade	Kube-proxy mode

2. On the Prometheus Instance's **Basic Info** page, click **Network** to view the cluster network:

Basic Info			
Name	Ô	Region	Guangzhou
Instance ID	ē.	AZ	Guangzhou Zone 4
Status	⊘ Running	Billing Mode	Pay as you go
Tag	Ø	Creation Time	2024/03/08 11:28:20

3. Compare the vpcid. If they are different, check if the VPCs are interconnected via CCN. If not, you need to associate the CCN to interconnect both VPCs or select **Create Public Network CLB Instance** when associating clusters. If CCN is interconnected but still unsuccessful, check if the CCN bandwidth limit is reached. If so, increase the CCN bandwidth limit.

Associate with CCN:

🔶 Details	of vpc-7	
Basic informa	ation Classiclink Monitoring	
Basic inform	nation	Associate with CCN
ID		CCN provides multi-point intranet interce customer IDCs. Learn more
Name	test	The current VPC is not associated with an
IPv4 CIDR	mary) 17 E)	
DNS	8 🖍	
Domain Name	e (j) - 🖍	
Tags	No tags found. 💋	

Select Create Public Network CLB Instance:

Associate Cluster	
(i) • Remaining IP(s)	of the subnet [ <b>1</b> ]: 160
Cluster Type	Standard cluster 🗸
Cross-VPC Association	Enable After enabling this option, you can use one instance to monitor clusters in different response on the second se
	✓ Create Public Network CLB Instance If the VPC where your instance resides is not interconnected with the cluster to be ass CLB instance because data cannot be collected otherwise. You don't need to do so if t
Cluster Region	Chengdu 🗸
	Tencent Cloud services in different regions cannot communicate with each other over you select a region closest to your end users to minimize access latency and improve region after you purchase the instance.
Cluster	Available clusters in the current region:

#### **Does the Security Group Allow Access?**

1. View the user cluster security group. For viewing methods, see User Cluster Security Group Viewing Method. Check if the rules meet the requirements.

2. If the user cluster is an independent cluster, view the Master&Etcd security group information. Click Node Management > Master&Etcd > Node Pool, click the Node Pool ID, hover over the Node ID, and then click Jump to CVM Instance Details Page. On the CVM Security groups page, you can view specific security group information:

-	-					
	e initial login name is root. If y Log in Shutdown			g the instance, check tł Terminate/Return	he password in <mark>Message Ce</mark> More actions <b>T</b>	enter. You ca
	Basic information E	ENI Public IP	Monitoring	Security groups	Operation logs	Run c
	Basic information E			Security groups	Operation logs Rule preview	Run co
			Sort		-	Run co Outbo

Check if the security group rules meet the requirements.

# Operation Guide Instance Creating Instance

Last updated : 2024-08-07 22:03:06

This document uses custom configuration as an example to describe how to create a TMP instance.

### Preparations

Before creating a TMP instance, you must complete the following operations: Sign up for a Tencent Cloud account and complete identity verification. Create a VPC in the target region and create a subnet in the target AZ in the VPC.

### Directions

- 1. Log in to the TMP console.
- 2. Click Create and configure the following information as prompted:

Туре	Required	Configuration Description
Region	Yes	Select a region based on the region of your Tencent Cloud service. The price may vary by region, and the real-time price displayed on the purchase page shall prevail. Tencent Cloud services in different regions cannot communicate with each other over the private network; for example, a service in the Guangzhou region cannot report data to a TMP instance in the Shanghai region over the private network. Once an instance is purchased, the region cannot be changed. Therefore, please select the region with caution.
AZ	Yes	Select as needed.
Network	Yes	It indicates a logically isolated network space in Tencent Cloud. A VPC consists of at least one subnet. The system will provide a default VPC and subnet for you in each region. If the existing VPCs/subnets don't meet your requirements, you can create new ones as instructed in Creating VPCs or Creating Subnets.
Product Specs	Yes	The price varies by product specs. For more information, please see Pricing.
Data Retention	Yes	The price varies by data retention period. For more information, please see

Period		Pricing.
Instance Name	Yes	Enter a custom name of the TMP instance.
Grafana/Grafana Password	Yes	Grafana is enabled by default. After the instance is successfully created, the system will generate a domain name accessible over the public network. The default account of the Grafana service is `admin`, and the password is user-defined.
Validity Period	Yes	Multiple validity periods are available for your choice.

### 3. After completing the configuration, click $\ensuremath{\textbf{Buy Now}}.$

Region and Ne	twork Config
Region	South China
	Guangzhou Shenzhen Finance
	Tencent Cloud services in different regions cannot communicate with each other over the private network. For example, the service in Guangzhou region cannot report data to TMP in Shanghai region over the private network it after purchasing the instance.
AZ	Guangzhou Zone 1       Guangzhou Zone 2       Guangzhou Zone 3       Guangzhou Zone 4       Guangzhou Zone 5       Guangzhou Zone 6       Guangzhou Zone 7
Network	vpc-rdalicw7   intl_test   10.0.0.0/16 V Subnet-9apu3jks   intl_test_1   10.0.0.0/24 V 🗘 Available IP(s) of the subnet: 252
	If the existing VPC/subnet does not meet your requirement, you can go to the console to create a VPC 2 or Create Subnet 2 Only services in the "Intl_test VPC" can report monitoring data. Please select the network with caution as you cannot change it after purchasing the instance.
Basic Instance	Config
Data Retention Period	15 days30 days45 days
Instance Name	example
Grafana	grafana-test-ai0efv5c   test0802 $\checkmark$ $\phi$
	If the existing Grafana instance does not meet your requirement, you can create one 🗹 in the console.
Tag (optional)	kkk ~ del ~ 删除
	+ Add
	If the existing tag/tag value does not meet your requirement, you can create one 🗹 in the console.
Terms of Agreement	Vive read and agree to Tencent Cloud Terms of Service, Tencent Cloud Prometheus Service Level Agreement, Billing Overview, and Payment Overdue

# Searching for Instance

Last updated : 2024-08-07 22:03:19

By default, TMP instances in the current region are displayed in the TMP console. To help you quickly find the instances in the current region, Tencent Cloud provides the search feature and allows you to filter instances by resource attributes such as instance ID, name, status, AZ, IPv4 address, and tag.

### Directions

Q

- 1. Log in to the TMP console.
- 2. Enter the conditions in the search box as needed and click

Tencent Managed Service	for Prometheus	🔇 Guangzhou 🔻						WeChat O
Create Edit Tag							Separate keywords with	
Instance ID/Name	Monitoring/St 🔻	AZ T	Network	Configuration	IPv4 Address	Billing Mode	Tag (key:value) (i)	Creation Tin
test	<b>II</b> ⊘ Running	Guangzhou Zone 1	Network: intl_test Subnet: intl_test_1	Data retention period: 15 day(s) Specs name: Shared Edition		- 🕑 -Expired		2021/11/24
intl test	Running	Guangzhou Zone 2	Network: Default-VPC Subnet: Default-Subnet	Data retention period: 15 day(s) Specs name: Shared Edition		Trial Edition 🥑	k1:v1	2021/11/10
Total items: 2							10 👻	/ page 🛛 🖌 🖣

3. You can filter instances by different conditions. Currently, the following dimensions are supported:

Filter Condition	Description
Instance ID	You can directly enter multiple instance IDs for quick filtering. Each instance ID supports only exact match-based filtering.
Instance name	You can enter only one instance names for filtering. Fuzzy match-based filtering is supported.
Instance status	You can enter multiple status values for filtering. You can also configure the list header for quick filtering.
AZ	You can enter multiple AZs for filtering. After the region is switched, the corresponding AZs in the region will be displayed. You can also configure the list header for quick filtering.



IPv4 address	You can enter multiple IPv4 addresses for filtering. Each IPv4 address supports only exact match-based filtering.
Тад	You can enter multiple tags to filter instances. You can also directly click a tag value in the instance list for filtering.

# **Renaming Instance**

Last updated : 2024-08-07 22:03:26

To help you quickly identify TMP instance names for easier instance management in the TMP console, Tencent Cloud allows you to name all instances and rename them at any time with immediate effect.

### Directions

1. Log in to the TMP console.

2. In the instance list, select the TMP instance to be renamed and click **More** > **Instance Configuration** > **Rename** on the right.

3. In the **Rename** window that pops up, enter the new instance name and click **OK**.

ou have selected	this instance:				
Instance ID/	Status	AZ	Network	Configuration	Billing Mod
		Guangzhou Zone 1	Network: intl_test Subnet: intl_test_1	Data retention period: 15 day(s)	-
stance Name	modify				

# **Terminating Instance**

Last updated : 2024-08-07 22:03:31

If you no longer use an instance, you can terminate it, and it will be suspended once terminated. You can reboot suspended instances according to different scenarios and needs.

### **Relevant Impact**

Once an instance is suspended, its data will be affected as follows:

IP: the corresponding IP address is retained but does not provide normal services.

Grafana: Grafana cannot be accessed at the corresponding domain name.

Data: data on the corresponding instance will be retained for a certain number of days. The specific number displayed in the confirmation information during instance termination in the console shall prevail.

### Directions

1. Log in to the TMP console.

2. In the instance list, select the TMP instance to be terminated and click **More** > **Instance Status** > **Terminate** on the right.

3. As termination is a high-risk operation, in the **Terminate** window that pops up, complete the termination steps as prompted and click **OK**.

ou have selected	this instance:				
Instance ID/	Status	AZ	Network	Configuration	Billing Mod
		Guangzhou Zone 1	Network: intl_test Subnet: intl_test_1	Data retention period: 15 day(s)	-
<ul> <li>Once t</li> </ul>	erminated, the ins	stance will be in the	e shutdown status for seven days, du		ned.
	he resources are t		-day unconditional refund (for one ir	nstance) will be returned to you e proportion of the cash and gi	

# **Rebooting Instance**

Last updated : 2024-08-07 22:03:36

You can reboot a suspended or abnormal instance in the console.

### Directions

1. Log in to the TMP console.

2. In the instance list, select the TMP instance to be rebooted and click **More** > **Instance Status** > **Reboot** on the right.

3. In the **Reboot Notes** window that pops up, click **OK**.

stance Renewa	I				
ou have selected t	nis instance:				
Instance ID/Name	Status	AZ	Network	Configuration	Billing Mode
		Guangzhou Zone 1	Network: intl_test Subnet: intl_test_1	Data retention period: 15 day(s)	-
lidity Period					
ito-Renewal 🗸	Auto-renew the devic	e every month when my ac	count has sufficient bal	ance	
es Ou	erying configuration fe				

#### Note:

If you reboot a running instance, the service will be interrupted during reboot. Please confirm the operation risks first.

# **Modifying Storage Period**

Last updated : 2024-08-07 22:03:43

TMP allows you to modify the data storage period after creating an instance. The longer the storage period, the higher the unit price of the instance. For billing details, see Pay-as-You-Go.

## Directions

1. Log in to the TMP console.

2. In the instance list, find the target TMP instance and click **More** > **Instance Configuration** > **Modify Storage Period** on the right.

Instance ID/Name	Monitoring/Status 🔻	AZ Y	Network	Configuration	IPv4 Address	Billing Mode	Tag (key:value) 🚯
test	Running	Singapore Zone 3		Data retention period: 15 day(s)		Pay as you go 🧭	
Total items: 1							

3. In the pop-up window, select the target storage period and click **OK**.

Instance ID/ Status AZ Network Configuration Billing Singapore Network:Default-VPC Data retention Pay as y Subnet:rs Subnet:rs	Mode
• • • • • • • • • • • • • • • • • • • •	
day(s)	/ou go
Data Storage Period 15 days 💌	

### Note:

After the modification, newly collected data will be billed based on the new storage period and unit price from 0:00 AM the next day.

Historical data is still stored based on the storage period configured before the modification.

# Viewing Instance's Basic Information

Last updated : 2024-07-31 10:44:36

To help you view the basic information of TMP instances, TMP allows you to select a desired instance on the instance list page to enter its management page and view its basic information.

# Directions

1. Log in to the TMP console.

2. In the instance list, select the TMP instance to be viewed and click its **instance ID** or **Manage** on the right.

						Follow WeChat O
Basic Info	Data Collection	Alarm Management	Recording Rule			
Basic Info						
Name	iou 🧷		Region	Guangzhou		Networl
Instance ID	Q		AZ	Guangzhou Zone 3		Subnet
Status	⊘ Running		Billing Mode	Pay as you go		IPv4 Ad
Tag	Ø		Creation Time	e 2024/04/08 16:43:50		
Grafana					Service Address	
Grafana Instai	nce Associate with TCM	IG			Token	***** <b>D</b>
Grafana Dat	a Source Configuration	Information			Remote Write Address	ht
HTTP URL	htt				Remote Read Address	ht
Basic auth use	er(APPID)	ē			HTTP API	h
Basic auth pas	ssword ***** 🗅				Pushgateway Address	

3. You can perform the following operations on the basic information page:

Rename the instance.

Edit instance tags.



Change the Grafana Admin password.

Upgrade preset Grafana dashboards.

# TKE Integration with TKE

Last updated : 2024-08-07 22:04:22

You can monitor the TKE service in business scenarios after integrating it. This document describes how to integrate the TKE service.

Tencent Kubernetes Engine (TKE) provides container-centric solutions based on native Kubernetes. It can solve environment-related issues in the process of development, testing, and Ops, helping you reduce costs and improve efficiency. Kubernetes is an open-source container orchestration tool developed by Google and has been used by Google for more than 15 years. As the de facto standard in the container field, Kubernetes can greatly simplify the management and deployment complexity of applications. By integrating with the TKE service, you can monitor the status of Kubernetes and the services running on it much more easily through Prometheus.

### Note

In order to ensure normal operation, existing instances will automatically update the component version when you edit collection configuration and associate new clusters. During the update process, data breakpoints may occur in the associated clusters.

## Directions

1. Log in to the TMP console.

2. In the TMP instance list, click the ID/Name of the newly created instance.

3. Go to the TMP management center and click Integrate with TKE on the left sidebar.

4. Perform the following operations on the cluster monitoring page:

Associate a cluster: Associate a cluster with a TMP instance as instructed in Associating with Cluster.

Configure data collection: Create a data collection rule to monitor your business data by adding the configuration in the console or via a YAML file. For more information, see Data Collection Configurations.

Streamline basic monitoring metrics: Select only the required metrics to avoid unnecessary fees as instructed Streamlining Monitoring Metrics.

5. After completing the above operations, you can view the monitoring data of your TKE service in Grafana.

Ø	Q Search dashboards by name	New Dashboard New Folder Import
Q	L≡ Sort (Default A-Z) ×	Filter by starred So Filter by tag So Filter by tag
+	Co Kubernetes	
88	API Server Kubernetes	api server cm-prometheus kubernetes readonly
Ø	Compute Resources / Cluster Kubernetes	cm-prometheus compute resources kubernetes readonly
¢ م	Compute Resources / Namespace (Pods) Kubernetes	cm-prometheus compute resources kubernetes readonly
\$ 0	Compute Resources / Node (Pods) Kubernetes	cm-prometheus compute resources kubernetes readonly
	Compute Resources / Pod Kubernetes	cm-prometheus compute resources kubernetes readonty
	Controller Manager Kubernetes	cm-prometheus controller kubernetes readonly
	CoreDNS Kubernetes	cm-prometheus coredns dns kubernetes readonly
	Kubelet Kubernetes	cm-prometheus kubelet kubernetes readonly
	Kubernetes Cluster Kubernetes	cluster cm-prometheus kubernetes readonly
	Networking / Cluster	cm-prometheus kubernetes networking readonly
	Networking / Namespace (Pods)	cm-prometheus kubernetes networking readonly
	Networking / Pod Kubernetes	cm-prometheus kubernetes networking readonly
	Nodes Kubernetes	cm-prometheus kubernetes node readonly
8	Prometheus	agent cm-prometheus prometheus readonly
0	Kubernetes Proxy	em-crometheus Indernetes provv Freedonbr

# **Integration Center**

Last updated : 2024-08-07 22:04:30

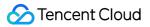
TMP integrates commonly used programming languages, middleware, big data, and infrastructure databases. It supports quick installation and custom installation. You only need to follow the instructions to monitor the corresponding components. It also provides out-of-the-box Grafana monitoring dashboards. The integration center covers three major monitoring scenarios of basic service monitoring, application layer monitoring, and TKE cluster monitoring, making it easier for you to integrate and use.

# List of Supported Services

Service Type	Service	Monitoring Metric	Quick Installation	Integration Guide
Big data	Elasticsearch	Cluster/Index/Node monitoring	Supported	Elasticsearch Exporter Integration
	Flink	Cluster/Job/Task monitoring	Not supported	Flink Integration
Development	CVM	The extended `cvm_sd_config` can be used to configure a CVM scrape task and collect Node Exporter or custom business metrics.	Supported	CVM Node Exporter
	Go	GC/Heap/Thread/Goroutine monitoring	Not supported	Go Application Integration
	JVM	Heap/Thread/GC/CPU/File monitoring	Not supported	JVM Integration
	Spring MVC	HTTP API/Exception/JVM monitoring	Not supported	Spring Boot Integration
Middleware	Kafka	Broker/Topic/Consumer group monitoring	Supported	Kafka Exporter Integration
	Consul	Consul monitoring	Supported	Consul Exporter Integration

	Etcd	Etcd monitoring	Not supported	-
	Istio	Istio monitoring	Not supported	-
Infrastructure	Kubernetes	API server/DNS/Workload/Network monitoring	Supported	Agent Management
Database	TencentDB for MongoDB	File count/Read and write performance/Network traffic monitoring	Supported	MongoDB Exporter Integration
	TencentDB for MySQL	Network/Connection count/Slow query monitoring	Supported	MySQL Exporter Integration
	TencentDB for PostgreSQL	CPU/Memory/Transaction/Lock/Read/Write monitoring	Supported	PostgreSQL Exporter Integration
	TencentDB for Redis	Memory utilization/Connection count/Command execution status monitoring	Supported	Redis Exporter Integration
	TencentDB for Memcached	Memcached monitoring	Supported	Memcached Exporter Integration
Inspection	Health check	Blackbox can be used to regularly test the connectivity of the target service, helping you stay up to date with the service health and discover exceptions in time.	Supported	Health Check
СМ	СМ	Tencent Cloud service monitoring	Supported	-
Custom	Scrape task	The native `static_config` can be used to configure a scrape task.	Supported	Scrape Configuration Description
Custom	CVM scrape task	The extended `cvm_sd_config` can be used to configure a CVM scrape task.	Supported	Scrape Configuration Description

## Directions



### **Quick installation**

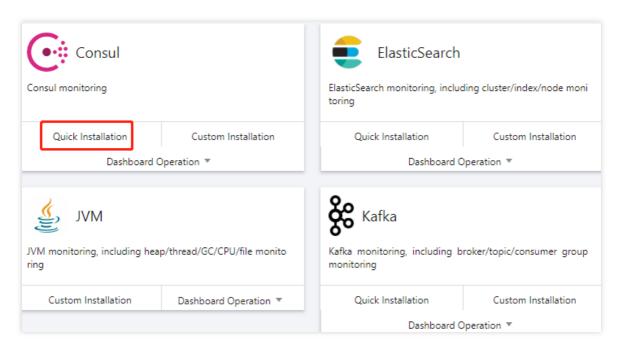
Some services support quick agent installation. For more information, see Integration Center > List of Supported Services.

1. Log in to the TMP console.

2. In the instance list, select the corresponding TMP instance.

3. Enter the instance details page and click Integration Center.

4. In the **Integration Center**, select the service that supports quick installation and click **Install** in the bottom-left corner.



5. On the **Integration List** page, enter the metric collection name and address and click **Save**. Below is a sample for Kafka:



ame *	Global unique name
Kafka instan	ce
address *	+ Add
tag 🛈	+ Add
Exporter con	ıfig
topic regular	Only collect topic match regular
group regular	Only collect group match regular

### **Custom installation**

- 1. Log in to the TMP console.
- 2. In the instance list, select the corresponding TMP instance.
- 3. Enter the instance details page and click Integration Center.
- 4. Select the target service in the integration center. You can click **Integration Guide** to view the integration guide.
- After successful integration, you can monitor the corresponding service in real time. You can also click

Install/Upgrade in Dashboard Operation to install or upgrade the Grafana dashboard for the service.

			Integrat	ion Center			
	Search for access mode	e by keyword					
	Category: All Midd	leware Big Data Application	Infrastructure Database				
Consul		ElasticSearch		🦂 Flink		Gola	
Consul monitoring		ElasticSearch monitoring, inclue toring	ding cluster/index/node moni	Flink monitoring, including clu	ister/job/task monitoring	Golang Runtime r outine monitoring	
Quick Installation	Custom Installation	Quick Installation	Custom Installation	Custom Installation	Dashboard Operation 🔻	Custom Insta	
Dashboard	Operation 🔻	Dashboard C	Operation 🔻				
MVL §		<b>čč</b> Kafka		Kubernetes		Men	
JVM monitoring, including hea ring	ap/thread/GC/CPU/file monito	Kafka monitoring, including b monitoring	roker/topic/consumer group	Kubernetes monitoring, inclu d/network monitoring	ding API server/DNS/workloa	Memcached mon	
Custom Installation	Dashboard Operation 🔻	Quick Installation	Custom Installation	Custom Installation	Dashboard Operation 🔻	Quick Instal	
		Dashboard C	Dashboard Operation 🔻				
MongoDB		MusqL MySQL		PostgreSQL		Re	
MongoDB instance monitoring d write performance/network t		MySQL instance monitoring, in count/slow query monitoring	ncluding network/connection	PostgreSQL instance monitori nsaction/lock/read/write mon	ng, including CPU/memory/tra itoring	Redis instance mo tion count/comma	
Quick Installation	Custom Installation	Quick Installation	Custom Installation	Quick Installation	Custom Installation	Quick Instal	
Dashboard Operation 💌		Dashboard Operation 🔻		Dashboard Operation 🔻			

# Data Multi-Write

Last updated : 2024-08-15 16:16:34

The data multi-write feature supports writing monitoring data to self-built Prometheus or other Prometheus monitoring instances.

#### Note:

Currently, only the multi-write of monitoring data collected by Integrated Container Service and Integrated Center is supported, and others are not supported yet.

# **Operation Guide**

1. Log in to TMP Console.

2. In the left sidebar, click TencentCloud Managed Service for Prometheus.

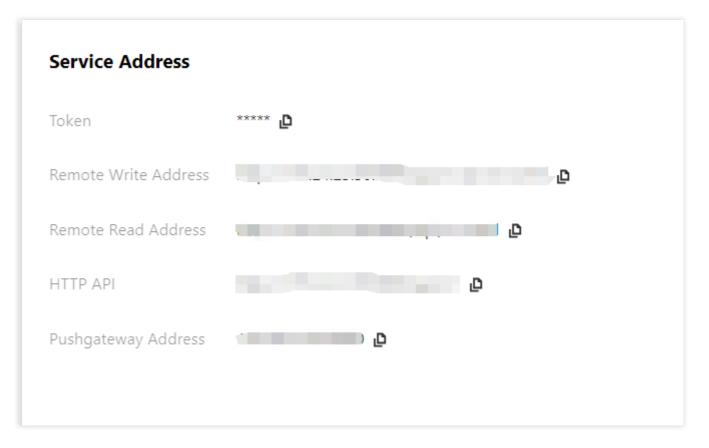
 Click the corresponding instance name to enter the instance details page, then click Data Collection > Data Multi-Write on the top navigation bar.

4. In the Data Multi-Write sub-window, click **Add** and see the following description to configure the data multi-write information.



← 11 11 11 11 11 11 11 11 11 11 11 11 11	
Basic Info Data Colle	Alarm Management Recording Rule instance diagnostics
Integrate with TKE In	tegration Center Data Multi-Write
	data to self-built Prometheus or other TencentCloud Managed Service for Prometheus instances. s only supported for monitoring data collected in "Integrate with TKE" and "Integration Center".
Prometheus Address	0
Security Authentication	N/A Basic authentication
Headers	Add
Relabel	<pre>1 # add label 2 #. target_label: key 3 # replacement: value 4 #discard metric 5 #. source_labels: [_name_] 6 # repex: kubelet*; 7 # action: drop</pre>
+ Add Configuration	
Save Cancel	

Prometheus address: Obtain the HTTP API address from the basic information of instances.



Security verification: Whether to enable security authentication. Once enabled, the username and password needs to be customized.

Prometheus Address		$(\mathbf{i})$	
Security Authentication	N/A	Basic authentication	
Headers	Add		
Username	Please er	nter the username	
Password	Please er	nter the passwori 🛛 🕲	
Relabel	1 2 3 4 5 6 7	<pre># add label #- target_label; key # replacement: value #discard metric #- source_labels: [name] # regex: kubelet+; # action: drop</pre>	
+ Add Configuration			
Save Cancel			

# Recording Rule Overview

Last updated : 2024-08-07 22:04:42

A recording rule allows you to calculate some commonly used or complex metrics in advance and then store the calculated data in new data metrics. In this way, querying the calculated data will be faster and easier than querying the original data. This is very suitable for dashboard scenarios and can solve the problems of complicated user configuration and slow query.

Recording rules exist in the form of rule group, and rules in the same group are executed sequentially at a certain interval. Rule names must conform to the corresponding Prometheus specification. Generally, a rule file is as follows:

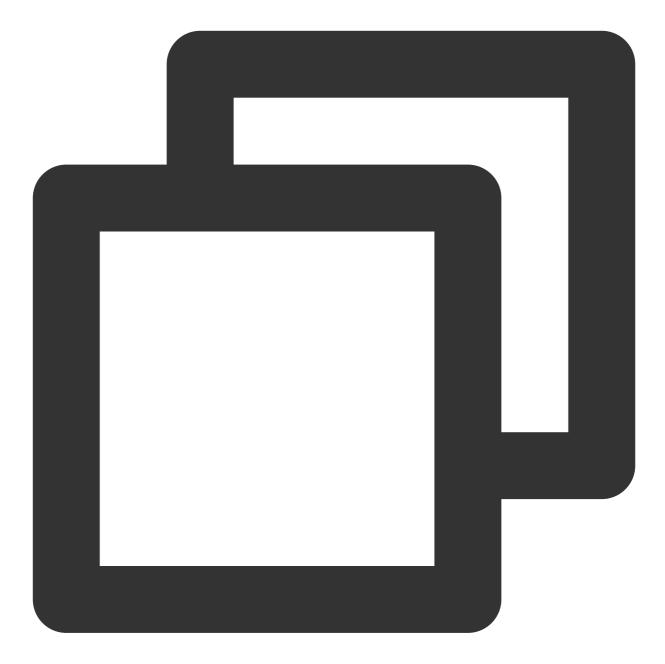




```
groups:
[ - <rule_group> ]
```

Below is a simple example of recording rule:





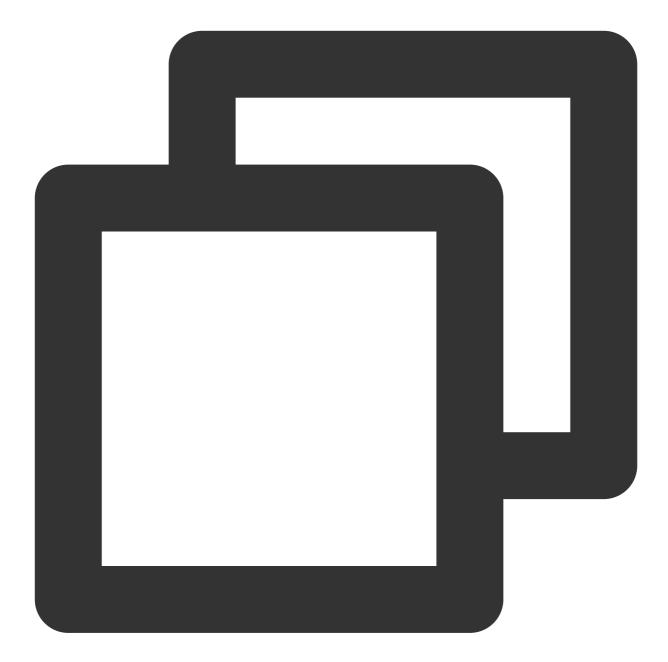
```
groups:
```

```
- name: example
```

```
rules:
```

- record: job:http\_inprogress\_requests:sum
  - expr: sum by (job) (http\_inprogress\_requests)

# Rule Group

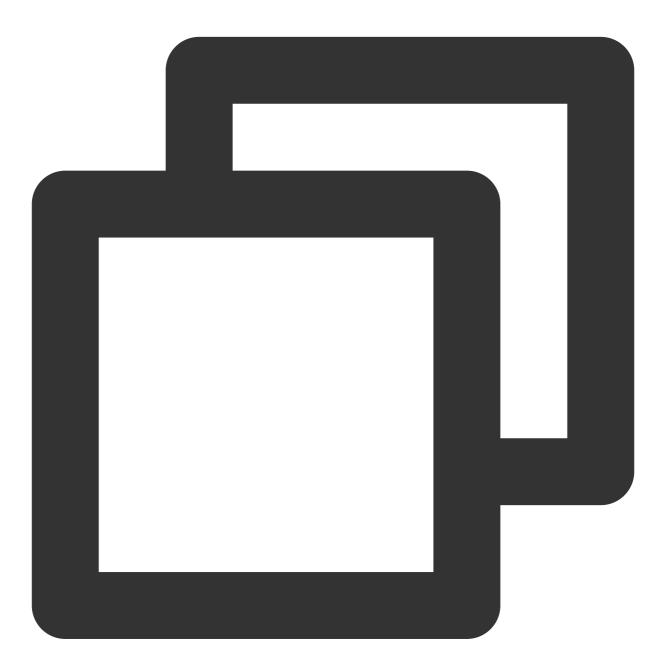


```
# A rule group name must be unique in the same file
name: <string>
# Rule detection interval
[ interval: <duration> | default = global.evaluation_interval ]
rules:
  [ - <rule> ... ]
```



## Rule

The recording rule syntax is as follows:



# The generated new metric name, which must be valid record: <string>

# PromQL expression. Each calculated data entry will be stored in the new metric na expr: <string>

# The label to be added or overwritten in the data to be stored

## 🕗 Tencent Cloud

```
labels:
  [ <labelname>: <labelvalue> ]
```

## **Recommended Name Format**

The recommended format for naming recording rules is level:metric:operations .

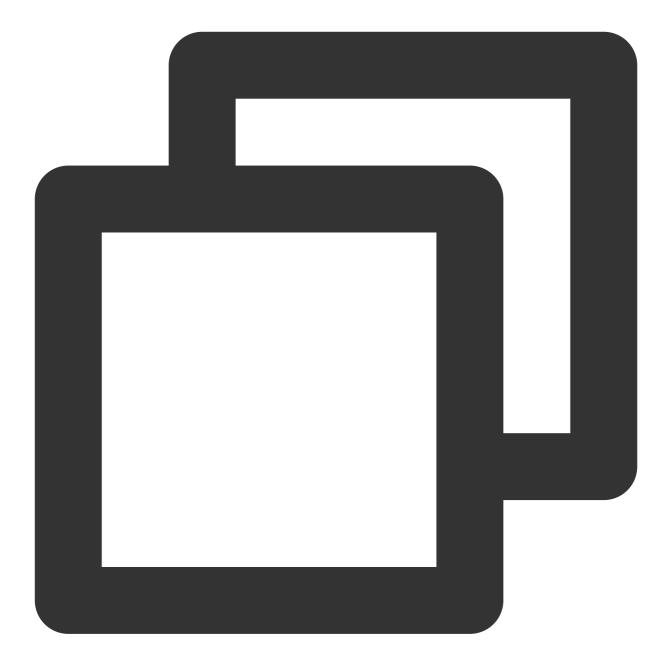
level: indicates the recording level and the output label of the rule.

metric: indicates the metric name.

operations: indicates the list of operations applied to the metric.

Example:





- record: instance\_path:requests:rate5m
  expr: rate(requests\_total{job="myjob"}[5m])
- record: path:requests:rate5m expr: sum without (instance)(instance\_path:requests:rate5m{job="myjob"})

# Rule Management

Last updated : 2024-08-07 22:04:47

## Overview

You can manage the TMP recording rules in the TMP console to avoid the hassle of having to modify the configuration file in native Prometheus.

# Preparations

- 1. Log in to the TMP console.
- 2. Create a TMP instance as instructed in Creating Instance.
- 3. Enter the TMP instance management page through the instance list.
- 4. Manage recording rules as instructed in Overview.

## Directions

## **Creating rule**

1. In the menu on the left of the instance management page, click **Recording Rule** > **Create** to enter the rule creation page, adjust the rule expression and the name of the new metric to be recorded according to your actual needs as shown below. For specific terms, please see Overview.

CreateRecording R	ule (RecordingRule)
<ul> <li>Please use the group at a time</li> </ul>	ne native Prometheus recording rule YAML configurations. Note: you can enter configurations for only one rule me.
Rule Group Name	Please enter the rule group name
YAML Configuration	<pre>1 name: example 2 rules: 3   - record: job:http_inprogress_requests:sum 4   expr: sum by (job) (http_inprogress_requests)</pre>
	OK Cancel

2. Click OK.

## Managing rule

In the rule list, you can temporarily **disable** rules or enable rules that are **not enabled**. Once disabled, a rule will stop working, and the collection of related recording metrics will also stop.

## **Deleting rule**

- 1. You can delete rules that are no longer used.
- 2. Select the rule to be deleted in the list and confirm in the pop-up window. Once deleted, a rule will stop working.

# List of Default Recording Rules

Last updated : 2024-08-07 22:04:52

Recording rules are created for associated clusters by default. For free metrics in pay-as-you-go mode, the following new metrics will be created after recording and will be billed normally. If you don't need them for data collection, you can disable the default recording rules when you associate the cluster for the first time, or later in the recording rule list on the recording page.

Metric	Preset Dashboard	A F T
:node_memory_MemAvailable_bytes:sum	Kubernetes / Compute Resources / Cluster	-
node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate	Kubernetes / Compute Resources / ClusterKubernetes /	-
node_namespace_pod_container:container_memory_working_set_bytes	Compute Resources / WorkloadKubernetes / Compute Resources /	-
node_namespace_pod_container:container_memory_rss	PodKubernetes / Compute Resources / Node (Pods)Kubernetes / Compute Resources / Namespace (Workloads)Kubernetes	-
node_namespace_pod_container:container_memory_cache		-
node_namespace_pod_container:container_memory_swap	/ Compute Resources / Namespace (Pods)Kubernetes / Networking /	-
namespace_workload_pod:kube_pod_owner:relabel	WorkloadKubernetes / Networking / Namespace (Workload)	-
namespace:kube_pod_container_resource_requests_memory_bytes:sum	-	K re
namespace:kube_pod_container_resource_requests_cpu_cores:sum	-	K re



# Instance Diagnosis

Last updated : 2024-08-07 22:04:57

# Background

To enhance the user experience with the Prometheus monitoring collection terminal, we now offer a new collection terminal architecture. The upgraded new architecture supports instance diagnosis, system health checks, and improves the resource utilization rate of the collection agent and the stability of metric collection. This document will guide users in upgrading the old collection architecture to the new one via the console instance diagnosis page and obtain detailed information about the current instance collection and storage for a better experience.

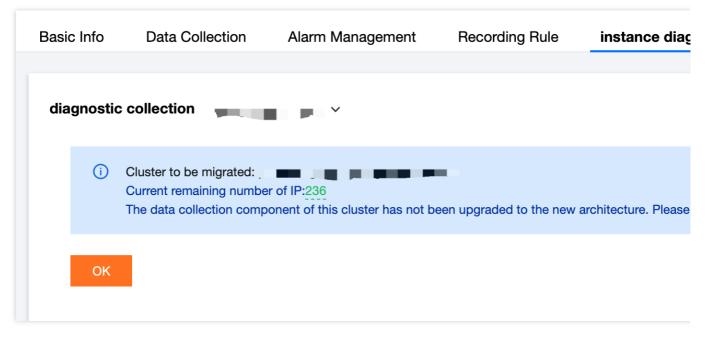
## Directions

## Upgrading to the New Architecture

- 1. Log in to TMP Console.
- 2. In the Prometheus instance list, click **Instance ID/Name**.

3. In the Prometheus management center, click **instance diagnostics** on the top navigation bar, and select the corresponding **collection cluster**.

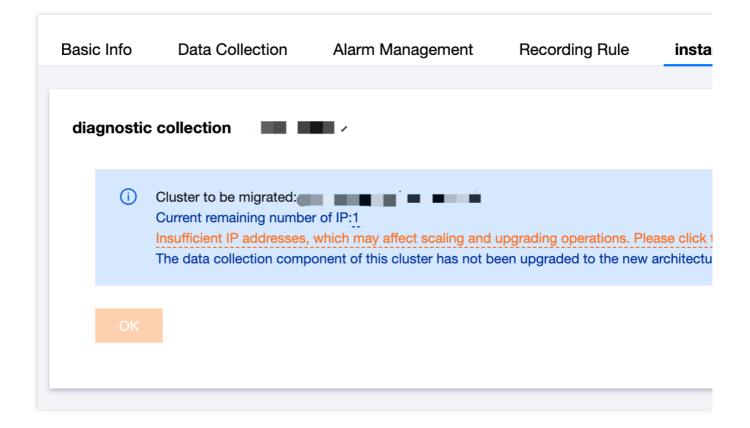
Click **OK** to upgrade to the new architecture.



### Note:

The upgrade process is expected to take 5 minutes and may experience a 1-2 minute metric interruption. If the number of IPs in the managed cluster of the instance is less than 10, a risk warning will appear. To prevent the

issue of insufficient IPs from causing component upgrade failures, please increase the available IPs as guided below.



## Adding Subnet

1. On the instance diagnosis page, click subnet information to enter the **Managed Cluster Subnet Management** page.

						Managed cluster subnet r	
asic Info	Data Collection	Alarm Manageme	nt	Recording Rule	instance diagnostics		
						(i) • The current remaining	ng number of IP addresses is: 1
liagnostic	c collection	~					
						• The current number	of remaining IP addresses in the
(j	Cluster to be migrated:					Please enable a new	<i>i</i> subnet
	Current remaining number of				and all the solution of the solution		
	Insufficient IP addresses, white The data collection component				ase click to add a subnet. architecture. Please upgrade it fir	Q	S
						Cultured ID	IP range
ОК						Subnet ID	IP range
						sub-test	
liagnostic	c storage 🤆					sub-test	
liagnostic	c storage 📿						
-	c storage 💭 of metric reporting(i)			The top 10 metric	cs based on the number of seri		
-	of metric reporting(i)			The top 10 metric			
Total rate	of metric reporting(i)				es_bytes_total		
Total rate 40.7 Chargeable Free metric	of metric reporting() 74 le metric rate 0.68	700	000	container_fs_write	es_bytes_total		
Total rate 40.7 Chargeable Free metric Instance se	of metric reporting() 74 le metric rate 0.68 c rate 40.07			container_fs_write container_fs_read kube_pod_status_	es_bytes_total		
Total rate of 40.7 Chargeable Free metric Instance se The storag	of metric reporting() 74 le metric rate 0.68 c rate 40.07 veries storage limit	s 100	000	container_fs_write container_fs_read kube_pod_status_ container_memory	s_bytes_total s_bytes_total phase y_working_set_bytes		
Total rate 40.7 Chargeable Free metric Instance se The storag The maxim	of metric reporting() 74 e metric rate 0.68 c rate 40.07 veries storage limit ge limit for a single metric series	s 100 ime 102	000 4	container_fs_write container_fs_read kube_pod_status_ container_memory container_cpu_us	s_bytes_total s_bytes_total _phase y_working_set_bytes age_seconds_total		
Total rate 40.7 Chargeable Free metric Instance se The storag The maxim Maximum I	of metric reporting() 74 le metric rate 0.68 c rate 40.07 reries storage limit ge limit for a single metric series num length of a metric label na	s 100 ime 102	000 4	container_fs_write container_fs_read kube_pod_status_ container_memory	s_bytes_total s_bytes_total _phase y_working_set_bytes age_seconds_total		
Total rate 40.7 Chargeable Free metric Instance se The storag The maxim Maximum I Maximum I	of metric reporting() 74 le metric rate 0.68 c rate 40.07 erries storage limit ge limit for a single metric series num length of a metric label nar length of Label values for metri	s 100 ime 102 rics 204 34	000 4	container_fs_write container_fs_read kube_pod_status_ container_memory container_cpu_us	s_bytes_total s_bytes_total uphase y_working_set_bytes age_seconds_total ue_bytes		
Total rate 40.7 Chargeable Free metric Instance se The storag The storag The maximum I Maximum I Maximum I The oldest	of metric reporting() 74 le metric rate 0.68 c rate 40.07 series storage limit ge limit for a single metric series num length of a metric label nar length of Label values for metri number of labels for metrics	s 100 ime 102 rics 204 34 stamps 5h	000 4 8	container_fs_write container_fs_read kube_pod_status_ container_memory container_cpu_us container_fs_usag container_fs_limit	s_bytes_total s_bytes_total uphase y_working_set_bytes age_seconds_total ue_bytes		
Total rate 40.7 Chargeable Free metric Instance se The storag The maxim Maximum I Maximum I The oldest maximum I	of metric reporting() 7 4 le metric rate 0.68 c rate 40.07 erries storage limit ge limit for a single metric series num length of a metric label nai length of Label values for metric number of labels for metrics trange allowed for metric times allowed range for metric times number of series per single qui	s 100 ime 102 rics 204 stamps 5h rtamps 10m iery 100	000 4 B 1 000	container_fs_write container_fs_read kube_pod_status_ container_memory container_cpu_us container_fs_usag container_fs_limit, kube_pod_contain	s_bytes_total s_bytes_total ,phase y_working_set_bytes age_seconds_total re_bytes _bytes		
Total rate 40.7 Chargeable Free metric Instance se The storag The maxim Maximum I Maximum I The oldest maximum I	of metric reporting() 74 e metric rate 0.68 c rate 40.07 erries storage limit ge limit for a single metric series num length of a metric label nai length of Label values for metri number of labels for metrics t range allowed for metric times allowed range for metric times	s 100 ime 102 rics 204 stamps 5h rtamps 10m iery 100 ime(i 200	000 4 B 1 000	container_fs_write container_fs_read kube_pod_status_ container_memory container_cpu_us container_fs_usag container_fs_limit, kube_pod_contain kube_pod_contain	s_bytes_total s_bytes_total phase y_working_set_bytes age_seconds_total pbytes _bytes _bytes age_status_restarts_total		

The page displays the subnets that have already been added to the managed cluster and those that have not in the current VPC. You can click **Enable Subnet** to enable subnets with sufficient remaining IPs as per your plan.
 If there are no available subnets, please Add Subnet first and then enable it on the current page.

## **Instance Diagnosis**

After the architecture is upgraded, the instance diagnosis page will include content on both collection and storage, helping users understand the operating status of Prometheus collection and storage to locate issues more quickly.

## **Diagnostic Collection**

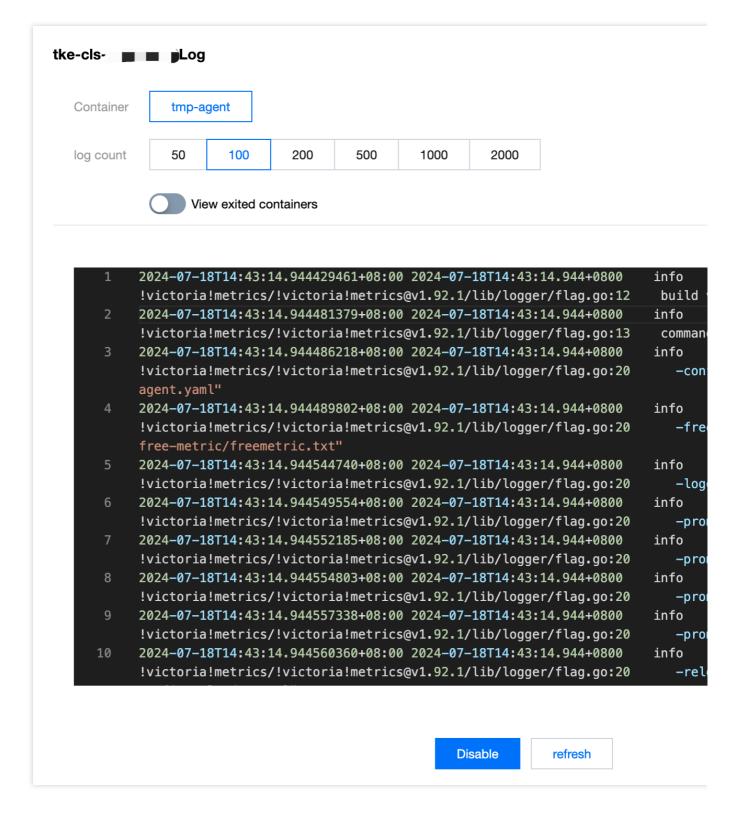
As shown in the following figure, collection diagnosis includes resource occupancy of the corresponding collection, collection configuration, target allocation status, target status, agent status, component version, and the collection architecture diagram.

Basic Info	Data Collection	Alarm Management	Recording Rule	instance diagnostics
diagnostic c	collection	~	data collectio	on architecture diagram 📿
resource utiliza	tion Number of Pods	<b>1/1</b> (1nuclear 2 G) (i)		
collection configuration	Ŀ		Managed clu collection co	uster for data omponents To assign a target tmp-agent
Target allocatio status	n Allocated5items	Not allocated0items	Available IP: Security Gro	: 240 collect
Target status	5 Up		Number of F	Pods 3/4 (3.50 core 5.25 G) (i)
Agent status	1 items			
Version	tmp-agent(v1.0	0.7)	user cluster	proxy-a
	proxy-agent(v1	.0.8)		
	proxy-server(v	1.0.2->v1.0.7)	Cluster Colle	ection Component Explanation 🗳 👘 collection t
	tmp-operator(	(1.1.0)		

#### **Resource Utilization**

Resource occupancy of the collection shard displays information including CPU, memory limit and occupancy, and inbound and outbound traffic. Click **View logs** to see the logs of the Pod, facilitating the review of running details and troubleshooting of exceptions.

asic Info Da	ta Collection Alarm Management	Recording Rule instance diagnostics	
		Search for Pod or If Q	
diagnostic colle	ction	Pod/IP	re
		data collection architecture diagram $\bigcirc$	
resource utilization	Number of Pods 1/1 (1nuclear 2 G) (i)	tke-cls-	1/1
collection configuration		Managed cluster for data collection components	
Target allocation status	Allocated5items Not allocated0items	Available IP: 240 Security Group sg	
Target status	5 Up	Number of Pods 3/4 (3.50 core 5.25 G) ①	
Agent status	1 items		
Version	tmp-agent(v1.0.7)	user cluster	
	proxy-agent(v1.0.8)		
	proxy-server(v1.0.2->v1.0.7)	Cluster Collection Component Explanation	
	tmp-operator(v1.1.0)		



### **Collection Configuration**

It displays the specific collection configuration of the current collection.

asic Info Da	ata Collection Alarm Management	Recording Rule instance diagnostics
		2 scrape_interval: 15s
		3 scrape_timeout: 10s
diagnostic colle	ction <sub>cls</sub> v	4 evaluation_interval: 1m
		data collection architecture diagram 🔿 5 🗸 external_labels:
	Number of Data d (d. (data bar 0.0)	6 cluster: cls-
resource utilization	Number of Pods 1/1 (1nuclear 2 G) (i)	7 cluster_type: tke
		Managed cluster for data 😤 tmp-operator-
collection	<b>.</b>	collection components $9 \lor -job_name: serviceMonitor/kube-system/kube 10 honor labels: true$
configuration		To assign a ta 11 honor timestamps: true
Target allocation	Allocated5items Not allocated0items	Available IP: 240 12 scrape_interval: 15s
status		Security Group sg-
Target status	E L In	Number of Pods 3/3 (2.50 core 3.25 G) (i) 14 metrics_path: /metrics
larget status	5 Up	15 scheme: http
A	4.16	16 follow_redirects: true
Agent status	1 items	17 enable_http2: true
		user cluster 18 relabel_configs:
Version	proxy-agent(v1.0.8)	19 V - source_labels: [job]
	tmp-agent(v1.0.7)	20 separator: ;
	tmp-operator(v1.1.0)	21     regex: (.*)       Cluster Collection Component Explanation 2     22     target_label:tmp_prometheus_job_nam
	,	22 replacement: \$1
	proxy-server( <u>v1.0.2-&gt;v1.0.7</u> )	24 action: replace

#### **Target Allocation Status**

It shows the URL of the collection target, the name of the collection job to which the target belongs, and which collection shard is currently collecting it.

						Target allocation stat	us C	
asic Info Da	ata Collection	Alarm Management	Recording Rule	instance	diagnostics	search URL	Q	
diagnostic colle	ction cls-	~				Job Name		URL
resource utilization	Number of Pods 1	/1 (1nuclear 2 G) (i)	data collect	tion architect	ure diagram <i>C</i>	cadvisor		ற https://kubernete
collection configuration	Ľ			cluster for data components	रु tmp-operator- To assign a ta	kube-proxy		₽ http:/
Target allocation status	Allocated5items	Not allocated0items	Available I					
Target status	5 Up		Security G Number of	roup <u>sc</u> f Pods <mark>3/3</mark> (2.50 (	x_ core 3.25 G) (i)	kubelet		ம https://kubernet
Agent status	1 items					serviceMonitor/kube-sys	stem/kube-state-	D http://
Version	tmp-agent(v1.0.7	7)	user cluste tke/cls	ər		metrics/0		凸 http://
	proxy-agent(v1.0	0.8)	INC/ CIS					
	tmp-operator(v1	.1.0)	Cluster Co	llection Compon	ent Explanation 2	serviceMonitor/kube-sys	stem/node-exporter/0	_ http:/.
	proxy-server(v1.	0.2->v1.0.7)						

#### **Target Status**

You can filter active collection targets based on the collection job, and obtain information about the corresponding target such as status, Labels, and Discovered Labels. The target status is "Healthy" in normal conditons, but if it is "abnormal" and there has already been the last scrape time, you can troubleshoot based on the error message on the

far right. Common issues may include the target itself being unhealthy, incorrect permission configuration, and network errors.

isic Info Da	ta Collection Alarm Management	Recording Rule instance diagnostics	search URL	Q		
liagnostic collee	ction ck		Scrape URL		Status T	Labels
		data collection architecture diagram $\bigcirc$				
esource utilization	Number of Pods 1/1 (1nuclear 2 G) (i)					pod:tke-kube-
ollection		Managed cluster for data 😤 tmp-operator-	凸 http://	)180/metrics	Healthy	service:tke-ku
configuration		collection components To assign a ta	incp <i>a</i> /	100/1101/00		container:kub
arget allocation	Allocated5items Not allocated0items	Available IP: 240 Security Group sg	Total items: 1			
arget status	5 Up	Number of Pods 3/3 (2.50 core 3.25 G) (i)				
gent status	1 items					
/ersion	tmp-agent(v1.0.7)	user cluster tke/cls-				
	proxy-agent(v1.0.8)					
	tmp-operator(v1.1.0)	Cluster Collection Component Explanation				

#### **Agent Status**

It shows the running status of the collection shard agent. Click **View logs** to see the logs of the Pod to understand running details and troubleshoot exceptions.

asic Info Da	ata Collection Alarm Mana	gement	Recording Rule	instance	diagnostics			
		-	-		_	Pod	idle time	Current
diagnostic colle	ction <sub>cls</sub> .					tke-cls-	2024-07-18 14:52:43	10017
resource utilization	Number of Pods 1/1 (1nuclear 2 G	) (j)	data collectio	on architect	ure diagram <i>C</i>			
collection configuration			Managed clu collection co	uster for data Imponents	रू tmp-operator- To assign a ta			
Target allocation status	Allocated5items Not allocated0ite	ms	Available IP: Security Gro	up sa				
Target status	5 Up				core 3.25 G) (i)			
Agent status	1 items							
Version	tmp-agent(v1.0.7)		user cluster tke/cls					
	proxy-agent(v1.0.8)							
	tmp-operator(v1.1.0)		Cluster Colle	ection Compor	nent Explanation			
	proxy-server(v1.0.2->v1.0.7)							

### **Component Version**



It shows the component version information of the current collection. On the **Component version** page, it displays IP quantity check and the current version, latest version, and component description, upgrade description of each component version.

					Component version	
asic Info Da	ata Collection Alarn	n Management	Recording Rule	instance diagnostics	Current remaining number	er of IP:240
diagnostic colle	ction cls-		data collectio	on architecture diagram 🖉	tmp-agent Latest	
resource utilization	Number of Pods 1/1 (1nu	clear 2 G) (i)			Component description	The tmp-agent, based on vmagent, is respo
collection configuration	<u>:</u>		Managed clu collection co	uster for data 😤 tmp-operator omponents To assign a ta	Current version	instances. v1.0.7
Target allocation status	Allocated5items Not alloc	cated0items	Available IP: Security Gro	240		
Target status	5 Up			Pods <b>3/3</b> (2.50 core 3.25 G) (i)	proxy-agent Latest	
Agent status	1 items				Component description	The proxy-agent connects to the proxy-serve access user clusters, enabling DNS resolution
/ersion	tmp-agent(v1.0.7) proxy-agent(v1.0.8)		user cluster tke/cls		Current version	v1.0.8
	tmp-operator(v1.1.0)		Cluster Coll	ection Component Explanation		
	proxy-server(v1.0.2->v1.	0.7)	Cluster Colle	ection component Explanation c	tmp-operator Latest	1
					Component description	The tmp-operator is responsible for the insta
	~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~					scheduler, and supports the tmp-agent to ru
diagnostic stora	age 🖓				Current version	scheduler, and supports the tmp-agent to ru v1.1.0
diagnostic stora Total rate of meti			The top 10 metri	cs based on the number of ser		
Total rate of metr	ric reporting() 4		The top 10 metric			v1.1.0
Total rate of met	ric reporting() 4 c rate 40.64		kube_resourceque		proxy-server Upgrad	v1.1.0
Total rate of metring the second seco	ric reporting() 4 rate 40.64 109.40	700000	kube_resourceque	ota operations_duration_seconds_bu	proxy-server Upgrad	v1.1.0
Total rate of metric 150.04 Chargeable metric Free metric rate	ric reporting() 4 rate 40.64 109.40	700000 100000	kubelet_runtime_c	ota operations_duration_seconds_bu	proxy-server Upgrac	v1.1.0
Total rate of metric 150.04 Chargeable metric Free metric rate Instance series stu The storage limit f	ric reporting() 4 rate 40.64 109.40 orage limit		kubelet_runtime_ kubelet_runtime_ kube_pod_status_ storage_operation	ota operations_duration_seconds_bu _phase n_duration_seconds_bucket	proxy-server Upgrac	v1.1.0 leable The proxy-server supports multipl path, layer-7 proxy through multip
Total rate of metr <b>150.04</b> Chargeable metric Free metric rate Instance series str The storage limit f The maximum len	ric reporting() 4 rate 40.64 109.40 orage limit for a single metric series	100000	kube_resourceque kubelet_runtime_o kube_pod_status_ storage_operation container_fs_read	ota operations_duration_seconds_bu _phase n_duration_seconds_bucket ls_bytes_total	proxy-server Upgrad	v1.1.0 teable The proxy-server supports multipl path, layer-7 proxy through multipl proceed with caution.
Total rate of metr <b>150.04</b> Chargeable metric Free metric rate Instance series str The storage limit f The maximum len Maximum length of	ric reporting) 4 rate 40.64 109.40 orage limit for a single metric series gth of a metric label name	100000 1024	kubelet_runtime_ kubelet_runtime_ kube_pod_status_ storage_operation	ota operations_duration_seconds_bu _phase n_duration_seconds_bucket ls_bytes_total	Component description	v1.1.0 The proxy-server supports multipl path, layer-7 proxy through multipl proceed with caution. v1.0.2 v1.0.7 t. Fixed the issue of being unable
Total rate of metr <b>1500.04</b> Chargeable metric Free metric rate Instance series str The storage limit f The maximum len Maximum length of Maximum number	ric reporting) 4 rate 40.64 109.40 orage limit for a single metric series gth of a metric label name of Label values for metrics	100000 1024 2048 34	kube_resourceque kubelet_runtime_c kube_pod_status storage_operation container_fs_read container_fs_write	ota operations_duration_seconds_bu _phase n_duration_seconds_bucket ls_bytes_total	proxy-server Upgrad Component description Current version Latest version	v1.1.0 The proxy-server supports multipl path, layer-7 proxy through multipl proceed with caution. v1.0.2 v1.0.7 1. Fixed the issue of being unable connection for scraping.
Total rate of metr <b>1500.04</b> Chargeable metric Free metric rate Instance series str The storage limit f The maximum length of Maximum number The oldest range a maximum allowed	ric reporting) 4 c rate 40.64 109.40 orage limit for a single metric series gth of a metric label name of Label values for metrics r of labels for metrics	100000 1024 2048 34 s 5h	kube_resourceque kubelet_runtime_c kube_pod_status_ storage_operation container_fs_read container_fs_write container_memor	ota operations_duration_seconds_bu _phase n_duration_seconds_bucket ls_bytes_total	proxy-server Upgrad Component description Current version Latest version	The proxy-server supports multiple path, layer-7 proxy through multiple proceed with caution. v1.0.2 v1.0.7 c content 1. Fixed the issue of being unable

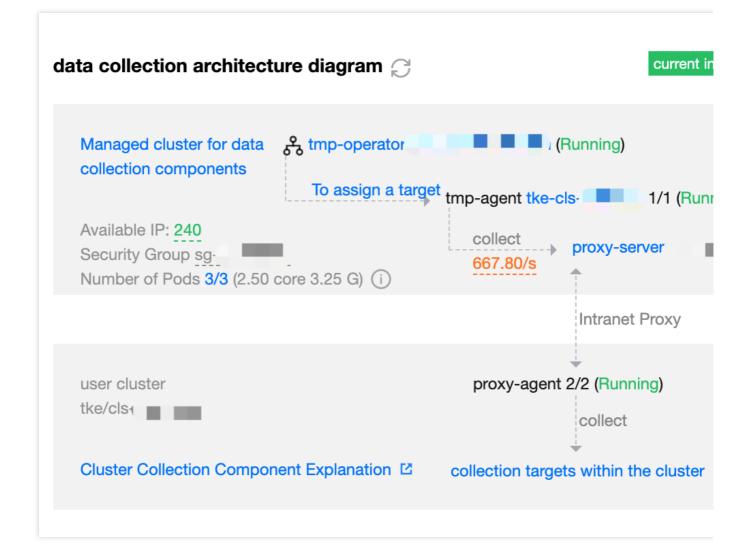
#### Note:

Please try to keep the collection component in the **latest version**, which can be upgraded through the **Upgrade** of the corresponding component.

Components tmp-operator, tmp-agent, and proxy-agent can be upgraded without impact under normal circumstances. During the upgrade of the proxy-server component, there will be collection breakpoints. The breakpoint duration is the time for eks to activate the component Pod, and it will affect the collection of the entire Prometheus instance (including the integration center and container clusters). Please operate cautiously.

#### **Data Collection Architecture Diagram**

The collection architecture diagram provides information about the current collection architecture.



#### Collection Component Managed Cluster:

Available IP count. When the number of IPs is insufficient, an **Insufficient IP Count** alert will appear. Click to enter the **Managed Cluster Subnet Management** page. For specific operations, see Adding Subnet. Managed cluster's security group and its pass-through requirements during normal operation. Some network issues during the collection may be caused by the security group not allowing pass-through. The number of Pods and resource utilization related to the current collection Running status of the collection scheduling component tmp-operator Collection target allocation status Running status of the collection shard component tmp-agent Metric collection rate and inbound bandwidth Running status of the proxy component proxy-server Metric write rate and outbound bandwidth Write target status, including data overwrite of the current instance's corresponding storage and user configuration.

#### User Cluster:

Status of the public network proxy CLB (if enabled) Running status of the proxy component proxy-agent; Status of the collection target within the cluster

### **Storage Diagnosis**

As shown in the figure, the storage diagnosis includes storage-related status and limitations.

diagnostic storage 📿 Total rate of metric reporting(i) The top 10 metrics based on the number kube\_resourcequota Chargeable metric rate 40.64 kubelet\_runtime\_operations\_duration\_secon Free metric rate 109.40 kube\_pod\_status\_phase Instance series storage limit 700000 The storage limit for a single metric series 100000 storage\_operation\_duration\_seconds\_bucke The maximum length of a metric label name 1024 container\_fs\_reads\_bytes\_total Maximum length of Label values for metrics 2048 container\_fs\_writes\_bytes\_total Maximum number of labels for metrics 34 container\_memory\_working\_set\_bytes The oldest range allowed for metric timestamps 5h 10m maximum allowed range for metric timestamps container\_cpu\_usage\_seconds\_total Maximum number of series per single query 100000 rest\_client\_request\_duration\_seconds\_buck maximum number of alarms per unit of time(i) 2000 node\_namespace\_pod\_container:container\_ maximum byte size limit for alarms per unit of time(i) 20971520 es

### **Parameter Description**

Parameter	Description
Total rate of metric reporting	Includes free metrics and paid metrics.
Instance series storage limit	The number of active series exceeding this value will cause the corresponding instance series discarded due to limit exceeded.
The storage limit for a single metric series	Different labels for the same metric name constitute different series. Exceeding this value will result in discard due to limit exceeded.
The maximum length of a metric label name	Exceeding this value will result in discard due to invalid length.



Maximum number of labels for metrics	Exceeding this value will result in discard due to invalid length.
The oldest range allowed for metric timestamps	Indicates the oldest timestamp acceptable in a single series (out of order not allowed).
maximum allowed range for metric timestamps	Indicates the latest timestamp acceptable in a single series (out of order not allowed).
Maximum number of series per single query	Indicates the maximum number of series involved in data query. It is recommended to shorten the range query time or use instant query in suitable scenarios.
Maximum number of alarms per unit of time	Indicates the maximum number of alarms triggered within 5 minutes.
Maximum byte size limit for alarms per unit of time	Indicates the total size limit of fields (such as lable and annotation) for alarms triggered within 5 minutes.
The top 10 metrics based on the number of series	Same metric name and different label keys are considered different series. Storage is subject to the upper limit for single metric series, and a large quantity can cause high cardinality issue.

# Alerting Rule Overview

Last updated : 2024-08-07 22:05:10

TMP allows you to define alert conditions based on Prometheus expressions. Once a metric reaches an alert condition, you will receive notifications through email and SMS, so you can take corresponding measures promptly. **Note:** 

Alerting in TMP combines the alarming capabilities of Cloud Monitor and the alerting capabilities of the open-source Prometheus ecosystem, making alerting more accurate and reasonable.

### Features

TMP supports Alertmanager features such as grouping and silencing to make alerts more reasonable.

You can group the metric data and regularly check and calculate the alerting rules to make alerting quicker and more convenient.

Alerting rules can be defined based on PromQL, making alerting more flexible.

Cloud Monitor's alarm notification templates can be used, which support multiple receiving channels such as email and SMS.

## Components

Term	Description
Rule name	Custom alerting rule name.
Alerting rule	It contains alert trigger condition and duration and should be defined through PromQL.
Alert object	Custom alert title.
Alert message	Custom alert content.
Label	A set of specified labels to be added to the alert.
Annotation	Custom additional alert message.
Notification template	It contains the template name, notification type, recipient, receiving channel. You can define the receiving channel and grouping method.



# Alerting Rule Description

Last updated : 2024-08-07 22:05:17

You can set alert conditions based on Prometheus expressions to monitor the service status in real time and receive prompt notifications when the service is exceptional.

## **Defining Alerting Rule**

Defining an alerting rule in TMP is very similar to defining a recording rule. Below is a sample alerting rule:





```
groups:
- name: example
rules:
- alert: HighRequestLatency
expr: job:request_latency_seconds:mean5m{job="myjob"} > 0.5
for: 10m
labels:
    severity: page
annotations:
    summary: High request latency
```

In an alerting rule file, you can define a set of relevant rules in the same group. In each group, you can define multiple alerting rules. A rule mainly consists of the following parts:

alert: alerting rule name.

expr: alert trigger condition based on a PromQL expression, which is used to calculate whether there is time series data meeting the condition.

for: assessment wait time, which is optional. It indicates how long a trigger condition can last before an alert is sent. New alerts generated during the wait time are in "Pending" status.

labels: custom labels, which are a set of specified labels to be added to alerts.

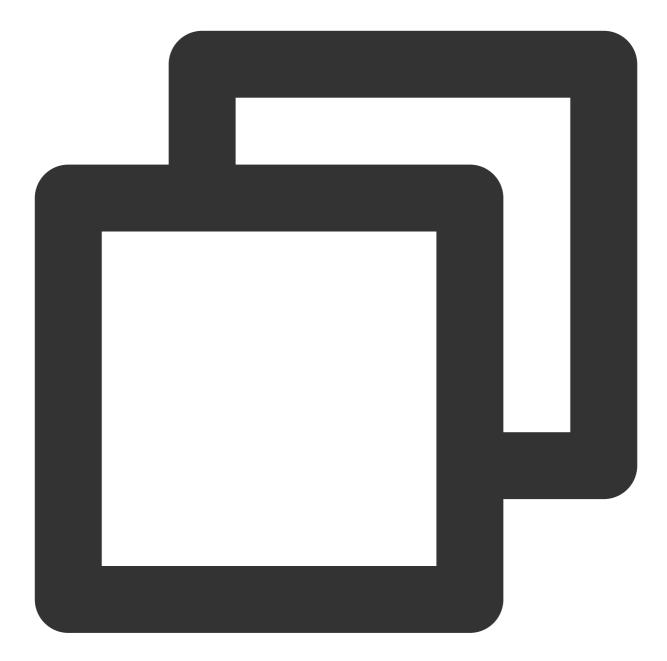
annotations: it is used to specify a set of additional information, such as text that describes alert details. It will be sent to Alertmanager as a parameter when an alert is generated.

## Template

Generally, annotations in an alerting rule file uses summary to describe the summary of alerts and description to describe alert details. In addition, Alertmanager UI will also display the alert information based on the two label values. To make the alert information more readable, TMP allows you to convert label values in labels and annotations into a template.

You can use the \$labels.<labelname> variable to access the value of the specified label on the current alert instance and use \$value to get the sample value calculated through the current PromQL expression.

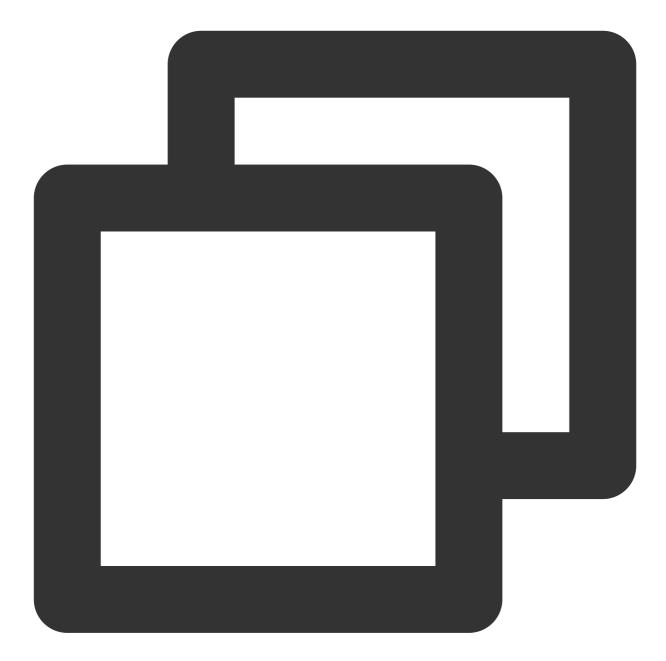




```
# To insert a firing element's label values:
{{ $labels.<labelname> }}
# To insert the numeric expression value of the firing element:
{{ $value }}
```

For example, you can use a template to optimize the readability of the content of summary and description :





```
groups:
- name: example
rules:
# Alert for any instance that is unreachable for >5 minutes.
- alert: InstanceDown
expr: up == 0
for: 5m
labels:
    severity: page
annotations:
```



```
summary: "Instance {{ $labels.instance }} down"
description: "{{ $labels.instance }} of job {{ $labels.job }} has been down f
# Alert for any instance that has a median request latency >1s.
- alert: APIHighRequestLatency
expr: api_http_request_latencies_second{quantile="0.5"} > 1
for: 10m
annotations:
   summary: "High request latency on {{ $labels.instance }}"
description: "{{ $labels.instance }} has a median request latency above 1s (c
```

# **Creating Alerting Rule**

Last updated : 2024-08-07 22:05:23

This document describes how to create an alerting rule in the TMP console. With such a rule, you will receive notifications when some metrics are exceptional, so that you can take corresponding measures promptly.

## Prerequisites

You have created a managed TKE cluster. You have created a TMP instance.

You have installed the Prometheus agent and other monitoring components.

### Directions

For more information on alerting rule, please see Alerting Rule Description.

1. Log in to the TMP console.

2. In the instance list, select the corresponding TMP instance and click **Alerting Rule** on the left.

3. On the alerting rule management page, click **Create** and configure the alerting rule information in the pop-up window.

**Rule Template Type**: select a rule template type. For more information, please see Alerting Rule Type Description. **Rule Name**: you can use the default rule name or customize it.

**PromQL-Based Rule**: you can use the default value or customize it. It indicates an alert trigger condition based on a PromQL expression, which is used to calculate whether there is time series data meeting the condition.

**Duration**: you can use the default value or customize it. It indicates how long a trigger condition can last before an alert is sent.

Alert Object: you can customize the alert title.

Alert Message: you can use the default value or customize it. It indicates the custom alert content.

**Advanced Configuration**: you can toggle it on for configuration, which contains configuration items for labels and annotations.

**Labels**: you can use the default value or customize it. It indicates a set of specified labels to be added to alerts. You can match the corresponding processing method based on the labels of a received alert.

**Annotations**: you can use the default value or customize it. It indicates the user-defined additional alert information. **Alert Notification**: you can customize the alert notification template, which contains the template name, notification type, recipient, and receiving channel. For more information, please see Notification Template.

Policy Template * Policy Name *	Please select a policy	template				-
Policy Name *						
	Please enter a policy	name				
PromQL-Based Rule *	Please enter the rule					
Duration	minu 🔻					
Alarm Object *	Please enter the alarr	n object				
Alarm Message *	Please enter the alarr	n message				
abels	Key: Please enter		Value:	Please enter	Save	
Annotations	Key: Please enter	· · · · · · · · · · · · · · · · · · ·	Value:	Please enter	Save	
Alarm Notification *	Select Template	Create 🗳				
	0 selected. 3 more can	be selected				
	Notification Templa	te Name		Included Open	ations	Operat
	The notification template list is empty. You can select some by clicking "Select Template".					
		The notification templa	ate list is	empty. You can select some by clicki	ng "Select Template".	

## **Disabling Alerting Rule**

Last updated : 2024-08-07 22:05:28

This document describes how to disable an alerting rule on the target instance in the TMP console.

## Directions

- 1. Log in to the TMP console.
- 2. In the instance list, select the corresponding TMP instance and click Alerting Rule on the left.
- 3. Find the alerting rule to be disabled and click **Disable** in the **Operation** column.
- 4. In the pop-up window, click **OK**.

Disable /	Alarm Policy
<b>(i)</b>	Are you sure you want to delete the selected alarm policy?
	OK Cancel

# Rule Type Description(old)

Last updated : 2024-08-07 22:05:33

TMP presets the master component, kubelet, resource use, workload, and node alert templates for TKE clusters.

### Kubernetes master component

The following metrics are provided for non-managed clusters:

Rule Name	Rule Expression	Duration	Desc
Error with client access to APIServer	(sum(rate(rest_client_requests_total{code=~"5"}[5m])) by (instance, job, cluster_id) / sum(rate(rest_client_requests_total[5m])) by (instance, job, cluster_id))> 0.01	15m	The erate of accentric the APIS abov
Imminent expiration of the client certificate for APIServer access	apiserver_client_certificate_expiration_seconds_count{job="apiserver"} > 0 and on(job) histogram_quantile(0.01, sum by (cluster_id, job, le) (rate(apiserver_client_certificate_expiration_seconds_bucket{job="apiserver"} [5m]))) < 86400	None	The certif for APIS acce expir hours
Recording API error	sum by(cluster_id, name, namespace) (increase(aggregator_unavailable_apiservice_count[5m])) > 2	None	The recor API repoi error last 5 minu
Low recording API availability	(1 - max by(name, namespace, cluster_id) (avg_over_time(aggregator_unavailable_apiservice[5m]))) * 100 < 90	5m	The availa of the recor API s in the



			minu belov
APIServer fault	absent(sum(up{job="apiserver"}) by (cluster_id) > 0)	5m	APIS disar from colle targe
Scheduler fault	absent(sum(up{job="kube-scheduler"}) by (cluster_id) > 0)	15m	The sche disar from colle targe
Controller manager fault	absent(sum(up{job="kube-controller-manager"}) by (cluster_id) > 0)	15m	The contr mana disar from colle targe

## Kubelet

Rule Name	Rule Expression	Duration	Des
Exceptional node status	kube_node_status_condition{job=~".*kube-state- metrics",condition="Ready",status="true"} == 0	15m	The stat exco for c min
Unreachable node	kube_node_spec_taint{job=~".*kube-state- metrics",key="node.kubernetes.io/unreachable",effect="NoSchedule"} == 1	15m	The unre and wor be sche aga
Too many Pods	count by(cluster_id, node) ((kube_pod_status_phase{job=~".*kube-state- metrics",phase="Running"} == 1) * on(instance,pod,namespace,cluster_id)	15m	The of P

running on node	group_left(node) topk by(instance,pod,namespace,cluster_id) (1, kube_pod_info{job=~".*kube-state-metrics"}))/max by(cluster_id, node) (kube_node_status_capacity_pods{job=~".*kube-state-metrics"} != 1) > 0.95		runr the clos upp
Node status fluctuation	sum(changes(kube_node_status_condition{status="true",condition="Ready"} [15m])) by (cluster_id, node) > 2	15m	The stat fluct betv norr exce
Imminent expiration of the kubelet client certificate	kubelet_certificate_manager_client_ttl_seconds < 86400	None	The cliei cert will 24 ł
Imminent expiration of the kubelet server certificate	kubelet_certificate_manager_server_ttl_seconds < 86400	None	The serv cert will 24 ł
Kubelet client certificate renewal error	increase(kubelet_certificate_manager_client_expiration_renew_errors[5m]) > 0	15m	An e occ whil rene kub cert
Kubelet server certificate renewal error	increase(kubelet_server_expiration_renew_errors[5m]) > 0	15m	An e occ whil rene kub serv cert
Time- Consuming PLEG	histogram_quantile(0.99, sum(rate(kubelet_pleg_relist_duration_seconds_bucket[5m])) by (cluster_id, instance, le) * on(instance, cluster_id) group_left(node) kubelet_node_name{job="kubelet"}) >= 10	5m	The perc PLE ope dura exco seco

Time- Consuming Pod start	histogram_quantile(0.99, sum(rate(kubelet_pod_worker_duration_seconds_bucket{job="kubelet"} [5m])) by (cluster_id, instance, le)) * on(cluster_id, instance) group_left(node) kubelet_node_name{job="kubelet"} > 60	15m	The perc Poc dura exco seco
Kubelet fault	absent(sum(up{job="kubelet"}) by (cluster_id) > 0)	15m	Kub disa fron colle targ

## Kubernetes Resource Use

Rule Name	Rule Expression	Duration	Description
Cluster CPU resource overload	<pre>sum by (cluster_id) (max by (cluster_id, namespace, pod, container) (kube_pod_container_resource_requests_cpu_cores{job=~".*kube- state-metrics"}) * on(cluster_id, namespace, pod) group_left() max by (cluster_id, namespace, pod) (kube_pod_status_phase{phase=~"Pending Running"} == 1))/sum by (cluster_id) (kube_node_status_allocatable_cpu_cores)&gt;(count by (cluster_id) (kube_node_status_allocatable_cpu_cores)-1) / count by (cluster_id) (kube_node_status_allocatable_cpu_cores)</pre>	5m	Too many CPU cores are applied for by Pods in the cluster, and no more failed nodes can be tolerated
Cluster memory resource overload	<pre>sum by (cluster_id) (max by (cluster_id, namespace, pod, container) (kube_pod_container_resource_requests_memory_bytes{job=~".*kube- state-metrics"}) * on(cluster_id, namespace, pod) group_left() max by (cluster_id, namespace, pod) (kube_pod_status_phase{phase=~"Pending Running"} == 1))/sum by (cluster_id) (kube_node_status_allocatable_memory_bytes) &gt; (count by (cluster_id) (kube_node_status_allocatable_memory_bytes)-1) / count by (cluster_id) (kube_node_status_allocatable_memory_bytes)</pre>	5m	Too much memory is applied for by Pods in the cluster, and no more failed nodes can be tolerated
Cluster CPU	sum by (cluster_id) (kube_resourcequota{job=~".*kube-state-metrics", type="hard", resource="cpu"})/sum by (cluster_id) (kube_node_status_allocatable_cpu_cores) > 1.5	5m	The CPU quota in the cluster



quota overload			exceeds the total number of allocable CPU cores
Cluster memory quota overload	sum by (cluster_id) (kube_resourcequota{job=~".*kube-state-metrics", type="hard", resource="memory"}) / sum by (cluster_id) (kube_node_status_allocatable_memory_bytes) > 1.5	5m	The memory quota in the cluster exceeds the total amount of allocable memory
Imminent runout of quota resources	sum by (cluster_id, namespace, resource) kube_resourcequota{job=~".*kube-state-metrics", type="used"} / sum by (cluster_id, namespace, resource) (kube_resourcequota{job=~".*kube- state-metrics", type="hard"} > 0) >= 0.9	15m	The quota resource utilization exceeds 90%
High proportion of restricted CPU execution cycles	<pre>sum(increase(container_cpu_cfs_throttled_periods_total{container!="", }[5m])) by (cluster_id, container, pod, namespace) /sum(increase(container_cpu_cfs_periods_total{}[5m])) by (cluster_id, container, pod, namespace) &gt; ( 25 / 100 )</pre>	15m	The proportion of restricted CPU execution cycles is high
High Pod CPU utilization	<pre>sum(rate(container_cpu_usage_seconds_total{job="kubelet", metrics_path="/metrics/cadvisor", image!="", container!="POD"}[1m])) by (cluster_id, namespace, pod, container) / sum(kube_pod_container_resource_limits_cpu_cores) by (cluster_id, namespace, pod, container) &gt; 0.75</pre>	15m	The Pod CPU utilization exceeds 75%
High Pod memory utilization	<pre>sum(rate(container_memory_working_set_bytes{job="kubelet", metrics_path="/metrics/cadvisor", image!="", container!="POD"}[1m])) by (cluster_id, namespace, pod, container) /sum(kube_pod_container_resource_limits_memory_bytes) by (cluster_id, namespace, pod, container) &gt; 0.75</pre>	15m	The Pod memory utilization exceeds 75%

## Kubernetes Workload



Rule Name	Rule Expression	Duration
Frequent Pod restarts	increase(kube_pod_container_status_restarts_total{job=~".*kube-state- metrics"}[5m]) > 0	15m
Exceptional Pod status	<pre>sum by (namespace, pod, cluster_id) (max by(namespace, pod, cluster_id) (kube_pod_status_phase{job=~".*kube-state-metrics", phase=~"Pending Unknown"}) * on(namespace, pod, cluster_id) group_left(owner_kind) topk by(namespace, pod) (1, max by(namespace, pod, owner_kind, cluster_id) (kube_pod_owner{owner_kind!="Job"}))) &gt; 0</pre>	15m
Exceptional container status	sum by (namespace, pod, container, cluster_id) (kube_pod_container_status_waiting_reason{job=~".*kube-state-metrics"}) > 0	1h
Deployment version mismatch	kube_deployment_status_observed_generation{job=~".*kube-state-metrics"} !=kube_deployment_metadata_generation{job=~".*kube-state-metrics"}	15m
Deployment replica quantity mismatch	(kube_deployment_spec_replicas{job=~".*kube-state-metrics"} != kube_deployment_status_replicas_available{job=~".*kube-state-metrics"}) and (changes(kube_deployment_status_replicas_updated{job=~".*kube-state- metrics"}[5m]) == 0)	15m
StatefulSet version mismatch	kube_statefulset_status_observed_generation{job=~".*kube-state-metrics"} != kube_statefulset_metadata_generation{job=~".*kube-state-metrics"}	15m
StatefulSet	(kube_statefulset_status_replicas_ready{job=~".*kube-state-metrics"} !=	15m



replica quantity mismatch	kube_statefulset_status_replicas{job=~".*kube-state-metrics"}) and ( changes(kube_statefulset_status_replicas_updated{job=~".*kube-state- metrics"}[5m]) == 0)	
Ineffective StatefulSet update	(maxwithout(revision)(kube_statefulset_status_current_revision{job=~".*kube- state-metrics"}unless kube_statefulset_status_update_revision{job=~".*kube- state-metrics"})*(kube_statefulset_replicas{job=~".*kube-state- metrics"}!=kube_statefulset_status_replicas_updated{job=~".*kube-state- metrics"})) and (changes(kube_statefulset_status_replicas_updated{job=~".*kube-state- metrics"}[5m])==0)	15m
Frozen DaemonSet change	<pre>((kube_daemonset_status_current_number_scheduled{job=~".*kube-state- metrics"}!=kube_daemonset_status_desired_number_scheduled{job=~".*kube- state-metrics"}) or (kube_daemonset_status_number_misscheduled{job=~".*kube-state- metrics"}!=0) or (kube_daemonset_updated_number_scheduled{job=~".*kube- state- metrics"}!=kube_daemonset_status_desired_number_scheduled{job=~".*kube- state-metrics"}) or (kube_daemonset_status_number_available{job=~".*kube- state- metrics"}!=kube_daemonset_status_desired_number_scheduled{job=~".*kube- state- metrics"}!=kube_daemonset_status_desired_number_scheduled{job=~".*kube- state- metrics"}!=kube_daemonset_status_desired_number_scheduled{job=~".*kube- state- metrics"}]==0)</pre>	15m
DaemonSet not scheduled on some nodes	kube_daemonset_status_desired_number_scheduled{job=~".*kube-state- metrics"} - kube_daemonset_status_current_number_scheduled{job=~".*kube- state-metrics"} > 0	10m
Faulty scheduling of DaemonSet on some nodes	kube_daemonset_status_number_misscheduled{job=~".*kube-state-metrics"} > 0	15m
Excessive Job execution	kube_job_spec_completions{job=~".*kube-state-metrics"} - kube_job_status_succeeded{job=~".*kube-state-metrics"} > 0	12h
Job execution failure	kube_job_failed{job=~".*kube-state-metrics"} > 0	15m
Mismatch between replica	(kube_hpa_status_desired_replicas{job=~".*kube-state-metrics"} != kube_hpa_status_current_replicas{job=~".*kube-state-metrics"}) and	15m



quantity and HPA	changes(kube_hpa_status_current_replicas[15m]) == 0	
Number of replicas reaching maximum value in HPA	kube_hpa_status_current_replicas{job=~".*kube-state-metrics"} == kube_hpa_spec_max_replicas{job=~".*kube-state-metrics"}	15m
Exceptional PersistentVolume status	kube_persistentvolume_status_phase{phase=~"Failed Pending",job=~".*kube- state-metrics"} > 0	15m

## Kubernetes Node

Rule Name	Rule Expression	Duration	Description
Imminent runout of filesystem space	(node_filesystem_avail_bytes{job="node- exporter",fstype!=""}/node_filesystem_size_bytes{job="node- exporter",fstype!=""}*100<15 and predict_linear(node_filesystem_avail_bytes{job="node- exporter",fstype!=""}[6h],4*60*60)<0 and node_filesystem_readonly{job="node- exporter",fstype!=""}==0)	1h	It is estimated that the filesystem space will be used up in 4 hours
High filesystem space utilization	(node_filesystem_avail_bytes{job="node- exporter",fstype!=""}/node_filesystem_size_bytes{job="node- exporter",fstype!=""}*100<5 and node_filesystem_readonly{job="node- exporter",fstype!=""}==0)	1h	The available filesystem space is below 5%
Imminent runout of filesystem inodes	<pre>(node_filesystem_files_free{job="node- exporter",fstype!=""}/node_filesystem_files{job="node- exporter",fstype!=""}*100&lt;20 and predict_linear(node_filesystem_files_free{job="node- exporter",fstype!=""}[6h],4*60*60)&lt;0 and node_filesystem_readonly{job="node- exporter",fstype!=""}==0)</pre>	1h	It is estimated that the filesystem inodes will be used up in 4 hours
High	(node_filesystem_files_free{job="node-	1h	The proportion of



filesystem inode utilization	exporter",fstype!=""}/node_filesystem_files{job="node- exporter",fstype!=""}*100<3 and node_filesystem_readonly{job="node- exporter",fstype!=""}==0)		available inodes is below 3%
Unstable network interface status	changes(node_network_up{job="node- exporter",device!~"veth.+"}[2m])	2m	The network interface status is unstable and frequently changes between "up" and "down"
Network interface data reception error	increase(node_network_receive_errs_total[2m]) > 10	1h	An error occurred while the network interface received data
Network interface data sending error	increase(node_network_transmit_errs_total[2m]) > 10	1h	An error occurred while the network interface sent data
Unsynced server clock	min_over_time(node_timex_sync_status[5m]) == 0	10m	The server time has not been synced recently. Please check whether NTP is correctly configured
Server clock skew	(node_timex_offset_seconds>0.05 and deriv(node_timex_offset_seconds[5m])>=0) or (node_timex_offset_seconds<-0.05 and deriv(node_timex_offset_seconds[5m])<=0)	10m	The server clock skew exceeds 300 seconds. Please check whether NTP is correctly configured

## Notification Template

Last updated : 2024-08-07 22:05:39

This document describes how to create a notification template in the Cloud Monitor alarming module.

## Use Cases

One template can be quickly reused for multiple policies, eliminating the need to repeatedly configure user notifications.

User notification methods can be configured in a more personalized way. For example, you can configure the alarm receiving channel as SMS/email by day and phone by night.

### Prerequisites

View notification templates: the sub-account must have the read permission of Cloud Monitor.

Create and edit notification templates: the sub-account must have the write permission of Cloud Monitor.

#### Note:

For more information on how to grant sub-accounts permissions, please see Cloud Access Management (CAM).

### **Use Limits**

Feature	Limit
User notification	Up to five items can be added
API callback	Up to three URLs accessible over the public network can be entered

### Directions

#### Creating notification template

1. Enter the Alarm Notification Template page in the Cloud Monitor console.

2. Click **Create** and enter relevant information in **Create Notification Template**.

Template Name: enter a custom template name.

Notification type:

Alarm triggered: a notification will be sent when an alarm is triggered.

Alarm cleared: a notification will be sent when an alarm is resolved.

User notification:

Recipient Object: you can choose a recipient group or recipient. If you need to create a group, please see Creating Alarm Recipient Group.

Notification Period: define the time period for receiving alarms.

Receiving Channel: four alarm channels are supported: email, SMS, WeChat, and phone. You can also set different channels and notification periods in different user dimensions.

Description of phone alarm settings:

Polling Times: the maximum number of dials for each polled recipient when there is no valid reach.

Polling Sequence: alarm calls will be dialed according to the order of the recipients. You can adjust the order of calling by dragging up and down recipients.

Polling Interval: time interval at which alarm calls will be dialed according to the order of the recipients.

Reach Notification: notifications will be to all recipients after successful reception of the call or calling all recipients. SMS messages are counted against the quota.

API Callback: you can enter a URL accessible over the public network as the callback API address, and Cloud Monitor will push alarm messages to it promptly. If the HTTP response returns code 200, the verification is successful. For more information on alarm callback fields, please see Alarm Callback.

#### Note:

After you save the callback URL, the system will automatically verify your URL once. The timeout threshold for this verification is 5 seconds. When an alarm policy created by the user is triggered or the alarm is resolved, the alarm messages will be pushed through the API callbacks. An alarm message can be pushed up to three times, and the timeout threshold for each request is 5 seconds.

When an alarm policy created by the user is triggered or the alarm is resolved, the alarm messages will be pushed through the API callbacks. API callbacks also support repeated alarms.

The outbound IP of the Cloud Monitor callback API is dynamically and randomly allocated, so no specific IP information can be provided to you, but the IP port is fixed at 80. We recommend you configure a weighted opening policy in the security group based on port 80.



Basic Info													
'emplate Vame *	example												
otification emplate 🛈	✓ Alarm Trigger	- 🔽 /	Alarm Recover	у									
otification	English				<b>*</b>								
	Recipient Object Notification Period Receiving Channel	User 00:00:00 ~ ✓ Email	<ul> <li>23:59:59</li> <li>SMS</li> </ul>	3						¢	Add User	Delete	
					Ade	d User Noti	ification						
l Callback <b>)</b>	Enter a URL acc	cessible over p	ublic network	s as the A	PI callbad	ck address (	(domain na	me or IP[:po	ort][/path]), e	.g. https://e	example.com:808	30/a Delete	View Usage Guide
					P	Add API Call	llback						
	i It supp	orts pushing to	o the WeCom	group rol	ootCome	and try it o	out 🖸						

#### Default notification template

The system automatically creates a default notification template for you as detailed below:

Feature	Default Configuration
Template name	Preset notification template
Notification type	Alarm trigger, alarm recovery
Alarm recipient	Root account admin
Notification period	00:00-23:59:59 (all day)
Receiving channel	Email, SMS

#### **Deleting template**

#### Note:

The default notification template cannot be deleted.

- 1. Find the name of the template to be deleted and click **Delete** in the **Operation** column.
- 2. In the pop-up window, click **OK**.

#### **Replicating template**

- 1. Find the name of the template to be replicated and click **Replicate** in the **Operation** column.
- 2. Modify the information in the redirected page or click **Complete** directly.

# Tag Examples

Last updated : 2024-08-07 22:05:51

#### Overview

A tag is a key-value pair provided by Tencent Cloud to identify a resource in the cloud.

You can use tags to classify TMP resources based on various factors such as service, usage, and owner. With tags, you can quickly sift through the resource pool and find the corresponding resources. The values of tag keys do not mean anything to Tencent Cloud semantically and will be parsed and matched strictly according to the string. Tencent Cloud will not use your set tags, which are used only for resource management. Below is a specific use case to show how a tag is used.

#### **Use Case Background**

A company owns 10 TMP instances in Tencent Cloud. Distributed in three departments (ecommerce, gaming, and entertainment), these instances are used to serve internal business lines such as marketing, game A, game B, and post-production. The OPS owners of the three departments are John, Jane, and Harry, respectively.

#### Setting Tag

To facilitate management, the company categorizes its TMP resources with tags and defines the following tag keyvalue pairs:

Tag Key	Tag Value
Department	Ecommerce, gaming, and entertainment
Business	Marketing, game A, game B, and post-production
OPS owner	John, Jane, and Harry

These tags are bound to TMP instances in the following way:

Instance ID	Department	Business	OPS Owner
prom-1jqwv1	Ecommerce	Marketing	Harry
prom-1jqwv12	Ecommerce	Marketing	Harry
prom-1jqwv13	Gaming	Game A	John
prom-1jqwv13	Gaming	Game B	John

prom-1jqwv14	Gaming	Game B	John
prom-1jqwv15	Gaming	Game B	Jane
prom-1jqwv16	Gaming	Game B	Jane
prom-1jqwv17	Gaming	Game B	Jane
prom-1jqwv18	Entertainment	Post-production	Harry
prom-1jqwv19	Entertainment	Post-production	Harry
prom-1jqwv110	Entertainment	Post-production	Harry

#### Using Tag

Filter out the TMP instances in the charge of Harry

Filter out the TMP instances where the OPS owner is "Harry". For detailed directions, please see Using Tag.

Filter out the TMP instances in the charge of Jane in the gaming department

Filter out the TMP instances where the department is "gaming" and OPS owner is "Jane". For detailed directions, please see Using Tag.

# Using Tag

Last updated : 2024-08-07 22:06:01

This document describes how to filter instance resources by tags in the TMP console.

## Directions

- 1. Log in to the TMP console.
- 2. Select the region at the top of the instance list page.
- 3. In the search box in the top-right corner of the instance list, click the blank space and select Tag.

v4 Addr value2 • ×	Tag:	Separate key	words with " "; press Er	nter to separate filter ta	ags	i	Q
tag1:value2		tag2	•	value2	•	×	Ş
kkkkkkkkvvvvvvv	4 Addr	+ Add					
OK Cancel		ОК					

- 4. Select the corresponding conditions in the tag filter selection box and click OK.
- 5. If you need to adjust the tag conditions, click the tag content after Tag: in the search box to edit it.
- 6. You can also directly click the corresponding tag value in the instance list to filter instances as shown below:

Create Edit Tag					Tag: tag2       value2         Separate keywords with " "; press Enter to separate filter tags				<b>13</b> (i) (1
Instance ID/Name	Monit 🔻	AZ 🕈	Network	Configura	IPv4 Addr	Billing Mo	Tag (key:value) 🛈	Creation T	Operation
prom 11	<b>II</b> ⊘ Running	Guangzhou Zone 2	Network: default_vpc Subnet: default_vpc_ subnet	Data retention period: 15 day(s) Specs name: Shared Edition	-	Trial Edition	tag2:value2 kkkkkk:vvvvvvv tag1:value1	2021/11/15 15:55:00	Manage More ▼
otal items: 1							10 🔻 / page 🛛 🖌 🔺	1 /1p	age 🕨

# **Editing Tag**

Last updated : 2024-08-07 22:06:06

This document describes how to edit the tags of an instance in the TMP console.

## Directions

#### Edit tag for a single instance

1. Log in to the TMP console.

2. On the instance management page, select the instance for which to edit tags and click More > Instance

Configuration > Edit Tag as shown below:

Instance ID/Name	Monit <b>T</b>	AZ T	Network	Configura	IPv4 Addr	Billing Mo	Tag (key:value)
prom11	<b>II</b> ⊘ Running	Guangzhou Zone 2	Network: default_vpc Subnet: default_vpc_ subnet	Data retention period: 15 day(s) Specs name: Shared Edition		Trial Edition	tag2:value2 kkkkkkk:vvvvvvv tag1:value1
							Modify Name

3. In the **Selected 1 resource** window that pops up, add, modify, or delete tags based on your actual needs.



Edit Tags				×
		s by category from differen ts, please go to <b>Manage Tag</b>		ons. If the existing
1 resource selected				
tag2	▼	value2	•	×
kkkkkk	▼	~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~	•	×
tag1	•	value1	•	×
Tag key	•	Tag value	•	×
+ Add				
		OK Cancel		

# Use Limits

Last updated : 2024-08-07 22:06:11

A tag is a key-value pair. You can set tags for TMP instances in the TMP console to manage them in a categorized manner. Then, you can easily filter and find desired resources with tags.

## Quantity Limit

One Tencent Cloud resource can have up to 50 tags.

## Tag Key Limit

It cannot begin with "qcloud", "tencent", or "project" as they are reserved by the system. It can contain only letters, digits, spaces, and certain special symbols  $(+, -, =, ., \_, ., ., ., ., ., ., ., ., ., .)$ . It can contain up to 255 characters.

## Tag Value Limit

It can contain only letters, digits, spaces, and certain special symbols  $(+, -, =, ., \_, :, /, @)$ . It can contain up to 127 characters.

# Access Control Overview

Last updated : 2024-08-07 22:06:26

If you have multiple users managing the TMP service, and they all share your Tencent Cloud account access key, you may face the following problems:

The risk of your key being compromised is high since multiple users are sharing it.

Your users might introduce security risks from maloperations due to the lack of user access control.

You can avoid the above problems by allowing different users to manage different services through sub-accounts. By default, sub-accounts have no permissions to use TMP. Therefore, you need to create a policy to grant different permissions to sub-accounts.

### Overview

Cloud Access Management (CAM) is a web-based Tencent Cloud service that helps you securely manage and control access permissions of your Tencent Cloud resources. Using CAM, you can create, manage, and terminate users (groups), and control the Tencent Cloud resources that can be used by the specified user through identity and policy management.

When using CAM, you can associate a policy with a user or user group to allow or forbid them to use specified resources to complete specified tasks. For more information on CAM policies, please see Element Reference. For more information on how to use CAM policies, please see Policy.

You can skip this section if you don't need to manage permissions of TMP resources for sub-accounts. This won't affect your understanding and use of the other sections of the document.

## **Getting Started**

A CAM policy must authorize or deny the use of one or more TMP operations. At the same time, it must specify the resources that can be used for the operations (which can be all resources or partial resources for certain operations). A policy can also include the conditions set for the manipulated resources.

Certain APIs of TMP don't support resource-level permissions, which means that for this type of API operations, you cannot specify a given resource for use when they are performed; instead, you must specify all resources for use.

# **Setting Policy**

Last updated : 2024-08-07 22:06:32

## Overview

Access policies can be used to grant access to TMP instances. They use JSON-based access policy syntax. You can authorize specified principals to perform specified operations on specified TMP resources through the access policy syntax.

The access policy syntax describes the basic elements and usage of the policy. For the description of the policy syntax, please see Permission.

## Elements in Access Policy

An access policy contains the following elements with basic meanings:

**statement**: it describes the details of one or more permissions. It contains a permission or permission set of multiple other elements such as effect, action, resource, and condition. One policy must and can have only one statement.

effect: it is required and describes the result of a statement. The result can be an "allow" or "explicit deny".

**action**: it is required and describes the allowed or denied action (operation). An operation can be an API (prefixed with "name") or a feature set (a set of specific APIs prefixed with "permid").

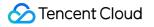
**resource**: it is required and describes the details of authorization. A resource is described in a six-segment format. Detailed resource definitions vary by product.

**condition**: it is optional and describes the condition for the policy to take effect. A condition consists of operator, action key, and action value. A condition value may contain information such as time and IP address. Some services allow you to specify additional values in a condition.

## **Element Usage**

#### Specifying effect

If access to a resource is not explicitly granted (allowed), then it is implicitly denied. It can also be explicitly denied, which ensures that users cannot access the resource even if they are granted the access permission by other policies. Below is an example of specifying the "allow" effect:



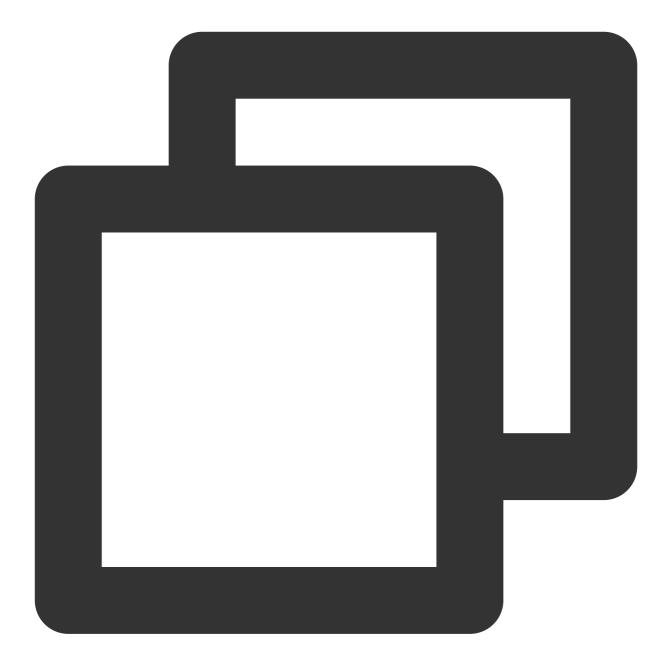


"effect" : "allow"

#### **Specifying action**

Cloud Monitor defines console operations that can be specified in a policy. The specified operations are divided into reading part of APIs (monitor:Describe\*) and all APIs (monitor:\*) according to the operation nature. Below is an example of specifying the allowed operations:



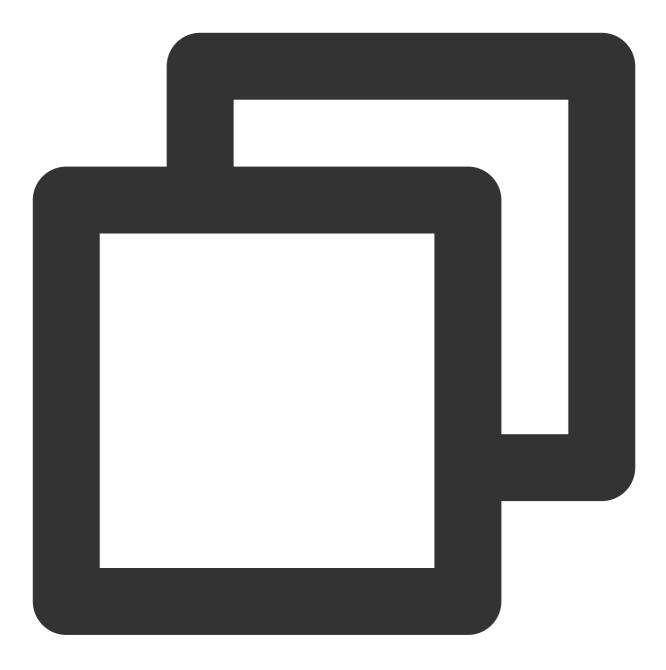


```
"action": [
    "name/monitor:Describe*"
]
```

#### Specifying resource

The resource element describes one or more operation objects, such as TMP resources. All resources can use the following six-segment format:





#### qcs:project\_id:service\_type:region:account:resource

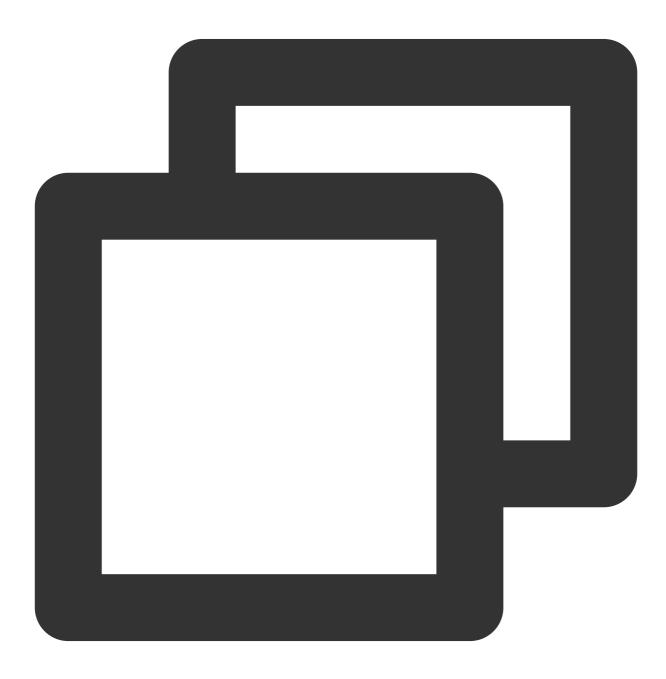
#### The parameters are as detailed below:

Parameter	Description	Required
qcs	Tencent Cloud service abbreviation, which indicates a service of Tencent Cloud	Yes
project_id	Project information, which is only used to enable compatibility with legacy CAM logic and generally can be left empty	No



service_type	Product abbreviation, which is monitor here	Yes
region	Region information	Yes
account	Root account information of the resource owner, i.e., root account ID in the format of uin/\${OwnerUin} , such as uin/10000000001	Yes
resource	Resource details prefixed with instance	Yes

Below is a sample six-segment TMP resource description:



"resource":["qcs::monitor:ap-guangzhou:uin/10000000001:prom-instance/prom-73jingds

#### **Specifying condition**

The access policy syntax allows you to specify the condition when granting permissions, which is mainly used to set tag authentication. The tag condition takes effect only for clusters bound with the tag. Below is a sample tag policy:



```
"condition": {
    "for_any_value:string_equal": {
        "qcs:tag": [
```

```
"testkey&testvalue"
]
}
}
```

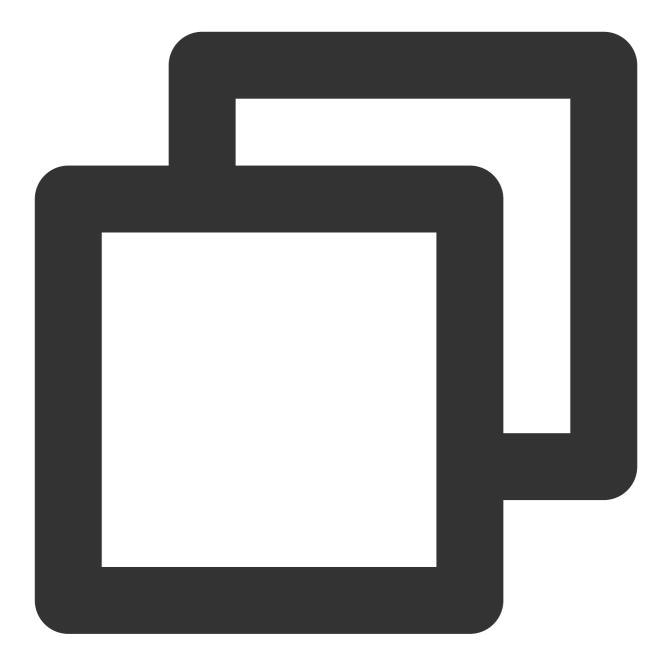
This statement means that the policy contains resources whose tag key is testkey and tag value is testvalue .

### Use Cases

#### Based on tag

In the following case, the policy is to allow access to the resource whose instance ID is prom-73jingds under UIN 125000000 and the resources whose tag key is testkey and tag value is testvalue (if this instance doesn't have the testkey& testvalue tag, it cannot be accessed).





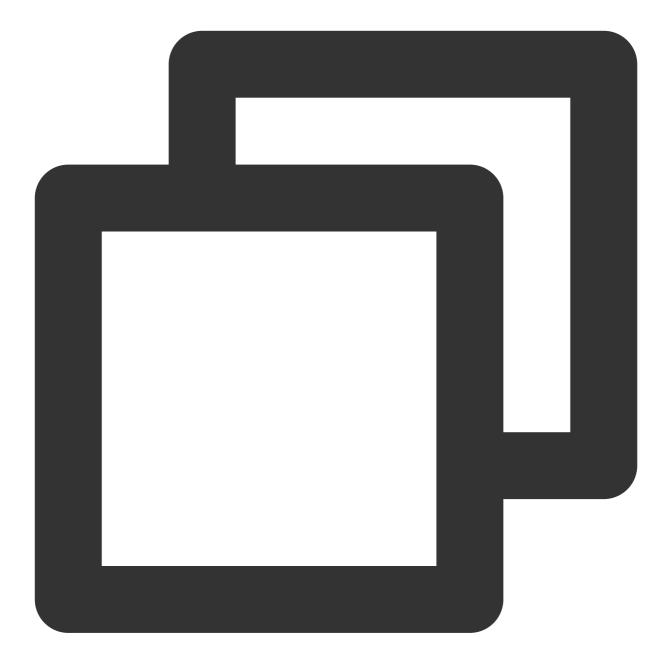
#### **Based on resource**

Grant the read and write permissions of specified resources based on resource ID. The root account ID is 1250000000:

Configure read-only access to resources in the Guangzhou region.

Sample: granting the sub-user read-only access to the instance1(prom-73jingds) and instance2(prom-65jidfafk) resources.





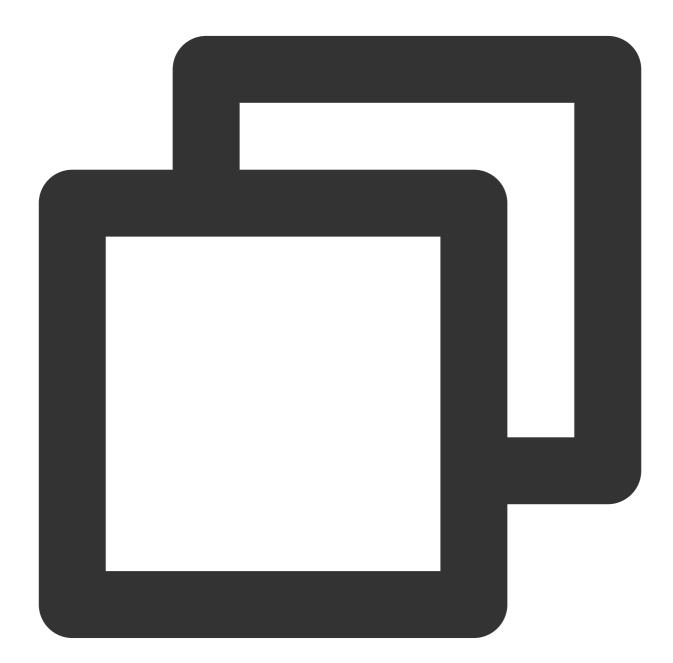
```
{
    "version": "2.0",
    "statement": [{
        "effect": "allow",
        "action": [
            "name/monitor:Describe*"
        ],
        "resource": [
            "qcs::monitor:ap-guangzhou:uin/125000000:prom-instance/prom-73jingds",
            "qcs::monitor:ap-guangzhou:uin/125000000:prom-instance/prom-65jidfafk"
        ],
```

#### 🔗 Tencent Cloud

```
"condition": []
}]
}
```

Configure partial read/write access to resources in the Guangzhou region.

Sample: granting the sub-user the permission to delete the instance1(prom-73jingds) resource



```
{
    "version": "2.0",
    "statement": [{
        "effect": "allow",
```

```
"resource": [
    "qcs::monitor:ap-guangzhou:uin/125000000:prom-instance/prom-73jingds"
],
"action": [
    "name/monitor:TerminatePrometheusInstances"
]
}]
}
```

# **Granting Policy**

Last updated : 2024-08-07 22:06:37

### Custom Policy for TMP

If preset policies cannot meet your needs, you can click Create Custom Policy to create custom policies.

icy							
<b>i</b> A	ssociate users or user groups with policies to g	rant permissions.					
Create	Custom Policy Delete			All Policies	Preset Policy	Custom Policy	Search by policy
	Policy Name	Service Type <b>Y</b>	Description			Last Mo	odified
	bsp	BatchCompute	bsp			2018-05	5-04 20:04:41
	QcloudCCNFullAccess	vpc	QcloudCCNFullAccess			2019-10	0-09 19:58:06
	FivAccess	fiv	FivAccess1			2020-01	1-06 15:35:35
	CaptchaAccess	captcha	CaptchaAccess			2019-08	3-02 14:25:17
	EMR_ADMIN	Elasticsearch MapReduce	EMR_ADMIN			2019-02	2-19 14:49:16
	EMR_OBSERVER	Elasticsearch MapReduce	EMR_OBSERVER			2019-09	9-10 22:34:27
	EMR_OPERATION	Elasticsearch MapReduce	EMR_OPERATION			2019-09	9-12 15:43:46
	QcloudFCRFullAccess	Collection Robot	QcloudFCRFullAccess,Can add,modify,view	and use all instances		2019-05	5-08 15:16:18
	FCR Administration	Collection Robot	Can add,modify,view and use all instances.			2019-04	4-08 16:26:54
	FCR Visit and Access Instances	Collection Robot	Can view and use all instances.			2019-04	4-08 16:26:45
						1	<b>0 ▼</b> / page 🛛

For the method of custom policy creation, please see Setting Policy.

### **Policy Authorization**

A configured policy can grant permissions by associating user groups or sub-users.

Policy							
0	Associate users or user groups with policies to grant p	permissions.					
Creat	te Custom Policy Delete			All Policies	Preset Policy	Custom Policy	
	Policy Name	Service Type 🔻	Description			Last Mc	dified
	bsp	BatchCompute	bsp			2018-05	-04 20:04:41
	QcloudCCNFullAccess	vpc	QcloudCCNFullAccess			2019-10	-09 19:58:06

### Resource Types Authorizable by Custom Policy

Resource-Level permission can be used to specify which resources a user can manipulate. TMP supports certain resource-level permissions. This means that for TMP operations that support resource-level permission, you can control the time when a user is allowed to perform operations or to use specified resources. The following table describes the types of resources that can be authorized in CAM.

Resource Type	Resource Description Method in Authorization Policy
TMP	<pre>qcs::monitor:\$region:\$account:prom-instance/*</pre>
	<pre>qcs::monitor:\$region:\$account:prom-instance/\$instanceId</pre>

The following table describes the TMP API operations that currently support resource-level permissions. When setting a policy, you can enter the API operation name in the action field to control the individual API. You can also use

\* as a wildcard to set the action .

#### List of APIs supporting resource-level authorization

API Operation	API Description
DescribePrometheusInstances	Lists all TMP instances of the user
TerminatePrometheusInstances	Terminates TMP instance
RecreatePrometheusInstance	Reboots TMP instance
ModifyPrometheusInstanceAttributes	Modifies TMP instance attributes
ChangeGrafanaAdminPassword	Changes Grafana admin Password
UpgradeGrafanaDashboard	Upgrades Grafana dashboard

DescribePrometheusKubeClusters	Lists TKE clusters that can be integrated with TMP
InstallPrometheusAgent	Installs Prometheus agent
UninstallPrometheusAgent	Uninstalls Prometheus agent
DescribeServiceDiscovery	Lists TMP scrape configurations
CreateServiceDiscovery	Creates TMP scrape configuration
UpdateServiceDiscovery	Updates TMP scrape configuration
DeleteServiceDiscovery	Deletes TMP scrape configuration
DescribePrometheusKubeBasicMonitor	Queries basic monitoring status
EnablePrometheusKubeBasicMonitor	Enables basic monitoring
DisablePrometheusKubeBasicMonitor	Disables basic monitoring
DescribePrometheusAgentRuntime	Gets the runtime status of Prometheus agent
DescribePrometheusJobTargets	Lists the status information of TMP metric scrape tasks
DescribeRecordingRules	Queries recording rules
CreateRecordingRule	Creates recording rule
UpdateRecordingRule	Updates recording rule
DeleteRecordingRules	Deletes recording rule
DescribeAlertRules	Queries alarming rules
DeleteAlertRules	Deletes alarming rule
UpdateAlertRuleState	Updates alarming rule status
CreateAlertRule	Creates alarming rule
UpdateAlertRule	Updates alarming rule

#### List of APIs not supporting resource-level authorization

For TMP API operations that don't support resource-level authorization, you can still authorize a user to perform them, but you must specify \* as the resource element in the policy statement.

API Operation	API Description	



CreatePrometheusInstance

Creates TMP instance

# Description of Role Permissions Related to Service Authorization

Last updated : 2024-08-07 22:06:42

When you use TMP, in order to use related Tencent Cloud resources, you will encounter a variety of scenarios that require service authorization. The CM\_QCSRole service role is mainly involved in the process of using TMP. This document describes the details, scenarios, and steps of each authorization policy by role. The preset policies associated with the CM\_QCSRole role by default include the following: QcloudAccessForCMRoleInPromHostingService: TKE permission required by TMP.

#### Use Cases

After you successfully create a TMP instance, you need to monitor the services running on TKE. In order to integrate the TKE service more conveniently, you need to access TKE-related APIs. In this case, your authorization is required before TKE can be normally accessed to install basic monitoring components and get their running status information. This role doesn't need to actively look for configuration. If its permission hasn't been granted, after you successfully create a TMP instance, the authorization page will automatically pop up when you enter the **Integrate with TKE** page for instance management.

### Authorization Steps

#### Authorizing by root account

1. After you successfully create a TMP instance, an authorization window will pop up when you access the **Integrate** with TKE page, and you need to authorize Cloud Monitor permissions as shown below:

2. Click Authorize Now in the window.

3. On the **CAM** > **Role Management** page, click **Grant**, and the system will prompt that the authorization is successful.

#### Note:

This authorization window will appear only once. If you have already authorized, it will not appear again.

#### Granting permissions to sub-account

After the root account completes the above authorization operations and successfully creates the CM\_QCSRole role, the sub-account doesn't have permission to access it. The sub-account must be granted the PassRole



permission by the root account before it can normally access TKE in TMP; otherwise, an error will be displayed when it accesses the TKE cluster list.

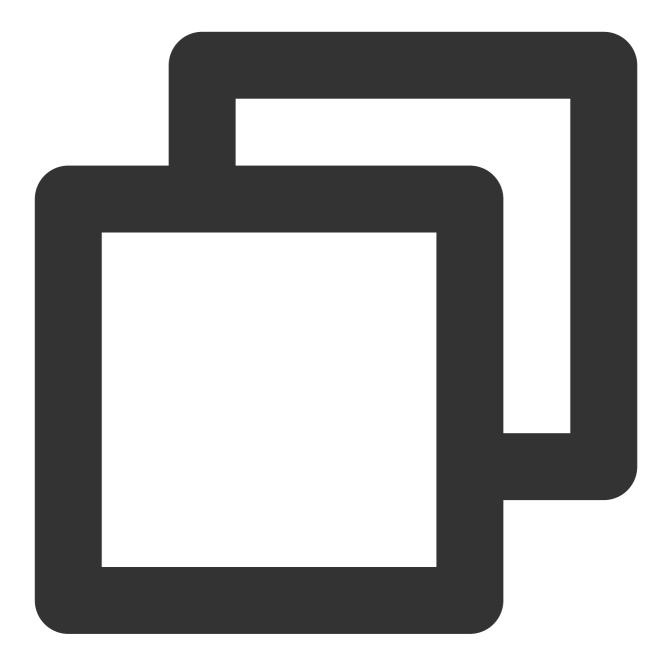
When granting the PassRole permission to your sub-account, please make sure that your sub-account has the following permissions:

Permission Description	Granted Policy
The sub-account needs to be granted access to CAM before granting the PassRole permission to the sub-account by the root account can take effect	QcloudCamReadOnlyAccess Or QcloudCamFullAcces
The Cloud Monitor policy depends on the Tencent Cloud service policy; therefore, before granting the PassRole permission to the sub- account, you need to make sure that the sub-account can normally access TKE resources	For more information, please see Permission Management

To ensure that the above permissions are granted successfully, please grant the cam:PassRole permission to the sub-account in the following steps.

1. Use the root account or a sub-account with administrative permissions to create the following custom policy:





```
{
    "version": "2.0",
    "statement": [
        {
            "effect": "allow",
            "action": "cam:PassRole",
            "resource": "qcs::cam::uin/${OwnerUin}:roleName/CM_QCSRole"
        }
    ]
}
```



2. After creation, associate the sub-account with the custom policy as instructed in CAM - Authorization Management. After granting the sub-account the cam:PassRole permission, access the **Integrate with TKE** page of the corresponding TMP instance, and an authorization window will pop up.

# Grafana

Last updated : 2024-08-07 22:06:47

TMP is highly integrated with TencentCloud Managed Service for Grafana (TCMG). One TCMG instance can be bound by multiple TMP instances at the same time to visualized TMP data in a unified way.

### Directions

TMP allows you to associate with a TCMG instance when creating a TMP instance. If you don't associate with the TCMG instance when purchasing the TMP service, you can follow the instructions below for binding.

#### Associating with a TCMG instance

- 1. Log in to the TMP console.
- 2. Find the corresponding TMP instance in the TMP instance list, and click **More>Grafana>Associate with TCMG** in the **Operation** column.
- 3. Select the TCMG instance in the pop-up window and click OK.

u nave selected	this instance:				
Instance ID/	Status	AZ	Network	Configuration	Billing Mode
	⊘ Running	Silicon Valley Zone 1	Netwo Subn	Data retention period: 15 day(s)	Pay as you go
	MD instance is as	ociated with a TCM	IG instance, the panel operations	in the TMP integration center will	
	ically apply to the	TCMG instance.			
U		TCMG instance.	7 *		

#### Note

You can only select the TCMG instance in the same VPC (private network) as the TMP instance. If there is no suitable TCMG instance, see Creating Instance to create one.

#### Disassociating from a TCMG instance

1. Log in to the TMP console.

2. Find the corresponding TMP instance in the TMP instance list, and click More>Grafana>Disassociate from

TCMG in the Operation column.

3. In the pop-up window, click **OK**.

You are disassociating this instance from <u>u</u> Once you confirm the disassociation, all the panels and data sources
related to this instance in g will be deleted. Are you sure you want to continue?

#### Logging in to the TCMG instance

1. Log in to the TMP console.

2. Find the corresponding TMP instance in the TMP instance list, and click the Grafana icon to the right of the instance ID/name.

Create Edit Tag	Pay-as-You-Go Instance Resour	rce Usage					Separate	eywords with ")"; press Enter to	separate fi
Instance ID/Name	Monitoring/Status <b>T</b>	AZ T	Network	Configuration	IPv4 Address	Billing Mode	Tag (key:value) 🚯	Creation Time	Ор
n-2605	ll ⊘ Running	Silicon Valley Zone 1	Networ Sub	Data retention period: 15 day(s)		Pay as you go 🕑	created_by;jliao bu:nasa	2023/05/08 13:30:56	Ma
ort 6	<b>II</b> ⊘ Running	Silicon Valley Zone 1	Networ Subne	Data retention period: 15 day(s)		Pay as you go 🥑	Team:myteam Role:test Region:na-siliconvalley	2023/01/18 08:01:43	Ma

3. On the TCMG login page, enter your account and password to log in.

#### Note

For more TCMG operations such as configuration and image rendering, see TencentCloud Managed Service for Grafana (TCMG).

# API Guide Overview

Last updated : 2024-08-07 22:06:58

### HTTP API

All stable HTTP APIs of Prometheus are under the path /api/v1. When you need to query monitoring data, you can request data through query APIs. To submit data, you can use the remote write protocol or Pushgateway.

### Supported APIs

API	Description	Authentication Required	Method
/api/v1/query	Query	Yes	GET/POST
/api/v1/query_range	Range query	Yes	GET/POST
/api/v1/series	Series query	Yes	GET/POST
/api/v1/labels	Labels query	Yes	GET/POST
/api/v1/label/ <label_name>/values</label_name>	Label value query	Yes	GET
/api/v1/prom/write	Data submission through remote write	Yes	remote write
Pushgateway	Data submission through Pushgateway	Yes	SDK

### Authentication Method

Authentication is enabled by default, so all APIs require authentication, and all authentication methods support bearer token and basic authentication.

#### Bearer token

A bearer token is generated as an instance is generated and can be queried in the console. For more information on bearer token, please see Bearer Authentication.

#### **Basic auth**

Basic auth is compatible with the native Prometheus query authentication method. The username is your APPID, and the password is the bearer token (generated when the instance is generated), which can be queried in the console. For more information on basic auth, please see Basic Authentication.

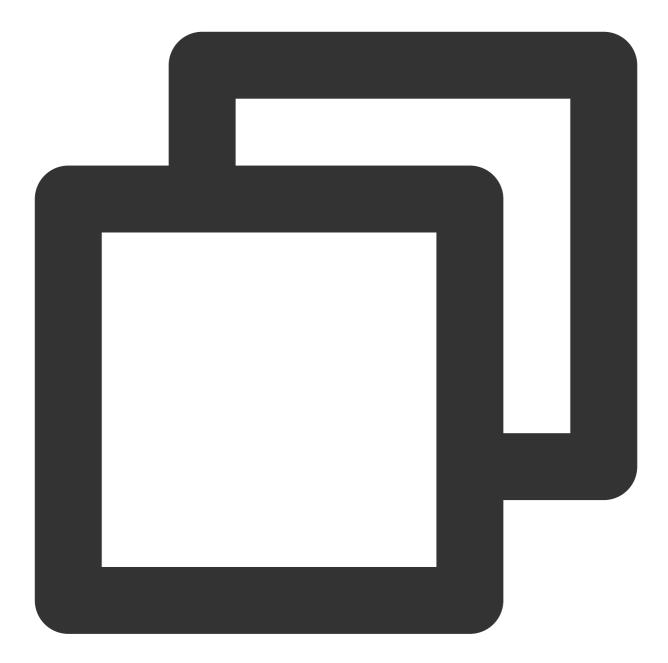
### Data Return Format

The response data of all APIs is in JSON format. Every successful request will return a status code of 2xx. An invalid request will return a piece of JSON data containing an error object and a status code as described below:

Status Code	Description
401	Authentication failed
400	A parameter was missing or incorrect
422	An invalid expression couldn't be specified (RFC4918)
503	The query was unavailable or canceled

Below is a response template for invalid requests:





```
{
   "status": "success" | "error",
   "data": <data>,
   // When the `status` is `error`, the following data will be returned
   "errorType": "<string>",
   "error": "<string>",
   // When there is a warning message during request execution, this field will be f
   "warnings": ["<string>"]
}
```



# Writing Data

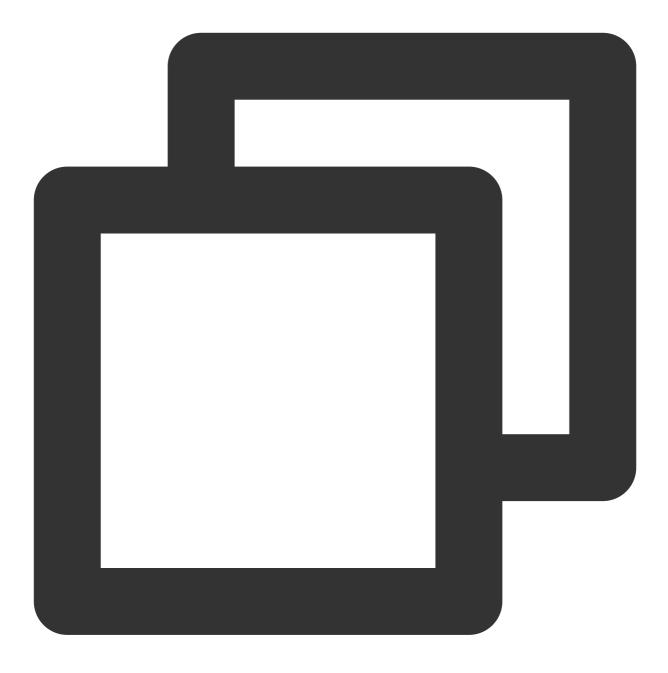
Last updated : 2024-08-07 22:07:03

### Overview

Similar to the scenario of data reporting by Flink jobs, you need to write data directly to Prometheus through APIs, as the lifecycle of these jobs may be very short, and it would be too late to wait for Prometheus to pull the data. To write data, you can directly use the remote write protocol or Pushgateway.

### **Remote Write**





#### POST /api/v1/prom/write

Remote write is a standard protocol of Prometheus. For more information, please see REMOTE WRITE TUNING. With remote write, you can write other Prometheus data in the VPC to TMP, which helps improve the data stability and make migration easier.

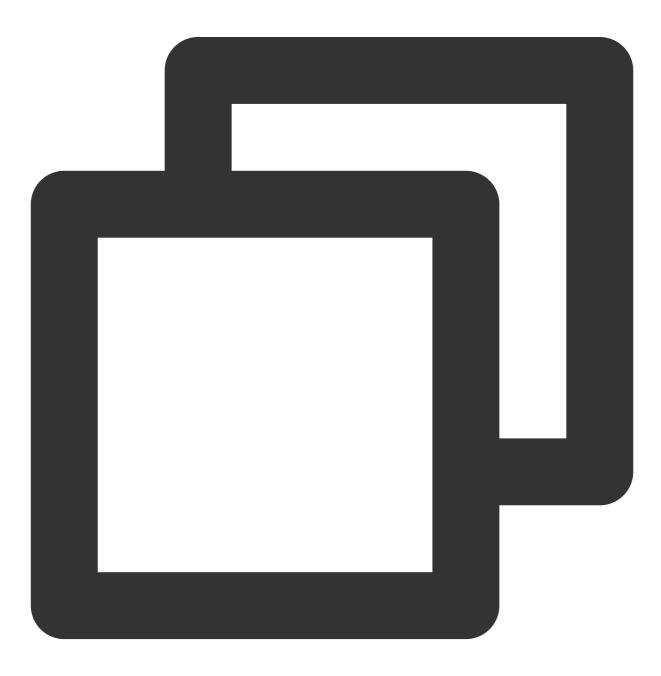
### Pushgateway

Although the pull method is recommended in Prometheus, you may still need to use the push method in some scenarios. For more information, please see WHEN TO USE THE PUSHGATEWAY.

TMP is natively integrated with the Pushgateway module, which can directly push data to it. Below is a simple example of pushing data in Go. You need to change the variables of <code>\$IP</code>, <code>\$PORT</code>, <code>\$APPID</code>, and <code>\$TOKEN</code> to the authentication information of your own instance, which can be queried in the console.

#### We strongly recommend you replace **\$INSTANCE** with an identifier of the current machine, such as

IP/Hostname . If this groupingKey is not set, when multiple machines report data together, the data entries will overwrite each other.





```
package main
import (
    "fmt"
    "time"
    "github.com/prometheus/client_golang/prometheus"
    "github.com/prometheus/client_golang/prometheus/push"
    "github.com/prometheus/common/expfmt"
)
var completionTime = prometheus.NewGauge(prometheus.GaugeOpts{
    Name: "db_backup_last_completion_timestamp_seconds",
    Help: "The timestamp of the last successful completion of a DB backup.",
})
func do() {
    completionTime.SetToCurrentTime()
}
func ExamplePusher_Push() {
    if err := push.New("http://$IP:$PORT", "db_backup").
        BasicAuth("$APPID", "$TOKEN").
        Collector(completionTime).
        Grouping("instance", "$INSTANCE").
        Grouping("db", "customers").
        Format(expfmt.FmtText).
        Push(); err != nil {
        fmt.Println("Could not push completion time to Pushgateway:", err)
    }
}
func main() {
    do()
    ticker := time.NewTicker(2 * time.Second)
    done := make(chan bool)
    for {
        select {
        case <-done:</pre>
            return
        case <-ticker.C:</pre>
            ExamplePusher Push()
        }
    }
}
```



#### Note:

You can customize the HTTP Client through the Client method for the object generated by push.New . We recommend you set an appropriate timeout period. In addition, if data is pushed, we recommend you use an async method for calls so as to avoid blocking the primary business process. For other reporting scenarios, please see Custom Monitoring.

# **Querying Monitoring Data**

Last updated : 2024-08-07 22:07:09

### Overview

When you need to query monitoring data, you can request data through query APIs.

### APPID/Token Acquisition Method

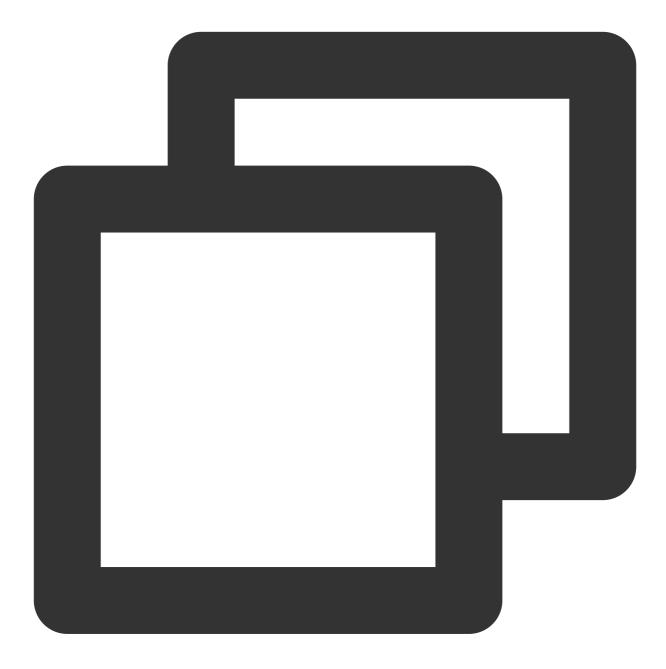
TMP usesAPPID+Tokento authenticate the access.

APPID can be obtained here.

Token can be obtained from the basic information of the corresponding TMP instance.

#### Query APIs





GET /api/v1/query
POST /api/v1/query

### **Query Parameters**

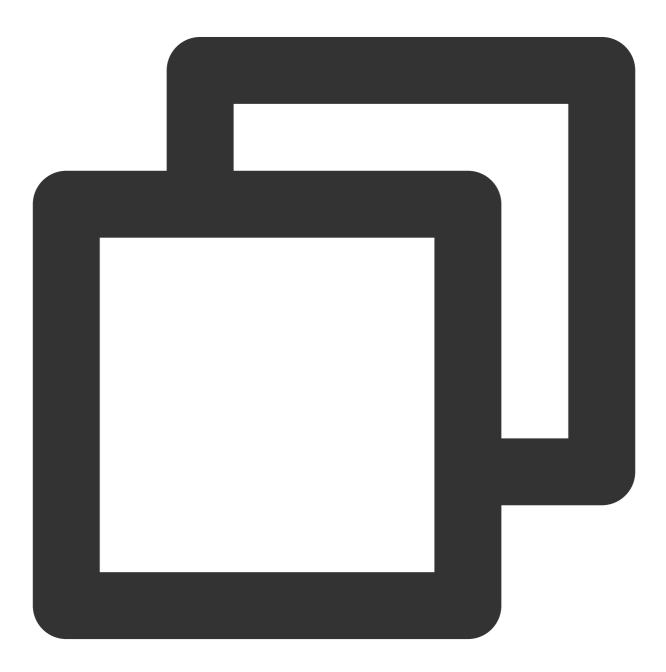
query=<string>: Prometheus: query expression.

time=<rfc3339 | unix\_timestamp>: timestamp, which is optional.

timeout=<duration>: detection timeout period, which is optional and specified by the \_\_query.timeout parameter by default.

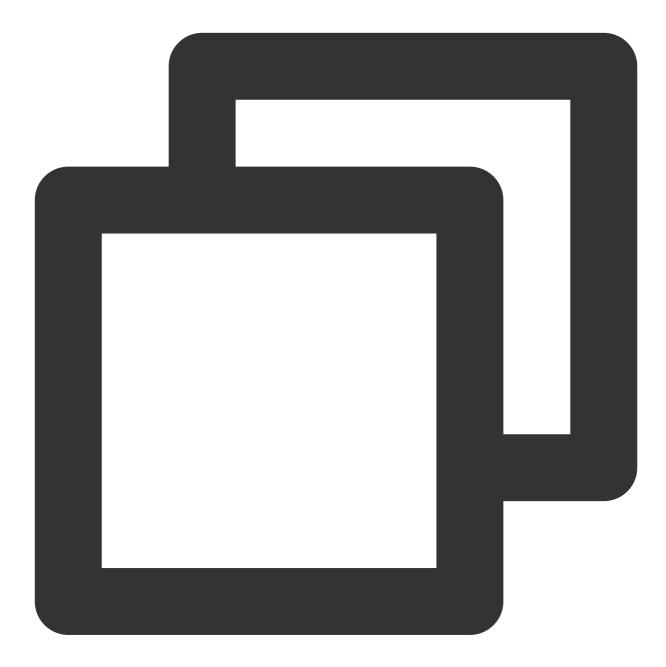
### Sample Simple Query

You can use the following sample to query data through an API. The query service address and authentication information can be viewed in the corresponding instance's information in the console:



curl -u "appid:token" 'http://IP:PORT/api/v1/query?query=up'

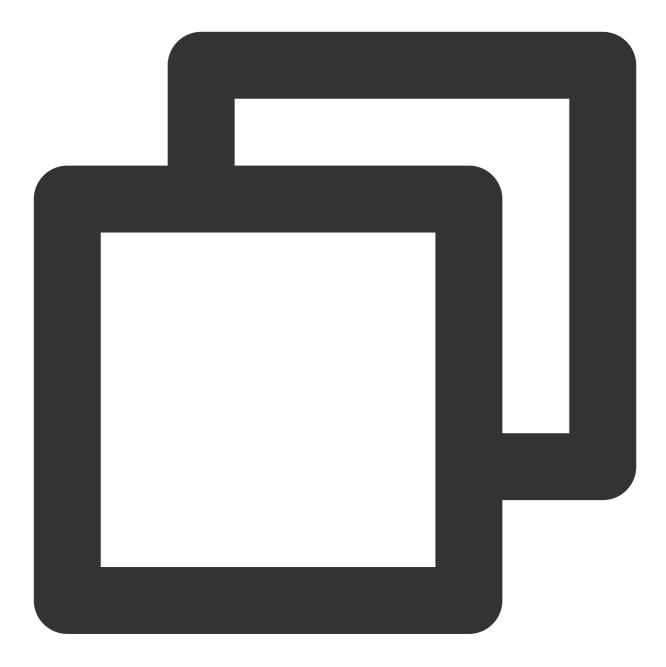
If the returned status code is 401, please check whether the authentication information is correct.



< HTTP/1.1 401 Unauthorized < Content-Length: 0

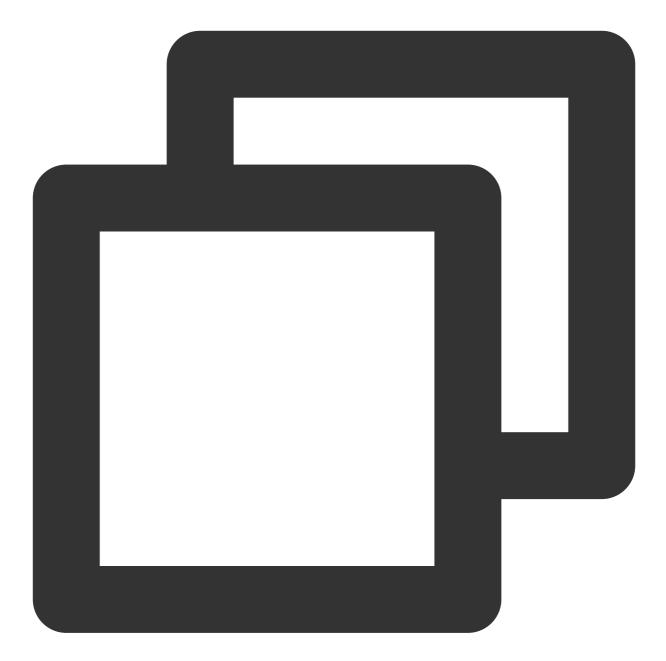
### Range Query





GET /api/v1/query\_range
POST /api/v1/query\_range

Querying data by time range is the most common query scenario. To do so, you can use the /api/v1/query\_range API as shown below:



```
},
            "values" : [
                [ 1435781430.781, "1" ],
                [ 1435781445.781, "1" ],
                [ 1435781460.781, "1" ]
            1
         },
         {
            "metric" : {
                "___name___" : "up",
                "job" : "node",
                "instance" : "localhost:9091"
            },
            "values" : [
                [ 1435781430.781, "0" ],
                [ 1435781445.781, "0" ],
                [ 1435781460.781, "1" ]
            ]
         }
      ]
  }
}
```

### Adding Data Source to Self-Built Grafana

You can add TMP as a data source in your self-built Grafana, and then you can view data in Grafana, provided that they are in the same VPC and can access each other over the network.

Enable the BasicAuth authentication method and enter the corresponding authentication information as shown below:

Data Sources / hosted-prometheus						
tki Settings 問 Dashboards						
Name 🛈	hosted-prometh	neus	Default			
НТТР						
URL 🔅	http://IP:POR	т				
Access	Server (defa	ult)	∽ Help	) >		
Whitelisted Cookies 🔅	Add Name		Add			
Auth						
Basic auth		With Credentials	0			
TLS Client Auth		With CA Cert				
Skip TLS Verify						
Forward OAuth Identity	0					
				1		
Basic Auth Details						
User	APPID					
Password	configured		Reset			

# TKE Metrics Free Metrics in Pay-as-You-Go Mode

Last updated : 2024-08-07 22:07:22

For pay-as-you-go instances with a storage period of more than 15 days, storage fees for their free metrics will be charged based on the excessive storage period.

Configuration File	Metric Name
node-exporter	node_boot_time_seconds
node-exporter	node_context_switches_total
node-exporter	node_cpu_seconds_total
node-exporter	node_disk_io_now
node-exporter	node_disk_io_time_seconds_total
node-exporter	node_disk_io_time_weighted_seconds_total
node-exporter	node_disk_read_bytes_total
node-exporter	node_disk_read_time_seconds_total
node-exporter	node_disk_reads_completed_total
node-exporter	node_disk_write_time_seconds_total
node-exporter	node_disk_writes_completed_total
node-exporter	node_disk_written_bytes_total
node-exporter	node_filefd_allocated
node-exporter	node_filesystem_avail_bytes
node-exporter	node_filesystem_free_bytes
node-exporter	node_filesystem_size_bytes
node-exporter	node_load1
node-exporter	node_load15
node-exporter	node_load5



node-exporter	node_memory_Buffers_bytes
node-exporter	node_memory_Cached_bytes
node-exporter	node_memory_MemAvailable_bytes
node-exporter	node_memory_MemFree_bytes
node-exporter	node_memory_MemTotal_bytes
node-exporter	node_netstat_TcpExt_ListenDrops
node-exporter	node_netstat_Tcp_ActiveOpens
node-exporter	node_netstat_Tcp_CurrEstab
node-exporter	node_netstat_Tcp_InSegs
node-exporter	node_netstat_Tcp_OutSegs
node-exporter	node_netstat_Tcp_PassiveOpens
node-exporter	node_network_receive_bytes_total
node-exporter	node_network_transmit_bytes_total
node-exporter	node_sockstat_TCP_alloc
node-exporter	node_sockstat_TCP_inuse
node-exporter	node_sockstat_TCP_tw
node-exporter	node_sockstat_UDP_inuse
node-exporter	node_sockstat_sockets_used
node-exporter	node_uname_info
cadvisor	container_cpu_usage_seconds_total
cadvisor	container_fs_limit_bytes
cadvisor	container_fs_reads_bytes_total
cadvisor	container_fs_usage_bytes
cadvisor	container_fs_writes_bytes_total
cadvisor	container_memory_working_set_bytes



cadvisor	container_network_receive_bytes_total
cadvisor	container_network_receive_packets_dropped_total
cadvisor	container_network_receive_packets_total
cadvisor	container_network_transmit_bytes_total
cadvisor	container_network_transmit_packets_dropped_total
cadvisor	container_network_transmit_packets_total
cadvisor	machine_cpu_cores
cadvisor	machine_memory_bytes
kubelet	kubelet_cgroup_manager_duration_seconds_count
kubelet	kubelet_node_config_error
kubelet	kubelet_node_name
kubelet	kubelet_pleg_relist_duration_seconds_bucket
kubelet	kubelet_pleg_relist_duration_seconds_count
kubelet	kubelet_pleg_relist_interval_seconds_bucket
kubelet	kubelet_pod_start_duration_seconds_count
kubelet	kubelet_pod_worker_duration_seconds_count
kubelet	kubelet_running_containers
kubelet	kubelet_running_pods
kubelet	kubelet_runtime_operations_duration_seconds_bucket
kubelet	kubelet_runtime_operations_errors_total
kubelet	kubelet_runtime_operations_total
kubelet	process_cpu_seconds_total
kubelet	process_resident_memory_bytes
kubelet	rest_client_request_duration_seconds_bucket
kubelet	rest_client_requests_total



kubelet	storage_operation_duration_seconds_bucket
kubelet	storage_operation_duration_seconds_count
kubelet	storage_operation_errors_total
kubelet	volume_manager_total_volumes
kube-state-metrics	kube_job_status_succeeded
kube-state-metrics	kube_job_status_failed
kube-state-metrics	kube_job_status_active
kube-state-metrics	kube_node_status_capacity_cpu_cores
kube-state-metrics	kube_node_status_capacity_memory_bytes
kube-state-metrics	kube_node_status_allocatable_cpu_cores
kube-state-metrics	kube_node_status_allocatable_memory_bytes
kube-state-metrics	kube_pod_info
kube-state-metrics	kube_pod_owner
kube-state-metrics	kube_pod_status_phase
kube-state-metrics	kube_pod_container_status_waiting
kube-state-metrics	kube_pod_container_status_running
kube-state-metrics	kube_pod_container_status_terminated
kube-state-metrics	kube_pod_container_status_restarts_total
kube-state-metrics	kube_pod_container_resource_requests_cpu_cores
kube-state-metrics	kube_pod_container_resource_requests_memory_bytes
kube-state-metrics	kube_pod_container_resource_limits_cpu_cores
kube-state-metrics	kube_pod_container_resource_limits_memory_bytes
kube-state-metrics	kube_replicaset_owner
kube-state-metrics	kube_statefulset_status_replicas
kube-controller-manager	rest_client_request_duration_seconds_bucket

kube-controller-manager	rest_client_requests_total
kube-controller-manager	workqueue_adds_total
kube-controller-manager	workqueue_depth
kube-controller-manager	workqueue_queue_duration_seconds_bucket
kube-apiserver	apiserver_current_inflight_requests
kube-apiserver	apiserver_current_inqueue_requests
kube-apiserver	apiserver_init_events_total
kube-apiserver	apiserver_longrunning_gauge
kube-apiserver	apiserver_registered_watchers
kube-apiserver	apiserver_request_duration_seconds_bucket
kube-apiserver	apiserver_request_duration_seconds_sum
kube-apiserver	apiserver_request_duration_seconds_count
kube-apiserver	apiserver_request_filter_duration_seconds_bucket
kube-apiserver	apiserver_request_filter_duration_seconds_sum
kube-apiserver	apiserver_request_filter_duration_seconds_count
kube-apiserver	apiserver_request_total
kube-apiserver	apiserver_requested_deprecated_apis
kube-apiserver	apiserver_response_sizes_bucket
kube-apiserver	apiserver_response_sizes_sum
kube-apiserver	apiserver_response_sizes_count
kube-apiserver	apiserver_selfrequest_total
kube-apiserver	apiserver_tls_handshake_errors_total
kube-apiserver	apiserver_watch_events_sizes
kube-apiserver	apiserver_watch_events_sizes_bucket
kube-apiserver	apiserver_watch_events_sizes_sum



kube-apiserver	apiserver_watch_events_sizes_count
kube-apiserver	apiserver_watch_events_total

## **Recommended Common Metrics for TKE**

Last updated : 2024-08-07 22:07:28

Based on the usage of most users, it is recommend that you configure the following commonly used TKE metrics. **Note** 

The following are paid metrics. For the billing method of metrics, see Billing Mode and Resource Usage.

Configuration File	Metric Name	Metric Description
kubelet	kubelet_running_container_count	kubelet_running_container Number of containers curre running
kubelet	kubelet_running_pod_count	kubelet_running_pod_cour Number of pods currently r
kube-state- metrics	kube_pod_container_info	Information about a contain pod.
kube-state- metrics	kube_deployment_status_replicas	The number of replicas per deployment.
kube-state- metrics	kube_deployment_labels	Kubernetes labels converte Prometheus labels.
kube-state- metrics	kube_pod_start_time	Start time in unix timestampod.
kube-state- metrics	kube_pod_status_ready	Describes whether the pod ready to serve requests.
kube-state- metrics	kube_node_info	Information about a cluster
kube-state- metrics	kube_node_status_condition	The condition of a cluster n
kube-state- metrics	kube_deployment_status_replicas_updated	The number of updated repper deployment.
kube-state- metrics	kube_deployment_status_replicas_available	The number of available re per deployment.
kube-state- metrics	kube_node_status_capacity_pods	The total pod resources of node.



kube-state- metrics	kube_pod_container_status_ready	Describes whether the con- readiness check succeede
kube-state- metrics	kube_deployment_spec_replicas	Number of desired pods for deployment.
kube-state- metrics	kube_pod_status_scheduled_time	Unix timestamp when pod ı into scheduled status
kube-state- metrics	kube_node_status_allocatable_pods	The pod resources of a noc are available for scheduling
kube-state- metrics	kube_pod_container_resource_limits	The number of requested li resource by a container.
kube-state- metrics	node_filefd_maximum	File descriptor statistics: maximum.
kube-state- metrics	kube_pod_container_resource_requests	The number of requested resource by a container.
kube-state- metrics	kube_namespace_labels	Kubernetes labels converte Prometheus labels.
kube-state- metrics	kube_deployment_status_replicas_unavailable	The number of unavailable replicas per deployment.
kube-state- metrics	kube_pod_created	Unix creation timestamp
kube-state- metrics	kube_pod_container_status_waiting_reason	Describes the reason the container is currently in wa state.
kube-state- metrics	kube_daemonset_status_desired_number_scheduled	The number of nodes that s be running the daemon poc
kube-state- metrics	kube_pod_restart_policy	Describes the restart policy use by this pod.
kube-state- metrics	kube_deployment_metadata_generation	Sequence number represe specific generation of the d state.
kube-state- metrics	kube_statefulset_status_update_revision	Indicates the version of the StatefulSet used to genera Pods in the sequence [repli updatedReplicas].



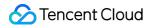
kube-state- metrics	kube_node_labels	Kubernetes labels converte Prometheus labels.
kube-state- metrics	kube_statefulset_replicas	Number of desired pods for StatefulSet.
kube-state- metrics	kube_statefulset_status_observed_generation	The generation observed b StatefulSet controller.
kube-state- metrics	kube_pod_container_status_last_terminated_reason	Describes the last reason the container was in terminated
kube-state- metrics	kube_replicaset_spec_replicas	Number of desired pods for ReplicaSet.
kube-state- metrics	kube_statefulset_created	Unix creation timestamp
kube-state- metrics	kube_statefulset_status_replicas_current	The number of current repl per StatefulSet.
kube-state- metrics	kube_statefulset_status_current_revision	Indicates the version of the StatefulSet used to genera Pods in the sequence [0].
kube-state- metrics	kube_statefulset_labels	Kubernetes labels converte Prometheus labels.
kube-state- metrics	kube_deployment_created	Unix creation timestamp
kube-state- metrics	kube_namespace_created	Unix creation timestamp
kube-state- metrics	kube_daemonset_status_number_ready	The number of nodes that s be running the daemon poc have one or more of the da pod running and ready.
kube-state- metrics	kube_deployment_status_observed_generation	The generation observed b deployment controller.
kube-state- metrics	kube_endpoint_info	Information about endpoint
kube-state- metrics	kube_statefulset_status_replicas_updated	The number of updated repper StatefulSet.



kube-state- metrics	kube_statefulset_metadata_generation	Sequence number represe specific generation of the d state for the StatefulSet.
kube-state- metrics	kube_secret_created	Unix creation timestamp
kube-state- metrics	kube_endpoint_address_not_ready	Number of addresses not rein endpoint
kube-state- metrics	kube_secret_type	Type about secret.
kube-state- metrics	kube_deployment_spec_paused	Whether the deployment is paused and will not be proceed by the deployment controlle
kube-state- metrics	kube_pod_container_status_terminated_reason	Describes the reason the container is currently in terminated state.
kube-state- metrics	kube_statefulset_status_replicas_ready	The number of ready replic StatefulSet.
kube-state- metrics	kube_endpoint_address_available	Number of addresses avail endpoint.
kube-state- metrics	kube_secret_info	Information about secret.
kube-state- metrics	kube_service_info	Information about service.
kube-state- metrics	kube_node_status_allocatable	The allocatable for different resources of a node that ar available for scheduling.
kube-state- metrics	kube_endpoint_labels	Kubernetes labels converte Prometheus labels.
kube-state- metrics	kube_deployment_status_condition	The current status conditio deployment.
kube-state- metrics	kube_endpoint_created	Unix creation timestamp
kube-state- metrics	kube_replicaset_labels	Kubernetes labels converte Prometheus labels.



kube-state- metrics	kube_replicaset_metadata_generation	Sequence number represe specific generation of the d state.
kube-state- metrics	kube_namespace_status_phase	kubernetes namespace sta phase.
kube-state- metrics	kube_service_created	Unix creation timestamp
kube-state- metrics	kube_configmap_created	Unix creation timestamp
kube-state- metrics	kube_secret_labels	Kubernetes labels converte Prometheus labels.
kube-state- metrics	kube_deployment_spec_strategy_rollingupdate_max_surge	Maximum number of replica can be scheduled above th desired number of replicas a rolling update of a deploy
kube-state- metrics	kube_configmap_metadata_resource_version	Resource version represen specific version of the confi
kube-state- metrics	kube_pod_labels	Kubernetes labels converte Prometheus labels.
kube-state- metrics	kube_replicaset_status_replicas	The number of replicas per ReplicaSet.
kube-state- metrics	kube_node_created	Unix creation timestamp
kube-state- metrics	kube_service_spec_type	Type about service.
kube-state- metrics	kube_secret_metadata_resource_version	Resource version represen specific version of secret.
kube-state- metrics	kube_configmap_info	Information about configma
kube-state- metrics	kube_replicaset_status_observed_generation	The generation observed b ReplicaSet controller.
kube-state- metrics	kube_service_labels	Kubernetes labels converte Prometheus labels.



kube-state- metrics	kube_replicaset_created	Unix creation timestamp
kube-state- metrics	kube_deployment_spec_strategy_rollingupdate_max_unavailable	Maximum number of unava replicas during a rolling upo a deployment.
kube-state- metrics	kube_replicaset_status_ready_replicas	The number of ready replic ReplicaSet.
kube-state- metrics	kube_replicaset_status_fully_labeled_replicas	The number of fully labeled replicas per ReplicaSet.
kube-state- metrics	kube_pod_status_scheduled	Describes the status of the scheduling process for the
kube-state- metrics	kube_storageclass_created	Unix creation timestamp
kube-state- metrics	kube_daemonset_status_number_misscheduled	The number of nodes runni daemon pod but are not supposed to.
kube-state- metrics	kube_storageclass_labels	Kubernetes labels converte Prometheus labels.
kube-state- metrics	kube_node_status_capacity	The capacity for different resources of a node.
kube-state- metrics	kube_daemonset_status_current_number_scheduled	The number of nodes runni least one daemon pod and supposed to.
kube-state- metrics	kube_storageclass_info	Information about storagec
kube-state- metrics	kube_node_spec_unschedulable	Whether a node can sched new pods.
kube-state- metrics	kube_daemonset_status_number_available	The number of nodes that s be running the daemon poc have one or more of the da pod running and available.
kube-state- metrics	kube_daemonset_labels	Kubernetes labels converte Prometheus labels.
kube-state-	kube_daemonset_created	Unix creation timestamp



metrics		
kube-state- metrics	kube_daemonset_status_number_unavailable	The number of nodes that s be running the daemon poor have none of the daemon p running and available
kube-state- metrics	kube_daemonset_metadata_generation	Sequence number represe specific generation of the d state.
kube-state- metrics	kube_mutatingwebhookconfiguration_info	Information about the MutatingWebhookConfigur
kube-state- metrics	kube_mutatingwebhookconfiguration_created	Unix creation timestamp.
kube-state- metrics	kube_mutatingwebhookconfiguration_metadata_resource_version	Resource version represen specific version of the MutatingWebhookConfigur
kube-state- metrics	kube_daemonset_updated_number_scheduled	The total number of nodes are running updated daemo
node- exporter	node_filesystem_files_free	Filesystem total free file not
node- exporter	node_filesystem_files	Filesystem total file nodes.
node- exporter	node_sockstat_UDP_mem_bytes	Number of UDP sockets in mem_bytes.
node- exporter	node_nf_conntrack_entries_limit	Maximum size of connectic tracking table.
node- exporter	node_memory_Shmem_bytes	Memory information field Shmem_bytes.
node- exporter	node_netstat_Tcp_RetransSegs	Statistic TcpRetransSegs.
node- exporter	node_sockstat_TCP_mem_bytes	Number of TCP sockets in mem_bytes.
node- exporter	node_network_info	Non-numeric data from /sys/class/net/ <iface></iface>



node- exporter	node_filesystem_readonly	Filesystem read-only status
node- exporter	node_exporter_build_info	A metric with a constant '1' labeled by version
node- exporter	node_network_iface_link_mode	iface_link_mode value of /sys/class/net/ <iface>.</iface>
node- exporter	node_network_receive_packets_total	Network device statistic receive_packets.
node- exporter	node_network_transmit_packets_total	Network device statistic transmit_packets.
node- exporter	node_memory_Mlocked_bytes	Memory information field Mlocked_bytes.
node- exporter	node_network_iface_id	iface_id value of /sys/class/net/ <iface>.</iface>
node- exporter	node_memory_WritebackTmp_bytes	Memory information field WritebackTmp_bytes.
node- exporter	kube_service_status_load_balancer_ingress	Service load balancer ingre status
node- exporter	node_vmstat_pgpgout	/proc/vmstat information fie pgpgout.
node- exporter	node_nf_conntrack_entries	Number of currently allocat flow entries for connection tracking.
node- exporter	node_memory_Inactive_file_bytes	Memory information field Inactive_file_bytes.
node- exporter	node_memory_SwapFree_bytes	Memory information field SwapFree_bytes.
node- exporter	node_sockstat_TCP_mem	Number of TCP sockets in mem.
node- exporter	node_memory_Slab_bytes	Memory information field Slab_bytes.
node- exporter	node_network_transmit_errs_total	Network device statistic transmit_errs.



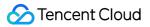
node- exporter	node_memory_Active_bytes	Memory information field Active_bytes.
node- exporter	node_procs_blocked	Number of processes block waiting for I/O to complete.
node- exporter	node_sockstat_UDP_mem	Number of UDP sockets in mem.
node- exporter	node_timex_maxerror_seconds	Maximum error in seconds.
node- exporter	node_memory_Inactive_bytes	Memory information field Inactive_bytes.
node- exporter	node_network_receive_errs_total	Network device statistic receive_errs.
node- exporter	node_memory_Unevictable_bytes	Memory information field Unevictable_bytes.
node- exporter	node_memory_KernelStack_bytes	Memory information field KernelStack_bytes.
node- exporter	node_procs_running	Number of processes in run state.
node- exporter	node_memory_SwapTotal_bytes	Memory information field SwapTotal_bytes.
node- exporter	node_netstat_lpExt_OutOctets	Statistic IpExtOutOctets.
node- exporter	node_memory_Active_file_bytes	Memory information field Active_file_bytes.
node- exporter	node_memory_SwapCached_bytes	Memory information field SwapCached_bytes.
node- exporter	node_netstat_lcmp_InMsgs	Statistic IcmpInMsgs.
node- exporter	node_forks_total	Total number of forks.
node- exporter	node_sockstat_RAW_inuse	Number of RAW sockets in inuse.



node- exporter	node_time_seconds	System time in seconds sir epoch (1970).
node- exporter	node_vmstat_pgpgin	/proc/vmstat information fie pgpgin.
node- exporter	node_memory_Mapped_bytes	Memory information field Mapped_bytes.
node- exporter	node_memory_SUnreclaim_bytes	Memory information field SUnreclaim_bytes.
node- exporter	node_memory_HardwareCorrupted_bytes	Memory information field HardwareCorrupted_bytes
node- exporter	node_memory_PageTables_bytes	Memory information field PageTables_bytes.
node- exporter	node_netstat_Udp6_InDatagrams	Statistic Udp6InDatagrams
node- exporter	node_netstat_lcmp_OutMsgs	Statistic IcmpOutMsgs.
node- exporter	node_netstat_Udp6_NoPorts	Statistic Udp6NoPorts.
node- exporter	node_memory_AnonPages_bytes	Memory information field AnonPages_bytes.
node- exporter	node_memory_Committed_AS_bytes	Memory information field Committed_AS_bytes.
node- exporter	node_netstat_TcpExt_ListenOverflows	Statistic TcpExtListenOver
node- exporter	node_netstat_UdpLite_InErrors	Statistic UdpLiteInErrors.
node- exporter	node_entropy_available_bits	Bits of available entropy.
node- exporter	node_memory_Inactive_anon_bytes	Memory information field Inactive_anon_bytes.
node- exporter	node_vmstat_pswpin	/proc/vmstat information fie pswpin.



node- exporter	node_memory_AnonHugePages_bytes	Memory information field AnonHugePages_bytes.
node- exporter	node_memory_SReclaimable_bytes	Memory information field SReclaimable_bytes.
node- exporter	node_netstat_lpExt_InOctets	Statistic lpExtInOctets.
node- exporter	node_netstat_Udp_NoPorts	Statistic UdpNoPorts.
node- exporter	node_timex_sync_status	Is clock synchronized to a server (1 = yes).
node- exporter	node_memory_CommitLimit_bytes	Memory information field CommitLimit_bytes.
node- exporter	node_memory_VmallocChunk_bytes	Memory information field VmallocChunk_bytes.
node- exporter	node_netstat_Udp_InDatagrams	Statistic UdpInDatagrams.
node- exporter	node_netstat_lcmp6_InErrors	Statistic lcmp6InErrors.
node- exporter	node_netstat_lcmp6_OutMsgs	Statistic lcmp6OutMsgs.
node- exporter	node_netstat_UdpLite6_InErrors	Statistic UdpLite6InErrors.
node- exporter	node_netstat_TcpExt_SyncookiesSent	Statistic TcpExtSyncookies
node- exporter	node_netstat_Tcp_InErrs	Statistic TcpInErrs.
node- exporter	node_intr_total	Total number of interrupts serviced.
node- exporter	node_timex_offset_seconds	Time offset in between loca system and reference clock
node- exporter	node_memory_Bounce_bytes	Memory information field Bounce_bytes.



node- exporter	node_memory_Writeback_bytes	Memory information field Writeback_bytes.
node- exporter	node_netstat_Udp_OutDatagrams	Statistic UdpOutDatagram
node- exporter	node_netstat_lcmp6_InMsgs	Statistic lcmp6InMsgs.
node- exporter	node_netstat_lp6_OutOctets	Statistic Ip6OutOctets.
node- exporter	node_netstat_lp_Forwarding	Statistic IpForwarding.
node- exporter	node_sockstat_TCP_orphan	Number of TCP sockets in orphan.
node- exporter	node_netstat_lp6_InOctets	Statistic lp6InOctets.
node- exporter	node_netstat_TcpExt_SyncookiesFailed	Statistic TcpExtSyncookiesFailed.
node- exporter	node_netstat_Udp_InErrors	Statistic UdpInErrors.
node- exporter	node_vmstat_pgmajfault	/proc/vmstat information fie pgmajfault.
node- exporter	node_network_transmit_drop_total	Network device statistic transmit_drop.
node- exporter	node_vmstat_pswpout	/proc/vmstat information fie pswpout.
node- exporter	node_network_up	Value is 1 if operstate is 'up
node- exporter	node_memory_NFS_Unstable_bytes	Memory information field NFS_Unstable_bytes.
node- exporter	node_memory_VmallocTotal_bytes	Memory information field VmallocTotal_bytes.
node- exporter	node_sockstat_FRAG_inuse	Number of FRAG sockets i inuse.



node- exporter	node_memory_Dirty_bytes	Memory information field Dirty_bytes.
node- exporter	node_netstat_Udp6_InErrors	Statistic Udp6InErrors.
node- exporter	node_netstat_TcpExt_SyncookiesRecv	Statistic TcpExtSyncookies
node- exporter	node_netstat_Udp6_OutDatagrams	Statistic Udp6OutDatagrar
node- exporter	node_memory_HugePages_Rsvd	Memory information field HugePages_Rsvd.
node- exporter	node_arp_entries	ARP entries by device
node- exporter	node_network_carrier	carrier value of /sys/class/net/ <iface>.</iface>
node- exporter	node_timex_pps_stability_exceeded_total	Pulse per second count of stability limit exceeded even
node- exporter	node_network_receive_compressed_total	Network device statistic receive_compressed.
node- exporter	node_network_transmit_carrier_total	Network device statistic transmit_carrier.
node- exporter	node_memory_DirectMap2M_bytes	Memory information field DirectMap2M_bytes.
node- exporter	node_memory_Hugepagesize_bytes	Memory information field Hugepagesize_bytes.
node- exporter	node_network_address_assign_type	address_assign_type value /sys/class/net/ <iface>.</iface>
node- exporter	node_network_receive_multicast_total	Network device statistic receive_multicast.
node- exporter	node_network_transmit_compressed_total	Network device statistic transmit_compressed.
node- exporter	node_memory_DirectMap4k_bytes	Memory information field DirectMap4k_bytes.
exporter	node_memory_Directiviap4k_bytes	DirectMap4k_bytes.



node- exporter	node_network_transmit_queue_length	transmit_queue_length valu/sys/class/net/ <iface>.</iface>
node- exporter	node_memory_HugePages_Free	Memory information field HugePages_Free.
node- exporter	node_network_receive_frame_total	Network device statistic receive_frame.
node- exporter	node_memory_HugePages_Total	Memory information field HugePages_Total.
node- exporter	node_network_flags	flags value of /sys/class/net/ <iface>.</iface>
node- exporter	node_network_receive_fifo_total	Network device statistic receive_fifo.
node- exporter	node_scrape_collector_duration_seconds	node_exporter: Duration of collector scrape.
node- exporter	node_network_speed_bytes	speed_bytes value of /sys/class/net/ <iface>.</iface>
node- exporter	node_sockstat_UDPLITE_inuse	Number of UDPLITE socke state inuse.
node- exporter	node_cpu_guest_seconds_total	Seconds the cpus spent in (VMs) for each mode.
node- exporter	node_filesystem_device_error	Whether an error occurred getting statistics for the give device.
node- exporter	node_scrape_collector_success	node_exporter: Whether a collector succeeded.
node- exporter	node_network_transmit_fifo_total	Network device statistic transmit_fifo.
node- exporter	node_vmstat_pgfault	/proc/vmstat information fie pgfault.
node- exporter	node_network_device_id	device_id value of /sys/class/net/ <iface>.</iface>
node- exporter	node_network_protocol_type	protocol_type value of /sys/class/net/ <iface>.</iface>



node- exporter	node_network_receive_drop_total	Network device statistic receive_drop.
node- exporter	node_timex_estimated_error_seconds	Estimated error in seconds
node- exporter	node_disk_writes_merged_total	The number of writes merg
node- exporter	node_network_transmit_colls_total	Network device statistic transmit_colls.
node- exporter	node_timex_tick_seconds	Seconds between clock tic
node- exporter	node_textfile_scrape_error	1 if there was an error oper reading a file
node- exporter	node_network_iface_link	iface_link value of /sys/class/net/ <iface>.</iface>
node- exporter	node_disk_reads_merged_total	The total number of reads merged.
node- exporter	node_timex_status	Value of the status array bi
node- exporter	node_netstat_lcmp_InErrors	Statistic IcmpInErrors.
node- exporter	node_memory_Active_anon_bytes	Memory information field Active_anon_bytes.
node- exporter	node_timex_pps_frequency_hertz	Pulse per second frequenc
node- exporter	node_network_mtu_bytes	mtu_bytes value of /sys/class/net/ <iface>.</iface>
node- exporter	node_timex_tai_offset_seconds	International Atomic Time ( offset.
node- exporter	node_timex_pps_jitter_total	Pulse per second count of j limit exceeded events.
node- exporter	node_timex_pps_jitter_seconds	Pulse per second jitter.



node- exporter	node_network_net_dev_group	net_dev_group value of /sys/class/net/ <iface>.</iface>
node- exporter	node_network_dormant	dormant value of /sys/class/net/ <iface>.</iface>
node- exporter	node_timex_pps_calibration_total	Pulse per second count of calibration intervals.
node- exporter	node_timex_pps_shift_seconds	Pulse per second interval duration.
node- exporter	node_timex_pps_error_total	Pulse per second count of calibration errors.
node- exporter	node_memory_VmallocUsed_bytes	Memory information field VmallocUsed_bytes.
node- exporter	node_timex_frequency_adjustment_ratio	Local clock frequency adjustment.
node- exporter	node_sockstat_FRAG_memory	Number of FRAG sockets i memory.
node- exporter	node_memory_HugePages_Surp	Memory information field HugePages_Surp.
node- exporter	node_timex_loop_time_constant	Phase-locked loop time cor
node- exporter	node_timex_pps_stability_hertz	Pulse per second stability

# **Recommended Common Metrics for TKE**

Last updated : 2024-08-07 22:07:38

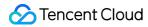
#### **Cluster Monitoring Overview**

Chart Name	Query Statement
CPU Requests Commitment	sum(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster"}) / sum(kube_i
CPU Limits Commitment	sum(kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster"}) / sum(kube_nod
Memory Utilisation	1 - sum(:node_memory_MemAvailable_bytes:sum{cluster="\$cluster"}) / sum(node_memory
Memory Requests Commitment	sum(kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster"}) / sum(ku
Memory Limits Commitment	sum(kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster"}) / sum(kube_
Node Count	count(kube_node_info{cluster="\$cluster"})
Pod Count	count(kube_pod_info{cluster="\$cluster"})



Node Request CPU Average Percent	avg(sum(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster"})by (node
Node Request Memory Average Percent	avg(sum(kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster"})by (node)/sum(kube_node_status_capacity_memory_bytes{cluster="\$cluster"})by(node))
API Server Success Request Percent	sum(irate(apiserver_request_total{cluster="\$cluster",code=~"20.*",verb=~"GET LIST"}[5m]] [5m]))
Namespace Overview	count(kube_pod_info{cluster="\$cluster"}) by (namespace)
	count(kube_service_info{cluster="\$cluster"}) by(namespace)
	count(kube_pod_container_info{cluster="\$cluster"}) by(namespace)
	count(kube_configmap_info{cluster="\$cluster"}) by(namespace)
	count(kube_secret_info{cluster="\$cluster"}) by(namespace)
	count(kube_deployment_created{cluster="\$cluster"}) by (namespace)
	count(kube_statefulset_created{cluster="\$cluster"}) by (namespace)
	count(kube_job_created{cluster="\$cluster"}) by (namespace)
	count(kube_cronjob_created{cluster="\$cluster"}) by (namespace)
	count(kube_pod_status_ready{cluster="\$cluster",condition="false"}==1) by(namespace) - ( by(namespace) or vector(0)) or count(kube_pod_status_ready{cluster="\$cluster",condition=

	count(kube_deployment_status_replicas_ready{cluster="\$cluster"} <kube_deployment_spe< td=""></kube_deployment_spe<>
	count(kube_statefulset_status_replicas_ready{cluster="\$cluster"} <kube_statefulset_replica< td=""></kube_statefulset_replica<>
	count(kube_daemonset_status_number_unavailable{cluster="\$cluster"}>0)by(namespace)
	<pre>count(kube_job_status_failed{cluster="\$cluster"} == 1) by (namespace)</pre>
	<pre>count(kube_daemonset_created{cluster="\$cluster"}) by (namespace)</pre>
	count(kube_persistentvolumeclaim_info{cluster="\$cluster"}) by (namespace)
CPU Usage	sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{clus
CPU Quota	<pre>sum(kube_pod_owner{cluster="\$cluster"}) by (namespace)</pre>
	count(avg(namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster"}) by (workload_pod:kube_pod_owner:relabel{cluster="\$cluster"})
	sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{clus
	<pre>sum(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster"}) by (namespa</pre>



	sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{clus sum(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster"}) by (namespa
	sum(kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster"}) by (namespace)
	sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{clus sum(kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster"}) by (namespace)
Memory Usage (working_set)	sum(container_memory_working_set_bytes{cluster="\$cluster", container!="", container!="F
	<pre>sum(kube_pod_owner{cluster="\$cluster"}) by (namespace)</pre>
	count(avg(namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster"}) by (wor
	sum(container_memory_rss{cluster="\$cluster", container!="", container!="POD"}) by (name
	sum(kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster"}) by (name
Memory Requests	sum(container_memory_rss{cluster="\$cluster", container!="", container!="POD"}) by (name sum(kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster"}) by (name
	sum(kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster"}) by (namesp
	sum(container_memory_rss{cluster="\$cluster", container!="", container!="POD"}) by (name sum(kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster"}) by (namespa
Node Memory Usage (Top 10)	sum(label_replace(topk(10, 1-(node_memory_MemAvailable_bytes{cluster="\$cluster"} / nc "instance", "(.*)"))by(node_ip)



Node CPU Usage (Top 10)	topk(10, sum(label_replace(1 - sum(rate(node_cpu_seconds_total{cluster="\$cluster",mode sum(rate(node_cpu_seconds_total{cluster="\$cluster"}[1m])) by (instance),"host_ip","\$1","in			
Node Disk Usage (Top 10)	topk(10, sum(label_replace(1- node_filesystem_free_bytes{cluster="\$cluster",mountpoint="/"}/node_filesystem_size_bytes (.*)"))by(host_ip))			
Node Network In (Top 10)	topk(10, sum(label_replace(max(irate(node_network_receive_bytes_total{cluster="\$cluster			
Node Network Out (Top 10)	topk(10, sum(label_replace(max(irate(node_network_transmit_bytes_total{cluster="\$cluste			
Node Sockets Count(Top 10)	topk(10, sum(label_replace(max(node_sockstat_TCP_alloc{cluster="\$cluster"}) by (insta			
Container Memory Usage(Top10)	topk(10, sum (container_memory_working_set_bytes{cluster="\$cluster",container !="",c			
Container Memory Usage/Limit(Top10)	topk(10, avg(container_memory_working_set_bytes{cluster="\$cluster",container!=""}/(cont namespace))			
Container CPU Usage(Top10)	topk(10, sum(rate(container_cpu_usage_seconds_total{cluster="\$cluster",container !="",cc			
0	topk(10, sum(irate(container_network_receive_bytes_total{cluster="\$cluster",image!="",coi			
Container Network	-topk(10, sum(irate(container_network_transmit_bytes_total{cluster="\$cluster",image!="",c			
Container Memory Usage/Limit (Top 10) topk(10, avg(container_memory_working_set_bytes{cluster="\$cluster",containation namespace))				
Container CPU Usage (Top 10)	topk(10, sum(irate(container_cpu_usage_seconds_total{cluster="\$cluster",container!="",cc			
Container Socket Count(Top 10)	topk(10, sum(container_sockets{cluster="\$cluster",container!=""}) by (container,pod,names			

## Cluster Namespace Dashboard

Chart Name	Query Statement
CPU Usage	sum(rate(container_cpu_usage_seconds_total{cluster="\$cluster",namespace=~"\$name
CPU Usage/Request(%)	sum(rate(container_cpu_usage_seconds_total{cluster="\$cluster",namespace=~"\$name [2m]))/sum(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster",nam
CPU Usage/Limit(%)	sum(rate(container_cpu_usage_seconds_total{cluster="\$cluster",namespace=~"\$name [2m]))/sum(kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster",namesp
CPU Request	sum(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster",namespace
CPU Limit	sum(kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster",namespace=~'
Cluster Available	sum(sum(kube_node_status_capacity{resource="cpu",cluster="\$cluster",namespace=~
StatefulSet Created	count(kube_statefulset_created{cluster="\$cluster",namespace="\$namespace"}) or on()
Pod Created	count(kube_pod_info{cluster="\$cluster",namespace="\$namespace"}) or on() vector(0)
Containers	count(kube_pod_container_info{cluster="\$cluster",namespace="\$namespace"}) or on()
DaemonSet Created	count(kube_daemonset_created{cluster="\$cluster",namespace="\$namespace"}) or on()
Job Created	count(kube_job_info{cluster="\$cluster",namespace="\$namespace"})or on() vector(0)
Job Active	count(kube_job_status_active{cluster="\$cluster",namespace="\$namespace"}==1)or on



Cron Job Created	count(kube_cronjob_created{cluster="\$cluster",namespace="\$namespace"}) or on() ver
Cron Job Active	count(kube_cronjob_status_active{cluster="\$cluster",namespace="\$namespace"}==1)
Unbound PVC	count(kube_persistentvolumeclaim_status_phase{phase!="Bound", cluster="\$cluster",n
PersistentVolumeClaim Created	count(kube_persistentvolumeclaim_info{cluster="\$cluster",namespace="\$namespace"})
Service Created	count(kube_service_info{cluster="\$cluster",namespace="\$namespace"}) or on() vector(
LoadBalancer Created	count(kube_service_spec_type{type="LoadBalancer", cluster="\$cluster",namespace="\$
Ingress Created	count(kube_ingress_info{cluster="\$cluster",namespace="\$namespace"})or on() vector((
ConfigMap Created	count(kube_configmap_info{cluster="\$cluster",namespace="\$namespace"})
Secret Created	count(kube_secret_info{cluster="\$cluster",namespace="\$namespace"}) or on() vector(0
PVC Storage Requests Total	sum(kube_persistentvolumeclaim_resource_requests_storage_bytes{cluster="\$cluster"
Pod NotReady	count(kube_pod_status_ready{condition="false", cluster="\$cluster",namespace="\$name by(namespace) or vector(0)) or count(kube_pod_status_ready{condition="false", cluster
Pod UnSchedulable	count(kube_pod_status_unschedulable{cluster="\$cluster",namespace="\$namespace"})
Deployment NotReady	count(sum(kube_deployment_status_replicas_ready{cluster="\$cluster",namespace="\$r vector(0)



Daemonset NotReady	count(kube_daemonset_status_number_unavailable{cluster="\$cluster",namespace="\$r
Job Failed	count(kube_job_status_failed{cluster="\$cluster",namespace="\$namespace"} == 1)
CPU Usage	sum(rate(container_cpu_usage_seconds_total{cluster="\$cluster",namespace=~"\$name
	sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{
	sum(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster", namespac
CPU Quota	sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{ sum(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster", namespac
	sum(kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster", namespace="
	sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{ sum(kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster", namespace="
Memory Usage	sum(container_memory_working_set_bytes{cluster="\$cluster",namespace=~"\$namespace
Memory Usage/Request(%)	sum(container_memory_working_set_bytes{cluster="\$cluster",namespace=~"\$namespa unit="byte", resource="memory"}) or on() vector(0)
Memory Usage/Limit(%)	sum(container_memory_working_set_bytes{cluster="\$cluster",namespace=~"\$namespa unit="byte", resource="memory"}) or on() vector(0)
Memory Request	sum(kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster",names
Memory Limit	sum(kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster",namespac



Cluster Available	sum(sum(kube_node_status_capacity{resource="memory"}) by (node) + sum(kube_noc
	sum(container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespa
Memory Usage (w/o cache)	scalar(kube_resourcequota{cluster="\$cluster", namespace="\$namespace", type="hard"
	scalar(kube_resourcequota{cluster="\$cluster", namespace="\$namespace", type="hard"
	sum(container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespa
	sum(kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster", names
	sum(container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespa sum(kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster",names
Memory Quota	sum(kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster", namespace
	sum(container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespa
	sum(container_memory_rss{cluster="\$cluster", namespace="\$namespace",container!='
	sum(container_memory_cache{cluster="\$cluster", namespace="\$namespace",containe
	sum(container_memory_swap{cluster="\$cluster", namespace="\$namespace",container
Containers	group by (image, container, container_id, pod) (kube_pod_container_info{cluster="\$clus
	group by (container, container_id, pod) (kube_pod_container_info{cluster="\$cluster",naı (kube_pod_container_status_running{cluster="\$cluster",namespace="\$namespace"})

group by (container, container\_id, pod) (kube\_pod\_container\_info{cluster="\$cluster",nai (kube\_pod\_container\_status\_restarts\_total{cluster="\$cluster",namespace="\$namespace"

group by (container, container\_id, pod) (kube\_pod\_container\_info{cluster="\$cluster",nai max(irate(container\_cpu\_usage\_seconds\_total{cluster="\$cluster",namespace="\$names

group by (container, container\_id, pod) (kube\_pod\_container\_info{cluster="\$cluster",nai (max(irate(container\_cpu\_usage\_seconds\_total{container!="",container!="POD",cluster (max(container\_spec\_cpu\_quota{container!="",container!="POD",cluster="\$cluster",nai

group by (container, container\_id, pod) (kube\_pod\_container\_info{cluster="\$cluster",naı (kube\_pod\_container\_resource\_requests\_cpu\_cores{cluster="\$cluster",namespace="\$

group by (container, container\_id, pod) (kube\_pod\_container\_info{cluster="\$cluster",nai (max(irate(container\_cpu\_usage\_seconds\_total{container!="",container!="POD",cluster (kube\_pod\_container\_resource\_requests\_cpu\_cores{cluster="\$cluster",namespace="\$i

group by (container, container\_id, pod) (kube\_pod\_container\_info{cluster="\$cluster",nai (kube\_pod\_container\_resource\_limits{resource="cpu",cluster="\$cluster",namespace="\$

group by (container, container\_id, pod) (kube\_pod\_container\_info{cluster="\$cluster",nai !="",container!="POD",cluster="\$cluster",namespace=~"\$namespace"}) by (pod,contain

group by (container, container\_id, pod) (kube\_pod\_container\_info{cluster="\$cluster",nai (max(container\_memory\_working\_set\_bytes{container!="",container!="POD",cluster="\$ !="",container!="POD",cluster="\$cluster",namespace=~"\$namespace"}) by (pod, contair

group by (container, container_id, pod) (kube_pod_container_info{cluster="\$cluster",nai (kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster",namespace
group by (container, container_id, pod) (kube_pod_container_info{cluster="\$cluster",nai !="",container!="POD",cluster="\$cluster",namespace=~"\$namespace"}) by (pod,contain

## API Server (Independent Cluster)

Chart Name	Query Statement
Availability > 99.000%	<pre>1 - (     (     sum by (cluster,cluster_type)     (increase(apiserver_request_duration_seconds_count{cluster="\$cluster"}[5m]))     -     sum by (cluster,cluster_type) (increase(apiserver_request_duration_seconds_bucket{le="1",     cluster="\$cluster"}[5m]))     )     +     sum by (cluster,cluster_type) (increase(apiserver_request_total{job="kube-apiserver",code=~"5     cluster="\$cluster"}[5m]) or vector(0))     )     /     sum by (cluster,cluster_type) (increase(apiserver_request_total{job="kube-apiserver",code=~"5     cluster="\$cluster"}[5m]) or vector(0))     )     /     sum by (cluster,cluster_type) (increase(apiserver_request_total{job="kube-apiserver",code=~"5 </pre>
ErrorBudget > 99.000%	100 * (1 - ( ( sum by (cluster,cluster_type) (increase(apiserver_request_duration_seconds_count{cluster="\$cluster"}[5m])) - sum by (cluster,cluster_type) (increase(apiserver_request_duration_seconds_bucket{le="1", cluster="\$cluster"}[5m])) )

	+ sum by (cluster,cluster_type) (increase(apiserver_request_total{job="kube-apiserver",code=~"5 cluster="\$cluster"}[5m]) or vector(0)) ) / sum by (cluster,cluster_type) (increase(apiserver_request_total{job="kube- apiserver",cluster="\$cluster"}[5m])) -0.990000)
Read Availability	<pre>1 - (     (     sum by (cluster,cluster_type) (increase(apiserver_request_duration_seconds_count{verb=~"LIST GET", cluster="\$cluster"}[5m])     -     sum by (cluster,cluster_type) (increase(apiserver_request_duration_seconds_bucket{verb=~"LIST GET",le="1", cluster="\$cluste [5m]))     )     +     sum by (cluster,cluster_type) (increase(apiserver_request_total{job="kube- apiserver",verb=~"LIST GET",code=~"5", cluster="\$cluster"}[5m]) or vector(0))     )     /     sum by (cluster,cluster_type) (increase(apiserver_request_total{job="kube- apiserver",verb=~"LIST GET", cluster="\$cluster"}[5m]))</pre>
Read SLI - Requests	sum by (code) (rate(apiserver_request_total{job="kube- apiserver",verb=~"LIST GET",cluster="\$cluster"}[5m]))
Read SLI - Errors	sum by (resource) (rate(apiserver_request_total{job="kube- apiserver",verb=~"LIST GET",code=~"5",cluster="\$cluster"}[5m]))/ sum by (resource) (rate(apiserver_request_total{job="kube- apiserver",verb=~"LIST GET",cluster="\$cluster"}[5m]))
Read SLI - Duration	histogram_quantile(0.99, sum by (le, resource,cluster,cluster_type) (rate(apiserver_request_duration_seconds_bucket{job="kube- apiserver",verb=~"LIST GET",cluster="\$cluster"}[5m]))) > 0
Write Availability	<pre>1 - (     (         sum by (cluster,cluster_type)     (increase(apiserver_request_duration_seconds_count{verb=~"POST PUT PATCH DELETE",     cluster="\$cluster"}[5m]))     -         sum by (cluster,cluster_type)     (increase(apiserver_request_duration_seconds_bucket{verb=~"POST PUT PATCH DELETE",le="</pre>



	<pre>cluster="\$cluster"}[5m]))   )   +   sum by (cluster,cluster_type) (increase(apiserver_request_total{job="kube- apiserver",verb=~"POST PUT PATCH DELETE",code=~"5", cluster="\$cluster"}[5m]) or vector(0))   )   /   sum by (cluster,cluster_type) (increase(apiserver_request_total{job="kube- apiserver",verb=~"POST PUT PATCH DELETE", cluster="\$cluster"}[5m]))</pre>
Write SLI - Requests	sum by (code) (rate(apiserver_request_total{job="kube- apiserver",verb=~"POST PUT PATCH DELETE",cluster="\$cluster"}[5m]))
Write SLI - Errors	sum by (resource) (rate(apiserver_request_total{job="kube- apiserver",verb=~"POST PUT PATCH DELETE",code=~"5",cluster="\$cluster"}[5m]))/ sum by (resource) (rate(apiserver_request_total{job="kube- apiserver",verb=~"POST PUT PATCH DELETE",cluster="\$cluster"}[5m]))
Write SLI - Duration	histogram_quantile(0.99, sum by (le, resource,cluster,cluster_type) (rate(apiserver_request_duration_seconds_bucket{job="kube- apiserver",verb=~"POST PUT PATCH DELETE",cluster="\$cluster"}[5m]))) > 0
Work Queue Add Rate	sum(rate(workqueue_adds_total{job="kube-apiserver", instance=~"\$instance", cluster=~"\$cluster"} [5m])) by (instance, name)
Work Queue Depth	sum(rate(workqueue_depth{job="kube-apiserver", instance=~"\$instance", cluster=~"\$cluster"}[5m]) by (instance, name)
Work Queue Latency	histogram_quantile(0.99, sum(rate(workqueue_queue_duration_seconds_bucket{job="kube- apiserver", instance=~"\$instance", cluster=~"\$cluster"}[5m])) by (instance, name, le))
Memory	process_resident_memory_bytes{job="kube-apiserver",instance=~"\$instance", cluster=~"\$cluster"}
CPU usage	rate(process_cpu_seconds_total{job="kube-apiserver",instance=~"\$instance", cluster=~"\$cluster"} [5m])

### Controller Manager (Independent Cluster)

Chart	Query Statement	Metrics Us
Name		



Up	sum(up{cluster=~"\$cluster",job="kube-controller-manager"})	up
Work Queue Add Rate	sum(rate(workqueue_adds_total{cluster=~"\$cluster",job="kube-controller-manager", instance=~"\$instance"}[5m])) by (instance, name)	workqueu
Work Queue Depth	sum(rate(workqueue_depth{cluster=~"\$cluster",job="kube-controller-manager", instance=~"\$instance"}[5m])) by (instance, name)	workqueu
Work Queue Latency	histogram_quantile(0.99, sum(rate(workqueue_queue_duration_seconds_bucket{cluster=~"\$cluster",job="kube- controller-manager", instance=~"\$instance"}[5m])) by (instance, name, le))	workqueu
	sum(rate(rest_client_requests_total{cluster=~"\$cluster",job="kube-controller-manager", instance=~"\$instance",code=~"2"}[5m]))	rest_client
Kube API	sum(rate(rest_client_requests_total{cluster=~"\$cluster",job="kube-controller-manager", instance=~"\$instance",code=~"3"}[5m]))	rest_client
Request	sum(rate(rest_client_requests_total{cluster=~"\$cluster",job="kube-controller-manager", instance=~"\$instance",code=~"4"}[5m]))	rest_client
	sum(rate(rest_client_requests_total{cluster=~"\$cluster",job="kube-controller-manager", instance=~"\$instance",code=~"5"}[5m]))	rest_client
Post Request Latency 99th Quantile	histogram_quantile(0.99, sum(rate(rest_client_request_duration_seconds_bucket{cluster=~"\$cluster",job="kube- controller-manager", instance=~"\$instance", verb="POST"}[5m])) by (verb, url, le))	rest_client
Get Request Latency 99th Quantile	histogram_quantile(0.99, sum(rate(rest_client_request_duration_seconds_bucket{cluster=~"\$cluster",job="kube- controller-manager", instance=~"\$instance", verb="GET"}[5m])) by (verb, url, le))	rest_client
Memory	process_resident_memory_bytes{cluster=~"\$cluster",job="kube-controller- manager",instance=~"\$instance"}	process_re
CPU usage	rate(process_cpu_seconds_total{cluster=~"\$cluster",job="kube-controller- manager",instance=~"\$instance"}[5m])	process_c



#### Kubelet

Chart Name	Query Statement
Up	sum(up{cluster="\$cluster", job="kubelet"})
Running Pods	sum(kubelet_running_pods{cluster="\$cluster", job="kubelet", instance=~"\$instance"})
Running Container	sum(kubelet_running_containers{cluster="\$cluster", job="kubelet", instance=~"\$instance"})
Actual Volume Count	sum(volume_manager_total_volumes{cluster="\$cluster", job="kubelet", instance=~"\$instance", state="actual_state_of_world"})
Desired Volume Count	sum(volume_manager_total_volumes{cluster="\$cluster", job="kubelet", instance=~"\$instance",state="desired_state_of_world"})
Config Error Count	sum(rate(kubelet_node_config_error{cluster="\$cluster", job="kubelet", instance=~"\$instance"}[5m]))
Operation Rate	sum(rate(kubelet_runtime_operations_total{cluster="\$cluster",job="kubelet",instance=~"\$instance"}[ (operation_type, instance)
Operation Error Rate	sum(rate(kubelet_runtime_operations_errors_total{cluster="\$cluster",job="kubelet",instance=~"\$inst (instance, operation_type)
Operation duration 99th quantile	histogram_quantile(0.99, sum(rate(kubelet_runtime_operations_duration_seconds_bucket{cluster="\$cluster",job="kubelet",in: [5m])) by (instance, operation_type, le))
Pod Start Rate	sum(rate(kubelet_pod_start_duration_seconds_count{cluster="\$cluster",job="kubelet",instance=~"\$ (instance)
	sum(rate(kubelet_pod_worker_duration_seconds_count{cluster="\$cluster",job="kubelet",instance=~ by (instance)
Pod Start Duration	histogram_quantile(0.99, sum(rate(kubelet_pod_start_duration_seconds_count{cluster="\$cluster",job="kubelet",instance=~"\$ (instance, le))
	histogram_quantile(0.99,



	<pre>sum(rate(kubelet_pod_worker_duration_seconds_bucket{cluster="\$cluster",job="kubelet",instance= by (instance, le))</pre>
Storage Operation Rate	sum(rate(storage_operation_duration_seconds_count{cluster="\$cluster",job="kubelet",instance=~"\$ (instance, operation_name, volume_plugin)
Storage Operation Error Rate	sum(rate(storage_operation_errors_total{cluster="\$cluster",job="kubelet",instance=~"\$instance"}[5m operation_name, volume_plugin)
Storage Operation Duration 99th quantile	histogram_quantile(0.99, sum(rate(storage_operation_duration_seconds_bucket{cluster="\$cluster", instance=~"\$instance"}[5m])) by (instance, operation_name, volume_plugin, le))
Cgroup manager operation rate	sum(rate(kubelet_cgroup_manager_duration_seconds_count{cluster="\$cluster", job="kubelet", insta [5m])) by (instance, operation_type)
Cgroup manager 99th quantile	histogram_quantile(0.99, sum(rate(kubelet_cgroup_manager_duration_seconds_bucket{cluster="\$c job="kubelet", instance=~"\$instance"}[5m])) by (instance, operation_type, le))
PLEG relist rate	sum(rate(kubelet_pleg_relist_duration_seconds_count{cluster="\$cluster", job="kubelet", instance=~' (instance)
PLEG relist interval	histogram_quantile(0.99, sum(rate(kubelet_pleg_relist_interval_seconds_bucket{cluster="\$cluster",job="kubelet",instance=~"; (instance, le))
PLEG relist duration	histogram_quantile(0.99, sum(rate(kubelet_pleg_relist_duration_seconds_bucket{cluster="\$cluster",job="kubelet",instance=~ (instance, le))
PPC Pata	sum(rate(rest_client_requests_total{cluster="\$cluster",job="kubelet", instance=~"\$instance",code=~"
RPC Rate	sum(rate(rest_client_requests_total{cluster="\$cluster",job="kubelet", instance=~"\$instance",code=~"
Request duration 99th quantile	histogram_quantile(0.99, sum(rate(rest_client_request_duration_seconds_bucket{cluster="\$cluster" instance=~"\$instance"}[5m])) by (instance, verb, url, le))
Memory	process_resident_memory_bytes{cluster="\$cluster",job="kubelet",instance=~"\$instance"}



CPU usage	rate(process_cpu_seconds_total{cluster="\$cluster",job="kubelet",instance=~"\$instance"}[5m])
Goroutines	go_goroutines{cluster="\$cluster",job="kubelet",instance=~"\$instance"}

# Proxy(Non-default Installation Component)

Chart Name	Query Statement
Up	sum(up{job="kube-proxy"})
Rules Sync Rate	sum(rate(kubeproxy_sync_proxy_rules_duration_seconds_count{job="kube-proxy", instance=~"\$instance"}[5m]))
Rule Sync Latency 99th Quantile	histogram_quantile(0.99,rate(kubeproxy_sync_proxy_rules_duration_seconds_bucket{job="kube proxy", instance=~"\$instance"}[5m]))
Network Programming Rate	sum(rate(kubeproxy_network_programming_duration_seconds_count{job="kube-proxy", instance=~"\$instance"}[5m]))
Network Programming Latency 99th Quantile	histogram_quantile(0.99, sum(rate(kubeproxy_network_programming_duration_seconds_bucket{job="kube-proxy", instance=~"\$instance"}[5m])) by (instance, le))
	sum(rate(rest_client_requests_total{job="kube-proxy", instance=~"\$instance",code=~"2"}[5m]))
Kube API	sum(rate(rest_client_requests_total{job="kube-proxy", instance=~"\$instance",code=~"3"}[5m]))
Request Rate	sum(rate(rest_client_requests_total{job="kube-proxy", instance=~"\$instance",code=~"4"}[5m]))
	sum(rate(rest_client_requests_total{job="kube-proxy", instance=~"\$instance",code=~"5"}[5m]))
Post Request Latency 99th Quantile	histogram_quantile(0.99, sum(rate(rest_client_request_duration_seconds_bucket{job="kube- proxy",instance=~"\$instance",verb="POST"}[5m])) by (verb, url, le))
Kube API	sum(rate(rest_client_requests_total{job="kube-proxy", instance=~"\$instance",code=~"2"}[5m]))

Request	sum(rate(rest_client_requests_total{job="kube-proxy", instance=~"\$instance",code=~"3"}[5m]))
Rate	sum(rate(rest_client_requests_total{job="kube-proxy", instance=~"\$instance",code=~"4"}[5m]))
	sum(rate(rest_client_requests_total{job="kube-proxy", instance=~"\$instance",code=~"5"}[5m]))
Post Request Latency 99th Quantile	histogram_quantile(0.99, sum(rate(rest_client_request_duration_seconds_bucket{job="kube- proxy",instance=~"\$instance",verb="POST"}[5m])) by (verb, url, le))
Get Request Latency 99th Quantile	histogram_quantile(0.99, sum(rate(rest_client_request_duration_seconds_bucket{job="kube- proxy", instance=~"\$instance", verb="GET"}[5m])) by (verb, url, le))
Memory	process_resident_memory_bytes{job="kube-proxy",instance=~"\$instance"}
CPU usage	rate(process_cpu_seconds_total{job="kube-proxy",instance=~"\$instance"}[5m])

# Scheduler(Independent Cluster)

Chart Name	Query Statement	Metrics Us
Up	sum(up{cluster=~"\$cluster", job="kube-scheduler"})	up
Kube API Request Rate	sum(rate(rest_client_requests_total{cluster=~"\$cluster",job="kube-scheduler", instance=~"\$instance",code=~"2"}[5m]))	rest_client
	sum(rate(rest_client_requests_total{cluster=~"\$cluster",job="kube-scheduler", instance=~"\$instance",code=~"3"}[5m]))	rest_client
	sum(rate(rest_client_requests_total{cluster=~"\$cluster",job="kube-scheduler", instance=~"\$instance",code=~"4"}[5m]))	rest_client
	sum(rate(rest_client_requests_total{cluster=~"\$cluster",job="kube-scheduler", instance=~"\$instance",code=~"5"}[5m]))	rest_client
Post Request Latency 99th Quantile	histogram_quantile(0.99, sum(rate(rest_client_request_duration_seconds_bucket{cluster=~"\$cluster",job="kube- scheduler", instance=~"\$instance", verb="POST"}[5m])) by (verb, url, le))	rest_client

Get Request Latency 99th Quantile	histogram_quantile(0.99, sum(rate(rest_client_request_duration_seconds_bucket{cluster=~"\$cluster",job="kube- scheduler", instance=~"\$instance", verb="GET"}[5m])) by (verb, url, le))	rest_client
Memory	process_resident_memory_bytes{cluster=~"\$cluster",job="kube-scheduler", instance=~"\$instance"}	process_re
CPU usage	rate(process_cpu_seconds_total{cluster=~"\$cluster",job="kube-scheduler", instance=~"\$instance"}[5m])	process_c

# Cluster Node Monitoring Details

Chart Name	Query Statement
Server resource	node_uname_info{job=~"\$job", cluster=~"\$cluster"} - 0
overview table	node_memory_MemTotal_bytes{job=~"\$job",cluster=~"\$cluster"} - 0
	count(node_cpu_seconds_total{job=~"\$job",mode='system',cluster=~"\$cluster"}) by (instance)
	sum(time() - node_boot_time_seconds{job=~"\$job",cluster=~"\$cluster"})by(instance)
	<pre>max((node_filesystem_size_bytes{job=~"\$job",cluster=~"\$cluster",fstype=~"ext.? xfs"}- node_filesystem_free_bytes{job=~"\$job",cluster=~"\$cluster",fstype=~"ext.? xfs"}) *100/(node_filesy {job=~"\$job",cluster=~"\$cluster",fstype=~"ext.? xfs"}+ (node_filesystem_size_bytes{job=~"\$job",cluster=~"\$cluster",fstype=~"ext.? xfs"}- node_filesystem_free_bytes{job=~"\$job",cluster=~"\$cluster",fstype=~"ext.? xfs"})))by(instance)</pre>
	(1 - avg(irate(node_cpu_seconds_total{job=~"\$job",mode="idle",cluster=~"\$cluster"}[5m])) by (insta
	(1 - (node_memory_MemAvailable_bytes{job=~"\$job",cluster=~"\$cluster"} / (node_memory_MemTotal_bytes{job=~"\$job",cluster=~"\$cluster"})))* 100



	node_load5{job=~"\$job",cluster=~"\$cluster"}
	max(irate(node_disk_written_bytes_total{job=~"\$job",cluster=~"\$cluster"}[5m])) by (instance)
	max(irate(node_network_receive_bytes_total{job=~"\$job",cluster=~"\$cluster"}[5m])*8) by (instance
	max(irate(node_network_transmit_bytes_total{job=~"\$job",cluster=~"\$cluster"}[5m])*8) by (instance
	node_load5{job=~"\$job",cluster=~"\$cluster"}
Overall total	count(node_cpu_seconds_total{job=~"\$job",cluster=~"\$cluster", mode='system'})
load and average CPU	sum(node_load5{job=~"\$job",cluster=~"\$cluster"})
utilization	avg(1 - avg(irate(node_cpu_seconds_total{job=~"\$job",mode="idle",cluster=~"\$cluster"}[5m])) by (i
Overall total memory and average memory utilization	sum(node_memory_MemTotal_bytes{job=~"\$job",cluster=~"\$cluster"})
	sum(node_memory_MemTotal_bytes{job=~"\$job",cluster=~"\$cluster"} - node_memory_MemAvailable_bytes{job=~"\$job",cluster=~"\$cluster"})
	(sum(node_memory_MemTotal_bytes{job=~"\$job",cluster=~"\$cluster"} - node_memory_MemAvailable_bytes{job=~"\$job",cluster=~"\$cluster"}) / sum(node_memory_MemTotal_bytes{job=~"\$job",cluster=~"\$cluster"}))*100
Overall total disk and average disk utilization	sum(avg(node_filesystem_size_bytes{job=~"\$job",cluster=~"\$cluster",fstype=~"xfs ext.*"})by(device
	sum(avg(node_filesystem_size_bytes{job=~"\$job",cluster=~"\$cluster",fstype=~"xfs ext.*"})by(device sum(avg(node_filesystem_free_bytes{job=~"\$job",cluster=~"\$cluster",fstype=~"xfs ext.*"})by(device



	(sum(avg(node_filesystem_size_bytes{job=~"\$job",cluster=~"\$cluster",fstype=~"xfs ext.*"})by(device sum(avg(node_filesystem_free_bytes{job=~"\$job",cluster=~"\$cluster",fstype=~"xfs ext.*"})by(device *100/(sum(avg(node_filesystem_avail_bytes{job=~"\$job",cluster=~"\$cluster",fstype=~"xfs ext.*"})by (sum(avg(node_filesystem_size_bytes{job=~"\$job",fstype=~"xfs ext.*"})by(device,instance)) - sum(avg(node_filesystem_free_bytes{job=~"\$job",cluster=~"\$cluster",fstype=~"xfs ext.*"})by(device
Running time	avg(time() - node_boot_time_seconds{instance=~"\$node",cluster=~"\$cluster"}) 75
CPU cores	count(node_cpu_seconds_total{cluster=~"\$cluster",instance=~"\$node", mode='system'})
Total memory	sum(node_memory_MemTotal_bytes{cluster=~"\$cluster",instance=~"\$node"})
Total CPU utilization	100 - (avg(irate(node_cpu_seconds_total{instance=~"\$node",mode="idle",cluster=~"\$cluster"}[5m]
Memory utilization	<pre>(1 - (node_memory_MemAvailable_bytes{instance=~"\$node",cluster=~"\$cluster"} / (node_memory_MemTotal_bytes{instance=~"\$node",cluster=~"\$cluster"})))* 100</pre>
Maximum partition utilization	(node_filesystem_size_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint= node_filesystem_free_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint=' /(node_filesystem_avail_bytes {cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoir (node_filesystem_size_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint= node_filesystem_free_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint=
CPU iowait	avg(irate(node_cpu_seconds_total{cluster=~"\$cluster",instance=~"\$node",mode="iowait"}[5m])) * 1
Available space for	node_filesystem_size_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint !
each partition	node_filesystem_avail_bytes {cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint
	(node_filesystem_size_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint node_filesystem_free_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint ! <sup>,</sup> *100/(node_filesystem_avail_bytes {cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint



	(node_filesystem_size_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint node_filesystem_free_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint !·
CPU	avg(irate(node_cpu_seconds_total{cluster=~"\$cluster",instance=~"\$node",mode="system"}[5m])) b
	avg(irate(node_cpu_seconds_total{cluster=~"\$cluster",instance=~"\$node",mode="user"}[5m])) by (
utilization	avg(irate(node_cpu_seconds_total{cluster=~"\$cluster",instance=~"\$node",mode="iowait"}[5m])) by
	(1 - avg(irate(node_cpu_seconds_total{cluster=~"\$cluster",instance=~"\$node",mode="idle"}[5m])) k
	node_memory_MemTotal_bytes{cluster=~"\$cluster",instance=~"\$node"}
Memory	node_memory_MemTotal_bytes{cluster=~"\$cluster",instance=~"\$node"} - node_memory_MemAvailable_bytes{cluster=~"\$cluster",instance=~"\$node"}
information	node_memory_MemAvailable_bytes{cluster=~"\$cluster",instance=~"\$node"}
	(1 - (node_memory_MemAvailable_bytes{cluster=~"\$cluster",instance=~"\$node"}/ (node_memory_MemTotal_bytes{cluster=~"\$cluster",instance=~"\$node"})))* 100
Network bandwidth usage per second System average load	irate(node_network_receive_bytes_total{cluster=~"\$cluster",instance=~'\$node',device=~"\$device"}
	irate(node_network_transmit_bytes_total{cluster=~"\$cluster",instance=~'\$node',device=~"\$device"
	node_load1{cluster=~"\$cluster",instance=~"\$node"}
	node_load5{cluster=~"\$cluster",instance=~"\$node"}
	node_load15{cluster=~"\$cluster",instance=~"\$node"}

	<pre>sum(count(node_cpu_seconds_total{cluster=~"\$cluster",instance=~"\$node", mode='system'}) by ( by(instance)</pre>
Disk read/write capacity per second	irate(node_disk_read_bytes_total{cluster=~"\$cluster",instance=~"\$node"}[5m])
	irate(node_disk_written_bytes_total{cluster=~"\$cluster",instance=~"\$node"}[5m])
Disk utilization	(node_filesystem_size_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint node_filesystem_free_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint !/ *100/(node_filesystem_avail_bytes {cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint (node_filesystem_size_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint node_filesystem_free_bytes{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.* xfs",mountpoint !/
	node_filesystem_files_free{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.? xfs"} / node_filesystem_files{cluster=~"\$cluster",instance=~'\$node',fstype=~"ext.? xfs"}
	irate(node_disk_reads_completed_total{cluster=~"\$cluster",instance=~"\$node"}[5m])
Disk read/write rate (IOPS)	irate(node_disk_writes_completed_total{cluster=~"\$cluster",instance=~"\$node"}[5m])
	node_disk_io_now{cluster=~"\$cluster",instance=~"\$node"}
I/O operation time percentage per second	irate(node_disk_io_time_seconds_total{cluster=~"\$cluster",instance=~"\$node"}[5m])
Time taken for each IO read/write	irate(node_disk_read_time_seconds_total{cluster=~"\$cluster",instance=~"\$node"}[5m]) / irate(node_disk_reads_completed_total{instance=~"\$node"}[5m])
	irate(node_disk_write_time_seconds_total{cluster=~"\$cluster",instance=~"\$node"}[5m]) / irate(node_disk_writes_completed_total{cluster=~"\$cluster",instance=~"\$node"}[5m])



	irate(node_disk_io_time_seconds_total{cluster=~"\$cluster",instance=~"\$node"}[5m])
	irate(node_disk_io_time_weighted_seconds_total{cluster=~"\$cluster",instance=~"\$node"}[5m])
	node_netstat_Tcp_CurrEstab{cluster=~"\$cluster",instance=~'\$node'}
	node_sockstat_TCP_tw{cluster=~"\$cluster",instance=~'\$node'}
	node_sockstat_sockets_used{cluster=~"\$cluster",instance=~'\$node'}
Network socket connection information	node_sockstat_UDP_inuse{cluster=~"\$cluster",instance=~'\$node'}
	node_sockstat_TCP_alloc{cluster=~"\$cluster",instance=~'\$node'}
	irate(node_netstat_Tcp_PassiveOpens{cluster=~"\$cluster",instance=~'\$node'}[5m])
	irate(node_netstat_Tcp_ActiveOpens{cluster=~"\$cluster",instance=~'\$node'}[5m])
	irate(node_netstat_Tcp_InSegs{cluster=~"\$cluster",instance=~'\$node'}[5m])
	irate(node_netstat_Tcp_OutSegs{cluster=~"\$cluster",instance=~'\$node'}[5m])
	irate(node_netstat_Tcp_RetransSegs{cluster=~"\$cluster",instance=~'\$node'}[5m])
Open file descriptors (left)/context switches per second (right)	node_filefd_allocated{cluster=~"\$cluster",instance=~"\$node"}
	irate(node_context_switches_total{cluster=~"\$cluster",instance=~"\$node"}[5m])
	(node_filefd_allocated{cluster=~"\$cluster",instance=~"\$node"}/node_filefd_maximum{cluster=~"\$clu *100



# Node Pod Monitoring

Query Statement
count(kube_pod_info{node=~"\$node"})
sum(kube_pod_container_resource_requests_memory_bytes{node=~"\$node"})by(node)
sum(kube_pod_container_resource_requests_cpu_cores{node=~"\$node"})by(node)
sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cluster node=~"\$node", container!="POD", container!=""}) by (pod)
sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cluster="\$cluster="}) by (pod)
sum(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster", node=~"\$node"}) by (pod)
sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cluster node=~"\$node", container!="POD", container!=""}) by (pod) / sum(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster", node=~"\$node"}) by (pod)
sum(kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster", node=~"\$node"}) by (pod)
sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cluster node=~"\$node", container!="POD", container!=""}) by (pod) / sum(kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster", node=~"\$node"}) by (pod)



Memory Usage	sum(node_namespace_pod_container:container_memory_working_set_bytes{cluster="\$cluster", node=~"\$node", container!="", container!="POD"}) by (pod)
	sum(node_namespace_pod_container:container_memory_working_set_bytes{cluster="\$cluster", node=~"\$node",container!="", container!="POD"}) by (pod)
	sum(kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster", node=~"\$node"}) by (pod)
Memory Quota	sum(node_namespace_pod_container:container_memory_working_set_bytes{cluster="\$cluster", node=~"\$node",container!="", container!="POD"}) by (pod) / sum(kube_pod_container_resource_requests_memory_bytes{node=~"\$node"}) by (pod)
	sum(kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster", node=~"\$node"}) by (por
	<pre>sum(node_namespace_pod_container:container_memory_working_set_bytes{cluster="\$cluster", node=~"\$node",container!="", container!="POD"}) by (pod) / sum(kube_pod_container_resource_limits_memory_bytes{node=~"\$node"}) by (pod)</pre>
Pod List	group (kube_pod_info{host_ip="\$node"})by(created_by_kind, created_by_name,host_network,pod_ip,pod,priority_class,namespace)
	min(kube_pod_info{host_ip="\$node"})by(pod) * on(pod) group_right() max(kube_pod_status_phase{}==1) by (pod, phase)
	min(kube_pod_info{host_ip="\$node"})by(pod) * on(pod) group_right() sum(container_memory_working_set_bytes) by (pod)
	min(kube_pod_info{host_ip="\$node"})by(pod) * on(pod) group_right() sum(rate(container_cpu_usage_seconds_total{image!=""}[5m])) by (pod)
	min(kube_pod_info{host_ip="\$node"})by(pod) * on(pod) group_right() max(time()-kube_pod_start_tim by (pod)

min(kube\_pod\_info{host\_ip="\$node"})by(pod) \* on(pod) max(kube\_pod\_status\_ready{condition="true" by (pod) or on() vector(0) min(kube\_pod\_info{host\_ip="\$node"})by(pod) \* on(pod) group\_right() max(rate(container\_network\_receive\_bytes\_total{image!=""}[5m])) by (pod) or on() vector(0) min(kube\_pod\_info{host\_ip="\$node"})by(pod) \* on(pod) group\_right() max(rate(container\_network\_transmit\_bytes\_total{image!=""}[5m])) by (pod) or on() vector(0) min(kube\_pod\_info{host\_ip="\$node"})by(pod) \* on(pod) group\_right() max(rate(container\_fs\_reads\_bytes\_total{container!="POD", container!=""}[5m])) by (pod) or on() vector(0) min(kube\_pod\_info{host\_ip="\$node"})by(pod) \* on(pod) group\_right() max(rate(container\_fs\_writes\_bytes\_total{container!="POD", container!=""}[5m])) by (pod) or on() vector(0)

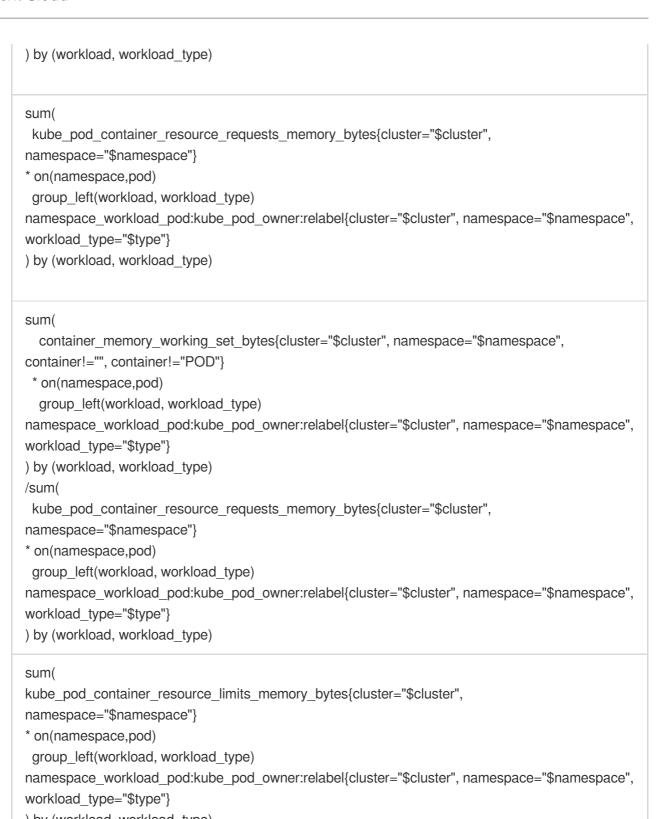
## Workload Monitoring Overview

Chart Name	Query Statement
CPU Usage	<pre>sum( node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cluster", namespace="\$namespace", container!="POD", container!=""} * on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload_type="\$type"} ) by (workload, workload_type) scalar(kube_resourcequota{cluster="\$cluster", namespace="\$namespace", type="hard",resource="requests.cpu"})</pre>



	scalar(kube_resourcequota{cluster="\$cluster", namespace="\$namespace", type="hard",resource="limits.cpu"})
CPU Quota	count(namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload_type="\$type"}) by (workload, workload_type)
	<pre>sum( node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cluster", namespace="\$namespace", container!="POD", container!=""} * on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload_type="\$type"} ) by (workload, workload_type)</pre>
	<pre>sum( kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster",     namespace="\$namespace"}  * on(namespace,pod)  group_left(workload, workload_type)  namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace",     workload_type="\$type"}  ) by (workload, workload_type)</pre>
	<pre>sum( node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cluster", namespace="\$namespace", container!="POD", container!=""} * on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload_type="\$type"} ) by (workload, workload_type)</pre>
	<pre>/sum( kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster", namespace="\$namespace"} * on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload_type="\$type"} ) by (workload, workload_type)</pre>
	<pre>sum( kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster", namespace="\$namespace"} * on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace", namespace="\$namespace", namespace="\$namespace", namespace="\$namespace", namespace="\$namespace", namespace="\$namespace", namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace="\$namespace", namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$n</pre>

	workload_type="\$type"} ) by (workload, workload_type)
	<pre>sum( node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cluster", namespace="\$namespace", container!="POD", container!=""} * on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload_type="\$type"} ) by (workload, workload_type) /sum( kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster", namespace="\$namespace"} * on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace"} * on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload_type="\$type"} ) by (workload, workload_type)</pre>
Memory Usage	<pre>sum(     container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace",     container!="", container!="POD"}     * on(namespace,pod)     group_left(workload, workload_type)     namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace",     workload_type="\$type"} ) by (workload, workload_type) scalar(kube_resourcequota{cluster="\$cluster", namespace="\$namespace",     "")</pre>
	type="hard",resource="requests.memory"}) scalar(kube_resourcequota{cluster="\$cluster", namespace="\$namespace", type="hard",resource="limits.memory"})
Memory Quota	count(namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload_type="\$type"}) by (workload, workload_type)
	<pre>sum( container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", container!="", container!="POD"} * on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload_type="\$type"}</pre>



```
) by (workload, workload_type)
```

#### sum(

container\_memory\_working\_set\_bytes{cluster="\$cluster", namespace="\$namespace", container!="", container!="POD"}

\* on(namespace,pod)

group\_left(workload, workload\_type)

namespace\_workload\_pod:kube\_pod\_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload\_type="\$type"}



) by (workload, workload_type)
/sum(
kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster",
namespace="\$namespace"}
* on(namespace,pod)
group_left(workload, workload_type)
namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace",
workload_type="\$type"}
) by (workload, workload_type)

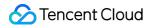
# Deployment

Chart Name	Query Statement
Age	time() - max(kube_deployment_created{cluster="\$cluster",namespace="\$namespace",deplo
Replicas(Pods)- Request	max(kube_deployment_spec_replicas{deployment="\$workload",cluster="\$cluster",namespa
Replicas(Pods)- Ready	max(kube_deployment_status_replicas_ready{deployment="\$workload",cluster="\$cluster",r
	max(kube_deployment_spec_replicas{deployment="\$workload",cluster="\$cluster",namespa
	max(kube_deployment_status_replicas{deployment="\$workload",cluster="\$cluster",namesp
Deplice Trend	min(kube_deployment_status_replicas_ready{deployment="\$workload",cluster="\$cluster",nappod)
Replica Trend	min(kube_deployment_status_replicas_available{deployment="\$workload",cluster="\$cluster (instance, pod)
	min(kube_deployment_status_replicas_updated{deployment="\$workload",cluster="\$cluster" (instance, pod)
	min(kube_deployment_status_replicas_unavailable{deployment="\$workload",cluster="\$clus (instance, pod)
CPU Usage	<pre>sum(     node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster=     container!="POD", container!=""}</pre>

	<pre>* on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe namespace="\$namespace", workload="\$workload", workload_type="deployment"} ) by (pod)</pre>
CPU Quota	<pre>sum(     node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster=     container!="POD", container!=""}     * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"}     ) by (pod)</pre>
	<pre>sum(     kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster", namespace="\$na  * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"}     ) by (pod)</pre>
	<pre>sum(     node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster=     container!="POD", container!=""}     * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"}     ) by (pod)</pre>
	<pre>/sum(    kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster", namespace="\$na  * on(namespace,pod)    group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe    namespace="\$namespace", workload="\$workload", workload_type="deployment"}    ) by (pod)</pre>
	<pre>sum(     kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster", namespace="\$name  * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"}     ) by (pod)</pre>
	<pre>sum(     node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster=     container!="POD", container!=""}     * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe</pre>



	<pre>namespace="\$namespace", workload="\$workload", workload_type="deployment"} ) by (pod) /sum(     kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster", namespace="\$name     * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"} ) by (pod)</pre>
CPU Limit-Total	<pre>sum(label_replace(     max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node),     "replicaset",     "\$1",     "created_by_name",     "(.+)" ) * on(replicaset) group_left()     max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De     (replicaset)     * on(pod) group_right() sum(kube_pod_container_resource_limits_cpu_cores{resource="cp     cluster="\$cluster",namespace="\$namespace"}) by (pod))</pre>
CPU Request- Total	<pre>sum(label_replace(     max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node),     "replicaset",     "\$1",     "created_by_name",     "(.+)" ) * on(replicaset) group_left()     max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De     (replicaset)     * on(pod) group_right() sum(kube_pod_container_resource_requests_cpu_cores{resource=     cluster="\$cluster",namespace="\$namespace"}) by (pod))</pre>
CPU Info	<pre>label_replace(   max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node),   "replicaset",   "\$1",   "created_by_name",   "(.+)" ) * on(replicaset) group_left()   max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De   (replicaset)</pre>

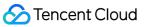


	* on(pod) group_right() max(rate(container_cpu_usage_seconds_total{cluster="\$cluster",namespace="\$namespace (pod, container)
	max(label_replace( max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node), "replicaset", "\$1",
	<pre>"created_by_name", "(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="Defination")</pre>
	<pre>(replicaset) * on(pod) group_right() max(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster",namespace="\$ by (pod, container))by(container)</pre>
	<pre>max(label_replace(     max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl     uid, pod, pod_ip, node),     "replicaset",     "\$1",     "created_by_name",</pre>
	<pre>"(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset) * on(pod) group_right() max(kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster",namespace="\$nam (pod, container))by(container)</pre>
CPU Usage/Limit	<pre>label_replace(   max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl   uid, pod, pod_ip, node),     "replicaset",     "\$1",     "created_by_name",     "(.+)"     ) * op(raplicaset) group_loft()</pre>
(%)	) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset) * on(pod) group_right() max(rate(container_cpu_usage_seconds_total{cluster="\$cluster",namespace="\$namespace (pod, container) / max by(container, pod) (kube_pod_container_resource_limits_cpu_cores{ cluster="\$cluster",namespace="\$namespace"})



CPU Usage/Request(%)	<pre>label_replace(     max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node),     "replicaset",     "\$1",     "created_by_name",     "(.+)" ) * on(replicaset) group_left()     max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De     (replicaset)     * on(pod) group_right()     max(rate(container_cpu_usage_seconds_total{cluster="\$cluster",namespace="\$namespace="\$namespace",owner_kind="De     (pod, container) / max by(container, pod) (kube_pod_container_resource_requests_cpu_col     cluster="\$cluster",namespace="\$namespace"})</pre>
CPU User Time(%)	<pre>avg(label_replace( max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node), "replicaset", "\$1", "created_by_name", "(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset) * on(pod) group_right() (max(rate(container_cpu_user_seconds_total{cluster="\$cluster",namespace="\$namespace" (pod,container) / max(rate(container_cpu_user_seconds_total{cluster="\$cluster",namespace="\$namespace" [5m])+rate(container_cpu_system_seconds_total{cluster="\$cluster",namespace="\$namespace="\$namespace" (pod,container))) by (pod,container)</pre>
Memory Usage	<pre>sum(     container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", co  * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"}     ) by (pod)</pre>
Memory Quota	<pre>sum(     container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", co  * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"}     ) by (pod)</pre>

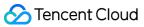
	<pre>sum(     kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster", namespace=     * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"}     ) by (pod)</pre>
	<pre>sum(     container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", co     * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"}     ) by (pod) /sum(     kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster", namespace=     * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"}     ) by (pod)</pre>
	<pre>sum(     kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster", namespace="\$n;  * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"}     ) by (pod)</pre>
	<pre>sum(     container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", co     * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"}     ) by (pod) /sum(     kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster", namespace="\$namespace="\$namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload_type] namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="deployment"}     ) by (pod)</pre>
Memory Limit- Total	<pre>sum(label_replace(     max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl     uid, pod, pod_ip, node),     "replicaset",     "\$1",     "created_by_name",</pre>



	"(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset) * on(pod) group_right() sum(container_spec_memory_limit_bytes{cluster="\$cluster",namesp
Memory Request- Total	<pre>sum(label_replace(     max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node),     "replicaset",     "\$1",     "created_by_name",     "(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De     (replicaset)     * on(pod) group_right() sum(kube_pod_container_resource_requests_memory_bytes{resource_resource="\$namespace"}) by (pod))</pre>
Memory Info	<pre>label_replace( max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node), "replicaset", "\$1", "created_by_name", "(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset) * on(pod) group_right() max by(container, pod) (container_memory_working_set_bytes{cluster container!="", image!="", container!="POD"})</pre>
	<pre>max(label_replace( max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node), "replicaset", "\$1", "created_by_name", "(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset) * on(pod) group_right() max by(container, pod) (kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster",namespace="\$</pre>



	<pre>max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node), "replicaset", "\$1", "created_by_name", "(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset) * on(pod) group_right() max by(container, pod) (kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster",namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$n</pre>
Memory Usage/Limit(%)	<pre>label_replace( max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node), "replicaset", "\$1", "created_by_name", "(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset) * on(pod) group_right() max by(container, pod) (container_memory_working_set_bytes{clus container!="", image!="", container!="POD"})/max by(container, pod) (kube_pod_container_resource_limits_memory_bytes{resource="memory", cluster="\$cluster"})</pre>
Memory Usage/Request(%)	<pre>label_replace(     max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node),     "replicaset",     "\$1",     "created_by_name",     "(.+)" ) * on(replicaset) group_left()     max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De     (replicaset)     * on(pod) group_right() max by(container, pod) (container_memory_working_set_bytes{clus     container!="", image!="", container!="POD"})/max by(container, pod)     (kube_pod_container_resource_requests_memory_bytes{resource="memory", cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="</pre>
Sockets	<pre>sum(label_replace(     max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl     uid, pod, pod_ip, node),     "replicaset",     "\$1",     "created_by_name",</pre>



	"(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset) * on(pod) group_right() sum(container_sockets{cluster="\$cluster",namespace="\$namespace
Network In	<pre>sum(label_replace(     max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl     uid, pod, pod_ip, node),     "replicaset",     "\$1",     "created_by_name",     "(.+)"     ) * on(replicaset) group_left()     max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De     (replicaset)     * on(pod) group_right() sum(rate(container_network_receive_bytes_total{cluster="\$cluster",     * cluster",     * on(pod) group_right() sum(rate(container_network_receive_bytes_total{cluster="\$cluster",     * cluster",     * cluster",     * cluster",     * cluster",     * cluster="\$cluster",     * cluster="\$cluster="\$cluster",     * cluster="\$cluster="\$cluster="\$cluster",     * cluster="\$cluster="\$cluster="\$cluster",     * cluster="\$cluster="\$cluster="\$cluster",     * cluster="\$cluster="\$cluster="\$cluster",     * cluster="\$cluster="\$cluster="\$cluster",     * cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster",     * cluster="\$cluster="\$cluster="\$cluster="\$cluster",     * cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster",     * cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster",     * cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster",     * cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$cluster",     * cluster="\$cluster="\$cluster="\$cluster="\$cluster="\$clust</pre>
Network Out	<pre>sum(label_replace(     max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node),     "replicaset",     "\$1",     "created_by_name",     "(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset)   * on(pod) group_right() sum(rate(container_network_transmit_bytes_total{cluster="\$cluster"}) </pre>
Network Errors	<pre>sum(label_replace(     max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl     uid, pod, pod_ip, node),     "replicaset",     "\$1",     "created_by_name",     "(.+)"     ) * on(replicaset) group_left()     max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De     (replicaset)     * on(pod) group_right() (sum(container_network_receive_errors_total{cluster="\$cluster",namespace="\$namespace="\$namespace="\$namespace",owner_kind="De     (replicaset)     * on(pod) group_right() (sum(container_network_receive_errors_total{cluster="\$cluster",namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$namespace="\$name</pre>
Network IO	label_replace( max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node),

	<pre>"replicaset", "\$1", "created_by_name", "(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset) * on(pod) group_right() max(rate(container_network_receive_bytes_total{cluster="\$cluster",</pre>
	<pre>label_replace(   max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node),   "replicaset",   "\$1",   "created_by_name",   "(.+)" ) * on(replicaset) group_left()   max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De   (replicaset)   * on(pod) group_right() max(rate(container_network_transmit_bytes_total{cluster="\$cluster"})</pre>
File System Read	<pre>label_replace( max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node), "replicaset", "\$1", "created_by_name", "(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset) * on(pod) group_right() max(rate(container_fs_reads_bytes_total{cluster="\$cluster",namesp container!=""}[5m])) by (pod,container)</pre>
File System Write	<pre>label_replace( max(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Repl uid, pod, pod_ip, node), "replicaset", "\$1", "created_by_name", "(.+)" ) * on(replicaset) group_left() max(kube_replicaset_owner{cluster="\$cluster",namespace="\$namespace",owner_kind="De (replicaset) * on(pod) group_right() max(rate(container_fs_writes_bytes_total{cluster="\$cluster",namespace container!=""}[5m])) by (pod,container)</pre>

## StatefulSet

Chart Name	Query Statement
Generation	max(kube_statefulset_metadata_generation{cluster="\$cluster",namespace="\$namespace",
Replicas(Pods)- Request	max(kube_statefulset_replicas{statefulset="\$workload",cluster="\$cluster",namespace="\$na
Replicas(Pods)- Ready	max(kube_statefulset_status_replicas_ready{statefulset="\$workload",cluster="\$cluster",nan
Age	time() - max(kube_statefulset_created{cluster="\$cluster",namespace="\$namespace",statefu
	max(kube_statefulset_replicas{statefulset="\$workload",cluster="\$cluster",namespace="\$na
	max(kube_statefulset_status_replicas{statefulset="\$workload",cluster="\$cluster",namespac pod)
Replica Trend	min(kube_statefulset_status_replicas_ready{statefulset="\$workload",cluster="\$cluster",nam (instance, pod)
	min(kube_statefulset_status_replicas_available{statefulset="\$workload",cluster="\$cluster",r (instance, pod)
	min(kube_statefulset_status_replicas_updated{statefulset="\$workload",cluster="\$cluster",n; (instance, pod)
CPU Usage	<pre>sum(     node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster=     namespace="\$namespace", container!="POD", container!=""}     * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="statefulset"}     ) by (pod)</pre>
CPU Quota	<pre>sum( node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster= namespace="\$namespace", container!="POD", container!=""} * on(namespace,pod)</pre>

group\_left(workload, workload\_type) namespace\_workload\_pod:kube\_pod\_owner:relabe namespace="\$namespace", workload="\$workload", workload\_type="statefulset"} ) by (pod)

#### sum(

kube\_pod\_container\_resource\_requests\_cpu\_cores{cluster="\$cluster", namespace="\$na
\* on(namespace,pod)

group\_left(workload, workload\_type) namespace\_workload\_pod:kube\_pod\_owner:relabe namespace="\$namespace", workload="\$workload", workload\_type="statefulset"} ) by (pod)

#### sum(

node\_namespace\_pod\_container:container\_cpu\_usage\_seconds\_total:sum\_rate{cluster= namespace="\$namespace", container!=""}

\* on(namespace,pod)

group\_left(workload, workload\_type) namespace\_workload\_pod:kube\_pod\_owner:relabe namespace="\$namespace", workload="\$workload", workload\_type="statefulset"}

### ) by (pod)

/sum(

kube\_pod\_container\_resource\_requests\_cpu\_cores{cluster="\$cluster", namespace="\$na
\* on(namespace,pod)

group\_left(workload, workload\_type) namespace\_workload\_pod:kube\_pod\_owner:relabe
namespace="\$namespace", workload="\$workload", workload\_type="statefulset"}
) by (pod)

#### sum(

kube\_pod\_container\_resource\_limits\_cpu\_cores{cluster="\$cluster", namespace="\$name
\* on(namespace,pod)

group\_left(workload, workload\_type) namespace\_workload\_pod:kube\_pod\_owner:relabe namespace="\$namespace", workload="\$workload", workload\_type="statefulset"} ) by (pod)

#### sum(

node\_namespace\_pod\_container:container\_cpu\_usage\_seconds\_total:sum\_rate{cluster= namespace="\$namespace", container!=""}

\* on(namespace,pod)

group\_left(workload, workload\_type) namespace\_workload\_pod:kube\_pod\_owner:relabe namespace="\$namespace", workload="\$workload", workload\_type="statefulset"} ) by (pod)

/sum(

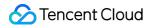
kube\_pod\_container\_resource\_limits\_cpu\_cores{cluster="\$cluster", namespace="\$name
\* on(namespace,pod)

group\_left(workload, workload\_type) namespace\_workload\_pod:kube\_pod\_owner:relabe namespace="\$namespace", workload="\$workload", workload\_type="statefulset"} ) by (pod)

CPU Limit-Total	<pre>sum(group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind= created_by_name="\$workload"}) by (pod) * on(pod) group_right() sum(kube_pod_container_resource_limits_cpu_cores{resource="cp cluster="\$cluster",namespace="\$namespace"}) by (pod))</pre>
CPU Request- Total	<pre>sum(group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind= created_by_name="\$workload"}) by (pod) * on(pod) group_right() sum(kube_pod_container_resource_requests_cpu_cores{resource= cluster="\$cluster",namespace="\$namespace"}) by (pod))</pre>
	group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Stat created_by_name="\$workload"}) by (pod) * on(pod) group_right() max(rate(container_cpu_usage_seconds_total{cluster="\$cluster",namespace="\$namespace [5m])) by (pod, container,image)
CPU Info	group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Stat created_by_name="\$workload"}) by (pod) * on(pod) group_right() max(kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster="\$cluster="}
	group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Stat created_by_name="\$workload"}) by (pod) * on(pod) group_right() max(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster",namespace="\$ container,image)
CPU Usage/Limit (%)	group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Stat created_by_name="\$workload"}) by (pod) * on(pod) group_right() max(rate(container_cpu_usage_seconds_total{cluster="\$cluster",namespace="\$namespace [5m])) by (pod, container) / max by(container, pod) (kube_pod_container_resource_limits_c; cluster="\$cluster",namespace="\$namespace"})
CPU Usage/Request(%)	group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Stat created_by_name="\$workload"}) by (pod) * on(pod) group_right() max(rate(container_cpu_usage_seconds_total{cluster="\$cluster",namespace="\$namespace [5m])) by (pod, container) / max by(container, pod) (kube_pod_container_resource_requests cluster="\$cluster",namespace="\$namespace"})
CPU User Time(%)	avg(group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind=' created_by_name="\$workload"}) by (pod)



	<pre>* on(pod) group_right() (max(rate(container_cpu_user_seconds_total{cluster="\$cluster",namespace="\$namespace' [5m])) by (pod, container,image) / max(rate(container_cpu_user_seconds_total{cluster="\$cluster",namespace="\$namespace" [5m])+rate(container_cpu_system_seconds_total{cluster="\$cluster",namespace="\$namespace" [5m])) by (pod,container,image))) by (pod,container,image)</pre>
Memory Usage	<pre>sum(     container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", co  * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="statefulset"}     ) by (pod)</pre>
Memory Quota	<pre>sum(     container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", co  * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="statefulset"}     ) by (pod)</pre>
	<pre>sum(    kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster", namespace=    * on(namespace,pod)    group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe    namespace="\$namespace", workload="\$workload", workload_type="statefulset"} ) by (pod)</pre>
	<pre>sum( container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", co * on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe namespace="\$namespace", workload="\$workload", workload_type="statefulset"} ) by (pod) /sum( kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster", namespace= * on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe namespace="\$namespace", workload="\$workload", workload_type="statefulset"} ) by (pod)</pre>
	sum( kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster", namespace="\$n; * on(namespace,pod) group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe



	namespace="\$namespace", workload="\$workload", workload_type="statefulset"} ) by (pod)
	<pre>sum(     container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", co     * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="statefulset"} ) by (pod) /sum(     kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster", namespace="\$ni     * on(namespace,pod)     group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$nimespace", workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload_type) namespace_workload_pod:kube_pod_owner:relabe     namespace="\$namespace", workload="\$workload", workload_type="statefulset"} ) by (pod)</pre>
Memory Limit- Total	<pre>sum(group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind= created_by_name="\$workload"}) by (pod) * on(pod) group_right() sum(container_spec_memory_limit_bytes{cluster="\$cluster",namespace="\$namespace",coi (pod))</pre>
Memory Request- Total	<pre>sum(group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind= created_by_name="\$workload"}) by (pod) * on(pod) group_right() sum(kube_pod_container_resource_requests_memory_bytes{resou cluster="\$cluster",namespace="\$namespace"}) by (pod))</pre>
	<pre>avg(group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind=' created_by_name="\$workload"}) by (pod) * on(pod) group_right() max by(container, pod, image) (container_memory_working_set_bytes{cluster="\$cluster",namespace="\$namespace", cont container!="POD"}))by (container, pod, image)</pre>
Memory Info	<pre>max(avg(group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_k created_by_name="\$workload"}) by (pod) * on(pod) group_right() max by(container, pod, image) (kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster",namespace="\$ pod))by(container)</pre>
	<pre>max(avg(group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_k created_by_name="\$workload"}) by (pod) * on(pod) group_right() max by(container, pod) (kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster",namespace="\$nam pod))by(container)</pre>
Memory	group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Stat

Usage/Limit(%)	<pre>created_by_name="\$workload"}) by (pod) * on(pod) group_right() max by(container, pod) (container_memory_working_set_bytes{cluster="\$cluster",namespace="\$namespace", cont container!="POD"})/max by(container, pod) (kube_pod_container_resource_limits_memory_ cluster="\$cluster",namespace="\$namespace"})</pre>
Memory Usage/Request(%)	group(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="Stat created_by_name="\$workload"}) by (pod) * on(pod) group_right() max by(container, pod) (container_memory_working_set_bytes{cluster="\$cluster",namespace="\$namespace", cont container!="POD"})/max by(container, pod) (kube_pod_container_resource_requests_mem_ cluster="\$cluster",namespace="\$namespace"})
Sockets	<pre>sum(sum(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="\$ created_by_name="\$workload"}) by (pod) * on(pod) group_right() sum(container_sockets{cluster="\$cluster",namespace="\$namespace"})</pre>
Network In	<pre>sum(sum(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="\$ created_by_name="\$workload"}) by (pod) * on(pod) group_right() sum(rate(container_network_receive_bytes_total{cluster="\$cluster", by (pod))</pre>
Network Out	<pre>sum(sum(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="\$ created_by_name="\$workload"}) by (pod) * on(pod) group_right() sum(rate(container_network_transmit_bytes_total{cluster="\$cluster" by (pod))</pre>
Network Errors	<pre>sum(sum(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="\$ created_by_name="\$workload"}) by (pod) * on(pod) group_right() (sum(container_network_receive_errors_total{cluster="\$cluster",nan sum(container_network_transmit_errors_total{cluster="\$cluster",namespace="\$namespace"</pre>
Network IO	<pre>sum(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="State created_by_name="\$workload"}) by (pod) * on(pod) group_right() max(rate(container_network_receive_bytes_total{cluster="\$cluster", by (pod)</pre>
	-sum(kube_pod_info{cluster="\$cluster",namespace="\$namespace",created_by_kind="State created_by_name="\$workload"}) by (pod) * on(pod) group_right() max(rate(container_network_transmit_bytes_total{cluster="\$cluster" by (pod)

## DaemonSet

Chart Name	Query Statement
CPU Usage	<pre>sum( node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cluster", namespace="\$namespace", container!="POD", container!=""}  * on(namespace,pod)  group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload="\$workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", ) by (pod)</pre>
CPU Quota	<pre>sum( node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cluster", namespace="\$namespace", container!="POD", container!=""}  * on(namespace,pod)  group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload="\$workload", workload_type="daemonset"} ) by (pod)</pre>
	<pre>sum(    kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster",    namespace="\$namespace"}    * on(namespace,pod)    group_left(workload, workload_type)    namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace",    workload="\$workload", workload_type="daemonset"} ) by (pod)</pre>
	<pre>sum( node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cluster", namespace="\$namespace", container!="POD", container!=""}  * on(namespace,pod)  group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload="\$workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload="\$workload", workload_type="daemonset"} ) by (pod) /sum(  kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster",</pre>

	<pre>namespace="\$namespace"}  * on(namespace,pod)  group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload="\$workload", workload_type="daemonset"} ) by (pod)</pre>
	<pre>sum(    kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster", namespace="\$namespace"}  * on(namespace,pod)  group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload="\$workload", workload_type="daemonset"} ) by (pod)</pre>
	sum(
	node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cluster", namespace="\$namespace", container!="POD", container!=""} * on(namespace,pod)
	group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload="\$workload", workload_type="daemonset"} ) by (pod)
	/sum( kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster", namespace="\$namespace"} * on(namespace,pod) group_left(workload, workload_type)
	namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload="\$workload", workload_type="daemonset"} ) by (pod)
Memory	<pre>sum(     container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace",     container!="", container!="POD"}     * on(namespace,pod)     group_loft(workload_workload_type)</pre>
Usage	group_left(workload, workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", workload="\$workload", workload_type="daemonset"} ) by (pod)
Memory Quota	sum( container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", container!="", container!="POD"} * on(namespace,pod) group_left(workload, workload_type)

-	ce_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", ="\$workload", workload_type="daemonset"}
namespace * on(name group_ namespace	d_container_resource_requests_memory_bytes{cluster="\$cluster", ce="\$namespace"} nespace,pod) left(workload, workload_type) ce_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", ="\$workload", workload_type="daemonset"}
container * on(nam group_ namespac	er_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", !="", container!="POD"} nespace,pod) left(workload, workload_type) ce_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", ="\$workload", workload_type="daemonset"}
namespace * on(namespace namespace	d_container_resource_requests_memory_bytes{cluster="\$cluster", ce="\$namespace"} nespace,pod) left(workload, workload_type) ce_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", ="\$workload", workload_type="daemonset"}
namespace * on(namespace group_ namespace	d_container_resource_limits_memory_bytes{cluster="\$cluster", ce="\$namespace"} nespace,pod) left(workload, workload_type) ce_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace", ="\$workload", workload_type="daemonset"}
container * on(nam group_	er_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", !="", container!="POD"} nespace,pod) left(workload, workload_type) ce_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace",

workload="\$workload", workload_type="daemonset"}
) by (pod)
/sum(
kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster",
namespace="\$namespace"}
* on(namespace,pod)
group_left(workload, workload_type)
namespace_workload_pod:kube_pod_owner:relabel{cluster="\$cluster", namespace="\$namespace",
workload="\$workload", workload_type="daemonset"}
) by (pod)

### **Chart Name Query Statement** time() - max(kube\_pod\_created{pod=~"\$pod",cluster="\$cluster",namespace="\$namespace"}) Age Restart Count-Last ceil(sum (increase(kube\_pod\_container\_status\_restarts\_total{pod=~"\$pod",cluster="\$cluster",name 1 Hour Requestssum(kube pod container resource requests cpu cores{pod=~"\$pod"}) or vector(0) CPU Requestssum(kube\_pod\_container\_resource\_requests\_memory\_bytes{pod=~"\$pod"}) or vector(0) Memory Limits-CPU sum(kube pod container resource limits cpu cores{pod=~"\$pod"}) or vector(0) Limitssum(kube pod container resource limits memory bytes{pod=~"\$pod"}) or vector(0) Memory Containers group by (image, container,pod) (kube\_pod\_container\_info{cluster="\$cluster",namespace="\$name: sum by (container,pod)(kube\_pod\_container\_resource\_requests\_cpu\_cores{cluster="\$cluster",nan sum by (container,pod)(kube\_pod\_container\_resource\_requests\_memory\_bytes{cluster="\$cluster"

## **Cluster Pod Monitoring**

pod=~"\$pod"})



	max by (container,pod)(kube_pod_container_status_running{cluster="\$cluster",namespace="\$nam
	sum by (container,pod)(kube_pod_container_resource_limits{resource="cpu",cluster="\$cluster",nai
	sum by (container,pod)(kube_pod_container_resource_limits{resource="memory",cluster="\$cluster pod=~"\$pod"})
	max by (container,pod)(kube_pod_container_status_restarts_total{cluster="\$cluster",namespace="
CPU Usage (%)	max(irate(container_cpu_usage_seconds_total{pod=~"\$pod",container!="",container!="POD",cluste [1m])) by (container,namespace,pod) / max(container_spec_cpu_quota{pod=~"\$pod",container!="",container!="POD",cluster="\$cluster",na (container,namespace,pod) or on() vector(0)
CPU Usage By Cores	max(irate(container_cpu_usage_seconds_total{pod=~"\$pod",container!="",container!="POD",cluste [1m])) by (pod,container,namespace)or on() vector(0)
CPU Load (10s)	max(container_cpu_load_average_10s{namespace=~"\$namespace", pod=~"\$pod", container!="", ( 1000)by(pod,container)
CPU Throttled Percent	max (rate (container_cpu_cfs_throttled_seconds_total{image!="", container!="", cluster="\$cluster",r pod=~"\$pod"}[1m])) by (container,pod) / max (rate (container_cpu_cfs_periods_total{image!="", cor cluster="\$cluster",namespace=~"\$namespace", pod=~"\$pod"}[1m])) by (container,pod) or on() vect
CPU Quota	<pre>sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cl pod="\$pod", container!="POD", container!=""}) by (container)</pre>
	sum(kube_pod_container_resource_requests_cpu_cores{cluster="\$cluster", namespace="\$names
	<pre>sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cl pod="\$pod", container!="POD", container!=""}) by (container) / sum(kube_pod_container_resource_ namespace="\$namespace", pod="\$pod"}) by (container)</pre>
	sum(kube_pod_container_resource_limits_cpu_cores{cluster="\$cluster", namespace="\$namespac
	sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{cluster="\$cl pod="\$pod", container!="POD", container!=""}) by (container) / sum(kube_pod_container_resource_ namespace="\$namespace", pod="\$pod"}) by (container)



Memory Usage (WSS)	max(container_memory_working_set_bytes{pod=~"\$pod",container !="",container!="POD",cluster= by (pod,namespace,container)
Memory Usage	max(container_memory_usage_bytes{pod=~"\$pod",container !="",container!="POD",cluster="\$clus (pod,namespace,container)
Memory Usage (RSS)	max(container_memory_rss{pod=~"\$pod",container !="",container!="POD",cluster="\$cluster",name (pod,namespace,container) or on() vector(0)
Memory Cache	max(container_memory_cache{pod=~"\$pod",container !="",container!="POD",cluster="\$cluster",na (pod,namespace,container)
Usage WSS/Limit (%)	(max(container_memory_working_set_bytes{pod=~"\$pod",container !="",container!="POD",cluster= by (pod,namespace,container)/ max(container_spec_memory_limit_bytes{pod=~"\$pod",container !="",container!="POD",cluster="\$cluster",namespace=~"\$namespace"}) by (pod,namespace, contai
Usage/Limit (%)	(max(container_memory_usage_bytes{pod=~"\$pod",container !="",container!="POD",cluster="\$clu (pod,namespace,container)/ max(container_spec_memory_limit_bytes{pod=~"\$pod",container !="",container!="POD",cluster="\$cluster",namespace=~"\$namespace"}) by (pod,namespace, contai
Usage RSS/Limit (%)	(max(container_memory_rss{pod=~"\$pod",container !="",container!="POD",cluster="\$cluster",name (pod,namespace,container)/ sum(container_spec_memory_limit_bytes{pod=~"\$pod",container !="",container!="POD",cluster="\$cluster",namespace=~"\$namespace"}) by (pod,namespace, contai
Memory Failcnt	max (increase(container_memory_failcnt{cluster="\$cluster",namespace=~"\$namespace", pod=~"\$r (pod,container)
Memory Quota	sum(container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", pod="\$ by (container)
	sum(kube_pod_container_resource_requests_memory_bytes{cluster="\$cluster", namespace="\$na
	sum(container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", pod="\$ by (container) / sum(kube_pod_container_resource_requests_memory_bytes{namespace="\$name
	sum(kube_pod_container_resource_limits_memory_bytes{cluster="\$cluster", namespace="\$name: (container)
	sum(container_memory_working_set_bytes{cluster="\$cluster", namespace="\$namespace", pod="\$ by (container) / sum(kube_pod_container_resource_limits_memory_bytes{namespace="\$namespa

Network Input	max (rate (container_network_receive_bytes_total{image!="",cluster="\$cluster",namespace=~"\$nai by(pod)
Network Output	max (rate (container_network_transmit_bytes_total{image!="",cluster="\$cluster",namespace=~"\$na [1m]))by(pod)
Network	max (increase (container_network_receive_packets_dropped_total{id!="/", cluster="\$cluster",name [1m])) by (pod,interface) / max (increase (container_network_receive_packets_total{id!="/", cluster= pod=~"\$pod"}[1m])) by (pod,interface)
Input Error (%)	max (increase (container_network_receive_errors_total{id!="/", cluster="\$cluster",namespace=~"\$r (pod,interface) / max (increase (container_network_receive_packets_total{id!="/", cluster="\$cluster pod=~"\$pod"}[1m])) by (pod,interface)
Network	max (increase (container_network_transmit_packets_dropped_total{id!="/", cluster="\$cluster",nam [1m])) by (pod,interface) / max (increase (container_network_transmit_packets_total{id!="/", cluster pod=~"\$pod"}[1m])) by (pod,interface)
Output Error (%)	max (increase (container_network_transmit_errors_total{id!="/", cluster="\$cluster",namespace=~"\$ (pod,interface) / max (increase (container_network_receive_packets_total{id!="/", cluster="\$cluster pod=~"\$pod"}[1m])) by (pod,interface)
File System Read	max (rate(container_fs_reads_bytes_total{cluster="\$cluster",namespace=~"\$namespace", pod=~"\$ (container,pod)
File System Write	max (rate(container_fs_writes_bytes_total{cluster="\$cluster",namespace=~"\$namespace", pod=~"\$ (container,pod)
Network Socket	max(container_sockets{cluster="\$cluster",namespace=~"\$namespace", pod=~"\$pod", container!="
Process Number	count(container_processes{cluster="\$cluster",namespace=~"\$namespace", pod=~"\$pod", containe

# Cluster Network Monitoring

Chart Name	Query Statement
Current Rate of	sort_desc(sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~".+"} (namespace))



Bytes Received	
Current Rate of Bytes Transmitted	sort_desc(sum(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~".+" (namespace))
	sort_desc(sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~".+"} (namespace))
	sort_desc(sum(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~".+" (namespace))
	sort_desc(avg(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~".+"}[ (namespace))
Current	sort_desc(avg(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~".+"} (namespace))
Status	sort_desc(sum(irate(container_network_receive_packets_total{cluster=~"\$cluster",namespace=~" (namespace))
	sort_desc(sum(irate(container_network_transmit_packets_total{cluster=~"\$cluster",namespace=~" (namespace))
	sort_desc(sum(irate(container_network_transmit_packets_total{cluster=~"\$cluster",namespace=~" (namespace))
	sort_desc(sum(irate(container_network_transmit_packets_dropped_total{cluster=~"\$cluster",name [5m])) by (namespace))
Average Rate of Bytes Received	sort_desc(avg(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~".+"}[ (namespace))
Average Rate of Bytes Transmitted	sort_desc(avg(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~".+"} (namespace))
Receive Bandwidth	sort_desc(sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~".+"} (namespace))
Transmit Bandwidth	sort_desc(sum(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~".+" (namespace))



Rate of Received Packets	sort_desc(sum(irate(container_network_receive_packets_total{cluster=~"\$cluster",namespace=~". (namespace))
Rate of Transmitted Packets	sort_desc(sum(irate(container_network_transmit_packets_total{cluster=~"\$cluster",namespace=~" (namespace))
Rate of Received Packets Dropped	sort_desc(sum(irate(container_network_receive_packets_dropped_total{cluster=~"\$cluster",names [5m])) by (namespace))
Rate of Transmitted Packets Dropped	sort_desc(sum(irate(container_network_transmit_packets_dropped_total{cluster=~"\$cluster",name [5m])) by (namespace))
Rate of TCP Retransmits out of all sent segments	sort_desc(sum(rate(node_netstat_Tcp_RetransSegs{cluster=~"\$cluster"}[5m]) / rate(node_netstat_Tcp_OutSegs{cluster=~"\$cluster"}[\$interval:\$resolution])) by (instance))
Rate of TCP SYN Retransmits out of all retransmits	sort_desc(sum(rate(node_netstat_TcpExt_TCPSynRetrans{cluster=~"\$cluster"}[\$interval:\$resolutic rate(node_netstat_Tcp_RetransSegs{cluster=~"\$cluster"}[\$interval:\$resolution])) by (instance))

## Namespace Pods Network Monitoring

Chart Name	Query Statement
Current Rate of Bytes Received	sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$namespace"]
Current Rate of	sum(irate(container_network_transmit_bytes_total{namespace=~"\$namespace"}[5m]))

Bytes Transmitted	
	sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$namespace"] (pod)
	sum(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~"\$namespace' (pod)
Current	sum(irate(container_network_receive_packets_total{cluster=~"\$cluster",namespace=~"\$namespac (pod)
Status	sum(irate(container_network_transmit_packets_total{cluster=~"\$cluster",namespace=~"\$namespace(pod)
	sum(irate(container_network_receive_packets_dropped_total{cluster=~"\$cluster",namespace=~"\$r [5m])) by (pod)
	sum(irate(container_network_transmit_packets_dropped_total{cluster=~"\$cluster",namespace=~"\$ [5m])) by (pod)
Receive Bandwidth	sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$namespace"] (pod)
Transmit Bandwidth	sum(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~"\$namespace' (pod)
Rate of Received Packets	sum(irate(container_network_receive_packets_total{cluster=~"\$cluster",namespace=~"\$namespac (pod)
Rate of Transmitted Packets	sum(irate(container_network_transmit_packets_total{cluster=~"\$cluster",namespace=~"\$namespa (pod)
Rate of Received Packets Dropped	sum(irate(container_network_receive_packets_dropped_total{cluster=~"\$cluster",namespace=~"\$r [5m])) by (pod)
Rate of Transmitted Packets Dropped	sum(irate(container_network_transmit_packets_dropped_total{cluster=~"\$cluster", namespace=~"\$namespace"}[5m])) by (pod)

# Namespace Workload Network Monitoring



Chart Name	Query Statement
Current Rate of Bytes Received	<pre>sort_desc(sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$na * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))</pre>
Current Rate of Bytes Transmitted	<pre>sort_desc(sum(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~"\$n * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))</pre>
Current Status	<pre>sort_desc(sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$na * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))</pre>
	sort_desc(sum(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~"\$n * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))
	sort_desc(avg(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$nai * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))
	sort_desc(avg(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~"\$na * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))
	<pre>sort_desc(sum(irate(container_network_receive_packets_total{cluster=~"\$cluster",namespace=~"\$ * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))</pre>
	sort_desc(sum(irate(container_network_transmit_packets_total{cluster=~"\$cluster",namespace=~"



	* on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))
	<pre>sort_desc(sum(irate(container_network_receive_packets_dropped_total{cluster=~"\$cluster",names [5m]) * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))</pre>
	<pre>sort_desc(sum(irate(container_network_transmit_packets_dropped_total{cluster=~"\$cluster",name [5m]) * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))</pre>
Average Rate of Bytes Received	sort_desc(avg(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$nal * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))
Average Rate of Bytes Transmitted	sort_desc(avg(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~"\$na * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))
Receive Bandwidth	<pre>sort_desc(sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$na * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))</pre>
Transmit Bandwidth	sort_desc(sum(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~"\$n * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))
Rate of Received Packets	sort_desc(sum(irate(container_network_receive_packets_total{cluster=~"\$cluster",namespace=~"\$ * on (namespace,pod) group_left(workload,workload_type)

	namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))
Rate of Transmitted Packets	<pre>sort_desc(sum(irate(container_network_transmit_packets_total{cluster=~"\$cluster",namespace=~" * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))</pre>
Rate of Received Packets Dropped	<pre>sort_desc(sum(irate(container_network_receive_packets_dropped_total{cluster=~"\$cluster",names [5m]) * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))</pre>
Rate of Transmitted Packets Dropped	<pre>sort_desc(sum(irate(container_network_transmit_packets_dropped_total{cluster=~"\$cluster",name [5m]) * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload_type="\$type"}) by (workload))</pre>

# Pod Network Monitoring

Chart Name	Query Statement
Current Rate of Bytes Received	sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$namespace", pod=~"\$pod"}[5m]))
Current Rate of Bytes Transmitted	sum(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~"\$namespace' pod=~"\$pod"}[5m]))
Receive Bandwidth	sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$namespace", pod=~"\$pod"}[5m])) by (pod)
Transmit Bandwidth	sum(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~"\$namespace' pod=~"\$pod"}[5m])) by (pod)
Rate of	sum(irate(container_network_receive_packets_total{cluster=~"\$cluster",namespace=~"\$namespac



Received Packets	pod=~"\$pod"}[5m])) by (pod)
Rate of Transmitted Packets	sum(irate(container_network_transmit_packets_total{cluster=~"\$cluster",namespace=~"\$namespar pod=~"\$pod"}[5m])) by (pod)
Rate of Received Packets Dropped	sum(irate(container_network_receive_packets_dropped_total{cluster=~"\$cluster",namespace=~"\$r pod=~"\$pod"}[5m])) by (pod)
Rate of Transmitted Packets Dropped	sum(irate(container_network_transmit_packets_dropped_total{cluster=~"\$cluster",namespace=~"\$ pod=~"\$pod"}[5m])) by (pod)

# Workload Network Monitoring

Chart Name	Query Statement
Current Rate of Bytes Received	<pre>sort_desc(sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$na * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload=~"\$workload_yod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace="\$namespace") by (pod))</pre>
Current Rate of Bytes Transmitted	<pre>sort_desc(sum(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~"\$n * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload=~"\$workload", workload_type="\$type"}) by (pod))</pre>
Average Rate of Bytes Received	sort_desc(avg(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$nal * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload=~"\$workload", workload_type="\$type"}) by (pod))
Average Rate of Bytes Transmitted	sort_desc(avg(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~"\$na * on (namespace,pod) group_left(workload,workload_type)



	namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespaceworkload=~"\$workload", workload_type="\$type"}) by (pod))
Receive Bandwidth	<pre>sort_desc(sum(irate(container_network_receive_bytes_total{cluster=~"\$cluster",namespace=~"\$na * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload=~"\$workload_yod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace="""</pre>
Transmit Bandwidth	<pre>sort_desc(sum(irate(container_network_transmit_bytes_total{cluster=~"\$cluster",namespace=~"\$n * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload=~"\$workload_type="\$type"}) by (pod))</pre>
Rate of Received Packets	<pre>sort_desc(sum(irate(container_network_receive_packets_total{cluster=~"\$cluster",namespace=~"\$ * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload=~"\$workload", workload_type="\$type"}) by (pod))</pre>
Rate of Transmitted Packets	sort_desc(sum(irate(container_network_transmit_packets_total{cluster=~"\$cluster",namespace=~" * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload=~"\$workload_type="\$type"}) by (pod))
Rate of Received Packets Dropped	<pre>sort_desc(sum(irate(container_network_receive_packets_dropped_total{cluster=~"\$cluster",names [5m]) * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload=~"\$workload", workload_type="\$type"}) by (pod))</pre>
Rate of Transmitted Packets Dropped	<pre>sort_desc(sum(irate(container_network_transmit_packets_dropped_total{cluster=~"\$cluster",name [5m]) * on (namespace,pod) group_left(workload,workload_type) namespace_workload_pod:kube_pod_owner:relabel{cluster=~"\$cluster",namespace=~"\$namespace workload=~"\$workload", workload_type="\$type"}) by (pod))</pre>

# PVC Storage Monitoring

Chart

Metrics Used

🔗 Tencent Cloud

Name		
	( sum without(instance, node) (kubelet_volume_stats_capacity_bytes{cluster="\$cluster", job="kubelet", namespace="\$namespace", persistentvolumeclaim="\$volume"})	kubelet_volume_stats_capacity_bytes
Volume Space Usage	sum without(instance, node) (kubelet_volume_stats_available_bytes{cluster="\$cluster", job="kubelet", namespace="\$namespace", persistentvolumeclaim="\$volume"}) )	kubelet_volume_stats_available_bytes
	sum without(instance, node) (kubelet_volume_stats_available_bytes{cluster="\$cluster", job="kubelet", namespace="\$namespace", persistentvolumeclaim="\$volume"})	kubelet_volume_stats_available_bytes
	( kubelet_volume_stats_capacity_bytes{cluster="\$cluster", job="kubelet", namespace="\$namespace", persistentvolumeclaim="\$volume"} -	kubelet_volume_stats_capacity_bytes
Volume Space Usage	kubelet_volume_stats_available_bytes{cluster="\$cluster", job="kubelet", namespace="\$namespace", persistentvolumeclaim="\$volume"} )	kubelet_volume_stats_available_bytes
	/ kubelet_volume_stats_capacity_bytes{cluster="\$cluster", job="kubelet", namespace="\$namespace", persistentvolumeclaim="\$volume"} * 100	kubelet_volume_stats_capacity_bytes
Volume inodes Usage	sum without(instance, node) (kubelet_volume_stats_inodes_used{cluster="\$cluster", job="kubelet", namespace="\$namespace", persistentvolumeclaim="\$volume"})	kubelet_volume_stats_inodes_used
	( sum without(instance, node) (kubelet_volume_stats_inodes{cluster="\$cluster", job="kubelet", namespace="\$namespace", persistentvolumeclaim="\$volume"}) -	kubelet_volume_stats_inodes

	sum without(instance, node) (kubelet_volume_stats_inodes_used{cluster="\$cluster", job="kubelet", namespace="\$namespace", persistentvolumeclaim="\$volume"}) )	kubelet_volume_stats_inodes_used
Volume	kubelet_volume_stats_inodes_used{cluster="\$cluster", job="kubelet", namespace="\$namespace", persistentvolumeclaim="\$volume"} /	kubelet_volume_stats_inodes_used
Usage	kubelet_volume_stats_inodes{cluster="\$cluster", job="kubelet", namespace="\$namespace", persistentvolumeclaim="\$volume"} * 100	kubelet_volume_stats_inodes

# Data Collection Configuration

Last updated : 2024-08-07 22:07:45

# Overview

This document describes how to configure monitoring collection items for the associated cluster.

## Prerequisites

Before configuring monitoring collection items, you need to perform the following operations:

The Prometheus monitoring instance has been successfully created.

The cluster to be monitored has been associated with the corresponding instance.

### Directions

### **Configuring Data Collection**

1. Log in to TMP Console.

2. On the instance list page, select the name of the instance that requires configuring data collection rules and enter its details page.

3. On the **Data Collection > Integration with TKE** page, click **Data Collection Configuration** on the right of the instance to enter the collection configuration list page.

4. Click **Metrics collection management** at the top of the page to pop up the Basic Monitoring Metrics Collection Management page.

If you need to collect all container chart metrics, you can click **One-click collection of preset chart metrics** in the pop-up window, and then click **OK** in the pop-up confirmation window to collect all preset chart metrics.

One-click collection of preset chart me	trics		Search by metric name	
Metric filtering Please select filt ✓				
Please select				~
— Metric Name	Real-Time Collec 🍸	Free 🝸	Collection com	Metric Collection S Before Filter(i)
✓ machine_memory_bytes	Collected	Yes	cadvisor	0.00Count/s
cadvisor_version_info	Uncollected	No	eks-network	0.27Count/s
container_blkio_device_usage_t	Uncollected	No	eks-network	48.80Count/s
container_cpu_cfs_periods_total	Uncollected	No	eks-network	0.80Count/s
container_cpu_cfs_throttled_peri	Uncollected	No	eks-network	0.80Count/s
✓ container_cpu_cfs_throttled_sec	Uncollected	No	eks-network	0.80Count/s
container_cpu_load_average_10s	Uncollected	No	eks-network	2.13Count/s
✓ container_cpu_system_seconds	Uncollected	No	eks-network	2.13Count/s

### Modify collection target

You are about to modify the collection status of 7 metrics. Please confirm before making the changes:

Metric Name	original state	Modified st	Free	Collection c	M Fi
cadvisor_version_info	Uncollected	Collected	No	eks-network	0.3
container_blkio_device_usage_t	Uncollected	Collected	No	eks-network	48
container_cpu_cfs_periods_total	Uncollected	Collected	No	eks-network	0.(
container_cpu_cfs_throttled_peri	Uncollected	Collected	No	eks-network	0.)
container_cpu_cfs_throttled_sec	Uncollected	Collected	No	eks-network	0.
container_cpu_load_average_10s	Uncollected	Collected	No	eks-network	2.1
container_cpu_system_seconds	Uncollected	Collected	No	eks-network	2.1
4					
		ОК	Cancel		

If only part of the metrics need to be collected, you can filter the metrics through, **monitoring charts** and **Recording Rule**, select the metrics you need, and then click **OK**. For details, see the following figure.

One-click collection of preset chart m	Search by metric name			
Please select filt V Please select filt V	1			~
Recording Rule Metric Name	Real-Time Collec 🝸	Free	Collection com	Metric Collection Sp Before Filter(i)
container_network_receive_byte	Collected	Yes	cadvisor	0.00Count/s
container_network_transmit_pac	Collected	Yes	cadvisor	0.00Count/s
✓ machine_cpu_cores	Collected	Yes	cadvisor	0.00Count/s
✓ machine_memory_bytes	Collected	Yes	cadvisor	0.00Count/s
cadvisor_version_info	Uncollected	No	eks-network	0.27Count/s
container_blkio_device_usage_t	Uncollected	No	eks-network	48.80Count/s
container_cpu_cfs_periods_total	Uncollected	No	eks-network	0.80Count/s

5. On the "Data Collection Configuration" page, click Customize Monitoring Configuration to add a new data collection configuration. TKE has preset some collection configuration files for collecting standard monitoring data.
You can configure new data collection rules to monitor your business data in the following two ways:

Adding Configuration in the Console

Adding Configuration via YAML File

#### **Monitoring Service**

#### 1. Click On this Page.

2. In the "Create Collection Configuration" pop-up window, fill in the configuration information. For details, see the following figure.

		×
Edit Mode	On this page Via YAML	
Monitoring Type	Service monitoring	
Name		
	It can contain up to 63 letters, digits, and hyphens (-). It must start with a letter and end with a digit or lowercase letter.	
Namespace	Please select	
Service	Please select 🗸	
servicePort	Please select 🗸	
metricsPath	/metrics	
	The default value is "/metrics". You can enter a custom value based your actual collection API.	
View Configuration File	Configuration File	
	If you have the requirement for special configurations, such as relabeling, please edit the configuration file.	
	OK Cancel	

#### 1. Click Via YAML.

2. In the pop-up window, select the monitoring type and fill in the corresponding configuration. You can complete the data collection configuration by submitting the corresponding YAML as per community usage.

Workload Monitoring: Corresponding configuration is PodMonitors.

**Service Monitoring**: Corresponding configuration is ServiceMonitors.

RawJobs Monitoring: Corresponding configuration is RawJobs.

6. Click **OK** to complete the configuration.

7. On the Data Collection Configuration page of the instance, view the status of the collection target. For details, see the following figure.

KUDEIC		110.93Count/!	s OCount/s	<mark>39.47</mark> Count/s	(1/1) up
		<mark>35.5</mark> 3Count/s	0Count/s	2.00Count/s	(1/1) up
Custom I Customize	Monitoring Monitoring Co	nfiguration			
Name	Туре	Metric Collection S	Billable Metric Colle ‡	Free metric collecti	Targets Template
	ServiceM	0Count/s	0/s	0.00Count/s	(0/0) No collection objects
	RawJobs	305.07Count/s	0/s	38.00Count/s	(1/1) up temp-cd7

**targets(1/1)** means (the actual number of targets scraped is 1 / the number of detected collection targets is 1). When the actual number of scraped targets equals the number of detected targets, it is displayed as up, indicating that the current scrape is normal. When the actual number of scraped targets is less than the number of detected targets, it is displayed as down, indicating that some endpoints failed to scrape.

Click the field value (1/1) in the above figure to view more details about the collection target. For details, see the following figure.

Please provide the sea	arch content	Q		
Endpoint	St T Labels		Metadata Last Scrape Time	Last Scrape I
p	Healthy ,		Expand details 2024-07-19 17:21:41	0.134

### **View Existing Configurations**

1. Log in to TMP Console.

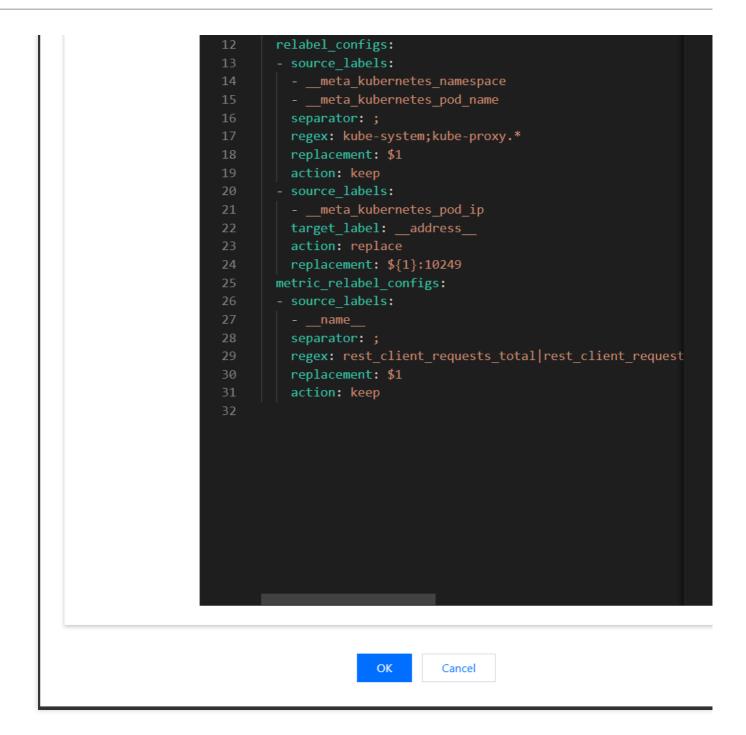
2. On the instance list page, select the name of the instance that requires configuring data collection rules and enter its details page.

3. On the **Data Collection > Integration TKE** page, click **Data Collection Configuration** on the right side of the instance to enter the collection configuration list page. Select **Basic Monitoring** or **Custom Monitoring**, then click **Edit** on the right.

4. In the pop-up edit window, view all monitoring objects currently configured in the YAML file. For details, see the following figure.

dit RawJobs		
Name	kube-proxy	
Configuration	1scrape_configs:2- job_name: kube-proxy3honor_timestamps: true4metrics_path: /metrics5kubernetes_sd_configs:6- role: pod7namespaces:8names:9- kube-system10tls_config:11insecure_skip_verify: true	





### **View Collection Targets**

1. Log in to TMP Console.

2. On the monitoring instance list page, select the instance name whose targets you want to view and enter the details page of the instance.

3. On the **Data Collection > Integration with TKE** page, click **Data Collection Configuration** on the right side of the instance.

v <sup>1</sup> in (勿識)							
isic Info Data Collection Alarm	Management Recon	ding Rule instance	diagnostics				
tegrate with TKE Integration Center	Data Multi-Write						
Associate Cluster Disassociate							
Cluster ID/Name	Agent Status	Region T	Cluster Type	Metric Collection Speed Before Filter	Billable Metric Collection Speed () \$	Free metric collection rate	Global Label(
m-70	Running	Guangzhou	Standard cluster	1010.33/s	46.8/s	164.53/s	cluster_type:

4. In the pop-up window, click the status under targets to jump to the details page of the current data.

Basic Monitoring				
Metric collection management	Metric Collection S	Billable Metric Colle 🛟	Free metric collecti	Targets
	148.27Count/s	41.93Count/s	26.20Count/s	(1/1) up
	105.47Count/s	4.87Count/s	20.86Count/s	(1/1) up
	305.07Count/s	0Count/s	38.00Count/s	(1/1) up
	<mark>0</mark> Count/s	<mark>0</mark> Count/s	<mark>0.00</mark> Count/s	(i No collec

ease provide the sea	arch content	Q		
Endpoint	St 了	Labels	Metadata Last Scrape Time	Last Scraj
۵ ۱	Healthy	n 0	Expand details 2024-07-19 17:26:06	0.004
otal items: 1			<b>10 ∨</b> / page	I

## **Related Operations**

### Mounting Files to a Collector

When configuring the collection item, if you need to provide some files for the configuration, such as a certificate, you can mount the files to the collector in the following way. The file update will be synchronized to the collector in real time.

#### prometheus.tke.tencent.cloud.com/scrape-mount = "true"

Add the above label to the configmap under the prom-xxx namespace. All keys will be mounted to the collector's path

/etc/prometheus/configmaps/[configmap-name]/ .

#### prometheus.tke.tencent.cloud.com/scrape-mount = "true"

Add the above label to the secret under the prom-xxx namespace. All keys will be mounted to the collector's path

/etc/prometheus/secrets/[secret-name]/ .

# **Configuring Necessary Monitoring Metrics**

Last updated : 2024-08-08 10:00:44

This document describes how to configure necessary collection metrics of TMP to avoid unnecessary expenses.

## Prerequisites

Before configuring monitoring collection items, you need to perform the following operations:

Create a Prometheus monitoring instance.

Associate the cluster to be monitored with the instance.

### **Configuring Necessary Metrics**

### **Configuring Necessary Metrics via Console**

TMP provides over 100 free basic monitoring metrics. For the complete metric list, see Free Metrics in Pay-as-You-Go Mode.

1. Log in to the TCOP console, and click Prometheus Monitoring in the left sidebar.

2. On the instance list page, select the name of the instance that requires configuring data collection rules and enter its details page.

3. On the **Data Collection** page, click **Integration with TKE**, find the target cluster, and click **Data Collection Configuration** in the operation bar, as shown below:

asic Info Data Collection Alarm Management Recording Rule							
Associate Cluster Disassociate	er Data Multi-Write Ag	ent Management					
Cluster ID/Name	Agent Status	Region T	Cluster Type 🝸	Metric Collection Speed Before Filter	Billable Metric Collection Speed () 💲	Free metric collection rate	
cls-nt2xlv8e I2 身 lucy-prom-勿制	Running	Guangzhou	Standard cluster	677.8/s	42.8/s	120.73/s	
Total items: 1							

4. In the data collection configuration window, click **Metric Details** in the operation bar to deselect unnecessary metrics, as shown below:

	(		ata Collection Configurat	ion
	Basic Monitoring			
	Metric collection management			
	Instance Type	Metric Collectio	n S Billable Metric Colle	‡ Free m
Free metric collection				
120.73/s	·system/node-exporter	98Count/s	4.47Count/s	19.40Count/s
	-system/kube-state-metrics	149.33Count/s	38.33Count/s	24.40Count/s

5. On the page below, you can view free and billed metrics. If a metric is selected, corresponding data will be collected. You can select metrics as needed. Only basic monitoring provides free metrics. For all free metrics, see Free Metrics in Pay-as-You-Go Mode. For billing of paid metrics, see TMP Pay-as-You-Go.

Basic Monitoring/kube-system/node-exporter			
Filter commonly used monitoring metrics with one click ba	sed on expertise	e analysis. For more informati	on, see Documentatic
— Metric Name	Free 🝸	Real-Time Coll 🍸	Metric Collection Before Filter(j)
go_gc_duration_seconds	No	Uncollected	0.33Count/s
go_gc_duration_seconds_count	No	Uncollected	0.07Count/s
go_gc_duration_seconds_sum	No	Uncollected	0.07Count/s
go_goroutines	No	Uncollected	0.07Count/s
✓ go_info	No	Uncollected	0.07Count/s
go_memstats_alloc_bytes	No	Uncollected	0.07Count/s
go_memstats_alloc_bytes_total	No	Uncollected	0.07Count/s
go_memstats_buck_hash_sys_bytes	No	Uncollected	0.07Count/s
		OK Cancel	

### **Configuring Necessary Metrics via YAML File**

The billing of TMP currently is based on the number of monitored data points. To minimize unnecessary expenses, it is recommended that you optimize the collection configuration to collect data of only necessary metrics, thereby reducing the overall data amount. For billing details and the usage of related cloud resources, see Resource Usage and Billing Overview.

The following part describes how to add filters to ServiceMonitor, PodMonitor, and native Job to configure necessary metrics.

1. Log in to the TCOP console, and click Prometheus Monitoring in the left sidebar.

2. On the instance list page, select the name of the instance that requires configuring data collection rules and enter its details page.

3. On the **Data Collection** page, click **Integration with TKE**, find the target cluster, and click **Data Collection Configuration** in the operation bar.

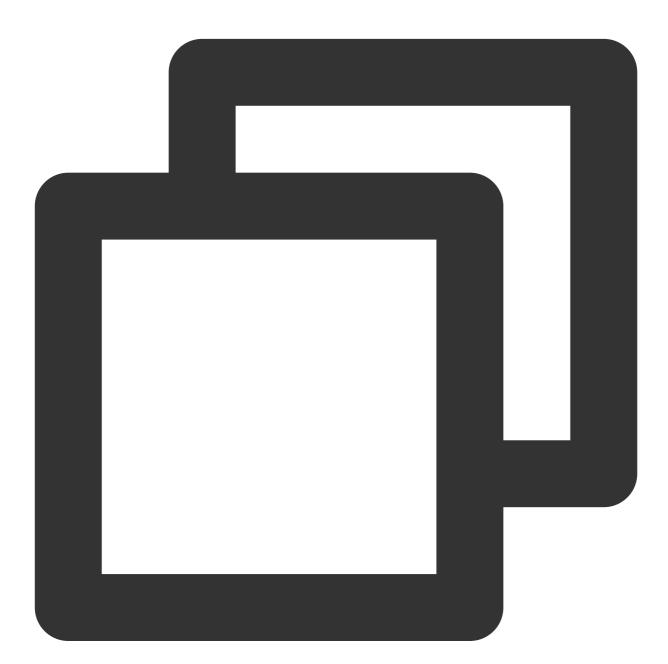


4. Find the target instance, and click **Editing** in the operation bar to view metric details.

ServiceMonitor and PodMonitor

Native Job

The filter configurations of ServiceMonitor and PodMonitor are the same. Take ServiceMonitor as an example. ServiceMonitor example:



apiVersion: monitoring.coreos.com/v1
kind: ServiceMonitor
metadata:
 labels:

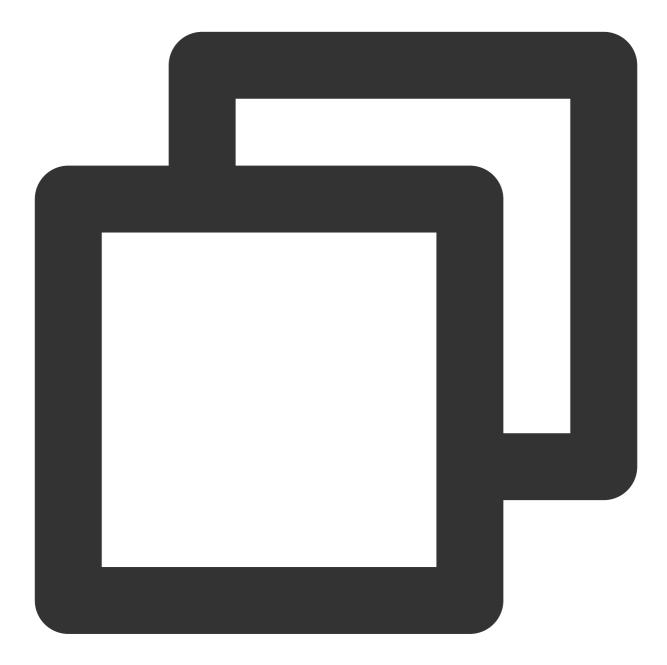
```
app.kubernetes.io/name: kube-state-metrics
    app.kubernetes.io/version: 1.9.7
 name: kube-state-metrics
 namespace: kube-system
spec:
 endpoints:
  - bearerTokenSecret:
      key: ""
    interval: 15s # This parameter indicates the collection interval. You can incre
   port: http-metrics
    scrapeTimeout: 15s # This parameter indicates the collection timeout. TMP requi
  jobLabel: app.kubernetes.io/name
 namespaceSelector: {}
  selector:
    matchLabels:
      app.kubernetes.io/name: kube-state-metrics
```

For example, if you want to collect data of metrics <a href="https://kube\_node\_info">kube\_node\_node\_role</a>, add the configuration with the <a href="https://metricRelabelings">metricRelabelings</a> field to the endpoints list of ServiceMonitor. Note that the field is

```
\ensuremath{\mathsf{metricRelabelings}} , not relabelings .
```

Example of adding metricRelabelings :





```
apiVersion: monitoring.coreos.com/v1
kind: ServiceMonitor
metadata:
   labels:
       app.kubernetes.io/name: kube-state-metrics
       app.kubernetes.io/version: 1.9.7
   name: kube-state-metrics
      namespace: kube-system
spec:
   endpoints:
      - bearerTokenSecret:
```

```
key: ""
interval: 15s # This parameter indicates the collection interval. You can incre
port: http-metrics
scrapeTimeout: 15s
# The following four lines are added:
metricRelabelings: # Each metric will undergo the following process.
- sourceLabels: ["__name__"] # The label name to be checked. __name__ refers to
regex: kube_node_info|kube_node_role # Check whether the label matches this r
action: keep # If the metric name meets the above conditions, the metric is
jobLabel: app.kubernetes.io/name
namespaceSelector: {}
```

If native Jobs of Prometheus are used, filter metrics using the method below. Job example:

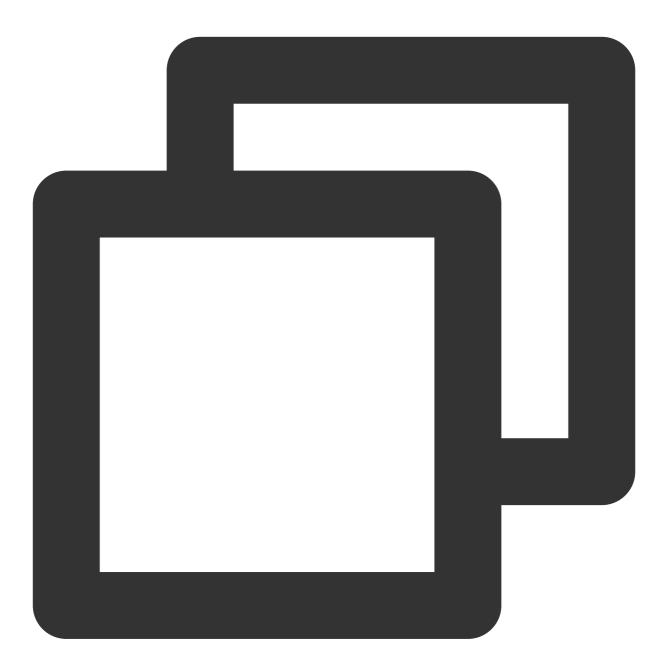




For example, if you want to collect data of metrics kube\_node\_info and kube\_node\_role only, add the configuration with the field metric\_relabel\_configs . Note: The field is metric\_relabel\_configs , not



```
relabel_configs .
Example of adding metric_relabel_configs :
```



- source\_labels: ["\_\_name\_\_"] # The label name to be checked. \_\_name\_\_ refers t regex: kube\_node\_info|kube\_node\_role # Check whether the label matches this r action: keep # If the metric name meets the above conditions, the metric is

5. Click **Confirm**.

#### **Skipping Data Collection**

#### **Skipping Data Collection from Entire Namespace**

TMP will manage all ServiceMonitors and PodMonitors in a cluster by default after the cluster is associated. If you want to skip data collection from a namespace, you can add the label tps-skip-monitor: "true" to the specified namespace. For label details, see Labels and Selectors.

#### **Skipping Data Collection**

TMP collects monitoring data by creating ServiceMonitor and PodMonitor types of CRD resources in a cluster. If you want to skip data collection from specified ServiceMonitor and PodMonitor resources, you can add the label tps-skip-monitor: "true" to these CRD resources. For label details, see Labels and Selectors.

# Adjusting Collection Interval

Last updated : 2024-08-08 10:02:32

This document describes how to adjust the **sampling interval** of TMP to avoid unnecessary expenses.

# Prerequisites

Before configuring monitoring collection items, you need to perform the following operations:

Create a Prometheus monitoring instance.

Associate the cluster to be monitored with the instance.

### Adjusting Collection Interval via TCOP Console

1. Log in to the TCOP console, and click Prometheus Monitoring in the left sidebar.

2. On the instance list page, select the name of the instance that requires configuring data collection rules and enter its details page.

3. On the **Data Collection** page, select **Integrate with TKE**, click **Data Collection Configuration** on the right side of the cluster, as shown in the following figure:

Basic Info Data Collection Alarm Management Recording Rule							
ntegrate with TKE	Integration Center Data Mult	ti-Write Agent Manag	ement				
Cluster ID/N	ame Age	nt Status	Region 🗑	Cluster Type 🗑	Metric Collection Speed Before Filter	Billable Metric Collection Speed () \$	Free metric collection rate
cls-nt2xlv8e lucy-prom-勿	C D Runs	ning	Guangzhou	Standard cluster	677.8/s	42.8/s	120.73/s

4. In the data collection configuration window, click **Edit** in the operation bar, as shown below:

			ita Collection Configurati	on
	Basic Monitoring			
	Metric collection management			
	Instance Type	Metric Collection	n S Billable Metric Colle	\$ Free met
Free metric collection				
120.73/s	-system/node-exporter	98Count/s	4.47Count/s	19.40Count/s
	-system/kube-state-metrics	149.33Count/s	38.33Count/s	24.40Count/s

5. Find the **interval** field, which represents the collection interval. You can increase its value to reduce data storage costs, as shown below:

lit ServiceMoni	tors
Name	kube-system/node-exporter
Configuration	1 apiVersion: monitoring.coreos.com/v1
configuration	2 kind: ServiceMonitor
	3 🗸 metadata:
	4 $\checkmark$ annotations:
	5 meta.helm.sh/release-name: tmp-scrape-component
	6 meta.helm.sh/release-namespace: kube-system
	7 creationTimestamp: "2024-07-31T03:17:52Z"
	8 generation: 2
	9 🗸 labels:
	10 app.kubernetes.io/managed-by: Helm
	11 app.kubernetes.io/name: node-exporter
	12 app.kubernetes.io/version: v0.18.1
	13 managedFields:
	14 V - apiVersion: monitoring.coreos.com/v1

15	ITETOSTYPE: FTETOSAT
16 🗸	fieldsV1:
17 🗸	f:metadata:
18 🗸	f:annotations:
19	
20	<pre>f:meta.helm.sh/release-name: {}</pre>
21	<pre>f:meta.helm.sh/release-namespace: {}</pre>
22 🗸	f:labels:
23	
24	<pre>f:app.kubernetes.io/managed-by: {}</pre>
25	<pre>f:app.kubernetes.io/name: {}</pre>
26	<pre>f:app.kubernetes.io/version: {}</pre>
27 🗸	f:spec:
28	.: {}
29	<pre>f:endpoints: {}</pre>
30	f:jobLabel: {}
31	<pre>f:namespaceSelector: {}</pre>
32	<pre>f:selector: {}</pre>
33	manager: controller
34	operation: Update
35	time: "2024-08-06T06:59:20Z"
36	name: node-exporter
37	namespace: kube-system
38	resourceVersion: "12692222780"
39	uid: feda3d44-ee96-4470-bbe1-67dd88571a6f
40 🗸 9	spec:
41	endpoints:
42 🗸	- bearerTokenSecret:
	OK Cancel

6. Click **Confirm** after completing the configuration.

### Adjusting Collection Interval via Integration Center

#### Note:

The Scraping Task application should have been installed and integrated. If it is not installed, click **One-click installation** on the right before you proceed with the following steps.

1. Log in to the TCOP console, and click Prometheus Monitoring in the left sidebar.

2. On the instance list page, select the name of the instance that requires configuring data collection rules and enter its details page.

3. On the **Data Collection** page, click **Integration Center**, find **Scrape Job**, and click **View Integrations** on the right, as shown below:

Easic Info	Alarm Management Recording Rule
Integrate with TKE Integratio	n Center Data Multi-Write Agent Management
	ethewa has integrated commonly used programming languages, middleware, big data products, and infrastructure databases, it allows you to install components through either the quick or custom installation mode and then monitor them. The integration center overs three major monitoring scenar d Monitor Application Check Middleware Big Data Database Alerting Others
Installed View All Integrations	
Application Name	Description
🔗 Cloud Monitor	TMP collects, stores, and visualizes the basic monitoring data of Tencent Cloud products.
🕒 Health Check	Blackbox can be used to regularly test the connectivity of the target service, helping you stay up to date with the service health and discover exceptions in time.
👙 Scrape Job	Scrape Job

4. On the page that appears, click any name to enter the collection tasks configuration page. Then, find the **scrape\_interval** field, which represents the collection interval. You can increase its value to reduce data storage costs, as shown below:

Edit		
Scrape Job		
Configuration *	job_name: prometheus scrape_interval: 30s static_configs: - targets: - 127.0.0.1:9090	
Estimated Collect	or Resource Occupation (): CPU-0.25 cores Memory-0.5 GiB	Billing Overview 🖸
Save	Cancel	

5. Click **Save** after completing the configuration.

# **Resource Usage and Billing Overview**

Last updated : 2024-08-07 22:07:52

When using the Tencent Managed Service for Prometheus (TMP) service, you may use resources such as **Tencent Kubernetes Engine (TKE) serverless clusters, TencentCloud Managed Service for Grafana (TCMG), and Cloud Load Balancer (CLB)**. This document describes the use cases and billing rules of these resources.

## **TKE Serverless Cluster**

### Use cases

You need to create a TKE serverless cluster if you use a TMP-associated cluster to monitor TKE.

A TKE serverless cluster will be automatically created for data collection when you install an integration plugin in the TMP integration center.

If both use cases are required for you, only one TKE serverless cluster will be created and shared. You can view the created clusters on the cluster list page.

Deploy Serverless container applications qu	uickly and enable in seconds, without the need to create K	8s clusters. <u>Container instance</u> 😢 is in beta te	st. You can get a 100 CNY voucher to join	it <u>Apply to join the beta</u> 😢		
ate						Separate filte
D/Name	Monitor	Kubernetes version	Type/State	Number of Pods	Resource volume	Operation
	-li	1.18.4	Elastic cluster(Running)	5	CPU:3.5-core MEM:5GIB	Configure alarm policy. View duster of
	.h	1.18.4	Elastic cluster(Running)	3	CPU:2-core MEM:2GIB	Configure alarm policy View cluster o
	-ti	1.18.4	Elastic cluster(Running)	11	CPU:9.25-core MEM:17.5GIB	Configure alarm policy. View cluster of
	-h	1.18.4	Elastic cluster(Idle)(j)	0	CPU:0-core MEM:0GIB	Configure alarm policy. View cluster o
	.h	1.18.4	Elastic cluster(idle)(	0	CPU:0-core MEM:0GIB	Configure alarm policy. View cluster c

### Note

The name of the TKE serverless cluster is the TMP **instance ID**, and the cluster description states that **For TMP use only. Do not modify or delete**.

Cluster name	prom- c 🖉
Cluster ID	cls-3sa8z1ko
Status	Running
K8s version	1.18.4
Deployment type	Elastic cluster
Region	Southeast Asia(Singapore)
Cluster network	
Container network	
Service CIDR block	
DNS Forward configuration	
Time created	2022-05-18 16:42:42
Tag	i -

### **Billing overview**

The billing mode is **pay-as-you-go**. For more information, see Product Pricing.

The TKE serverless cluster automatically scales according to the monitoring size. The relationship between the monitoring size and the TKE serverless cluster cost is shown below:

Reported Instantaneous Series	Estimated TKE Serverless Cluster Resources Required	List Price/Day
< 500,000	1.25 cores, 1.6 GiB	0.35 USD
1 million	0.5 cores, 1.5 GiB*2	1.46 USD
5 million	1 core, 3 GiB*3	2.93 USD
20 million	1 core, 6 GiB*5	7.98 USD
30 million	1 core, 6 GiB*8	12.77 USD

#### Sample TKE serverless cluster costs are as follows:

If the TKE serverless cluster used for a newly initialized TMP instance consumes 1.25 CPU cores and 1.5 GiB memory, then the estimated list price per day will be 0.0319 24 + 0.0132 24 = 1.0824 USD.

# TCMG

## Use cases

When creating a TMP instance, you need to associate it with a TCMG instance in the same region for the visual display of monitoring data collected by TMP. For billing information, see Billing Overview.

## CLB

When you use a TMP-associated cluster to monitor TKE, a private network CLB instance will be created under your account for network connectivity between the collector and the cluster.

If you associate an edge cluster or a cluster with no network connection, a public network CLB instance will be created for network connectivity.

To access the TCMG service over the public network, you need to create a public network CLB instance.

These CLB resources will be charged. You can view the resource information of the created public network CLB instances in the CLB console.

Resources are billed based on the actual usage. For billing details, see Network Pricing.

## **Resource Termination**

If you terminate TMP instances in the TMP console, all relevant resources will also be terminated. Tencent Cloud does not repossess TMP instances proactively. If you no longer use TMP, you need to delete the instances promptly to avoid extra charges. For instance termination directions, see Terminating Instance.

# Practical Tutorial Migration from Self-Built Prometheus

Last updated : 2024-08-07 22:08:03

## Overview

You can quickly migrate from your self-built Prometheus service to TMP.

## Directions

Prometheus itself supports remote write to an external storage; therefore, you can add a remote write configuration pointing to TMP in the configuration file of your self-built Prometheus. The steps are as follows:

1. Get the remote write address and token of TMP through the basic information of the instance as follows:

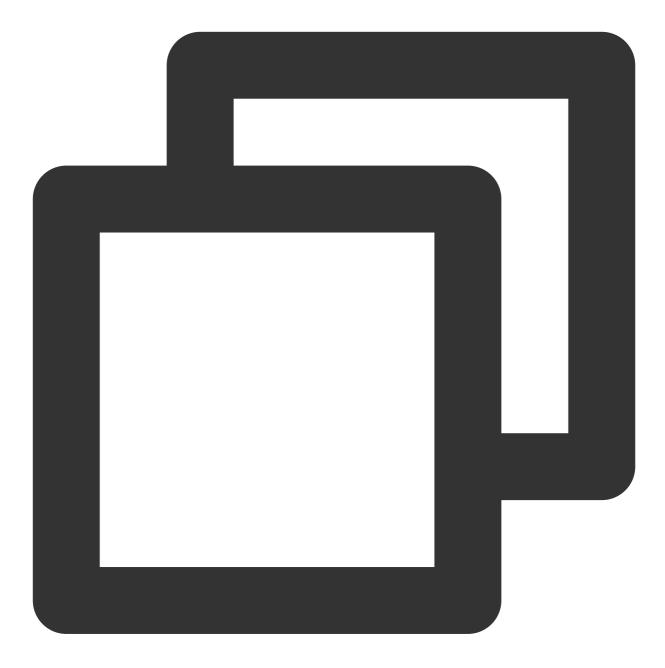
Ba	asic Info	
	Basic Info	
	Name	prom11 🧪
	Instance ID	L L L L L L L L L L L L L L L L L L L
	Status	
	Region	Guangzhou
	AZ	Guangzhou Zone 2
	Network	default_vpc
	Subnet	default_vpc_subnet
	Tag	

IPv4 Address		à							
Grafana Status		Disabled							
Billing Mode	Trial Edit	ion							
Creation Time	2021/11	/15 15:55:00	0						
Service Addr	'ess								
Service Addr		*****				1			
Service Addr Token		***** Г				]			
	1	***** <b>©</b> http		pi/v1/prom,	/write <b>Г</b>				
Token	ddress			pi/v1/prom, pi/v1 <b>T</b>	/write <b>Г</b>				

2. Modify prometheus.yml and restart Prometheus. The specific configuration is as follows. For more information on the remote write configuration parameters, please see remote\_write.

Circent Cloud





#### remote\_write:

- name: cm\_prometheus # Remote write name
url: http://ip:port/api/v1/prom/write # Get the remote write address from the
remote\_timeout: 30s # Set according to the actual situation
bearer\_token: k32\*\*\*\*trR # Get the token information from the basic informatio

3. Open the Grafana service that comes with TMP and use Explore to verify whether the data is written successfully as shown below. You can also customize Grafana monitoring dashboards.

	@ Explore	Prometheus-1	~							Split 🕘	Last 1 hour 🗸	Q
	Metrics ~ 0	order_service_	order_queue_	_size								Ste
	+ Add query	🕤 Query history										
	^ Graph											
]	1.3											
	1.2 1.1											
	1.0 0.9											
	0.8											
	0.7 12:	50 12:5	55 1:	3:00 1	3:05	13:10	13:15	13:	20 13	:25	13:30	13:
	- order_servic	e_order_queue_size{ii	nstance="localhos	t:2112", job="go_der	no*, type="ma	ke_order*}						
	^ Table											
										type		
	2021-05-13 13	46:40	andan aamilar	e_order_queue_si		ost:2112		go_demo		make_c	and an	

4. You can also use Prometheus APIs for self-built visualization. For more information, please see Monitoring Data Query.

# Custom Integration with CVM

Last updated : 2024-08-07 22:08:08

This document describes how to integrate CVM with TMP.

# Purchasing a TMP Instance

#### Note:

The purchased TMP instance must be in the same VPC as the monitored CVM instance for network connectivity. 1. Log in to the TMP console and click **Create** to purchase a TMP instance.

Cloud Monitor	Tencent Managed Service	for Prometheus	🛇 Singapore 🔻	
Hanitor Overview	Create Edit Tag			
🕒 Dashboard 🛛 👻		Marilanian (Status V	AZ T	Network
田 Instance Group	Instance ID/Name	Monitoring/Status <b>T</b>	AZ '	Network
Alarm Management	test	ll ⊘ Running	Singapore Zone 3	Network:D Subnet: <sub>rs</sub>
🛆 Alarm List				
△ Alarm ~ Configuration	Total items: 1			
() Trigger Condition Template				
Notification Template				
Cloud Native Monitor				
Managed Service for Prometheus				
Managed Service for Grafana				

2. On the purchase page, select the target instance specification and network. Make sure that the TMP and CVM instances have the same VPC IP range so that data can be collected. Select the instance specification based on your reported data volume.



Tencent I	Managed Ser	vice for Pro	metheus Return to product detail	s page	∃ Pro
Billing Mode	Pay-as-you-go				
Region and N	letwork Config				
Region	Asia Pacific	Europe and North A	America		
	Singapore	Tokyo			
	Tencent Cloud services in diffe it after purchasing the instance		cate with each other over the private network. For examp	le, the service in Guangzhou region cannot report o	ata to TMP in Shanghai region over the private ne
AZ	Singapore Zone 1	Sold of Singapore Zone 2 Sin	ngapore Zone 3 Singapore Zone 4		
Network	Select a VPC		∽ N/A		1
	If the existing VPC/subnet doe	as not meet your requirement, y	you can go to the console to create a VPC Z		
Basic Instance		45 days			
Instance Name	Please enter the instanc	e name			
Grafana	Please select a Grafana	a instance	~ ¢		
	If the existing Grafana instanc	e does not meet your requirem	rent, you can create one 🗹 in the console.		
Tag (optional)	Tag key	✓ Tag	g value 🗸 🗸	Delet e	
			+ Add		
	If the existing tag/tag value do	es not meet your requirement,	you can create one 🗹 in the console.		
Terms of Agreement	I've read and agree to	Tencent Cloud Terms of S	Service, Tencent Cloud Prometheus Service Leve	Agreement, Billing Overview, and Paymen	Overdue

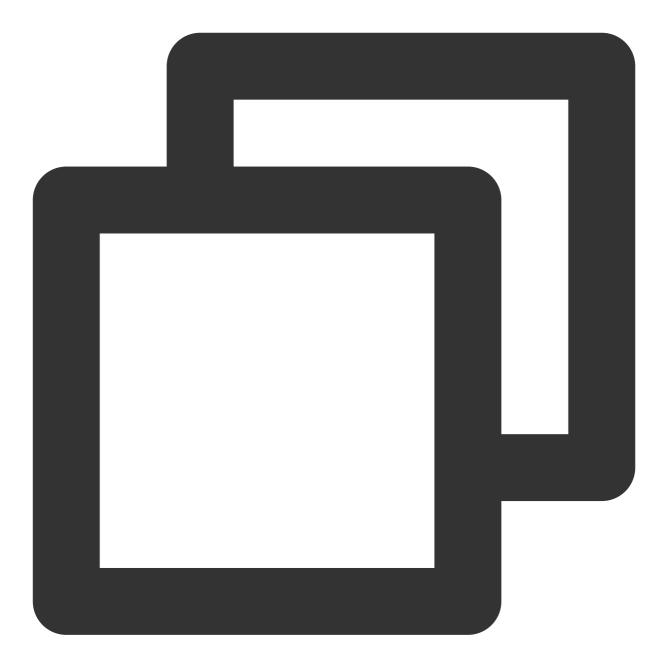
3. Click **Buy Now** and make the payment.

## Integrating CVM Basic Metrics

1. Download and install Node Exporter.

Download and install Node Exporter (used to collect basic metric data) in the target CVM instance. Click here or run the following command for download:





wget https://github.com/prometheus/node\_exporter/releases/download/v1.3.1/node\_expo

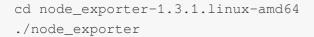
The file directory is as follows:

rw-rr	1	3434	3434	11357	Aug	6	2021	LICENSE
rwxr-xr-x	1	3434	3434	18494215	Aug	6	2021	node_exporter
- rw-rr	1	3434	3434	463	Aug	6	2021	NOTICE

- 2. Run Node Exporter to collect basic monitoring data.
- 2.1 Go to the target folder and run Node Exporter.





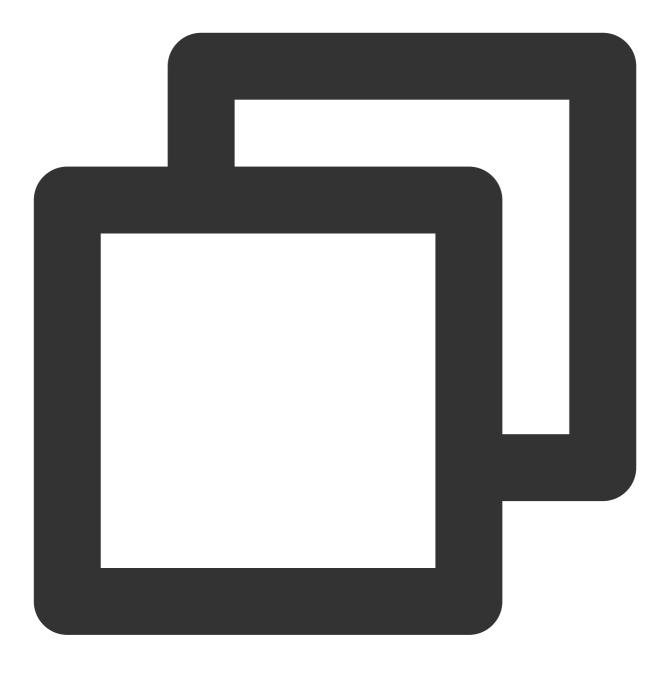


If the following result is displayed, basic monitoring data has been collected successfully.

rw-rr 1 3434 3434 463 Aug 6 2021 NOTICE
root@VM-0-7-centos node_exporter-1.2.2.linux-amd64]# ./node_exporter
.evel=info ts=2022-02-11T07:15:26.555Z caller=node_exporter.go:182 msg="Starting node_exporter" version="(version=1.2.2
n=26645363b486e12be40af7ce4fc91e731a33104e)"
.evel=info ts=2022-02-11T07:15:26.555Z caller=node_exporter.go:183 msg="Build context" build_context="(go=go1.16.7, use
late=20210806-13:44:18)"
.evel=warn ts=2022-02-11T07:15:26.555Z caller=node_exporter.go:185 msg="Node Exporter is running as root user. This exp
run as unpriviledged user, root is not required."
.evel=info ts=2022-02-11T07:15:26.555Z caller=filesystem_common.go:110 collector=filesystem msg="Parsed flagcollectc
nts-exclude" flag=^/(dev proc sys var/lib/docker/.+)(\$ /)
.evel=info ts=2022-02-11T07:15:26.555Z caller=filesystem_common.go:112 collector=filesystem msg="Parsed flagcollecto
exclude" flag=^(autofs binfmt_misc bpf cgroup2? configfs debugfs devpts devtmpfs fusectl hugetlbfs iso9660 mqueue nsfs
store rpc_pipefs securityfs selinuxfs squashfs sysfs tracefs)\$
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:108 msg="Enabled collectors"
evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=arp
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=bcache
evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=bonding
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=btrfs
evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=conntrack
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=cpu
evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=cpufreq
evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=diskstats
evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=edac
evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=entropy
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=fibrechannel
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=filefd
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=filesystem
evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=hwmon
.evel=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=infiniband

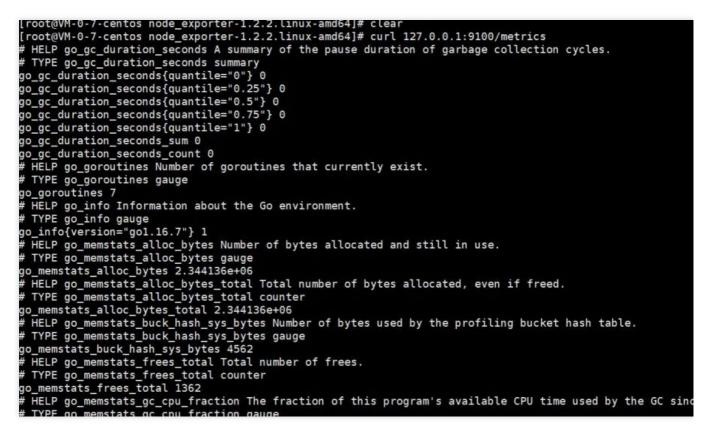
2.2 Run the following command to expose the basic monitoring data to port 9100:





curl 127.0.0.1:9100/metrics

You can see the following metric monitoring data that is exposed after the command is executed.

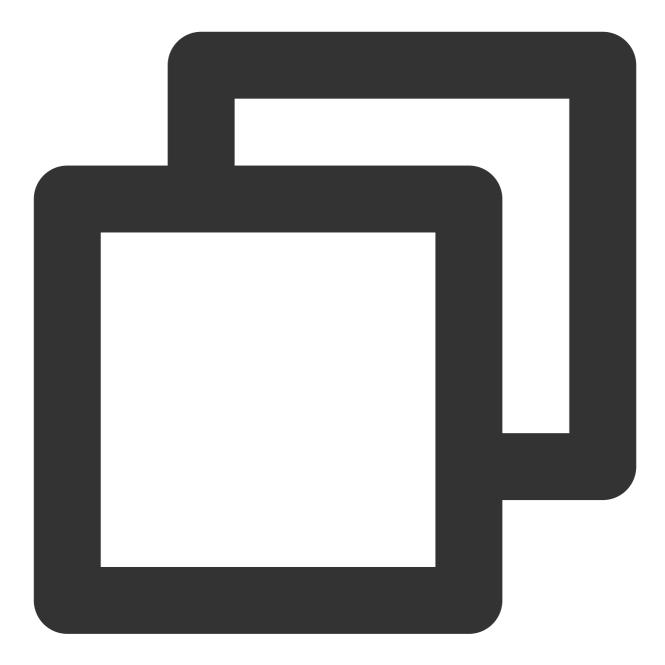


3. Add a scrape task.

Log in to the TMP console, select Integration Center > CVM, and configure the information in Task Configuration as prompted.

Below is a sample configuration of a scrape task:





```
job_name: example-job-name
metrics_path: /metrics
cvm_sd_configs:
- region: ap-guangzhou
ports:
- 9100
filters:
- name: tag:Sample tag key
values:
- Sample tag value
relabel_configs:
```

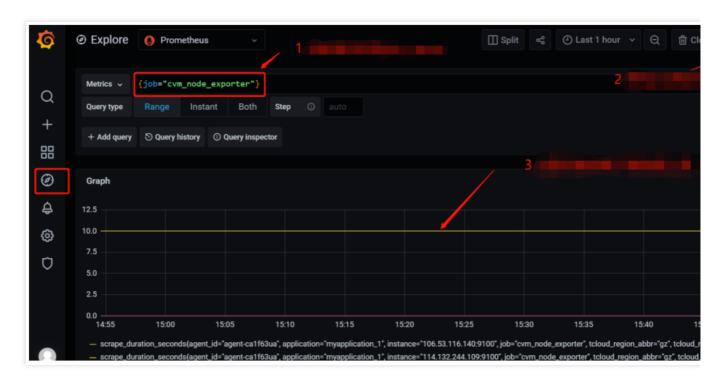


```
source_labels: [__meta_cvm_instance_state]
regex: RUNNING
action: keep
regex: __meta_cvm_tag_(.*)
replacement: $1
action: labelmap
source_labels: [__meta_cvm_region]
target_label: region
action: replace
```

4. Check whether data is reported successfully.

Log in to the TMP console and click the Grafana icon to enter Grafana.

Search for {job="cvm\_node\_exporter"} in **Explore** to see whether there is data, and if so, data is reported successfully.



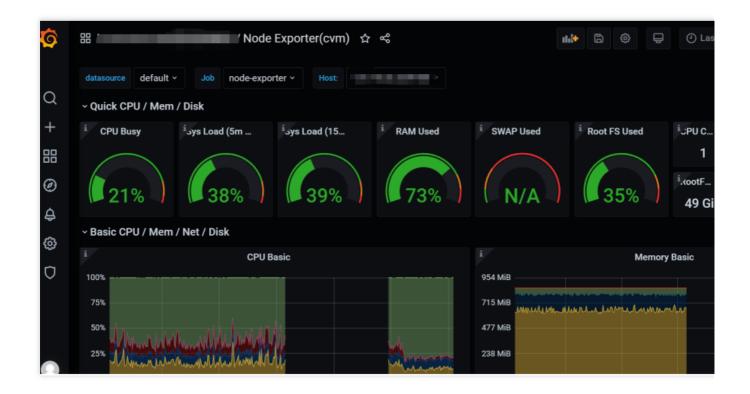
5. Configure the dashboard page: Every product has some existing JSON files that can be directly imported into the dashboard.

5.1 **Download a dashboard file**: Go to the **Dashboard** page, search for **node\_exporter**, and select the latest dashboard for download.

Node	e Exporter Full	v rfraile				
		, mane				
	dated: 3 days ago					
Start wit	th Grafana Cloud and the p	ew FRFE tier, Inclu	ides 10K series	Prometheus or Graphi	te Metrics and 50gb Loki Log	e.
Overview	evisions Reviews					
Overview R	evisions Reviews					
Overview R	levisions Reviews					
Overview R	levisions Reviews					Get this das
Overview R	levisions Reviews		-			Get this das
Overview R	levisions Reviews		-			1860
Overview R	levisions Reviews		-			
Overview R	levisions Reviews		-			1860
	evisions Reviews	s node exporter g	graphed.			1860

5.2 Import a JSON file into the dashboard: Log in to the TMP console, select Basic Info > Grafana Address to enter Grafana. In the Grafana console, select Create > Import and upload the dashboard file in Upload JSON file.

<b>Q</b>	Import dashboard from file or Grafana.com
Q	Options
+	Name Node Exporter Full
	Folder
	General
Ø	Unique identifier (uid)
¢	The unique identifier (uid) of a dashboard can be used for uniquely identify a dashboard between multiple Grafana installs. The uid allows having consistent URL's for accessing dashboards so changing the title of a dashboard will not break any bookmarked links to
ŝ	that dashboard.
$\Box$	rYdddlPWk1
	Prometheus
	Import Cancel



## Integrating CVM Metrics at the Business Layer

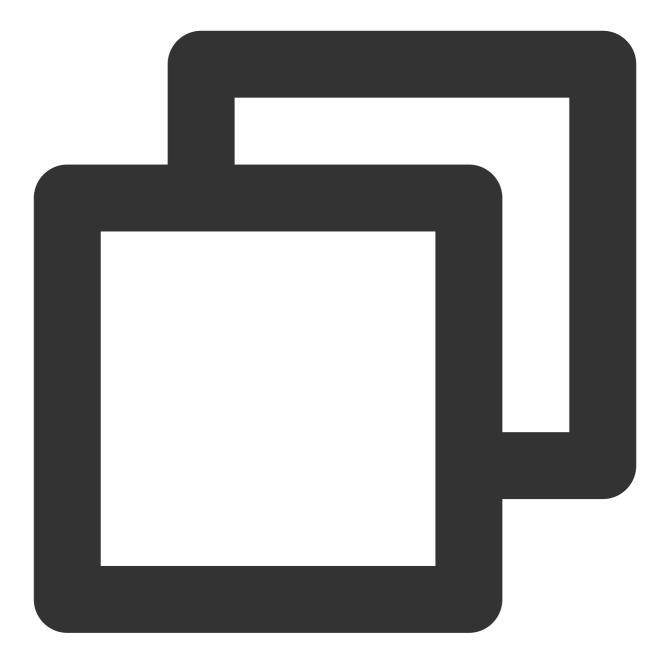
Prometheus provides four metric types for different monitoring scenarios: Counter, Gauge, Histogram, and Summary. The Prometheus community provides SDKs for multiple programing languages, which are basically similar in usage and mainly differ in the syntax. The following uses Go as an example to describe how to report custom monitoring metrics. For detailed directions of other metric types, see Custom Monitoring.

### Counter

A metric in Counter type increases monotonically and will be reset after service restart. You can use counters to monitor the numbers of requests, exceptions, user logins, orders, etc.

1. You can use a counter to monitor the number of orders as follows:



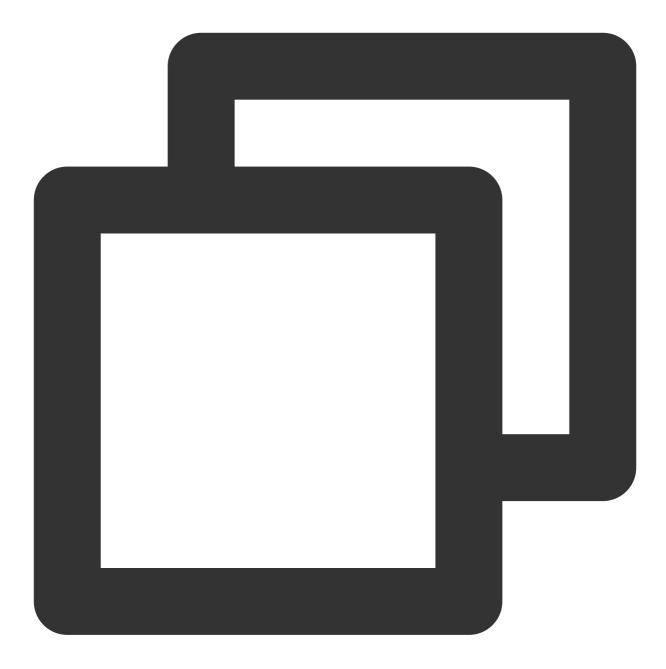


```
package order
import (
  "github.com/prometheus/client_golang/prometheus"
  "github.com/prometheus/client_golang/prometheus/promauto"
)
// Define the counter object to be monitored
var (
  opsProcessed = promauto.NewCounterVec(prometheus.CounterOpts{
     Name: "order_service_processed_orders_total",
```

```
Help: "The total number of processed orders",
}, []string{"status"}) // Processing status
)
// Process the order
func makeOrder() {
   opsProcessed.WithLabelValues("success").Inc() // Success
   // opsProcessed.WithLabelValues("fail").Inc() // Failure
   // Order placement business logic
}
```

For example, you can use the <code>rate()</code> function to get the order increase rate:



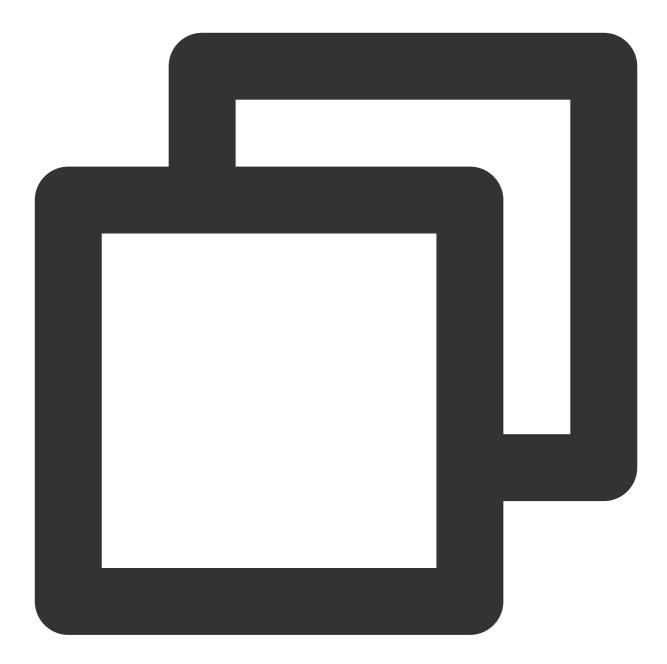


rate(order\_service\_processed\_orders\_total[5m])

#### 2. Expose Prometheus metrics:

Use promhttp.Handler() to expose the metric tracking data to the HTTP service.





```
package main
import (
  "net/http"
  "github.com/prometheus/client_golang/prometheus/promhttp"
)
func main() {
   // Business code
```

```
// Expose Prometheus metrics in the HTTP service
http.Handle("/metrics", promhttp.Handler())
// Business code
```

3. Collect data:

}

After the tracking of custom metrics for your business is completed and the application is released, you can use Prometheus to collect the monitoring metric data. After the collection is completed, wait a few minutes and then you can view the business metric monitoring data in Grafana integrated in TMP.

+ Add query       O Query history         - Graph         1.1       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -       -	м	etrics ~ order_s	ervice_order_	queue_size								Step	
1.3 1.2 1.1 1.0 0.9 0.8 0.7 12:50 12:55 13:00 13:05 13:10 13:15 13:20 13:25 13:30 13:35 13:40 - order_service_order_queue_size(instance='localhost:2112', job='go_demo', type='make_order')		Add query 🕤 Que	ry history										
1.3 1.2 1.1 1.0 0.9 0.8 0.7 12:50 12:55 13:00 13:05 13:10 13:15 13:20 13:25 13:30 13:35 13:40 - order_service_order_queue_size(instance='localhost:2112', job='go_demo', type='make_order') ^ Table													
1.2       1.1         1.1       1.1         1.0       1.1         0.9       1.1         0.8       1.1         0.7       12:50         12:55       13:00         13:05       13:15         13:15       13:20         13:25       13:30         13:35       13:40         - order_service_order_queue_size(instance="localhost:2112", job="go_demo", type="make_order")		Graph											
1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1.1       1	1.3												
1.0	1.2												
0.9 0.8 0.7 12:50 12:55 13:00 13:05 13:10 13:15 13:20 13:25 13:30 13:35 13:40 - order_service_order_queue_size(instance="localhost:2112", job="go_demo", type="make_order") - Table	1.1												
0.8 0.7 12:50 12:55 13:00 13:05 13:10 13:15 13:20 13:25 13:30 13:35 13:40 - order_service_order_queue_size(instance="localhost:2112", job="go_demo", type="make_order") - Table	1.0												
0.7 12:50 12:55 13:00 13:05 13:10 13:15 13:20 13:25 13:30 13:35 13:40 - order_service_order_queue_size{instance="localhost:2112", job="go_demo", type="make_order"} ^ Table	0.9												
12:50 12:55 13:00 13:05 13:10 13:15 13:20 13:25 13:30 13:35 13:40 — order_service_order_queue_size(instance="localhost:2112", job="go_demo", type="make_order") ^ Table													
^ Table	0.7	12:50	12:55	13:00	13:05	13	:10 13	15	13:20	13:25	13:30	13:35	13:40
∧ Table		<ul> <li>order_service_order_c</li> </ul>	ueue_size{instance='	'localhost:2112", job="go_	_demo", ty	ype="make_order	r"}						
Timename instance job type		Table											
Time instance job type													
											ре		

# **TKE Monitoring**

Last updated : 2024-08-07 22:08:14

# Background

As we all know, Prometheus is the best monitoring tool for container scenarios. However, self-building Prometheus is too expensive for small and medium-sized enterprises with limited Ops manpower, and it is likely to hit performance bottlenecks for large enterprises with rapid business development. Therefore, using Prometheus managed on the cloud has become the first choice for more and more cloud companies. For how to use the managed Prometheus to monitor TKE, see Tencent Kubernetes Engine (TKE).

## Directions

### Step 1. Purchasing an instance

1. Log in to the TMP console.

2. Click **Create**, select the purchase region, storage duration and select the Grafana instance to be associated based on your needs. If there is no Grafana instance, see <u>Creating Instance</u> to create one. You need to create an instance and complete the purchase.

3. After completing the configuration, click **Buy Now**. For more information on billing rules, see Pay-as-You-Go Description.

### Step 2. Integrate with TKE

1. After creating the instance, click the **ID**/**Name** of the target instance in the instance list to enter the instance details page.

2. On the left sidebar, click Integrate with TKE > Associate Cluster.

3. Select the cluster that needs to be associated in the pop-up window. A total of 4 types of clusters are supported: standard cluster, elastic cluster, registered cluster, and edge cluster. The clusters can be across VPCs. If different VPCs are not interconnected, you need to create a public network CLB instance.

Cluster type	General cluster 💌
Cross-VPC association	☑ Enable When it is enabled, you can monitor clusters under different VPCs in different regions in the same PROM instance.
	Create public CLB
	You must select "Create public CLB" if the VPC of your instance does not interconnect with the network of the desired cluster.
Region	· · · · · · · · · · · · · · · · · · ·
	Tencent Cloud resources in different regions cannot communicate via private network. The region cannot be changed after purchase. Please choose a region close to your end-users to minimize access latency and improve download speed.
Cluster	The following clusters are available for the current region.3/3 loaded1 item selected
	Separate filters with carriage return Q Node ID/N Type VPC Status
	Node ID/Na Type VPC Status
	Running
	Press and hold Shift key to select more
	Please reserve at least 0.5-core 100M for each cluster.
iobal label	Enable
	The key name can contain up to 63 characters. It supports letters, numbers, and "_", "_" cannot be placed at the beginning. A prefix is supported. Learn more 🗹 The label key value can only include letters, numbers and separators ("-", "_", ","). It must start and end with letters and numbers.

After associating the cluster, you can manually configure metrics for collection on Cluster Monitoring > Data
 Collection Configuration, view the default free basic collection metrics, and add or reduce the metrics as needed.

Basic monitoring Custom monitoring							
Instance type							
kube-system/kube-state-metrics							
cadvisor	0/sec				(1/1) up		
eks-network	Basic monitoring/kube-system/kube-state-metrics					×	
	Filter common monitoring metrics quickly. These metrics are metrics. For more information, see Metric description 🗳 .	expert recomme	endations provided by TMP	based on the analysis of user	Enter the metric nar	Q,	
	Metric name	Free o Y	Real-time coll 🗡	Metric collection rate before filtering ()	Metric collection rate *		
	kube_node_status_allocatable_memory_bytes	Yes	Collected	0.07/sec	0.07/sec	^	
	kube_pod_owner	Yes	Collected	1.6/sec	1.6/sec		
	kube_replicaset_owner	Yes	Collected	0.6/sec	0.6/sec		
	kube_validatingwebhookconfiguration_metadata_reso	No	Not collected	0.07/sec	0/sec		
	kube_job_owner	No	Not collected	0.27/sec	0/sec		
	kube_statefulset_status_update_revision	No	Not collected	0.07/sec	0/sec		
	kube_deployment_status_replicas_updated	No	Not collected	0.6/sec	0/sec		
	kube_hpa_spec_min_replicas	No	Not collected	0.07/sec	0/sec	*	
		Co	nfirm Cancel				

## Step 3. View monitoring data in Grafana

1. Click the Grafana icon to the right of the instance in the instance list to enter the Grafana service platform.

2. In the dashboard search list, TKE-related monitoring panels are preset by default, and click a panel name.

\$ Search dashboards by name	×
© Recent	
D tps	~
Kubemetes / API server(	kubernetes-mixin
Kubernetes / Compute Resources / Cluster	kubernetes-mixin
Kubernetes / Compute Resources / Namespace (Pode)	kubernetes-mixin
Kubernetes / Compute Resources / Namespace (Workloads)	kubernetes-mixin
Kubernetes / Compute Resources / Node (Pods)	kubernetes-mixin
Kubernetes / Compute Resources / Pod	kubernetes-mixin
Kubernetes / Compute Resources / Workload	kubernetes-mixin
Kubernetes / Controller Manager	kubernetes-mixin
Kubernetes / Kubelet	kubernetes-mixin
Kubernetes / Networking / Cluster	kubernetes-mixin

Enter the panel page, and you can view the preset monitoring data charts.

datasource v cluste node					
~ CPU Usage					
	CP	U Usage			
0.00500					
0.00500					
0.00300					
0.00200					
0 16:05 16:10 16:15 16:20	16:25 16:30	16:35 16:40	16:45 16:5	0 1655	17:00
16:05 16:10 16:15 16:20	16:25 16:30	16:15 16:40	16:45 16:5	0 16:55	17:00
~ CPU Quota					
	0	U Quota ~			
	CPU Usage	CPU Requests	CPU Requests %	CPU Limits	CPU Limits %
kvass-operator-585b8d67cd-wfkt2	0.00	0.50	0.10%		
proxy-agent-5b9f8485f7-xf2ww	0.00		0.16%		
tke-kube-state-metrics-0	0.00				
coredns-5b8b5c9954-wzglb	0.00		0.56%		0.56%
coredns-5b8b5c9954-n4bt2	0.00		0.58%		0.58%
~ Memory Usage					
	Memory Us	age (w/o cache)			
191 MB					

## Step 4. Configure the alert policy

On the Prometheus instance details page, click **Alerting Rule**, and you can select a preset template type without manual configuration. For alert notifications, you can select existing notification templates on the TCOP to quickly configure alerts.



÷	Alerting Rule / Create						User G
Basic Info	Rule Template	Please select a policy template					
instance	Kule lemplate	Please select a policy template					Ŷ
Monitoring	Rule Name *		۹				
Agent		MySQL	·				
anagement	PromQL-Based Rule •	Kubelet	*				
grate with TKE		Kubernetes Masters	*				
		Kubernetes Nodes	*				
egration Center		Kubernetes Resources	•				
ording Rule	Duration	Kubernetes Workload	*				
erting Rule	Alert Notification Cycle 🕢	HealthCheck	*				
		Redis	÷				
ert Manager	Alert Object *	CVM	•				
	Alert Message *	ClickHouse Hease enter the alert message	>				
	Labels	Key: Please enter	Value:	Please enter	Save		
	Annotations	Maria Lanca da	Value:		Save		
	Annotations	Key: Please enter	value.	Please enter	Save		
	Alert Notification •	Select Template Create 🖬					
		0 selected. 3 more can be selected					
		Notification Template Name				Included Operations	Operation
				Т	he notification template list is empty. You can sel	ect some by clicking "Select Template".	
	Save Cancel						

# Enabling Public Network Access for TKE Serverless Cluster

Last updated : 2024-08-07 22:08:22

## Overview

TMP is integrated with CM. When creating an integration, if you select COS, you need to enable public network access for the TKE Serverless cluster of the target CM exporter, as COS doesn't support private network access.

## Directions

1. Log in to the TMP console.

2. Click the target instance to enter the instance management page. Then, click **Integration Center** > **Integration List**.

3. Click Log in the Operation column of the line where Type is CM.

Integration Center				Sca	n code plus technical exc	change group 🖳	Integration Cent
Integration Center	Integration list						
Targets							
name		type	instance informa	Operating status	Acquisition rate	Targets	oper
cloud		Cloud monitoring	MySQL(CDB)	O deployed	15.98 /sec	(1/1)up	Indic delet
cloudcloud		Cloud monitoring	CLB(Public),NAT	O deployed	7.38 /sec	(1/1)up	Indic delet
kafka		Cloud monitoring	Cafka	O deployed	0.88 /sec	(1/1)up	Indic delet
test		Cloud monitoring	CVM,Ckafka,Mon	O deployed	17.27 /sec	(1/1)up	Indic delet

4. On the topbar, switch to the Pod management page. Click the instance name to enter the cluster page.

5. On the Basic Info page, click Container Network.

🕗 Tencent Cloud

<b>Basic information</b>	
Basic information	
Cluster name	bear 🎤
Cluster ID	
Status	Running
K8s version	1.20.6
Deployment type	Elastic cluster
Region	
Cluster network	vp
Container network	subne
Service CIDR block	192.168.0.0/20
DNS Forward configuration	Learn more 🛂
Time created	2022-05-24 10:36:35
Tag	<i>I</i> *
Description	N/A 🎤

6. On the topbar of the container network page, switch to the **Routing Policy** tab. Click the route table link ( rtbxxx in the list) to enter the route table page.



← Details o			
Basic information Routing r	ules ACL rules		
Routing rules			
Destination	hange route table Next hop type	Next hop	Notes
	LOCAL	Local Local	Delivered by defa

7. On the route table page, click **Create Routing Policy**.

**Destination**: Enter 0.0.0/0.

Next Hop Type: Select NAT Gateway.

Next Hop: Select the target gateway. If there is no gateway, create one as instructed in Getting Started.

Add a route				
Destination	Next hop type	Next hop	Notes	Operation
0.0.0/0	NAT gateway	▼ nat Create a NAT gateway	¥	0
+ New line				
		Create	ise	

8. Click Create. After the creation is successful, public network access is enabled for TKE Serverless.

# Connecting TMP to Local Grafana

Last updated : 2024-08-07 22:08:28

If you need to view the relevant data of TMP in the local Grafana system, you can use the HTTP API provided by TMP to do so. This document describes how to connect TMP data to local Grafana.

## Directions

### Step 1: Obtain the HTTP API provided by TMP

- 1. Log in to the TMP console.
- 2. Click the corresponding pay-as-you-go instance to enter the basic information page of TMP.

3. Get the HTTP API address in the service address module. If you need to improve the security of Grafana data read, you can obtain the authentication token of the TMP instance and fill it in as instructed in step 2.

Grafana Data Source Configuration Information				
HTTP URL	http://			
Basic auth user(APPID)	12. 62 🖬			
Basic auth password	***** 🖻			

### Step 2. Add data source to local Grafana

- 1. Log in to the local Grafana system with admin account.
- 2. Select **Configuration > Data Sources** on the left sidebar (non-admins cannot view this menu).
- 3. On the Data Sources page, click Add data source.
- 4. On the Add data source page, select Prometheus.

	d data source se a data source type	~	Data source deleted	×
Q Filter by na	ame or type			← Cancel
Time series of	databases			
	Prometheus Open source time series database & alerting Core			[͡ <sup>7]</sup> Learn more
<b>,</b>	Graphite Open source time series database Core			
$\bigcirc$	InfluxDB Open source time series database Core			
	OpenTSDB Open source time series database Core			

5. On the **Settings** tab, enter a custom name in the **Name** field, and paste the **HTTP API** obtained in step 1 in the URL field.

6. Toggle on **Basic Auth** in the **Auth** module. In the **Basic Auth Details** module, set **User** same as **Basic auth user** and **Password** as **Basic auth password** obtained in the step 1.

7. Click Save & test to complete the settings.

HTTP				
URL	0	http://10.ſ∎ <sup>*</sup> :9090		
Access		Server (default)		Help >
Allowed cookies	(	New tag (enter key to add)		
Timeout	<b>(</b> )	Timeout in seconds		
Auth				
Basic auth		With Credentials	0	
TLS Client Auth		With CA Cert	3	
Skip TLS Verify				
Forward OAuth Identity	(i)			
Basic Auth Details				
User	1251			
Password	config	gured	Reset	

Step 3. Verify whether the connection is successful



Follow the steps below to verify whether TMP is successfully connected to the local Grafana:

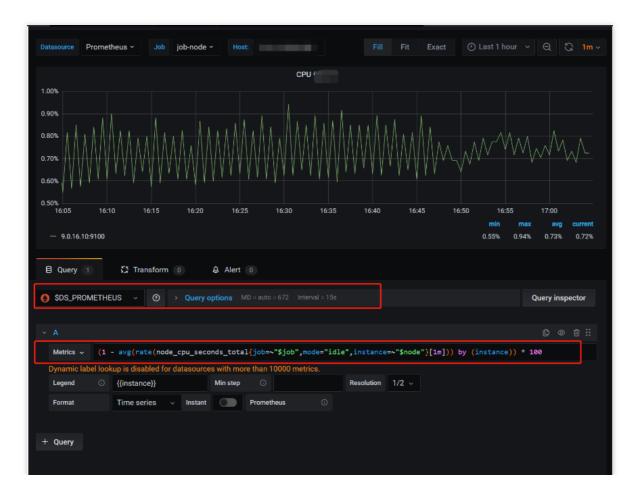
1. Log in to your local Grafana system.

2. On the left sidebar, select + > Create.

3. On the New dashboard page, click Add a new panel.

4. On the **Edit Panel** page, select the data source added in step 2 in the drop-down box on the **Query** page, enter the metric name in the **Metrics** field in the A module and press Enter.

5. If the chart of the corresponding metric can be displayed, the operation is successful. Otherwise, check whether the API address or token entered is correct, and whether the data source has TMP data.



# Terraform Terraform Overview

Last updated : 2023-12-29 16:06:50

Terraform is an Infrastructure as Code (IaC) tool that allows creation of reusable, maintainable infrastructure code. You can manage and update your infrastructure using codes, thereby simplifying the deployment and management of cloud infrastructure. With abundant resource types, powerful modules, and flexible customization options, Terraform makes it easy to create and manage a variety of resources, such as CVM and CLB.

### **Terraform Features and Advantages**

Terraform is a powerful open-source tool for automatically deploying and orchestrating cloud infrastructure.

Terraform enables you to operate, construct, define, deploy, and preview cloud infrastructure on Tencent Cloud using simple template language commands.

Terraform is remarkably user-friendly. It uses the simple HCL or JSON configuration language to define infrastructure, requiring virtually no code writing.

With variables, modules, and data source features, Terraform's configuration files are easy to use, reusable, and maintainable.

Terraform can provide plan reports at different stages, so users can get to know the actual operations executions. Terraform tracks the actual configuration and status information of infrastructure through state files and supports rollback operations.

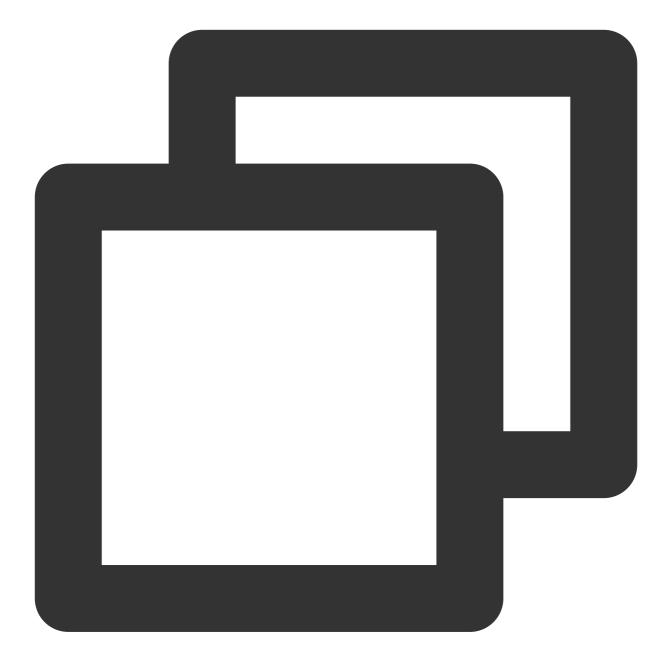
Note: For details about the application scenarios of Terraform, see Use Cases.

### **Terraform Resources**

In Terraform, resources mainly fall into two categories: Resource and Data Source.

#### Resource

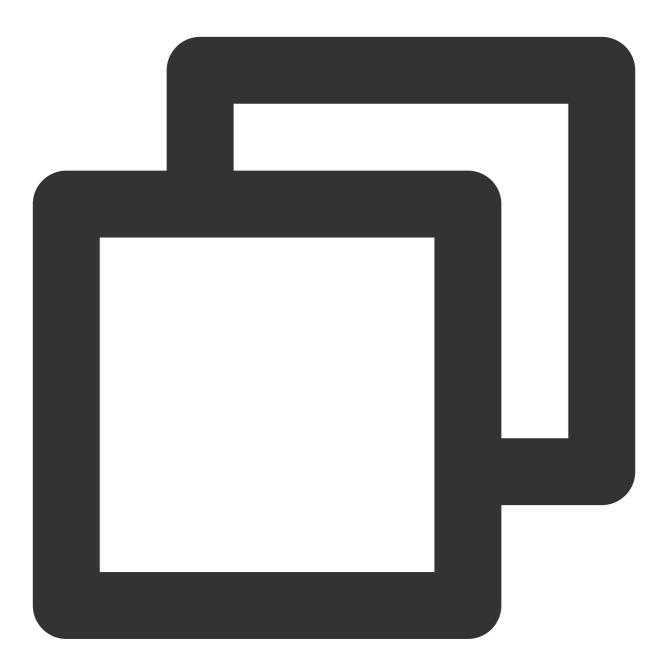
Resource enables the creation and management of new cloud infrastructure components. Through resource definitions, Terraform can be used for creating, modifying, and deleting a diverse range of resources supported by cloud service providers, such as creating a new CVM instance.



```
#Resource of cvm chc_assist_vpc.
# Reference document: https://registry.terraform.io/providers/tencentcloudstack/ten
resource "tencentcloud_cvm_chc_assist_vpc" "chc_assist_vpc" {
    chc_id = "chc-xxxxx"
}
```

### Data Source

A data source is used to access and retrieve information of already existing resources and subsequently implementing it as an input or attribute within the configuration. A data source can be used to retrieve the configuration and attribute information from an existing cloud infrastructure, making it available for reference and usage within a Terraform configuration file. The following example shows that it is used to query an instance of a created cvm chc\_denied\_actions:



```
#Using this data source to extract detailed information from cvm chc_denied_actions
# Reference document: https://registry.terraform.io/providers/tencentcloudstack/ten
data "tencentcloud_cvm_chc_denied_actions" "chc_denied_actions" {
    chc_ids = ["chc-xxxxx"]
}
```

For information about the Resource of the TCOP Prometheus version, see Usage with Terraform.

### **Terraform Tools**

#### Terraform Command Line Interface (CLI):

Terraform CLI is a command line interface designed to create, manage, and manage infrastructure. It brings a suite of commands and options, enabling users to perform various operations, including initializing configuration, creating plans, and applying changes. Terraform CLI also incorporates several sub-commands associated with additional Terraform functionality, such as configuration validation and code formatting. Employing Terraform CLI enables users to interact with Terraform configuration files and manage the infrastructure effectively.

#### **Terraform Provider:**

Each cloud provider provides its proprietary Terraform Provider, which is a plugin used to incorporate the features of the cloud provider into Terraform. Each Terraform Provider is tasked with interacting with the APIs of the corresponding cloud service provider, defining available resources and data sources, and performing operations of creating, modifying, or deleting these resources. By deploying the appropriate provider, you can directly use the resources and features of the cloud provider within Terraform configurations.

#### Note:

The Terraform CLI serves as the command line interface to manage Terraform, while Terraform Provider is the plugin incorporating the features of diverse cloud providers into Terraform. These two components work collaboratively, enabling users to use Terraform to create and manage infrastructure. For more information about Terraform, see the usage document of Terraform.

### **Usage Advantages**

Multi-Cloud Solution: Applicable to multi-cloud solutions, Terraform can deploy analogous infrastructure to Tencent Cloud and local data centers. Due to Terraform's portability, developers can manage resources from distinct cloud providers simultaneously using the same tools and configuration files. Such flexibility allows users to switch to other cloud platforms at any time without the need to modify their configuration files.

Automated Management: Terraform enables the creation of configuration file templates, defining, provisioning, and configuring resources in a repeatable and predictable manner, reducing deployment and management errors due to human factors. Using Terraform allows users to deploy the same template multiple times, creating identical development, test, and production environments. Additionally, Terraform automates the process of creating and managing infrastructure, thereby enhancing deployment speed and efficiency.

Infrastructure as Code: The Infrastructure as Code (IaC) method of Terraform allows for the management of resources through code, providing convenience in version control, collaboration, and code reuse. Using Terraform, you can save the state of your infrastructure to track system changes and share these configurations with others. This code-based management approach not only simplifies deployment but also ensures consistency, security, and maintainability. Visual Interface: Terraform's graphical interface allows for the viewing and management of infrastructure resources. This enables users to view and understand their infrastructures and interdependencies of infrastructures more directly. Through this visual interface, users can locate and rectify potential configuration problems faster and easier.

Community Support: Developed by HashiCorp, Terraform is backed by a large community. This community comprises professionals across various industries and experience levels capable of sharing best practices and solutions, offering supports like forums and email lists.

### Using TCOP Prometheus Managed Service via Terraform

#### Prometheus supports managing the following resources via Terraform:

Name	Description
tencentcloud_monitor_tmp_instance	Prometheus Instances
tencentcloud_monitor_tmp_alert_rule	Prometheus Alarm Rules
tencentcloud_monitor_tmp_recording_rule	Prometheus Recording Rules
tencentcloud_monitor_tmp_cvm_agent	Prometheus Cvm Agent
tencentcloud_monitor_tmp_scrape_job	Prometheus cvm_agent Capture Tasks
tencentcloud_monitor_tmp_exporter_integration	Prometheus Integration Center Resources
tencentcloud_monitor_tmp_manage_grafana_attachment	Prometheus Association with Grafana
tencentcloud_monitor_tmp_tke_cluster_agent	Prometheus Association with Containers
tencentcloud_monitor_tmp_tke_config	Prometheus Cluster Collection Configuration

# Managing Prometheus Instances Using Terraform

Last updated : 2023-12-29 16:08:36

### Prerequisites

### **Installing Terraform**

For detailed instructions on installing Terraform, see Installing Terraform.

### Note:

The installed version of Terraform must not be below v1.6.3. You can verify the version of Terraform installed by using the command terraform --version.

### **Configuring Tencent Cloud Account Information**

Before using Terraform for the first time, go to API Key to apply for a SecretID and SecretKey of the security credentials. If you already have valid credentials, skip this step.

1. Log in to the CAM Console and choose Access Key > API Keys from the left sidebar.

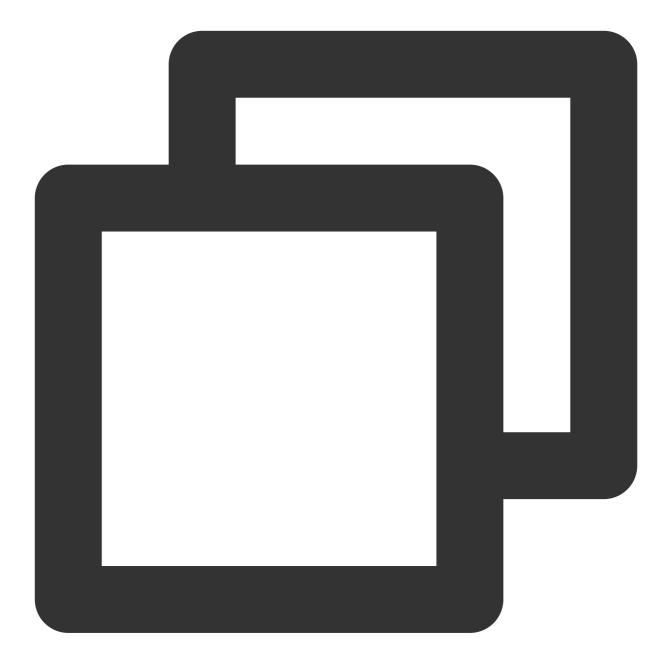
2. On the API Keys page, click Create Key to create a SecretId/SecretKey pair.

### There are two methods for configuring Tencent Cloud account information:

Authentication by Static Credentials

Create a provider.tf file in the user directory with the following content. Replace my-secret-id and mysecret-key with your SecretId and SecretKey, respectively.



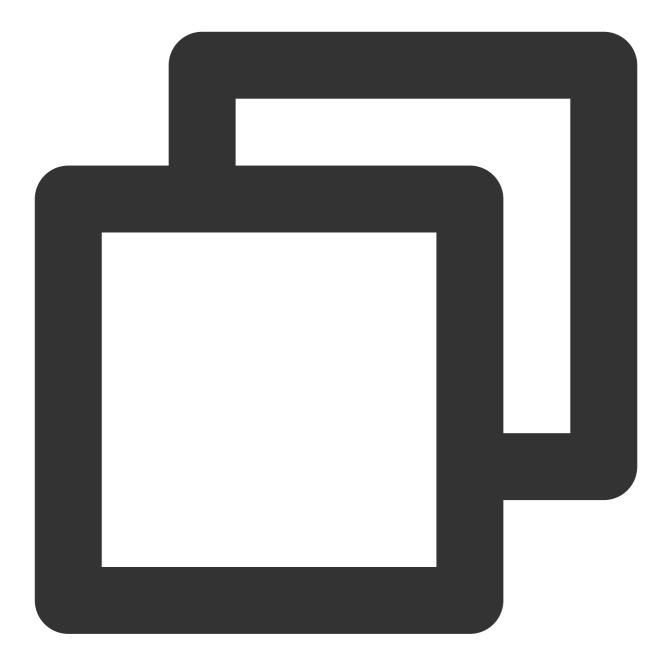


```
provider "tencentcloud" {
   secret_id = "my-secret-id"
   secret_key = "my-secret-key"
}
```

### Authentication by Environment Variables

Configure the computer environment variables or cloud environment variables by running the following command. Replace YOUR\_SECRET\_ID and YOUR\_SECRET\_KEY with your SecretId and SecretKey, respectively.



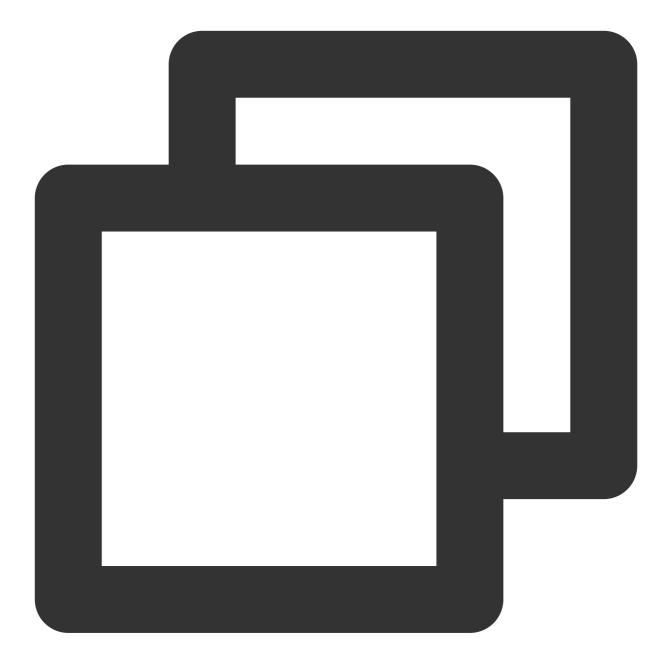


export TENCENTCLOUD\_SECRET\_ID=YOUR\_SECRET\_ID
export TENCENTCLOUD\_SECRET\_KEY=YOUR\_SECRET\_KEY

### **Creating Prometheus Instances**

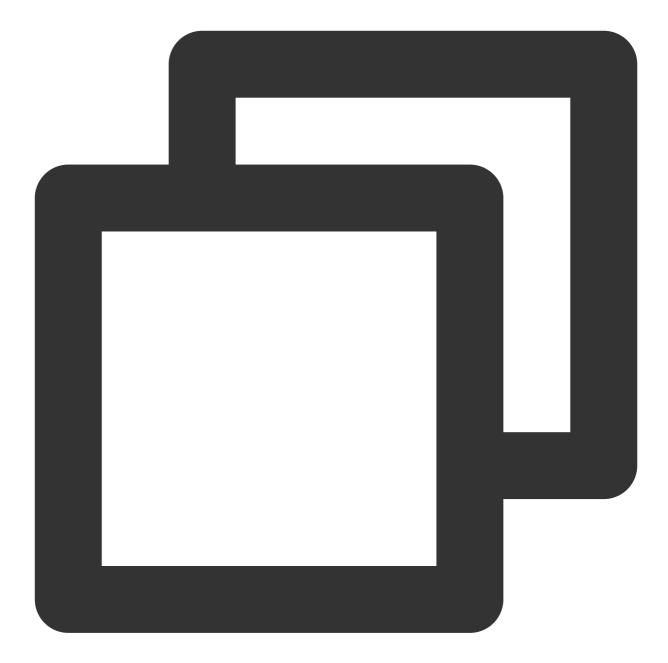
1. Create the configuration file for Terraform. Create a file named main.tf in the project directory and open it with an appropriate editor. Add the Tencent Cloud provider configuration to the main.tf or provider.tf file.





```
provider "tencentcloud" {
  region = "your_region"
  secret_id = "your_secret_id"
  secret_key = "your_secret_key"
}
```

2. Defining Tencent Cloud Resources: Use the resources provided by Tencent Cloud in the main.tf file to define the resources you want to create and manage.



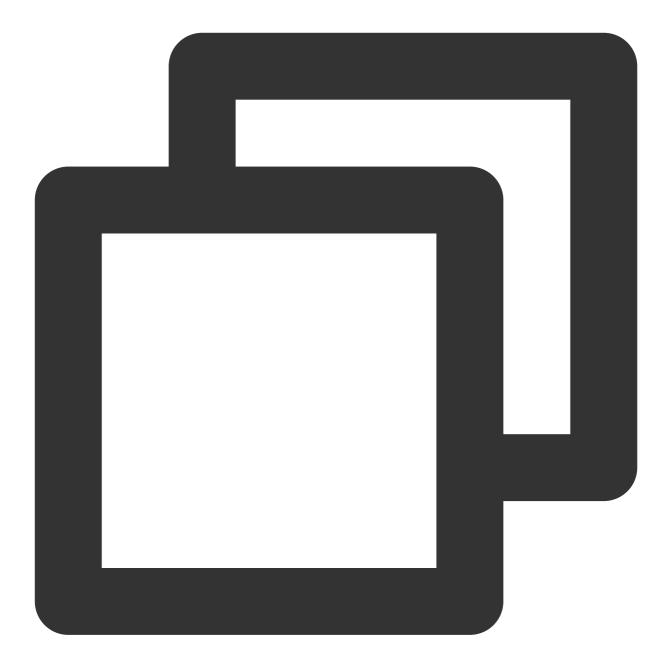
```
#Specifying Provider Configuration Information
terraform {
  required_providers {
    tencentcloud = {
        source = "tencentcloudstack/tencentcloud"
     }
   }
#Creating a Prometheus Instance in the Mumbai Region
```

```
🔗 Tencent Cloud
```

```
resource "tencentcloud_monitor_tmp_instance" "foo" {
 instance_name = "tf-tmp-instance-sjtest"
                   = "vpc-0n42dxzs"
 vpc_id
 subnet_id = "subnet-es8rv1kx"
 data_retention_time = 30
                    = "ap-mumbai-1"
 zone
 tags = {
   "createdBy" = "terraform"
 }
}
#Creating a Grafana Instance in the Mumbai Region (Optional)
resource "tencentcloud_monitor_grafana_instance" "foo" {
 instance_name = "tf-grfana-cstest"
 vpc_id
                      = "vpc-0n42dxzs"
 subnet_ids
               = ["subnet-es8rv1kx"]
 grafana_init_password = "1234567890"
 enable_internet
                     = false
 is_destroy
                      = true
 tags = \{
   "createdBy" = "test"
  }
}
#Binding Grafana and Prometheus Instances (Optional)
resource "tencentcloud_monitor_tmp_manage_grafana_attachment" "foo" {
 grafana_id = tencentcloud_monitor_grafana_instance.foo.id
 instance_id = tencentcloud_monitor_tmp_instance.foo.id
}
```

3. Initialize the Terraform runtime environment by running the following command:

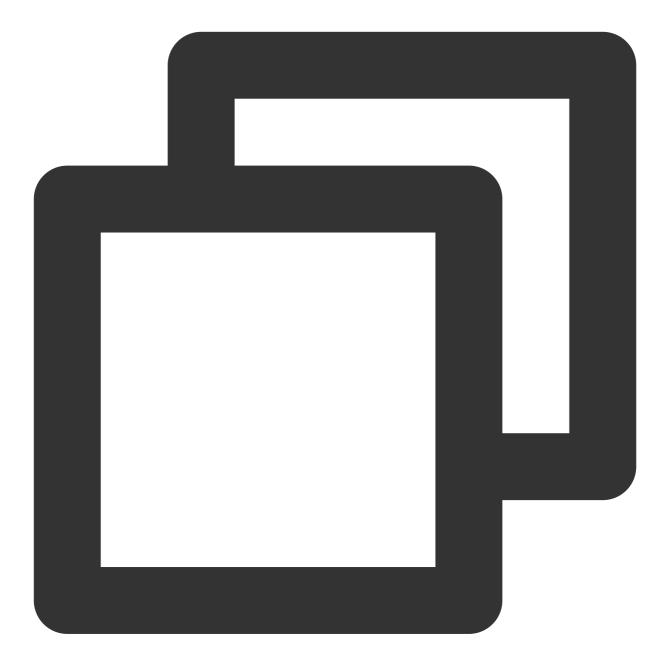




terraform init

Expected output information:





Initializing the backend...

Initializing provider plugins...

- Reusing previous version of tencentcloudstack/tencentcloud from the dependency lo
- Using previously-installed tencentcloudstack/tencentcloud v1.81.32

Terraform has been successfully initialized!

You may now begin working with Terraform. Try running "terraform plan" to see any changes that are required for your infrastructure. All Terraform commands should now work.



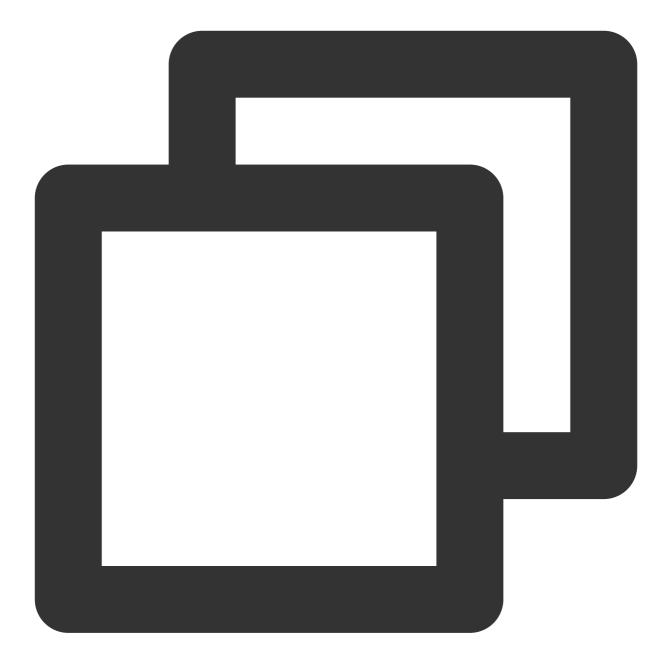
If you ever set or change modules or backend configuration for Terraform, rerun this command to reinitialize your working directory. If you forget, other commands will detect it and remind you to do so if necessary.

4. To generate resource planning, run the following command:



terraform plan

Expected output information:



tencentcloud\_vpc.vpc: Refreshing state... [id=vpc-3csjp7k8]
tencentcloud\_monitor\_tmp\_instance.foo: Refreshing state... [id=prom-4wcvt7p1]

Terraform used the selected providers to generate the following execution plan. Res
following symbols:
 + create

Terraform will perform the following actions:

```
# tencentcloud_monitor_tmp_instance.foo will be created
```

```
+ resource "tencentcloud_monitor_tmp_instance" "foo" {
```

```
+ api_root_path = (known after apply)
+ data_retention_time = 30
+ id = (known after apply)
+ instance_name = "tf-tmp-instance-sjtest"
+ ipv4_address = (known after apply)
+ proxy_address = (known after apply)
+ remote_write = (known after apply)
+ subnet_id = "subnet-es8rv1kx"
+ tags = {
    + "createdBy" = "terraform"
    }
+ vpc_id = "vpc-0n42dxzs"
+ zone = "ap-mumbai-1"
}
Plan: 1 to add, 0 to change, 0 to destroy.
```

5. To create an instance, run the following command:





terraform apply

Expected output information:



tencentcloud\_vpc.vpc: Refreshing state... [id=vpc-3csjp7k8]
tencentcloud\_monitor\_tmp\_instance.foo: Refreshing state... [id=prom-4wcvt7p1]

Terraform used the selected providers to generate the following execution plan. Res
following symbols:
 + create

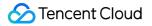
Terraform will perform the following actions:

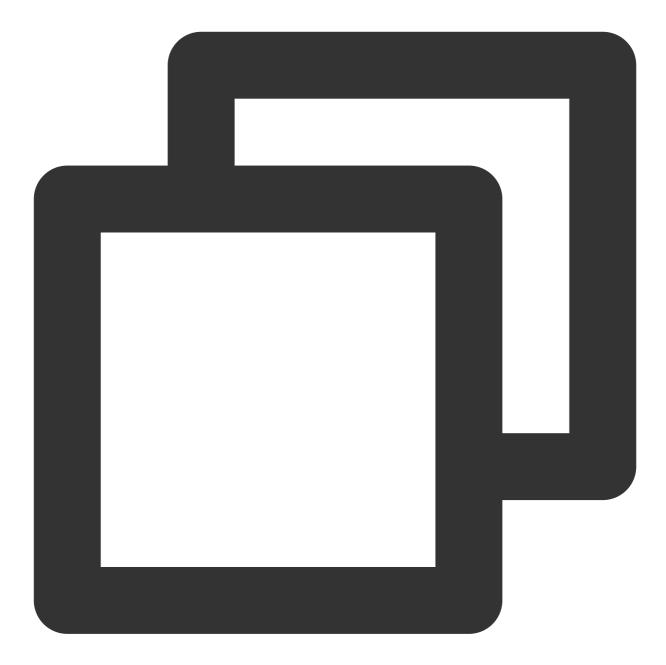
```
# tencentcloud_monitor_tmp_instance.foo will be created
```

```
+ resource "tencentcloud_monitor_tmp_instance" "foo" {
```

```
+ api_root_path = (known after apply)
+ data_retention_time = 30
+ id = (known after apply)
+ instance_name = "tf-tmp-instance-sjtest"
+ ipv4_address = (known after apply)
+ proxy_address = (known after apply)
+ remote_write = (known after apply)
+ subnet_id = "subnet-es8rv1kx"
+ tags = {
    + "createdBy" = "terraform"
    }
+ vpc_id = "vpc-0n42dxzs"
+ zone = "ap-mumbai-1"
}
Plan: 1 to add, 0 to change, 0 to destroy.
```

Enter "yes" to the following information to proceed:



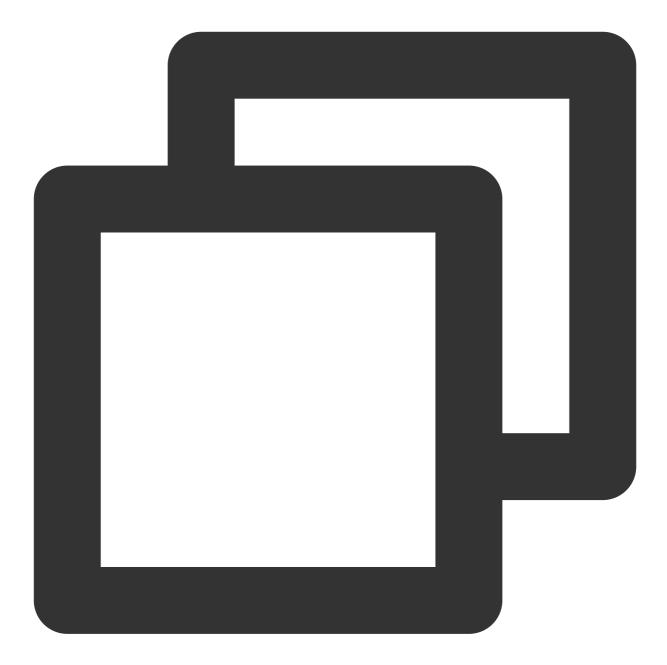


Do you want to perform these actions? Terraform will perform the actions described above. Only 'yes' will be accepted to approve.

Enter a value:

If the following information appear, the creation is successful:





tencentcloud\_monitor\_tmp\_instance.foo: Creating...
tencentcloud\_monitor\_tmp\_instance.foo: Still creating... [10s elapsed]
tencentcloud\_monitor\_tmp\_instance.foo: Creation complete after 12s [id=prom-8dyb6in

Apply complete! Resources: 1 added, 0 changed, 0 destroyed.

### Viewing the Status of Prometheus Instances

Log in to the TCOP, and choose **Managed Service for Prometheus** from the left-hand navigation bar. Existing instances are displayed in the Prometheus instance list.

### **Deleting Prometheus Instances**

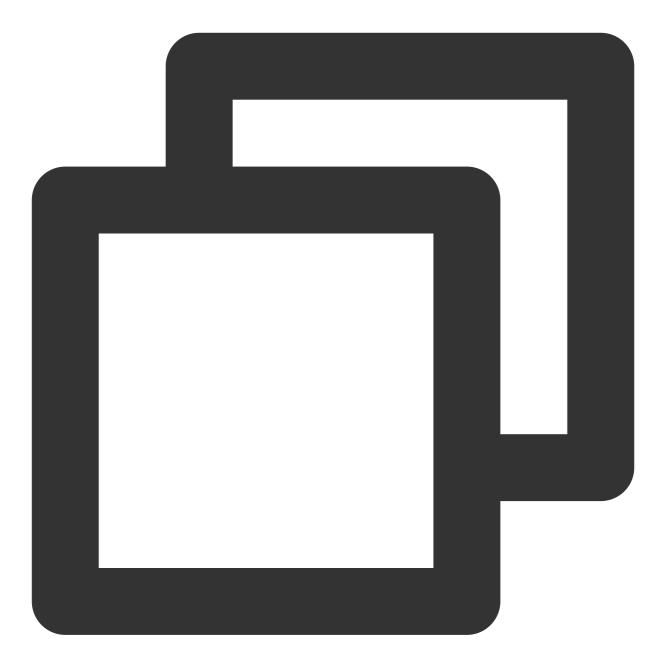
To delete a Prometheus instance, run the following command:



terraform destroy



If the following information appears, you are prompted to enter "yes" for confirmation:



tencentcloud\_monitor\_tmp\_instance.foo: Refreshing state... [id=prom-8dyb6iny]

Terraform used the selected providers to generate the following execution plan. Res following symbols:

- destroy

Terraform will perform the following actions:

# tencentcloud\_monitor\_tmp\_instance.foo will be destroyed



```
- resource "tencentcloud_monitor_tmp_instance" "foo" {
     - api_root_path = "http://10.0.0.34:9090/api/v1" -> null
      - data retention time = 30 -> null
                          = "prom-8dyb6iny" -> null
      - id
     - instance_name
                          = "tf-tmp-instance-sjtest" -> null
                          = "10.0.0.34" -> null

    ipv4_address

                          = "10.0.0.34:9090" -> null
     - proxy_address
      - remote_write
                          = "http://10.0.0.34:9090/api/v1/prom/write" -> null
                          = "subnet-es8rv1kx" -> null
     - subnet_id
      - tags
                           = {
         - "createdBy" = "terraform"
       } -> null
                          = "vpc-0n42dxzs" -> null
     - vpc_id
     - zone
                          = "ap-mumbai-1" -> null
    }
Plan: 0 to add, 0 to change, 1 to destroy.
Do you really want to destroy all resources?
 Terraform will destroy all your managed infrastructure, as shown above.
 There is no undo. Only 'yes' will be accepted to confirm.
 Enter a value: yes
tencentcloud_monitor_tmp_instance.foo: Destroying... [id=prom-8dyb6iny]
tencentcloud_monitor_tmp_instance.foo: Destruction complete after 6s
Destroy complete! Resources: 1 destroyed.
```

#### Note:

When Destroy complete! Resources: 1 destroyed. is displayed, the instance has been successfully deleted.

## Managing the Integration Center of Prometheus Instances Using Terraform

Last updated : 2023-12-29 16:09:44

### Prerequisites

### **Installing Terraform**

For detailed instructions on installing Terraform, see Installing Terraform.

#### Note:

The installed version of Terraform must not be below v1.6.3. You can verify the version of Terraform installed by using the command terraform --version.

### **Configuring Tencent Cloud Account Information**

Before using Terraform for the first time, go to API Key to apply for a SecretID and SecretKey of the security credentials. If you already have valid credentials, skip this step.

1. Log in to the CAM Console and choose Access Key > API Keys from the left sidebar.

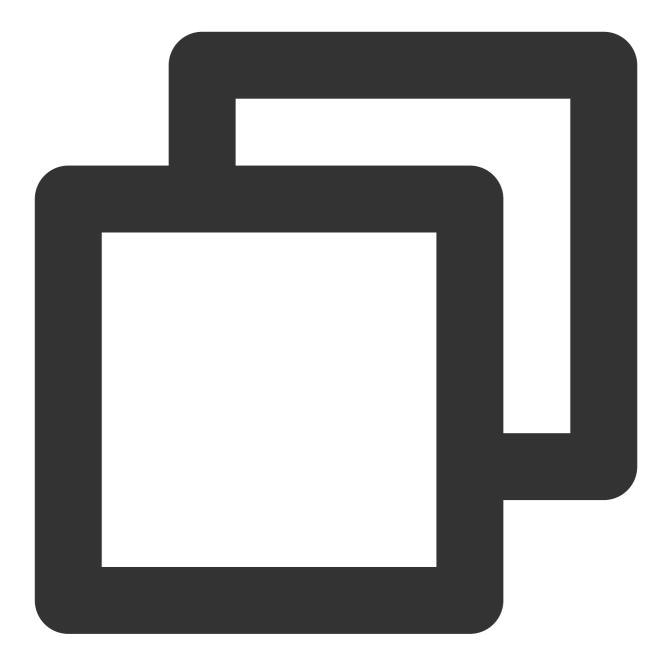
2. On the API Keys page, click Create Key to create a SecretId/SecretKey pair.

### There are two methods for configuring Tencent Cloud account information:

Authentication by Static Credentials

Create a provider.tf file in the user directory with the following content. Replace my-secret-id and mysecret-key with your SecretId and SecretKey, respectively.



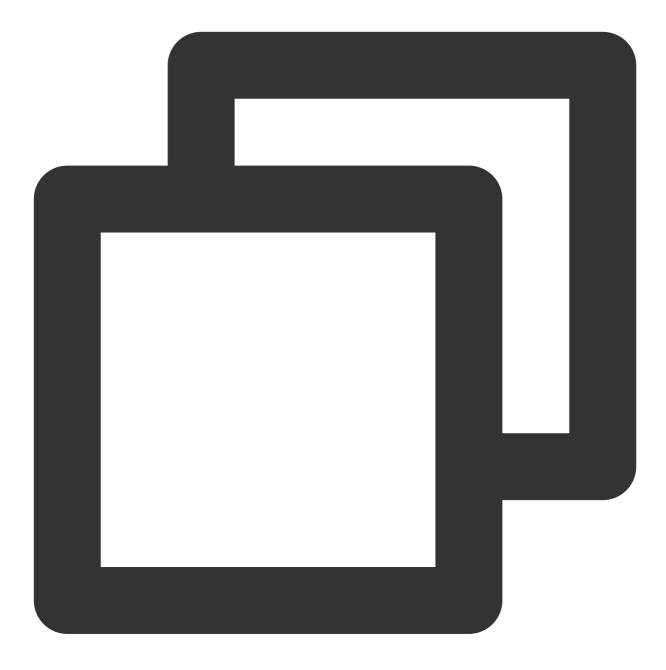


```
provider "tencentcloud" {
   secret_id = "my-secret-id"
   secret_key = "my-secret-key"
}
```

### Authentication by Environment Variables

Configure the computer environment variables or cloud environment variables by running the following command. Replace YOUR\_SECRET\_ID and YOUR\_SECRET\_KEY with your SecretId and SecretKey, respectively.

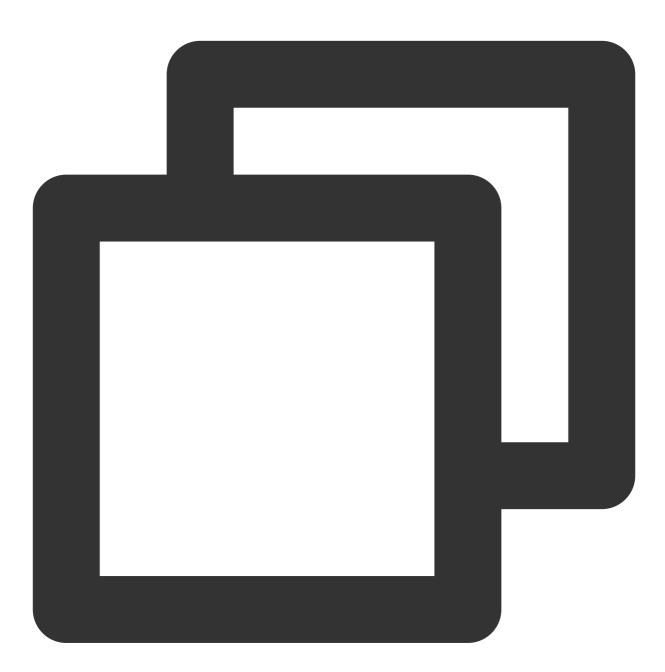




export TENCENTCLOUD\_SECRET\_ID=YOUR\_SECRET\_ID
export TENCENTCLOUD\_SECRET\_KEY=YOUR\_SECRET\_KEY

# Incorporating Component Integration from the Integration Center of a Prometheus Instance

1. Creating a new Terraform configuration file: Create a new directory and a file named main.tf in the directory. The following is the configuration information:



```
#Specifying Provider Configuration Information
terraform {
  required_providers {
    tencentcloud = {
       source = "tencentcloudstack/tencentcloud"
    }
```

```
}
#Prometheus Managing Cloud Monitoring Integration
##black-box Integration
resource "tencentcloud_monitor_tmp_exporter_integration" "tmpExporterIntegration" {
  instance_id = tencentcloud_monitor_tmp_instance.foo.id
            = "blackbox-exporter"
 kind
            = "{\\"name\\":\\"balck-box-tf-test\\", \\"kind\\":\\"blackbox-exporte
 content
 kube_type
             = 1
 cluster_id = ""
}
##Cloud Monitoring Plugin Integration
resource "tencentcloud_monitor_tmp_exporter_integration" "tmpExporterMointor" {
  instance_id = tencentcloud_monitor_tmp_instance.foo.id
 kind
            = "qcloud-exporter"
  content
            = "{\\"name\\":\\"tf-test-cjtest\\",\\"kind\\":\\"qcloud-exporter\\",
 kube_type = 1
 cluster_id = ""
}
```

### Note:

The fields to be configured when you create the Integration Center component of the Prometheus instance are as follows:

cluster\_id: (Required, string) Cluster ID.

content: (Required, string) Integration configuration (internal parameters within content can be replaced as needed). instance\_id: (Required, string) Instance ID.

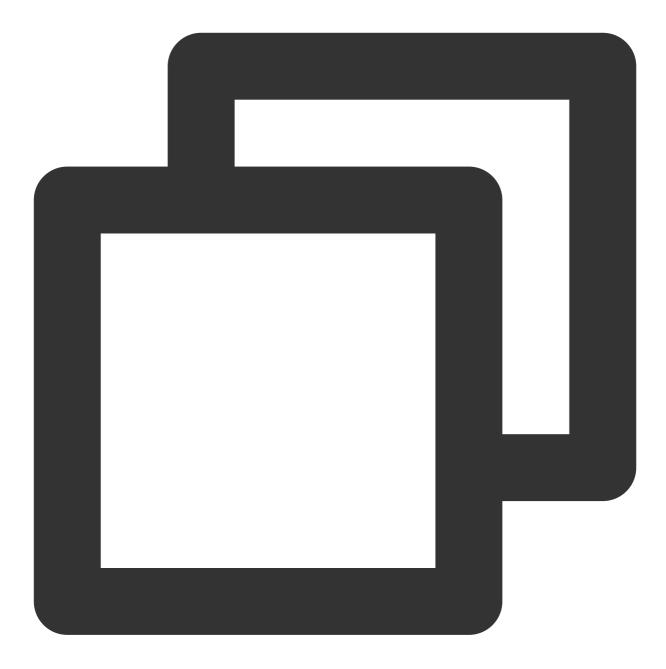
kind: (Required, string) Type.

kube\_type: (Required, integer) Integration configuration.

For more detailed parameters, see Tencent Cloud Integration with GitHub or Terraform Tencent Cloud Provider.

2. To initialize the Terraform runtime environment, run the following command:

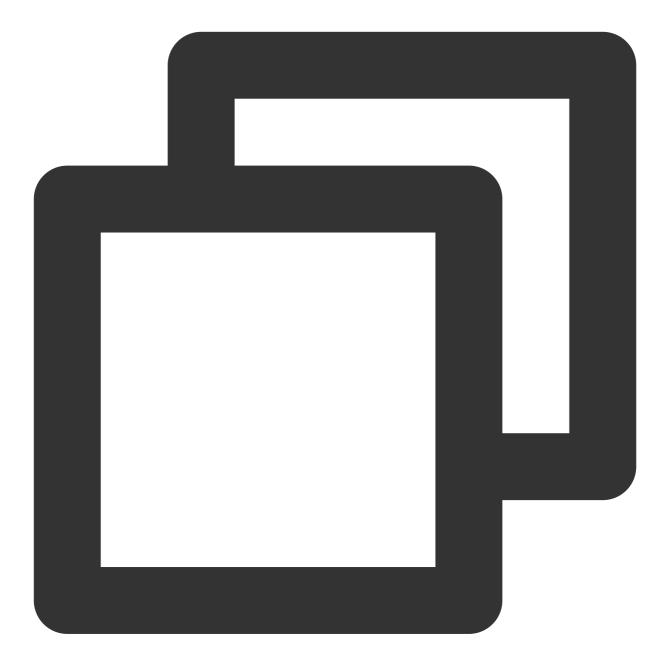




terraform init

Expected output:





Initializing the backend...

Initializing provider plugins...

- Reusing previous version of tencentcloudstack/tencentcloud from the dependency lo
- Using previously-installed tencentcloudstack/tencentcloud v1.81.32

Terraform has been successfully initialized!

You may now begin working with Terraform. Try running "terraform plan" to see any changes that are required for your infrastructure. All Terraform commands should now work.



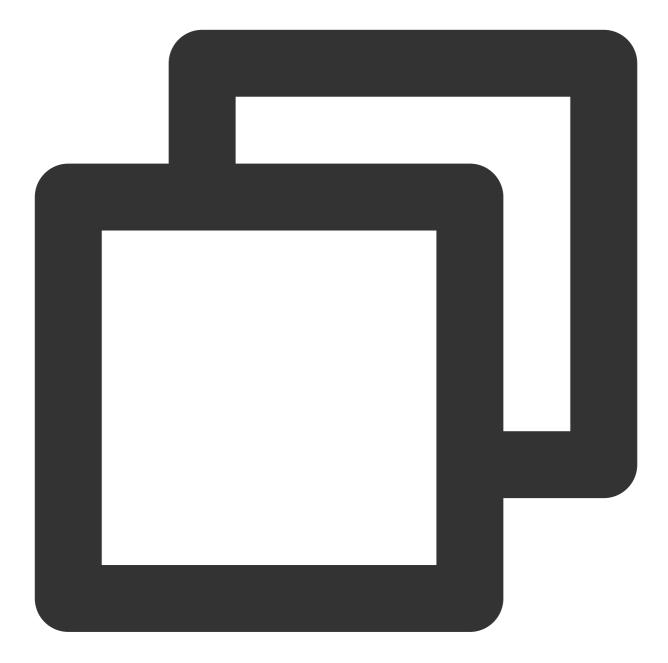
If you ever set or change modules or backend configuration for Terraform, rerun this command to reinitialize your working directory. If you forget, other commands will detect it and remind you to do so if necessary.

3. To generate a resource plan, run the following command:



terraform plan

Expected output:



Terraform used the selected providers to generate the following execution plan. Res following symbols:

+ create

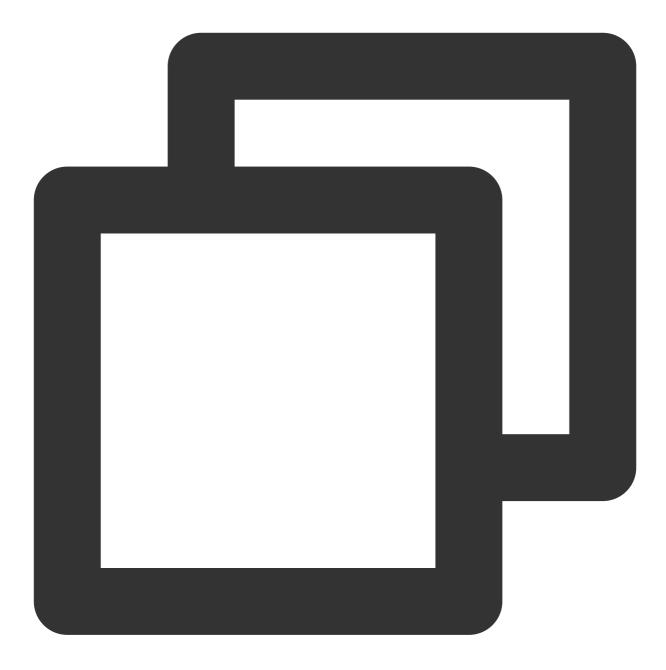
Terraform will perform the following actions:

```
+ id = (known after apply)
     + instance_id = (known after apply)
     + kind = "blackbox-exporter"
     + kube_type = 1
   }
  # tencentcloud_monitor_tmp_exporter_integration.tmpExporterMointor will be create
 + resource "tencentcloud_monitor_tmp_exporter_integration" "tmpExporterMointor" {
     + content = jsonencode(
          XXX...
       )
     + id
             = (known after apply)
     + instance_id = (known after apply)
     + kind = "qcloud-exporter"
     + kube_type = 1
   }
  # tencentcloud_monitor_tmp_instance.foo will be created
  + resource "tencentcloud_monitor_tmp_instance" "foo" {
     + api_root_path = (known after apply)
     + data_retention_time = 30
     + id
                        = (known after apply)
                         = "tf-tmp-instance-sjtest"
     + instance_name
     + ipv4_address
                         = (known after apply)
     + proxy_address
                         = (known after apply)
     + remote_write
                         = (known after apply)
     + subnet_id
                         = "subnet-es8rv1kx"
     + tags
                         = {
         + "createdBy" = "terraform"
       }
     + vpc_id
                         = "vpc-0n42dxzs"
     + zone
                         = "ap-mumbai-1"
   }
Plan: 3 to add, 0 to change, 0 to destroy.
```

Note: You didn't use the -out option to save this plan, so Terraform can't guarante you run "terraform apply" now.

4. To create component integration of the Integration Center, run the following command:

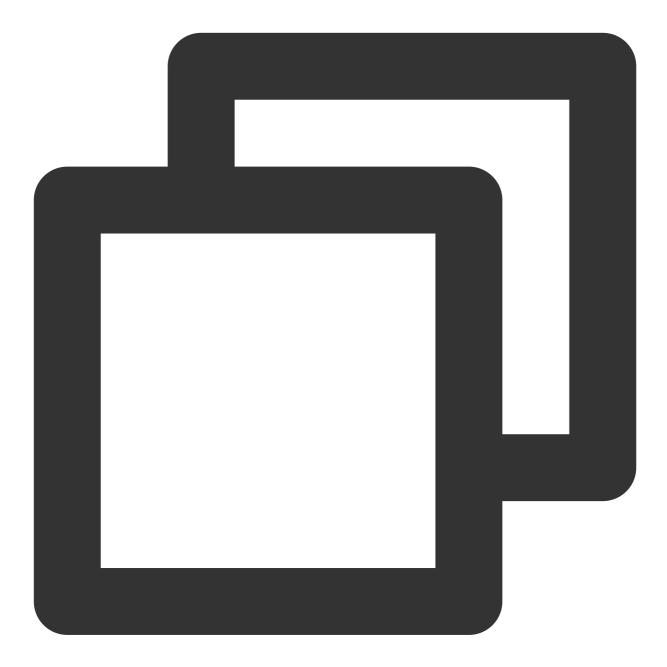




terraform apply

Expected output:





Terraform used the selected providers to generate the following execution plan. Res following symbols:

+ create

Terraform will perform the following actions: XXX...

Plan: 3 to add, 0 to change, 0 to destroy.

Do you want to perform these actions?

```
Terraform will perform the actions described above.
Only 'yes' will be accepted to approve.
Enter a value: yes
tencentcloud_monitor_tmp_instance.foo: Creating...
tencentcloud_monitor_tmp_instance.foo: Still creating... [10s elapsed]
...
Apply complete! Resources: 3 added, 0 changed, 0 destroyed.
```

#### Note:

When the message "Apply complete! Resources: 3 added, 0 changed, 0 destroyed." is displayed, the creation is successful.

### Viewing the Status of Prometheus Instances

Log in to the TCOP, and choose **Managed Service for Prometheus** from the left-hand navigation bar. Existing instances are displayed in the Prometheus instance list.

### Deleting the Component Integration in the Prometheus Instance Integration Center

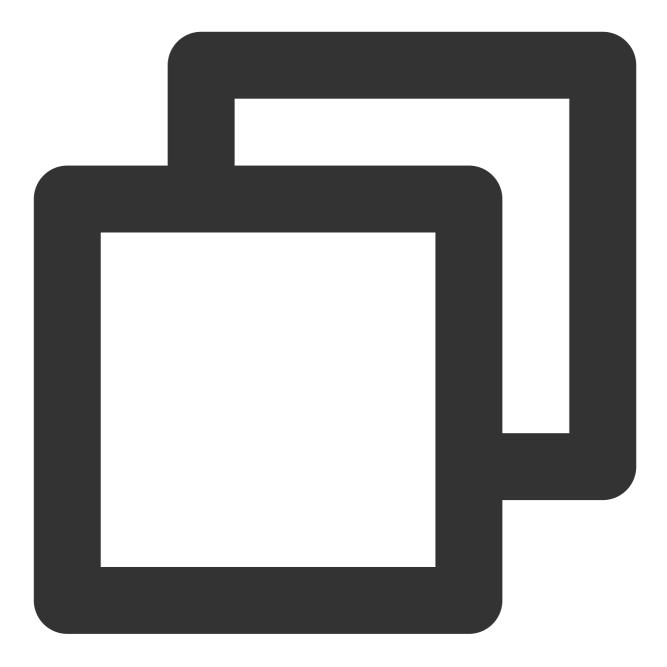
To terminate the resources, run the following command:





terraform destroy





tencentcloud\_monitor\_tmp\_instance.foo: Refreshing state... [id=prom-8dyb6iny]

Terraform used the selected providers to generate the following execution plan. Res following symbols:

- destroy

Terraform will perform the following actions: XXX...

Plan: 0 to add, 0 to change, 1 to destroy.



Do you really want to destroy all resources? Terraform will destroy all your managed infrastructure, as shown above. There is no undo. Only 'yes' will be accepted to confirm.

Enter a value: yes

tencentcloud\_monitor\_tmp\_instance.foo: Destroying... [id=prom-8dyb6iny]
tencentcloud\_monitor\_tmp\_instance.foo: Destruction complete after 6s

Destroy complete! Resources: 1 destroyed.

#### Note:

When Destroy complete! Resources: 1 destroyed. is displayed, the instance has been successfully deleted.

# Collecting Container Monitoring Data Using Terraform

Last updated : 2023-12-29 16:10:38

## Prerequisites

#### **Installing Terraform**

For detailed instructions on installing Terraform, see Installing Terraform.

#### Note:

The installed version of Terraform must not be below v1.6.3. You can verify the version of Terraform installed by using the command terraform --version.

#### **Configuring Tencent Cloud Account Information**

Before using Terraform for the first time, go to API Key to apply for a SecretID and SecretKey of the security credentials. If you already have valid credentials, skip this step.

1. Log in to the CAM Console and choose Access Key > API Keys from the left sidebar.

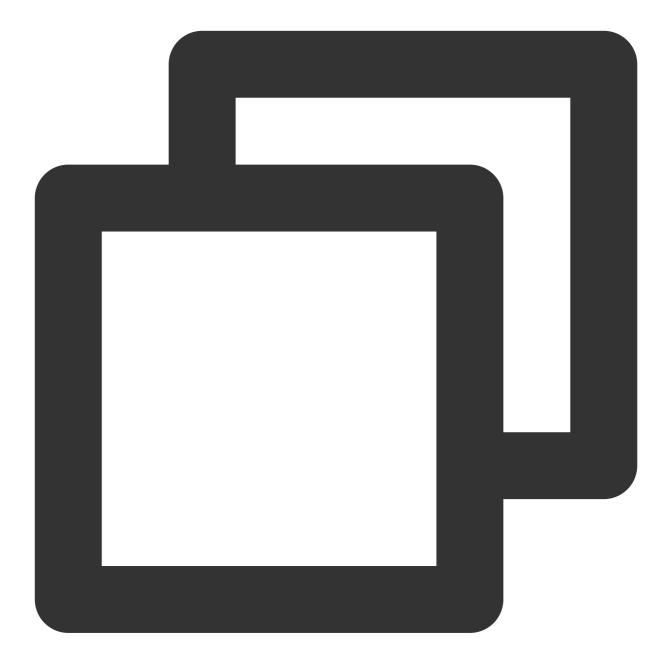
2. On the API Keys page, click Create Key to create a SecretId/SecretKey pair.

#### There are two methods for configuring Tencent Cloud account information:

Authentication by Static Credentials

Create a provider.tf file in the user directory with the following content. Replace my-secret-id and mysecret-key with your SecretId and SecretKey, respectively.



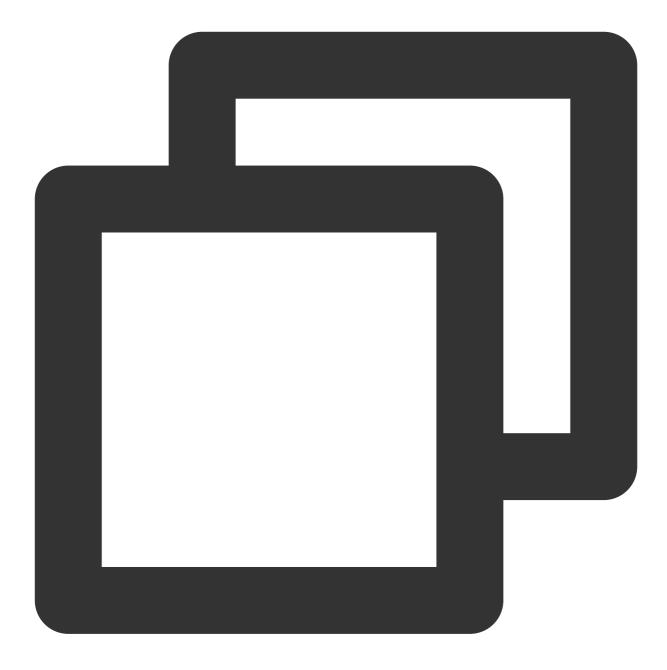


```
provider "tencentcloud" {
   secret_id = "my-secret-id"
   secret_key = "my-secret-key"
}
```

#### Authentication by Environment Variables

Set the computer variables or cloud environment variables by running the command below. Replace YOUR\_SECRET\_ID and YOUR\_SECRET\_KEY with your SecretId and SecretKey, respectively.



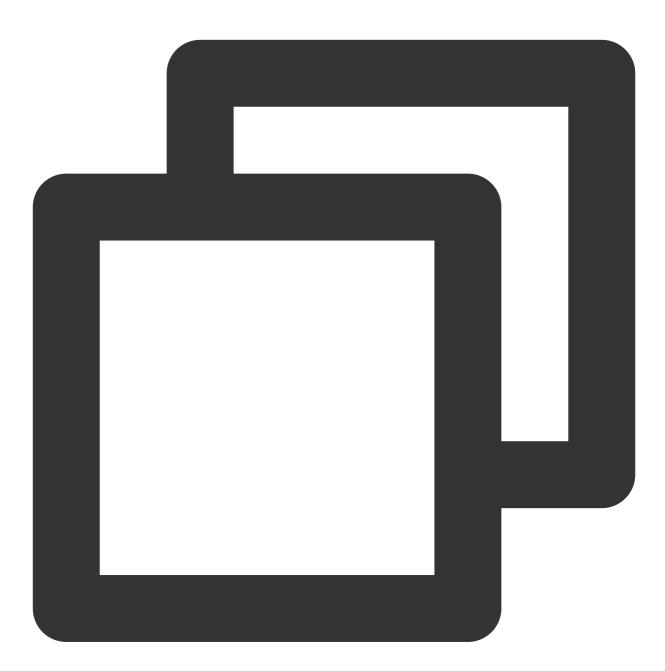


export TENCENTCLOUD\_SECRET\_ID=YOUR\_SECRET\_ID
export TENCENTCLOUD\_SECRET\_KEY=YOUR\_SECRET\_KEY

## Adding Container Monitoring Data Collection to Prometheus Instances



1. Create a new Terraform configuration file: Create a new directory and a file named main.tf in the directory. The following shows the configuration:



```
#Specifying Provider Configuration Information
terraform {
  required_providers {
    tencentcloud = {
       source = "tencentcloudstack/tencentcloud"
    }
```

```
}
}
#Prometheus Managing Cloud Monitoring Integration
#Prometheus Management of Container Clusters (Collection Containers)
resource "tencentcloud_monitor_tmp_tke_cluster_agent" "foo" {
    instance_id = tencentcloud_monitor_tmp_instance.foo.id
    agents {
        region = "ap-mumbai"
        cluster_type = "eks"
        cluster_id = "cls-1uary7z2" #ID of the cluster that needs to be associated
        enable_external = false
    }
}
```

#### Note:

To create the Integration Center component of a Prometheus instance, the corresponding fields must be configured as follows:

instance\_id: (Required, string, ForceNew) Instance ID.

agents: (Required, list) List of agents.

cluster\_id: (Required, string) Identifier for the cluster ID.

region: (Required, string) Region restrictions.

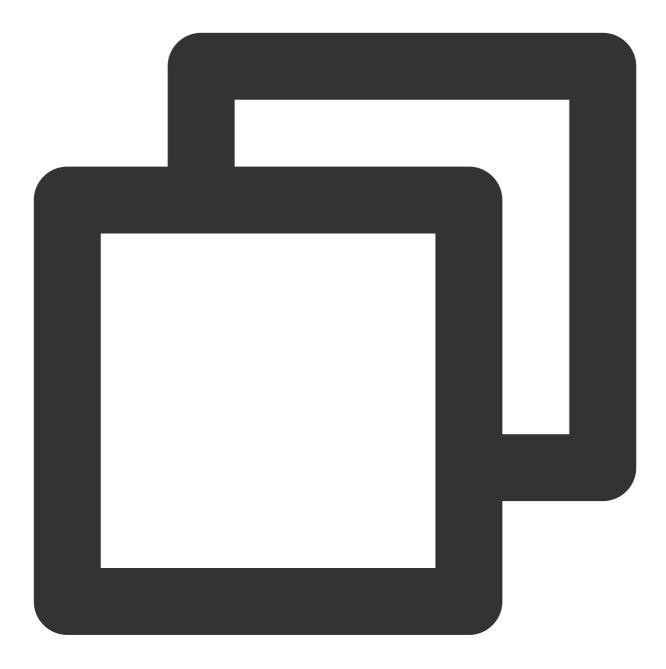
enable\_external: (Required, Bool) Whether to enable public network CLB.

cluster\_type: (Required, string) Cluster type.

For more details about parameters, see **Detailed Parameters**. You can also see Tencent Cloud Provider on Github and Terraform Tencent Cloud Provider.

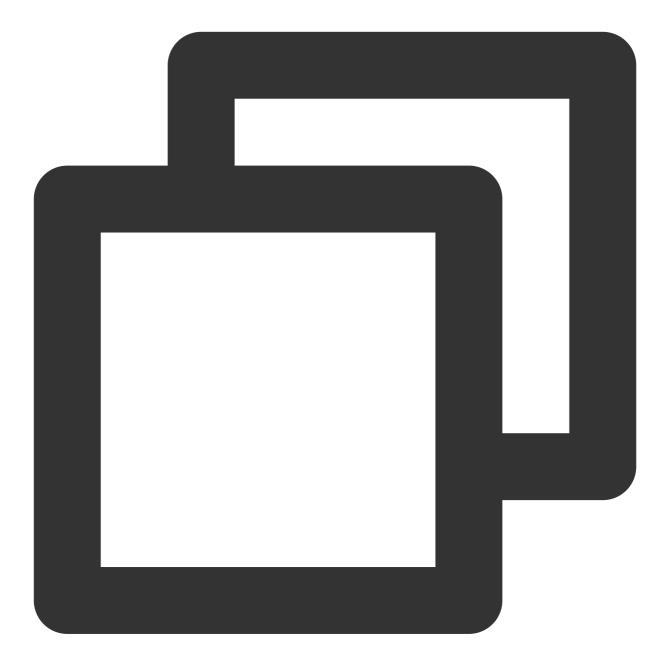
2. To initialize the Terraform runtime environment, run the following command:





terraform init





Initializing the backend...

Initializing provider plugins...

- Reusing previous version of tencentcloudstack/tencentcloud from the dependency lo
- Using previously-installed tencentcloudstack/tencentcloud v1.81.32

Terraform has been successfully initialized!

You may now begin working with Terraform. Try running "terraform plan" to see any changes that are required for your infrastructure. All Terraform commands should now work.



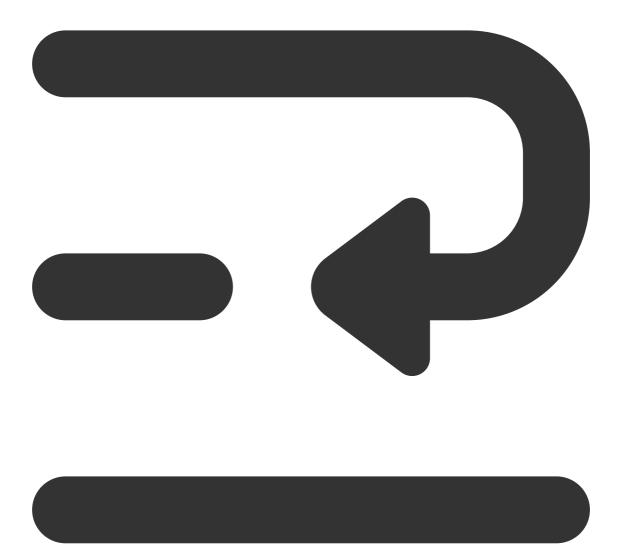
If you ever set or change modules or backend configuration for Terraform, rerun this command to reinitialize your working directory. If you forget, other commands will detect it and remind you to do so if necessary.

3. To generate a resource plan, run the following command:

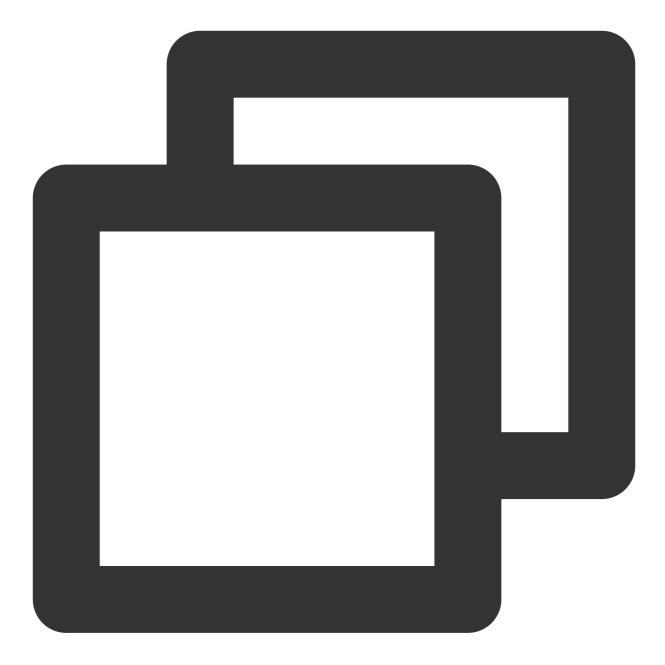


terraform plan









tencentcloud\_monitor\_tmp\_instance.foo: Refreshing state... [id=prom-jh0zntj2]
tencentcloud\_monitor\_tmp\_exporter\_integration.tmpExporterIntegration: Refreshing st
tencentcloud\_monitor\_tmp\_exporter\_integration.tmpExporterMointor: Refreshing state.

Terraform used the selected providers to generate the following execution plan. Res following symbols:

+ create

Terraform will perform the following actions:

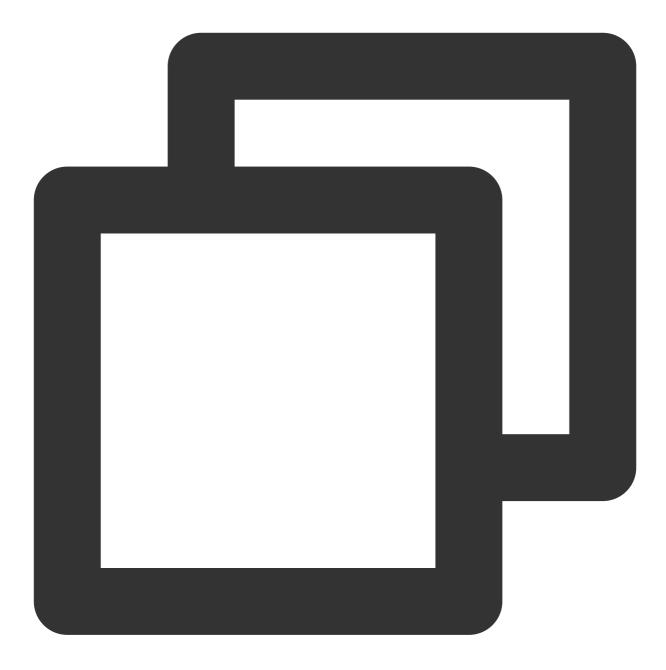
# tencentcloud\_monitor\_tmp\_tke\_cluster\_agent.foo will be created

```
+ resource "tencentcloud_monitor_tmp_tke_cluster_agent" "foo" {
    + id = (known after apply)
    + instance_id = "prom-jh0zntj2"
    + agents {
        + cluster_id = "cls-1uary7z2"
        + cluster_name = (known after apply)
        + cluster_type = "eks"
        + enable_external = false
        + region = "ap-mumbai"
        + status = (known after apply)
        }
    }
Plan: 1 to add, 0 to change, 0 to destroy.
Note: You didn't use the -out option to save this plan, so Terraform can't guarante
```

4. To create component integration of the Integration Center, run the following command:

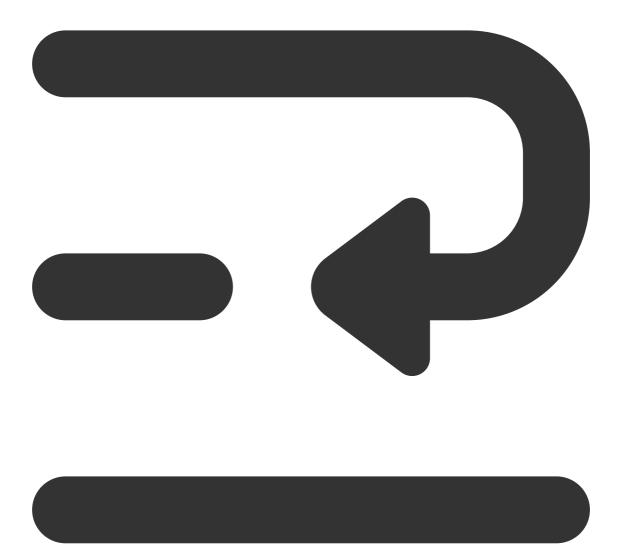
you run "terraform apply" now.

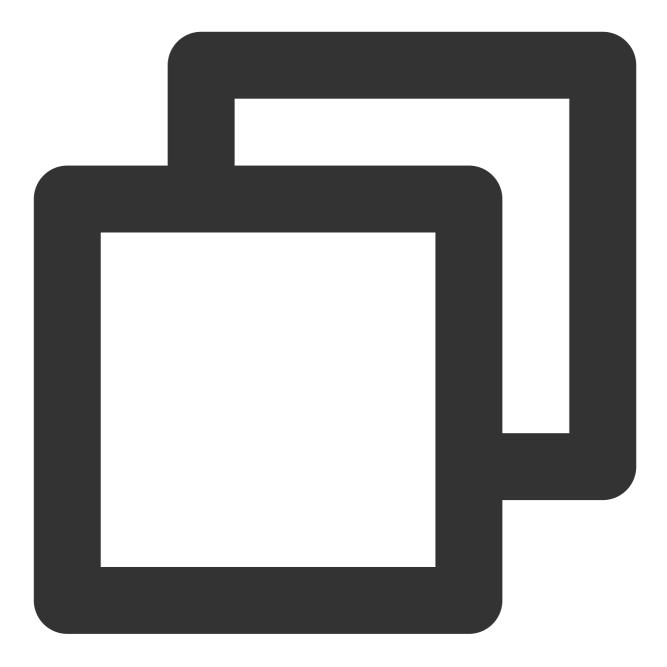




terraform apply







tencentcloud\_monitor\_tmp\_instance.foo: Refreshing state... [id=prom-jh0zntj2]
tencentcloud\_monitor\_tmp\_exporter\_integration.tmpExporterIntegration: Refreshing st
tencentcloud\_monitor\_tmp\_exporter\_integration.tmpExporterMointor: Refreshing state.

Terraform used the selected providers to generate the following execution plan. Res following symbols:

+ create

Terraform will perform the following actions:

# tencentcloud\_monitor\_tmp\_tke\_cluster\_agent.foo will be created

```
+ resource "tencentcloud_monitor_tmp_tke_cluster_agent" "foo" {
     + id = (known after apply)
     + instance_id = "prom-jh0zntj2"
     + agents {
         + cluster_id = "cls-1uary7z2"
         + cluster_name = (known after apply)
         + cluster_type = "eks"
         + enable_external = false
         + region
                         = "ap-mumbai"
         + status = (known after apply)
       }
    }
Plan: 1 to add, 0 to change, 0 to destroy.
Do you want to perform these actions?
 Terraform will perform the actions described above.
 Only 'yes' will be accepted to approve.
 Enter a value: yes
tencentcloud_monitor_tmp_tke_cluster_agent.foo: Creating...
Instance content...
Apply complete! Resources: 1 added, 0 changed, 0 destroyed.
```

## Viewing the Status of the Prometheus Instance

Log in to the TCOP, and choose **Managed Service for Prometheus** from the left-hand navigation bar. Existing instances are displayed in the Prometheus instance list.

## Deleting the Component Integration in the Prometheus Instance Integration Center

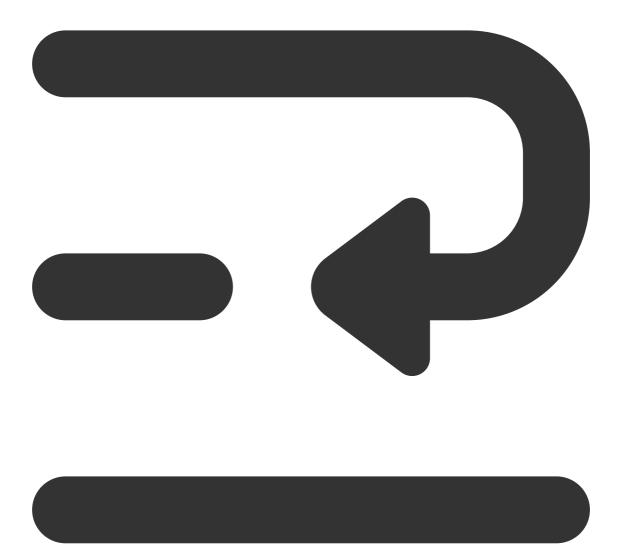
To terminate the resources, run the following command:

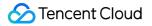


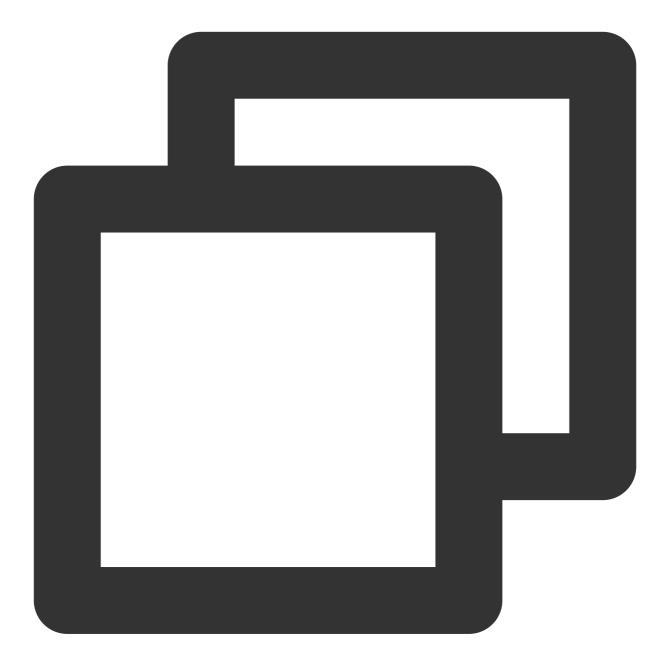


terraform destroy









tencentcloud\_monitor\_tmp\_instance.foo: Refreshing state... [id=prom-8dyb6iny]

Terraform used the selected providers to generate the following execution plan. Res following symbols:

- destroy

Terraform will perform the following actions: Instance content...

Plan: 0 to add, 0 to change, 1 to destroy.



Do you really want to destroy all resources? Terraform will destroy all your managed infrastructure, as shown above. There is no undo. Only 'yes' will be accepted to confirm.

Enter a value: yes

tencentcloud\_monitor\_tmp\_instance.foo: Destroying... [id=prom-8dyb6iny]
tencentcloud\_monitor\_tmp\_instance.foo: Destruction complete after 6s

Destroy complete! Resources: 1 destroyed.

#### Note:

If Destroy complete! Resources: (numbers of existing instances) destroyed. is displayed, the instance has been deleted.

## **Configuring Alarm Policies Using Terraform**

Last updated : 2023-12-29 16:11:24

## Prerequisites

#### Installing Terraform

For detailed instructions on installing Terraform, see Installing Terraform.

#### Note:

The installed version of Terraform must not be below v1.6.3. You can verify the version of Terraform installed by using the command terraform --version.

#### **Configuring Tencent Cloud Account Information**

Before using Terraform for the first time, go to API Key to apply for a SecretID and SecretKey of the security credentials. If you already have valid credentials, skip this step.

1. Log in to the CAM Console and choose Access Key > API Keys from the left sidebar.

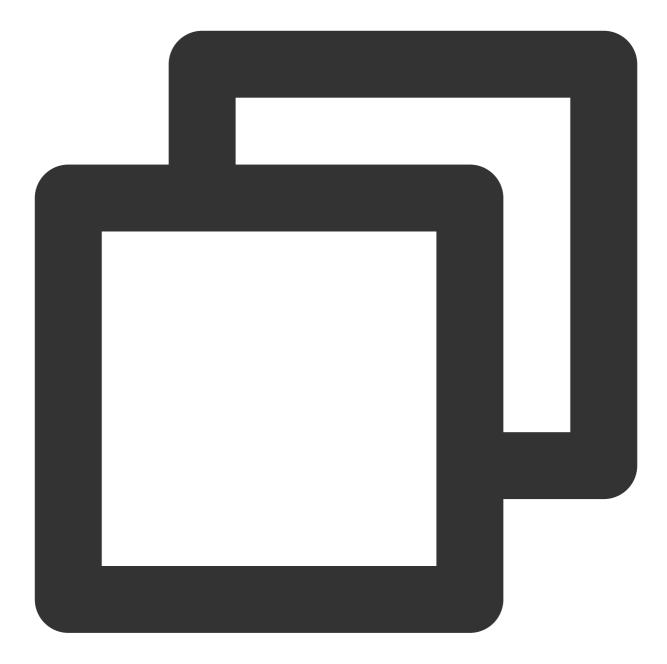
2. On the API Keys page, click Create Key to create a SecretId/SecretKey pair.

#### There are two methods for configuring Tencent Cloud account information:

Authentication by Static Credentials

Create a provider.tf file in the user directory with the following content. Replace my-secret-id and my-secret-key with your SecretId and SecretKey, respectively.

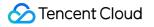


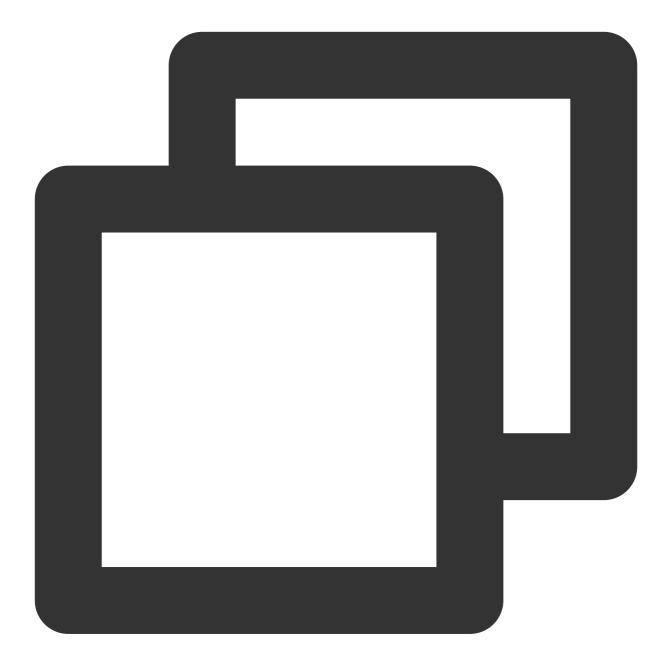


```
provider "tencentcloud" {
   secret_id = "my-secret-id"
   secret_key = "my-secret-key"
}
```

#### Authentication by Environment Variables

Configure the computer environment variables or cloud environment variables by running the following command. Replace YOUR\_SECRET\_ID and YOUR\_SECRET\_KEY with your SecretId and SecretKey, respectively.



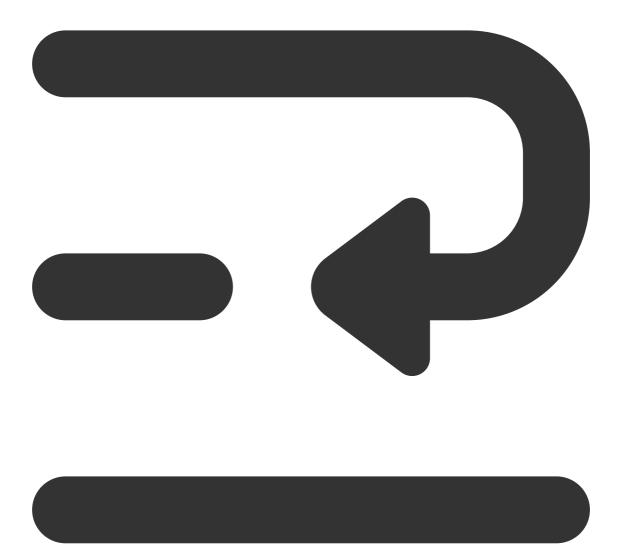


export TENCENTCLOUD\_SECRET\_ID=YOUR\_SECRET\_ID
export TENCENTCLOUD\_SECRET\_KEY=YOUR\_SECRET\_KEY

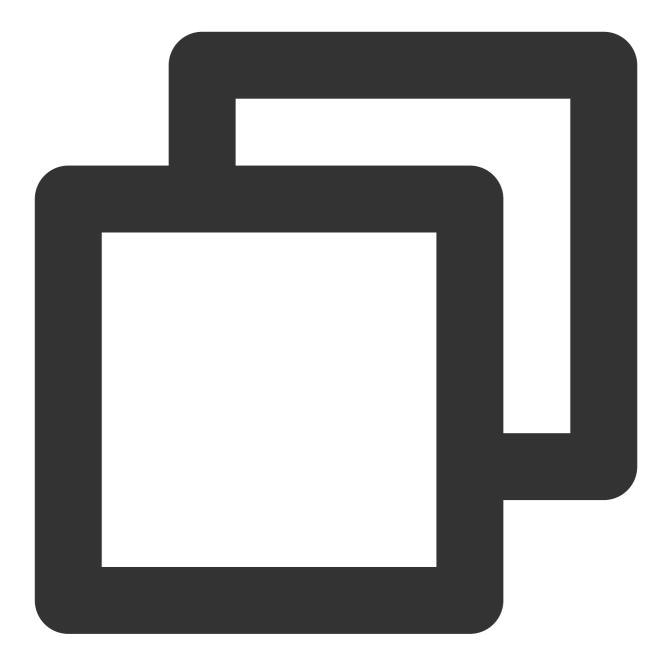
### Adding a Terraform Alarm Policy to a Prometheus Instance

1. Create a new Terraform configuration file. Create a new directory and a file named main.tf in the directory. The following is the configuration information:









```
#Specifying Provider Configuration Information
terraform {
   required_providers {
     tencentcloud = {
        source = "tencentcloudstack/tencentcloud"
     }
   }
   #Configuring Alarm for Prometheus (Cloud Monitoring Side Configuration)
```



```
resource "tencentcloud_monitor_tmp_alert_rule" "foo" {
 duration = "2m"
 expr
             = "avg by (instance) (mysql_global_status_threads_connected) / avg by
 instance_id = tencentcloud_monitor_tmp_instance.foo.id
 receivers = ["notice-zmjsavnp"] # Notifications templates can be created here usi
 rule_name = "MySQL Excessive Connections--tf-Cloud Monitor Test"
  rule_state = 2
 type = "MySQL/Excessive MySQL Connections"
 annotations {
   key = "description"
   value = "Too many MySQL connections, instance: {{$labels.instance}}, current va
  }
  annotations {
   key = "summary"
   value = "MySQL has too many connections(>80%)"
  }
 labels {
   key = "severity"
   value = "warning"
  }
}
```

#### Note:

The fields of the configuration are as follows:

duration: Rule persistence duration.

expr: Alarm expression.

instance\_id: Instance ID.

receivers: List of alarm recipients.

rule\_name: Alarm rule name.

rule\_state: Alarm rule status.

type: Alarm rule type.

For more detailed parameters, see Tencent Cloud Integration to GitHub or Terraform Tencent Cloud Provider.

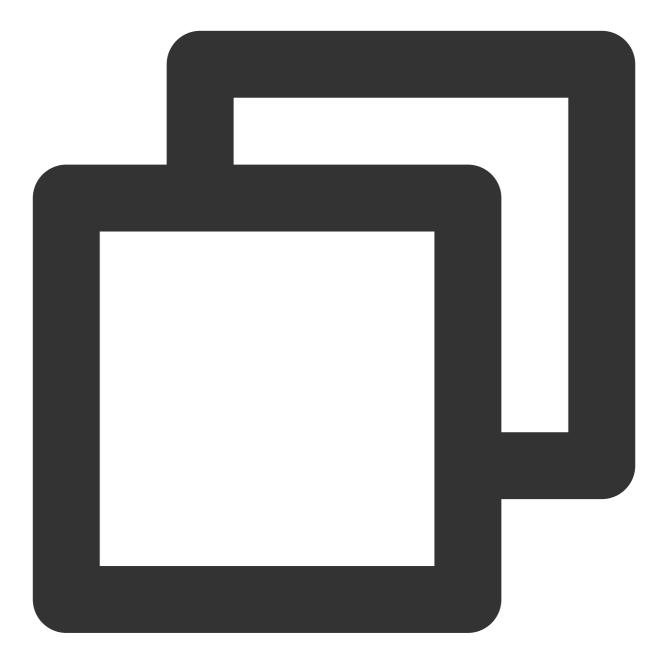
2. To initialize the Terraform runtime environment, run the following command:





terraform init





Initializing the backend...

Initializing provider plugins...

- Reusing previous version of tencentcloudstack/tencentcloud from the dependency lo

- Using previously-installed tencentcloudstack/tencentcloud v1.81.32

Terraform has been successfully initialized!

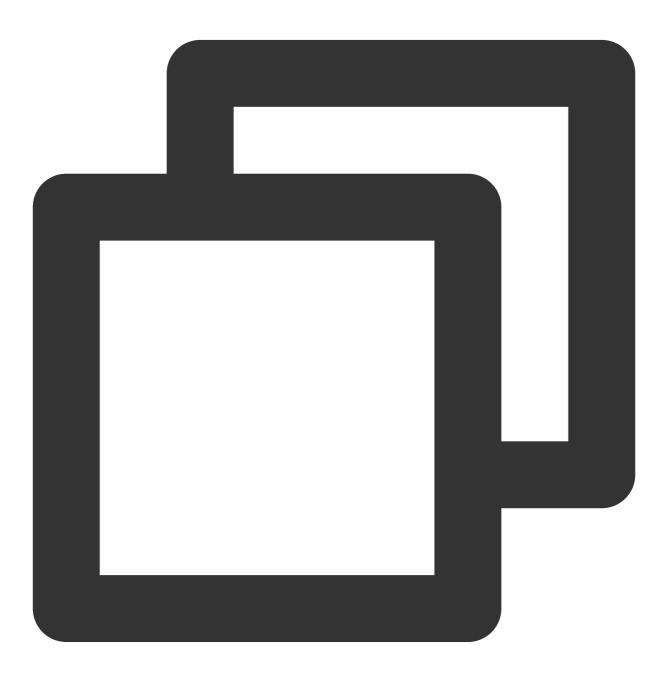
You may now begin working with Terraform. Try running "terraform plan" to see any changes that are required for your infrastructure. All Terraform commands



should now work.

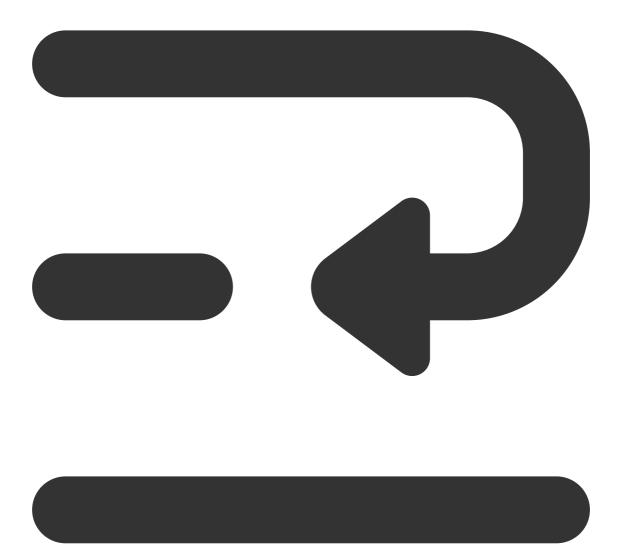
If you ever set or change modules or backend configuration for Terraform, rerun this command to reinitialize your working directory. If you forget, other commands will detect it and remind you to do so if necessary.

3. To generate a resource plan, run the following command:

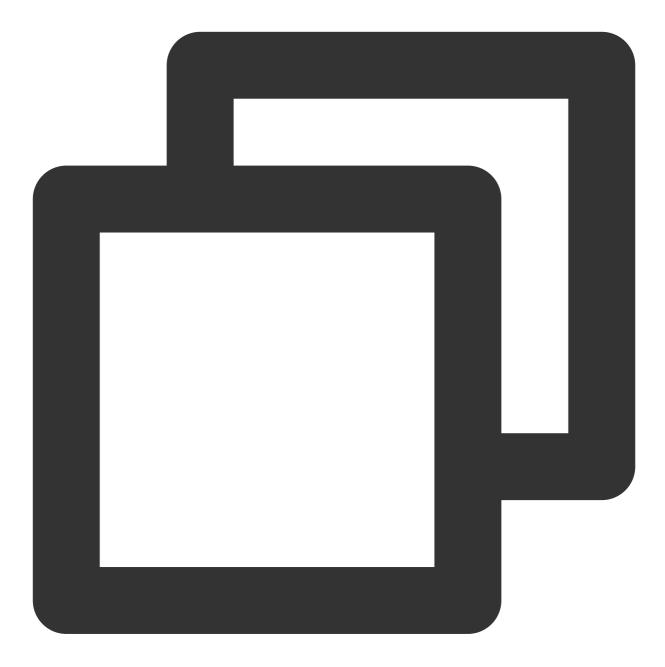


terraform plan









tencentcloud\_monitor\_tmp\_instance.foo: Refreshing state... [id=prom-jh0zntj2]
tencentcloud\_monitor\_tmp\_exporter\_integration.tmpExporterIntegration: Refreshing st
tencentcloud\_monitor\_tmp\_tke\_cluster\_agent.foo: Refreshing state... [id=prom-jh0znt
tencentcloud\_monitor\_tmp\_exporter\_integration.tmpExporterMointor: Refreshing state.

Terraform used the selected providers to generate the following execution plan. Res following symbols:

+ create

Terraform will perform the following actions:

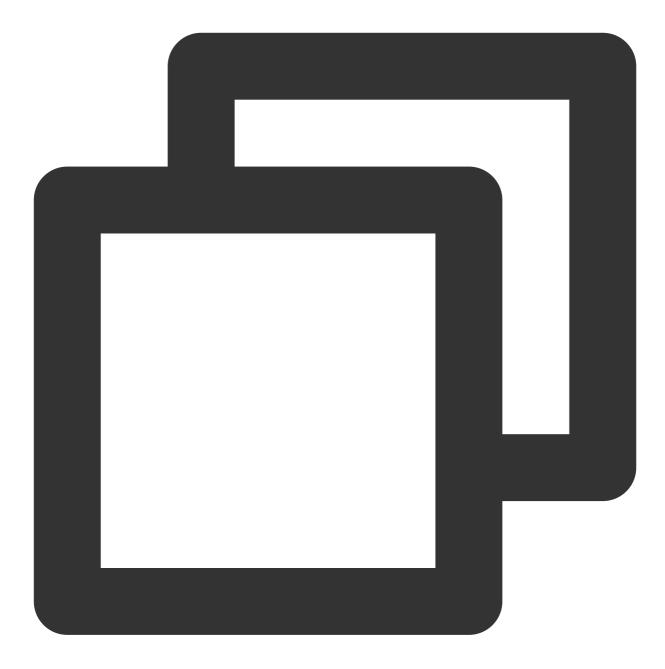
```
Stencent Cloud
```

```
# tencentcloud_monitor_tmp_alert_rule.foo will be created
  + resource "tencentcloud_monitor_tmp_alert_rule" "foo" {
     + duration = "2m"
     + expr
                  = "avg by (instance) (mysql_global_status_threads_connected) /
     + id
                 = (known after apply)
     + instance_id = "prom-jh0zntj2"
     + receivers = [
         + "notice-zmjsavnp",
       1
     + rule_name = "Excessive MySQL connections--tf-Cloud monitoring test"
     + rule_state = 2
              = "MySQL/Excessive MySQL Connections"
     + type
     + annotations {
         + key = "description"
         + value = "Excessive MySQL connections, instance: {{$labels.instance}}, C
       }
     + annotations {
         + key = "summary"
         + value = "Excessive MySQL connections (>80%)"
       }
     + labels {
        + key = "severity"
         + value = "warning"
       }
    }
Plan: 1 to add, 0 to change, 0 to destroy.
```

Note: You didn't use the -out option to save this plan, so Terraform can't guarante you run "terraform apply" now.

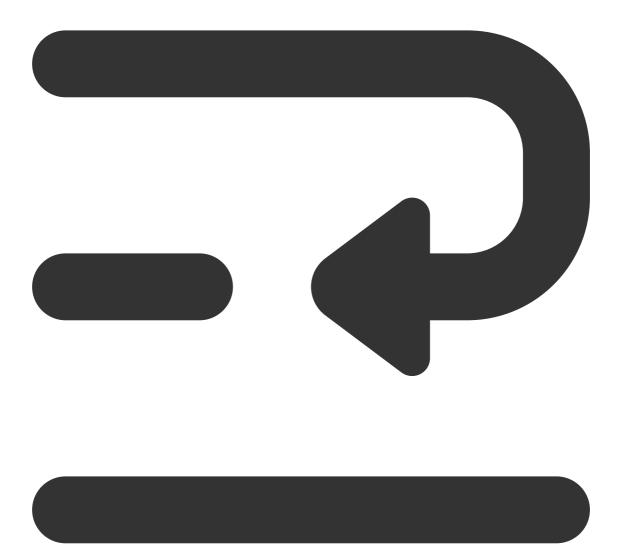
4. To create component integration of the Integration Center, run the following command:

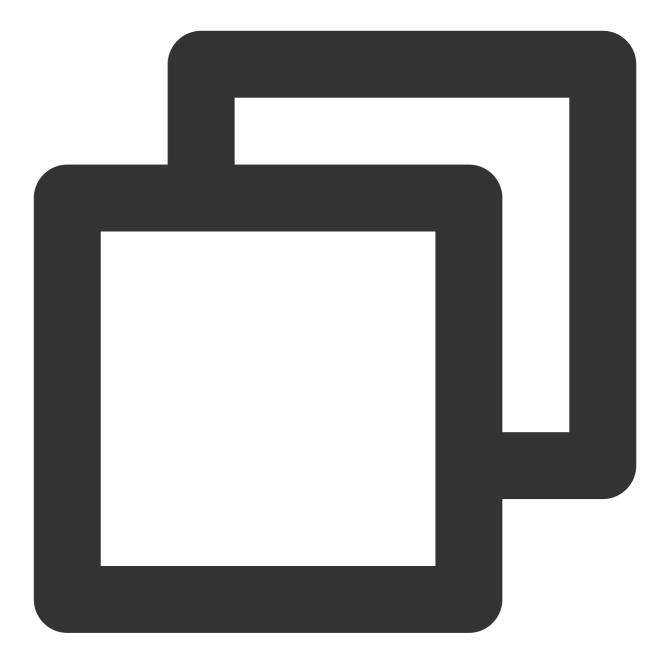




terraform apply







tencentcloud\_monitor\_tmp\_instance.foo: Refreshing state... [id=prom-jh0zntj2]
tencentcloud\_monitor\_tmp\_exporter\_integration.tmpExporterIntegration: Refreshing st
tencentcloud\_monitor\_tmp\_exporter\_integration.tmpExporterMointor: Refreshing state.
tencentcloud\_monitor\_tmp\_tke\_cluster\_agent.foo: Refreshing state... [id=prom-jh0znt

Terraform used the selected providers to generate the following execution plan. Res following symbols:

+ create

Terraform will perform the following actions: Instance content...

```
Plan: 1 to add, 0 to change, 0 to destroy.
Do you want to perform these actions?
Terraform will perform the actions described above.
Only 'yes' will be accepted to approve.
Enter a value: yes
tencentcloud_monitor_tmp_alert_rule.foo: Creating...
tencentcloud_monitor_tmp_alert_rule.foo: Creation complete after 2s [id=prom-jh0znt
Apply complete! Resources: 1 added, 0 changed, 0 destroyed.
```

### Viewing the Status of the Prometheus Instance

Log in to the TCOP, and choose **Managed Service for Prometheus** from the left-hand navigation bar. Existing instances are displayed in the Prometheus instance list.

### Deleting the Component Integration in the Prometheus Instance Integration Center

To terminate the resources, run the following command:

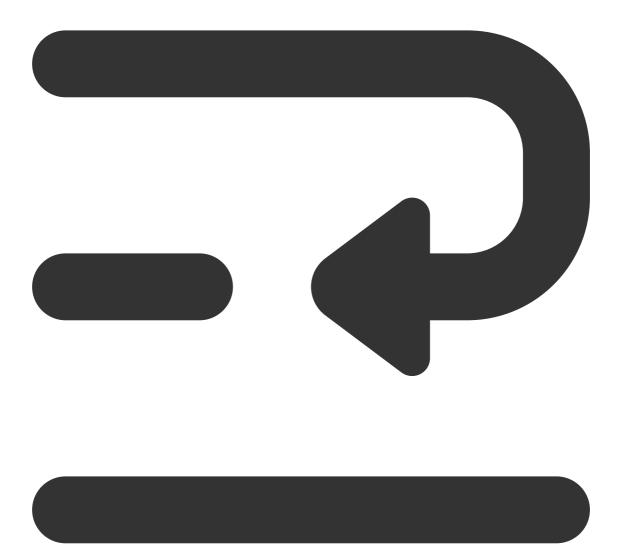




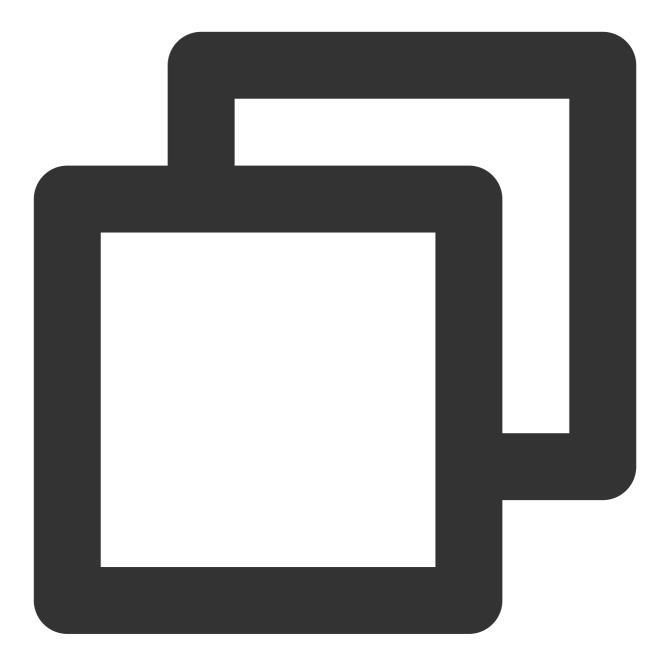
terraform destroy

Expected output:









```
tencentcloud_monitor_tmp_instance.foo: Refreshing state... [id=prom-8dyb6iny]
```

Terraform used the selected providers to generate the following execution plan. Res following symbols: - destroy

Terraform will perform the following actions: Instance content...

```
Do you really want to destroy all resources?
Terraform will destroy all your managed infrastructure, as shown above.
```

#### ठ Tencent Cloud

There is no undo. Only 'yes' will be accepted to confirm. Enter a value: yes tencentcloud\_monitor\_tmp\_instance.foo: Destroying... [id=prom-8dyb6iny] tencentcloud\_monitor\_tmp\_instance.foo: Destruction complete after 6s Destroy complete! Resources: 1 destroyed.

#### Note:

If Destroy complete! Resources: (numbers of existing instances) destroyed. is displayed, the instance has been deleted.

## FAQs Basic Questions

Last updated : 2024-08-15 17:56:35

#### How to use Prometheus's native default UI and other features?

TencentCloud Managed Service for Prometheus (TMP) is different from the open source stand-alone Prometheus. TMP is of structure with separated collection and storage and does not provide native default UI features. Use Tencent Cloud Managed Service for Grafana as an alternative for querying features.

# When the alarm is recovered, the \$value in the notification template is incorrect. How to handle it?

The \$value during the alarm recovering is the value that meets the alarm expression condition for the last time. The value that does not meet the condition cannot be obtained. Regarding the design, the alarm expression is a whole. The alarm expression has no difference from the PromQL query in the normal scene. If the series in the query result meets the duration, the alarm will be triggered. When the next query result does not contain a series, the alarm corresponding to the series will be recovered. Prometheus cannot deconstruct and interpret the alarm expression by itself, because some expressions themselves do not contain comparison relationships such as thresholds, for example: a and b and 123456789.

#### How to view the alarm records?

Prometheus itself does not have the concept of alarm records. Due to special reasons, it cannot fully support the feature of alarm records, but the related alarms and status can be viewed through the ALERTS or ALERTS\_FOR\_STATE metrics. In addition, all Prometheus alarms are currently delivered to the Tencent Cloud Observability Platform Alarm Service, and then sent to the user side through various notification channels of the Platform. The alarm records feature provided by the Tencent Cloud Observability Platform can make up for this defect to a certain extent. Due to conceptual differences in design, it cannot be used as a complete alternative but can be used for troubleshooting and other requirements. You can view records through the Alarm Records Feature.

#### Is the alarm repeat interval inconsistent with the configured duration?

Prometheus alarm repeat interval (repeat\_interval) is not set for a single alarm but for the entire group. When there is a new alarm or alarms are recovered in the group, the entire group will be notified, and then notifications will be sent at a repeated interval. Currently, TencentCloud Managed Service for Prometheus is grouped by a single alarm policy. For example, a single alarm policy may be for the restart of instance A/B/C. When A/B restarts, the alarm is triggered and then notifications will be sent at a repeated interval. After a period of time, C also restarts, and the interval will be interrupted and recovered. At this time, it will be notified to the user taking A/B/C as a Group. Currently, due to implementation limitations, the alarms delivered based on the Tencent Cloud Observability Platform cannot be

grouped and aggregated as a single message to notify users. If necessary, you can use a custom Alertmanager or Tencent Cloud Managed Service for Grafana.

#### Why doesn't the rate/irate function produce any data when the original metrics exist?

The rate/irate function requires at least two data points to calculate, so the time range calculated by rate/irate must cover at least two data points. Considering the possibility of data point loss due to network exception, the official recommendation for the time range is **four times the collection interval**.

#### Why does the rate/irate function calculate a very large outlier?

rate/irate functions can only be used for Counter type metrics, which are defined as strictly increasing numbers. Querying through Prometheus will process the problem of Counter being reset to 0, such as server restart. Normally, this will not affect the calculation results unless there is a problem of data point disorder. For example, the disorder of two second-level data points 9999 and 10000 results in an outlier of (10000+9999)-10000 = 9999 (normally it should be 1). The typical scene of this situation is as follows:

There are multiple collection components collecting the same metric and repeatedly reporting it to the same Prometheus instance, which may cause disorder problems, resulting in a large outlier calculated after the reset processing logic of Prometheus. Solution: If the collection component collects the same metric, it does not meet the expectations. It is recommended to troubleshoot and solve the problem of repeated collection. If the collection solution is designed to ensure the high availability of the collection component, different collection components need to add replica or other tags to distinguish all collected metrics.

Report directly to TencentCloud Managed Service for Prometheus through pushgateway and other methods and no timestamp information is included. TencentCloud Managed Service for Prometheus side will store the current time as the timestamp of the metric. If the network jitters, the arrival time of the data points will become disordered or the timestamp information granted by different processes/threads processing different data points will be disordered, resulting in calculation errors. The solution is to include the timestamp information when reporting. TencentCloud Managed Service for Prometheus built-in pushgateway reporting method is only a supplement to remotewrite and does not realize the complete feature. It is recommended to use Agent(remotewrite) for collection and reporting when it is not necessary.

#### Why is the interval of data points returned by the query different from that of capturing?

The interval of data points returned by querying through Prometheus is determined by the query parameter interval/step. Each data point is filled in and aligned strictly according to the interval/step, and there is no one-to-one correspondence with the capturing interval. If there is too much data loss or the capturing interval is too large, data will not be filled in, and the storage side will not store any information related to the capturing configuration. Users need to process their capturing configuration.

#### Why does the query return extra data for the last five minutes?

By default, Prometheus will fill in data for certain queries. Even if there is only one data point in the last five minutes, it may return multiple data points in the past five minutes (according to the step/interval parameters of the query). This is the default behavior of open source Prometheus and cannot be adjusted currently. Generally, it does not affect normal use.

#### How do time-related functions in PromQL handle local time zones?

All time in Prometheus is in UTC time. There are no exceptions, and the concept of time zone is not considered in the design, which may cause  $day_of_*$  series of functions to be inconvenient to use, and the official support for them can not be realized in the short term. The temporary solution for the domestic time zone can be as follows:  $day_of_week(vector(time() + 8*3600))$ .

#### Can Prometheus-related use limits be adjusted?

The use limits related to Prometheus are adjustable. In most cases, exceeding these use limits may lead to degradation of user experience and performance. Therefore, adjusting the limits upward cannot guarantee the query and write performance benchmarks before adjustments. The service level agreement may no longer apply. Users need to have certain psychological expectations for potential related risks.

#### From the chart, it seems that the alarm meets the duration but is not actually triggered?

The basic principle of Prometheus alarm detection is to use Instant Query to query data every minute, and trigger alarms if the alarm condition is (continuously) met. In some cases, Prometheus's Range Query will fill in the missing data, and the continuous timeline on the chart may be discontinuous under the execution logic of the alarm component. In addition, due to the characteristics of alarm scheduled inspection, the time when the alarm is detected may be slightly delayed. You can check the status of the alarm by querying the ALERTS or ALERTS\_FOR\_STATE metrics.

### Integration with TKE Cluster

Last updated : 2024-08-07 22:08:34

#### What should I do if I can't access the TKE cluster over the private network?

When installing an agent, you need to access TKE over the private network. If **private network access** is not enabled for the corresponding TKE cluster, the installation will fail. You can solve this problem by following the steps below:

1. Log in to the TKE console and select a container cluster in the corresponding region.

2. Enable **Private Network Access** under **Basic Information** > **Cluster APIServer Information**.

#### What should I do if the status of all kube-proxy collection targets is "Down"?

In TKE, the launch parameter <u>--metrics-bind-address</u> is not specified for kube-proxy, and the default listening address of the metrics service is 127.0.0.1; therefore, the agent cannot pull metrics by Pod IP. You can solve this problem by following the steps below:

1. Log in to the TKE console and select a container cluster in the corresponding region.

2. Go to **Basic Information** > **Cluster APIServer Information** > **Connecting to Kubernetes cluster through KubectI** to configure KubectI as instructed.

3. Run kubectl edit ds kube-proxy -n kube-system and add the launch parameter --metricsbind-address=0.0.0.0 in spec.template.spec.containers.args .

## What should I do if the status of all component collection targets on the master node in a dedicated TKE cluster is "Down"?

The inbound rule of the default security group of the master node in a dedicated TKE cluster does not allow access to the metrics ports of some components. You can solve this problem by following the steps below:

1. Log in to the Security Group console and select the corresponding region.

2. Enter tke-master-security-for-<tke cluster id> in the security group search box. For example, if the cluster ID is cls-xxx , then enter tke-master-security-for-cls-xxx .

3. Click the returned security group ID to enter the Edit Inbound Rule window.

4. The protocol port column of the rule to be edited should include TCP: 60001, 60002 . Select the rules one by one and add ports 10249, 10252, 10251, 9100, and 9153 for the following purposes respectively:

10249: kube-proxy metrics port

10252: kube-controller-manager metrics port

10251: kube-scheduler metrics port

9100: node-exporter metrics port

9153: core-dns metrics port

## **Product Consulting**

Last updated : 2024-08-07 22:08:40

#### Does TMP support customizing the reported data?

Yes. TMP supports customizing the reported metric monitoring data in multiple languages and displaying such data in the integrated Grafana dashboards.

#### How does TMP collect monitoring data?

TMP can pull data through the Prometheus agent or write data through Pushgateway. It is fully compatible with opensource Prometheus in terms of collection method.

#### Is there a separate exporter for each instance?

Each MySQL or Kafka instance has one exporter, while multiple Redis instances can share the same exporter.

#### What is the difference between TMP and a self-built Prometheus service?

TMP is fully compatible with the open-source Prometheus ecosystem and is integrated with CM to help you quickly set up a monitoring system (for custom monitoring, component monitoring, basic monitoring, etc.). It supports Grafana, with preset common monitoring dashboards. It also supports diverse exporters, with preset common alarm templates. It well solves various problems with open-source Prometheus, such as complicated HA setup, poor performance, and labor-intensive Ops.

#### Which Tencent Cloud products does TMP support?

TMP supports CVM, TencentDB for MongoDB, TencentDB for MySQL, TencentDB for PostgreSQL, TencentDB for Redis, ES, and TKE. For more information, see Integration Center.

#### Can I integrate other data sources on the Grafana page of TMP?

No. We recommend you integrate them by using Grafana.

#### Can I pull monitoring data in Prometheus pull mode?

No, you can only pull data through APIs. For detailed directions, see Querying Monitoring Data.

#### What is the series storage limit on pay-as-you-go TMP instances?

Each instance can have up to 4.5 million series. If you need to adjust it, submit a ticket for application.

#### Does TMP support integration with EKS?

Yes. Pay-as-you-go TMP instances support general clusters, EKS elastic clusters, edge clusters, and registered clusters.

## Use and Technology

Last updated : 2024-08-07 22:08:46

#### How do I migrate from self-built Prometheus to TMP?

In the configuration file of self-built Prometheus, add a remote write configuration pointing to TMP for migration. For more information, see Migration from Self-Built Prometheus.

#### Can I batch import/export dashboards into/from Grafana?

You can do so through APIs as instructed in HTTP API reference.

#### What should I do if one exporter has too much data?

You can use Prometheus in this scenario, but we recommend you not expose too many metrics in the exporter, as the exporter itself and the Prometheus agent consume a lot of memory. Only expose key metrics, and do not use parameters such as user ID and URL in metric labels.

#### How do I use TMP to monitor clusters in two different VPCs?

You can interconnect the VPCs of the two clusters through CCN and then integrate the clusters into Prometheus.
 Install the agent in one of the clusters, and then use the Nginx proxy to remotely write the target address to the TMP instance.

#### How do I integrate native container services with TMP?

You can't integrate them directly. However, you can create data and add a remote write configuration pointing to TMP in the configuration file of your self-built Prometheus. For detailed directions, see Migration from Self-Built Prometheus.

## How do I set the parameters of the alarm cycle when creating an alarm rule through Prometheus APIs?

Add key=\_interval\_ value=1m/5m/10m/15m/30m/60m to the new Labels parameter.

#### Will data get lost during instance rebooting?

Data will not get lost during instance rebooting, as it is stored in TencentDB for CTSDB. However, because data cannot be reported normally during rebooting, data breakpoints may occur.

#### Is it normal to see a rise in generated data after the TMP instance is rebooted?

Retries will be performed after instance rebooting, so it is normal if the data volume seems to fluctuate greatly in a short period of time.

# Can I batch add MySQL instances when integrating MySQL with Prometheus in the integration center?

No. You can only add instances one by one.

#### Where can I view the security group of the EKS cluster created by Prometheus?

Go to **Security** > **Security Group** in the VPC console and search by the current TMP instance ID.

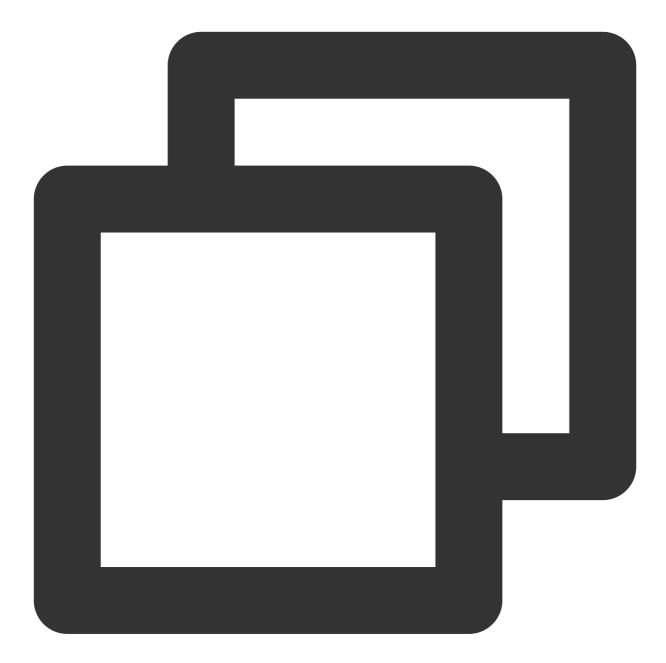
#### Can I configure multiple scrape tasks for a PodMonitor?

Yes, but you should ensure that the tasks have the same port name and label.

## (TMP Pushgateway usage) If the clients of multiple instances push the same job with different label metrics, metrics will be overwritten. What should I do?

You can use groupingkey to solve this problem. Below is the sample code:





if err := push.New("http://\$IP:\$PORT", "db\_backup").
 BasicAuth("\$APPID", "\$TOKEN").
 Collector(completionTime).
 Grouping("instance", "\$INSTANCE").