


TencentCloud Managed Service for Prometheus Integration Guide Product Documentation



Copyright Notice

©2013-2025 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

Contents

Integration Guide

- Scrape Configuration Description

- Custom Monitoring

- EMR Integration

 - EMR Metric Collection Component of Prometheus

- 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

 - Apache Exporter Integration

- Health Check

- Instructions for Installing Components in the TKE Cluster

- Cloud Monitoring

- Read Cloud-Hosted Prometheus Instance Data via Remote Read

- Agent Self-Service Access

- Pushgateway Integration

- Security Group Open Description

Integration Guide

Scrape Configuration Description

Last updated : 2024-01-29 15:55:08

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
[ 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_config> ... ]

# 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>	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/document/api/213/15692#.E5.9C.B0.E5.9F.9F.E5.88.97.E8
.A1.A8.
region: <string>

# Custom endpoint.
[ endpoint: <string> ]

# Credential information for accessing TencentCloud API. If it is not set, the
values of the `TENCENT_CLOUD_SECRET_ID` and `TENCENT_CLOUD_SECRET_KEY`
environment variables will be used.
# Leave it empty if you use a CVM scrape task in **Integration Center** for
configuration.
[ 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, visit https://www.tencentcloud.com/document/product/213/33258.
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 not specified, the value of job in the corresponding metric label will be
`<namespace>/<name>`
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
scrape task
spec:
  # Enter the target Pod label. PodMonitor will use the corresponding value as
the job label value
  # If Pod YAML configuration is to be viewed, use the value in
`pod.metadata.labels`
  # If `Deployment/Daemonset/Statefulset` is to be viewed, use
`spec.template.metadata.labels`
  [ 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 multiple Endpoints
  podMetricsEndpoints:
    [ - <endpoint_config> ... ] # For more information, see the endpoint
description below
  # Select the namespace where the Pod to be monitored resides. If it is not
specified, all namespaces will be selected
  [ 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 information, see [LabelSelector v1 meta](https://v1-
17.docs.kubernetes.io/docs/reference/generated/kubernetes-
api/v1.17/#labelselector-v1-meta)
  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
Prometheus exporter in the Pod YAML configuration file
      path: /metrics # Enter the value of the corresponding path of the
Prometheus exporter. If it is not specified, it will be `/metrics` by default
      relabelings:
        - action: replace
          sourceLabels:
            - instance
          regex: (.*?)
          targetLabel: instance
          replacement: 'crs-xxxxxx' # Change it to the corresponding Redis
instance ID
        - 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
resides
        matchNames:
          - redis-test
      selector: # Enter the label value of the Pod to be monitored to locate
the target Pod
        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 not specified, the value of job in the corresponding metric label will be
the Service name
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
scrape task
spec:
  # Enter the target Pod label (metadata/labels). ServiceMonitor will use the
corresponding value as the job label value
  [ 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 multiple Endpoints
  endpoints:
    [ - <endpoint_config> ... ] # For more information, see the endpoint
description below
  # Select the namespace where the Pod to be monitored resides. If it is not
specified, all namespaces will be selected
  [ 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 information, see [LabelSelector v1 meta](https://v1-
17.docs.kubernetes.io/docs/reference/generated/kubernetes-
api/v1.17/#labelselector-v1-meta)
  selector:
    [ matchExpressions: array ]
      [ example: - {key: tier, operator: In, values: [cache]} ]
    [ matchLabels: object ]
      [ example: k8s-app: redis-exporter ]

```

Sample

```
apiVersion: monitoring.coreos.com/v1
```

```

kind: ServiceMonitor
metadata:
  name: go-demo      # Enter a unique name
  namespace: cm-prometheus # The namespace is fixed. Do not change it
spec:
  endpoints:
    - interval: 30s
      # Enter the name of the corresponding port of the Prometheus exporter in
the Service YAML configuration file
      port: 8080-8080-tcp
      # Enter the value of the corresponding path of the Prometheus exporter.
If it is not specified, it will be `/metrics` by default
      path: /metrics
      relabelings:
        # ** There must be a label named `application`. Here, suppose that K8s
has a label named `app`
        # 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
Service
      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: 80. Corresponding values are as follows:
# ServiceMonitor: `Service>spec/ports/name`
# PodMonitor description:
#   If Pod YAML configuration is to be viewed, use the value in
`pod.spec.containers.ports.name`
# If `Deployment/Daemonset/Statefulset` is to be viewed, use
`spec.template.spec.containers.ports.name`
[ 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
configure the corresponding TLS parameters
[ tlsConfig: TLSConfig ]
# Read the value of the bearer token through the corresponding file and add it
to the header of the scrape task
[ bearerTokenFile: string ]
# You can use the corresponding K8s secret key to read the bearer token. Note
that the secret namespace must be the same with that of the
PodMonitor/ServiceMonitor
[ bearerTokenSecret: string ]
# Solve the conflict between the scraped label and the label added to
Prometheus on the backend
# true: Retain the scraped label and ignore the label conflicting with
Prometheus on the backend
# false: Add `exported_<original-label>` before the scraped label to add the
label on the Prometheus backend
[ 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 value for `username/password`. Note that the secret namespace must be the
same as that of the PodMonitor/ServiceMonitor
[ basicAuth: BasicAuth ]
# Use a proxy service to scrape metrics on the target and enter the
corresponding proxy service address
[ proxyUrl: string ]
# After scraping the data, change the label on the target through the
relabeling mechanism and run multiple relabeling rules in sequence
# 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 to change the label value and run multiple relabeling rules in
sequence
# 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 taken values are concatenated and separated with the symbol defined in
`separator`
# The corresponding configuration item for PodMonitor/ServiceMonitor is
`sourceLabels`
[ source_labels: '[' <labelname> [, ...] ']' ]
# Define the separator symbol for concatenating the labels to be relabeled.
Default value: `;`
[ separator: <string> | default = ; ]

# If `action` is `replace` or `hashmod`, you need to use the `target_label` to
specify the corresponding label name
# The corresponding configuration item for PodMonitor/ServiceMonitor is
`targetLabel`
[ 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 is used if `action` is `hashmod`
[ modulus: <int> ]

# If `action` is `replace`, use `replacement` to define the expression to be
replaced after regular match. You can replace it based on regex
[ replacement: <string> | default = $1 ]

# Perform an action based on the value matched by the regex. Valid values of
`action` are as follows (the default value is `replace`):
# replace: Replace the matched value with that defined in `replacement` if the
regex has any match and use `target_label` to set the value and add the
corresponding label
# 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 value specified by `modulus` and add a label with the name specified by
`target_label`
# labelmap: Use `replacement` to replace the corresponding label name if the
regex has any match
# 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 : 2024-01-29 15:55:07

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](#)

[C](#)

[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, `http_request_total` 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 programming 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 `rate()` 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 duration
    defer opsProcessCost.WithLabelValues("success").Observe((float64)(time.Now().UnixNano() - start) / 1e9)

    // 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_count

```

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())
}

```



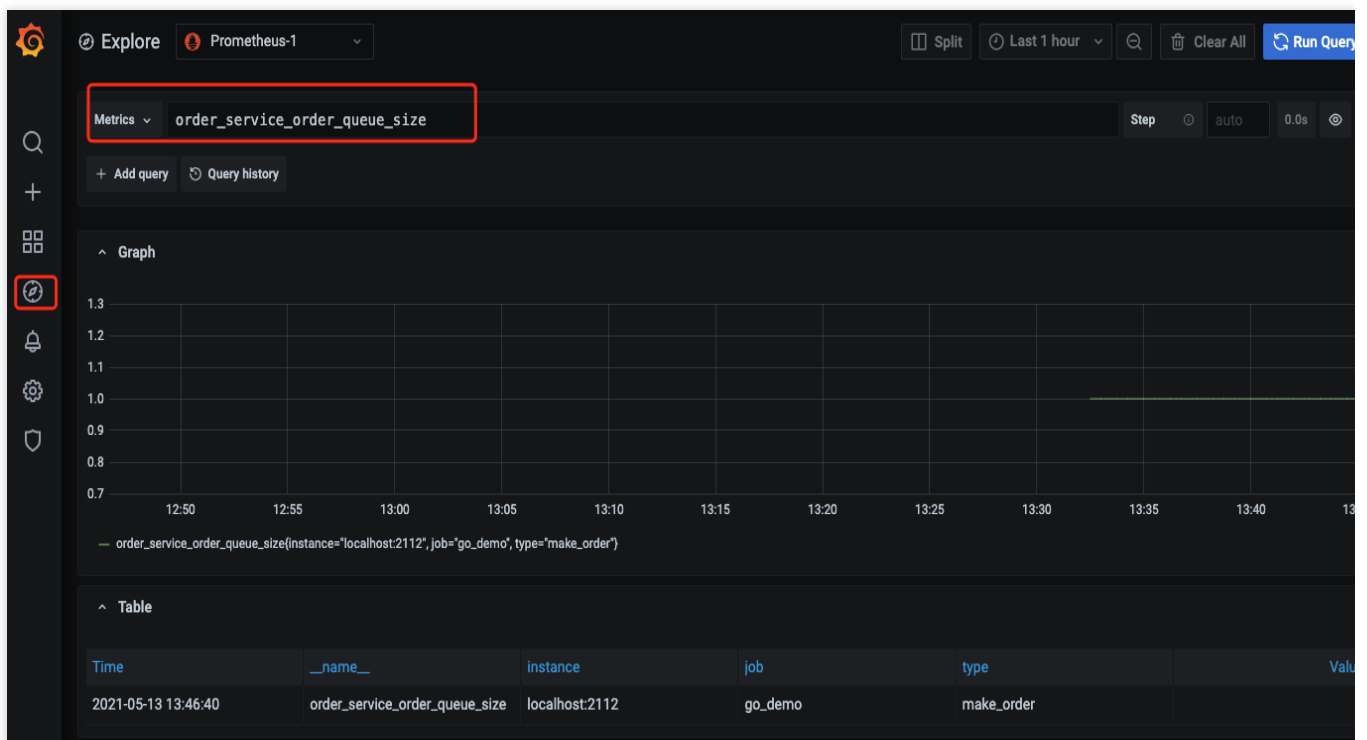
```
// Business code  
}
```

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

EMR Metric Collection Component of Prometheus

Last updated : 2024-10-29 11:44:46

Overview

During the use of [Tencent Cloud Elastic MapReduce](#) (EMR), you need to report EMR monitoring metrics to TencentCloud Managed Service for Prometheus (TMP). This document will guide you on how to quickly collect EMR monitoring metrics.

Prerequisites

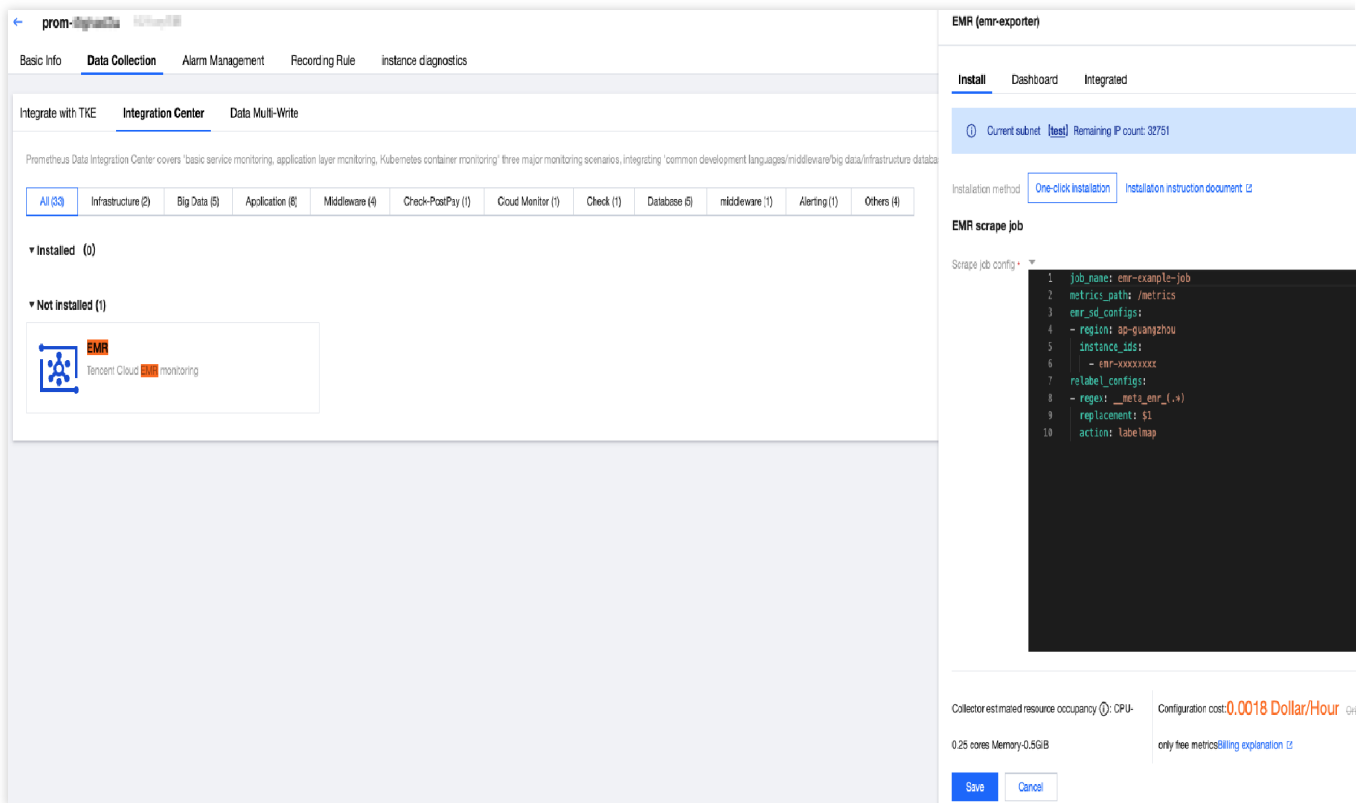
You have used EMR and enabled the Prometheus Exporter feature.

Use the same region and Virtual Private Cloud (VPC) as EMR to purchase a Tencent Cloud [Prometheus monitoring instance](#). You can check the [regions supported by TMP](#).

Directions

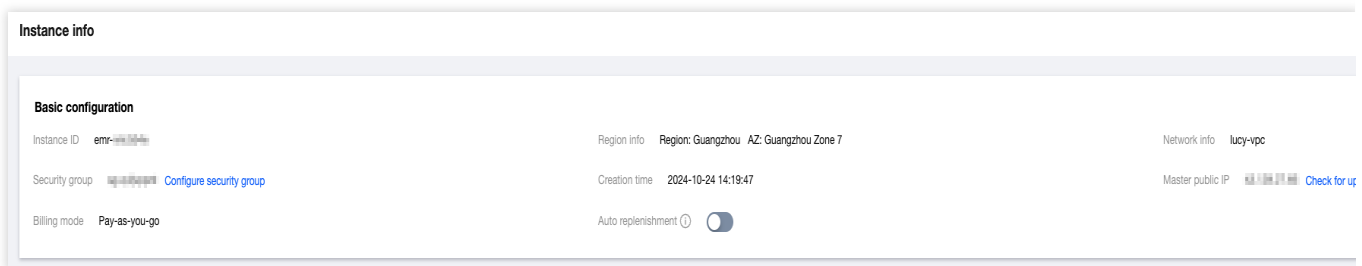
1. log in to [TCOP](#).
2. In the left menu bar, click **Managed Service for Prometheus**.
3. In the Prometheus instance list, select the corresponding Prometheus instance.
4. Enter the instance details page, click **Data Collection > Integration Center**.
5. Search for **EMR** in the integration center, and click it to pop up an installation window. Then, confirm the information and click **Save**.

Search for the required CAM policy as needed, and click to complete policy association.



6. Log in to the [EMR Console](#), click Cluster ID/Name > **Instance info**, and obtain the region where the EMR cluster is located and the EMR instance ID of the cluster.

Search for the required CAM policy as needed, and click to complete policy association.



7. Fill in the task configuration (in YAML format) in EMR. Then, fill in the task name, region where the EMR cluster is located, and EMR instance ID in the red box of the following figure.

Note:

For the format of the region, see the region description in [Service Region](#), for example, `ap-guangzhou`.

Multiple instance IDs are supported.

For the relabel_configs configuration, see [Capture Configuration Instructions](#).

Search for the required CAM policy as needed, and click to complete policy association.

EMR (emr-exporter)

[Install](#)[Dashboard](#)[Integrated](#)

① Current subnet **[lucy-subnet-4]** Remaining IP count: 190

Installation method

[One-click installation](#)[Installation instruction document](#)

EMR scrape job

Scrape job config

```
1 job_name: emr-example-job
2 metrics_path: /metrics
3 emr_sd_configs:
4 - region: ap-guangzhou
5   instance_ids:
6     - emr-xxxxxxx
7 relabel_configs:
8 - regex: __meta_emr__(.*)
9   replacement: $1
10  action: labelmap
```

Collector estimated resource occupancy ①: CPU-

0.25 cores Memory-0.5GiB

Configuration cost: **0.0018 Dollar/Hour** Original price: 0.0065 Dollar/Hour No charge for colleconly free metrics [Billing explanation](#)[Save](#)[Cancel](#)

Supported Metrics

TMP supports all EMR metrics. For a detailed metric list, see [EMR Cluster Monitoring Metrics](#).

Java Application Integration

Spring Boot Integration

Last updated : 2024-01-29 15:29:42

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>
<dependency>
  <groupId>io.micrometer</groupId>
  <artifactId>micrometer-registry-prometheus</artifactId>
</dependency>
```

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

The maximum length of 40 characters, 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

Label =

[New variable](#)

Can only contain letters, numbers and separators ("_", "-", ".", ":"), and must start and end with letters and numbers

Namespaces

type ☒ Deployment (Scalable Deployment Pod)
☐ DaemonSet (Run Pod on each host)
☐ StatefulSet (operating Pod with stateful set)
☐ CronJob (run regularly according to Cron's plan)
☐ Job (single task)

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 mounted to the specified path of the container. [Guidelines for use](#)

Instance content container

name

Up to 63 characters, can only contain lowercase letters, numbers and separators ("-"), and cannot start or end with separators

Mirror image [Select mirror](#)

Mirror version (Tag)

Image pull strategy

Access Settings (Service)

Service ☒ Enable

Service access method ☒ Access only within the cluster ☐ Host port access ☐ Public network LB access ☐ Intranet LB access [show to choose](#)

That is, the ClusterIP type will provide an entry that can be accessed by other services or containers in the cluster. It supports the TCP/UDP protocol. Database services such as Mysql can be accessed with cluster to ensure service network isolation.

☐ Headless Service [?](#) (Headless Service only supports selection during creation, and does not support changing the access method after creation)

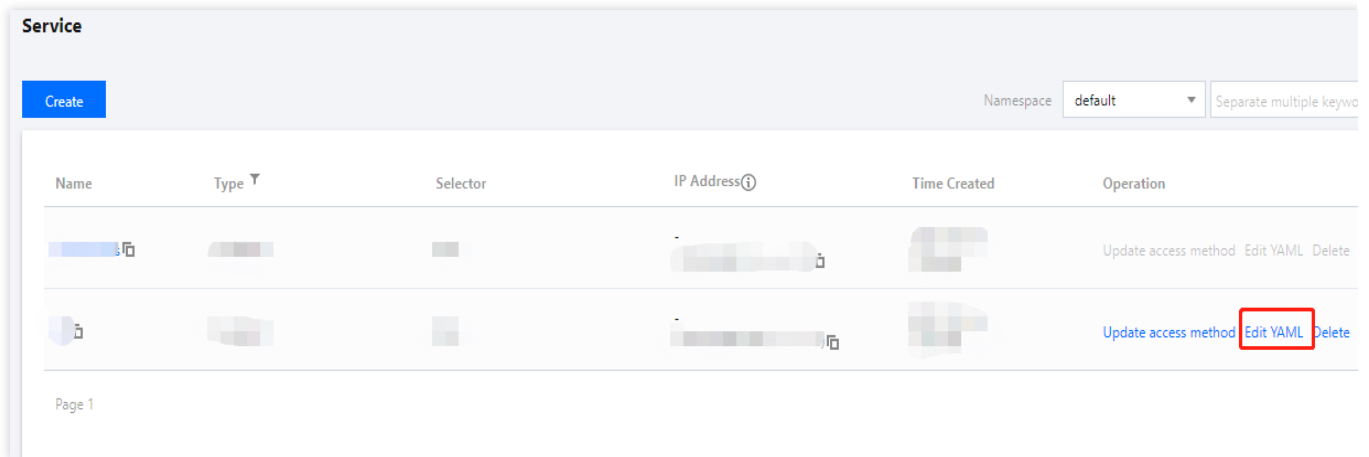
Port Mapping

protocol ^①	Container port ^②	Service port ^③
TCP	The port that the application in the	It is recommended to be consistent

[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:



```

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**.

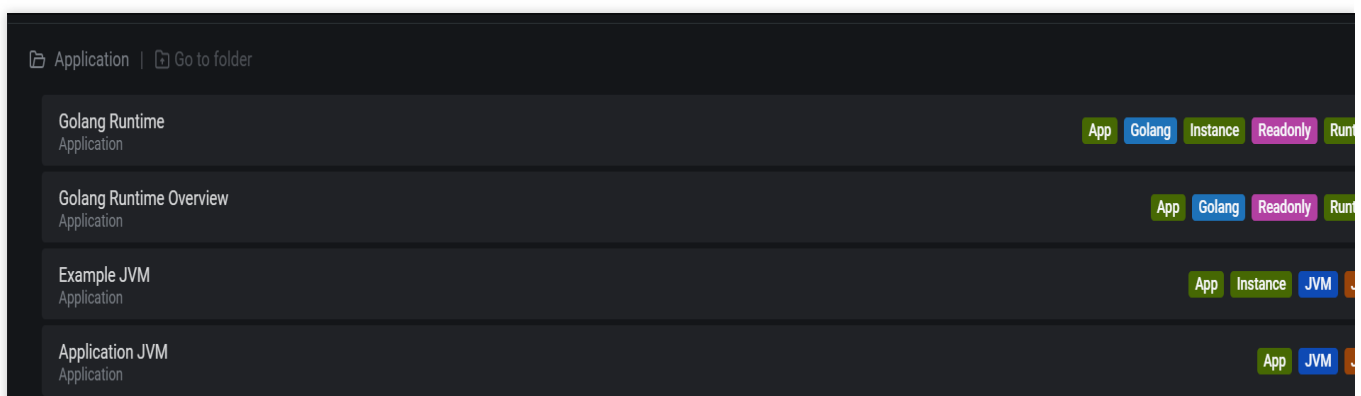
Spring MVC application: monitoring data of MVC status, such as the request latency, number of requests, success rate, and exception distribution.

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 : 2024-01-29 15:55:08

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 `System.gc()` 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]
```

Below is a sample:

```
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

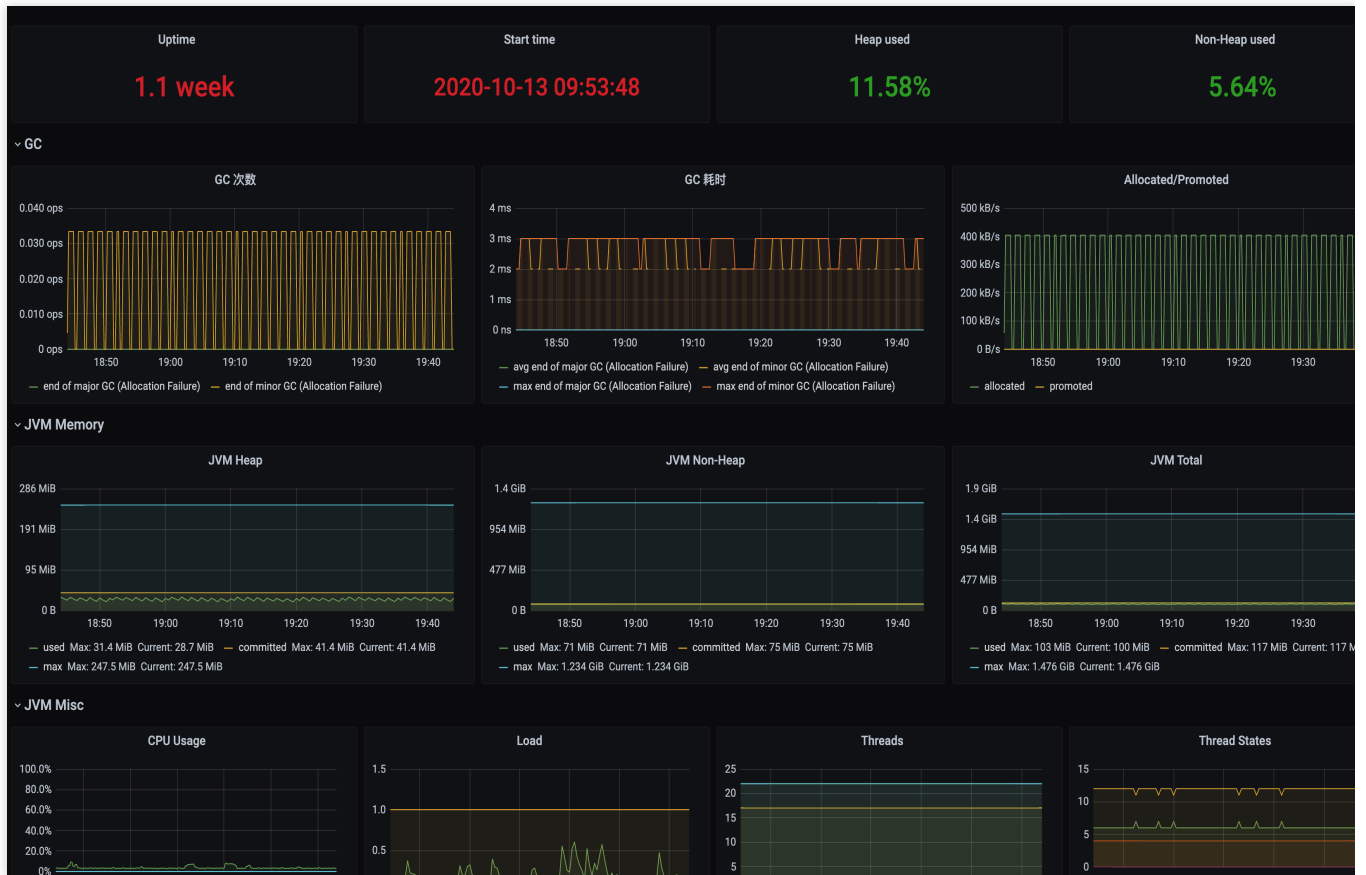
1. Log in to the [TMP console](#) and select the target TMP instance to enter the management page.
2. Click a **cluster ID** in the TKE cluster list to enter the **Integrate with TKE** page.
3. In **Scrape Configuration**, add **Pod Monitor** to define a Prometheus scrape task. Below is a sample YAML 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.

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 使用率	Heap 大小	最大线程数
	13.04 hour	0.02%	0	0 s	1.03%	38.78 MiB	6
	26.89 s	0.00%	-	-	0.20%	7.58 MiB	6



Go Application Integration

Last updated : 2025-03-17 10:55:18

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.

```
apiVersion: monitoring.coreos.com/v1
kind: ServiceMonitor
metadata:
  name: go-demo      # Enter a unique name
  namespace: cm-prometheus # The namespace is fixed. Do not change it
spec:
  endpoints:
  - interval: 30s
    # Enter the name of the corresponding port of the Prometheus exporter in the
    port: 2112
    # Enter the value of the corresponding path of the Prometheus exporter. If it
    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-of-the-box integration features of TMP. For more advanced usage, please see [ServiceMonitor](#) or [PodMonitor](#).

Viewing monitoring information

1. In the [TMP instance](#) list, find the corresponding TMP instance, click



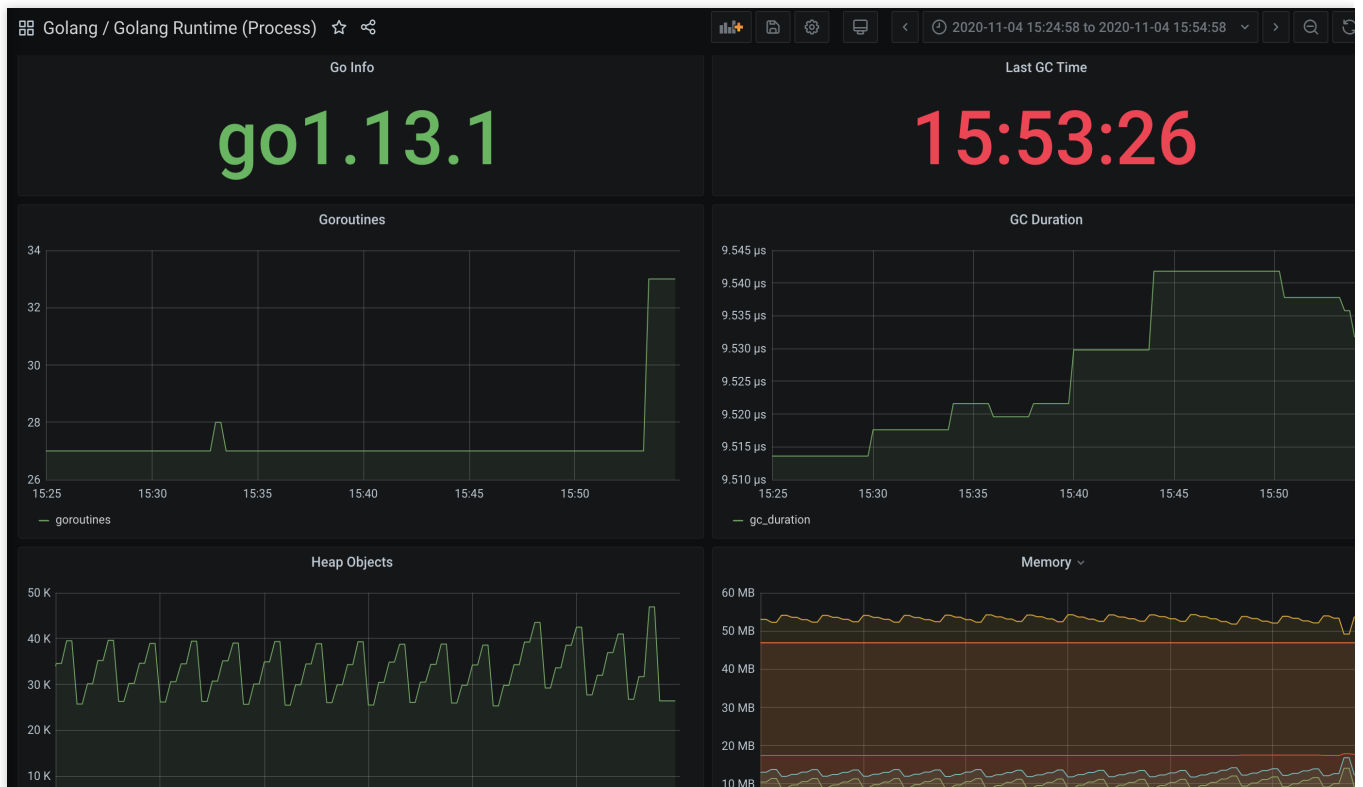
on the right of the instance ID to open your Grafana page, and enter your account and password to access the Grafana visual dashboard operation section.

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.

Instance ^	Uptime	CPU Usage	Memory(RSS) v	Threads	Goroutines	GC Duration	Heap Objects
10902	5.20 day	0.00	62.07 MiB	18	33	20.52 μs	73.30 K
900	5.20 day	0.38	2.02 GiB	39	57	188.75 μs	14.56 Mil
9002	5.20 day	0.00	59.36 MiB	16	33	21.10 μs	46.92 K
900	5.20 day	0.34	2.25 GiB	39	56	492.40 μs	15.71 Mil



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 may 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 : 2024-01-29 15:55:07

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.
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 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.
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 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
```

Note:

The Elasticsearch connection string is in the format of `<proto>://<user>:<password>@<host>:<port>` , 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: {}
```

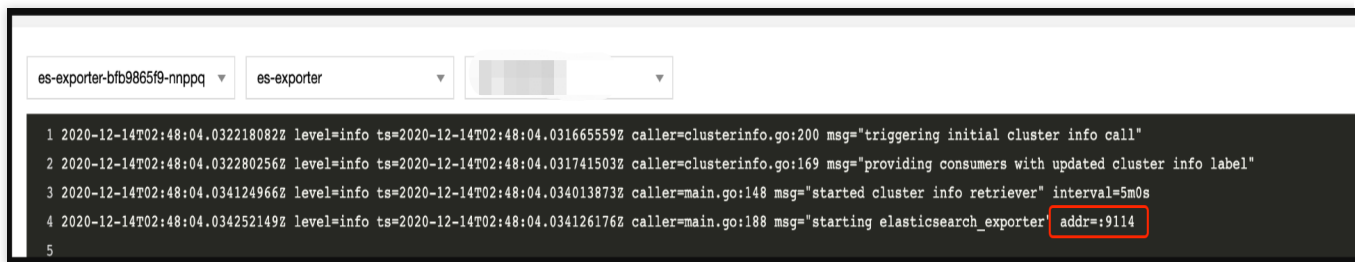
```
terminationGracePeriodSeconds: 30
```

Note:

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:



```
es-exporter-bfb9865f9-nppq  es-exporter  [REDACTED]
1 2020-12-14T02:48:04.032218082Z level=info ts=2020-12-14T02:48:04.031665559Z caller=clusterinfo.go:200 msg="triggering initial cluster info call"
2 2020-12-14T02:48:04.032280256Z level=info ts=2020-12-14T02:48:04.031741503Z caller=clusterinfo.go:169 msg="providing consumers with updated cluster info label"
3 2020-12-14T02:48:04.034124966Z level=info ts=2020-12-14T02:48:04.034013873Z caller=main.go:148 msg="started cluster info retriever" interval=5m0s
4 2020-12-14T02:48:04.034252149Z level=info ts=2020-12-14T02:48:04.034126176Z caller=main.go:188 msg="starting elasticsearch_exporter" addr:9114
5
```

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",cluster="demo"} 2.0102643e+07
elasticsearch_breakers_estimated_size_bytes{breaker="accounting",cluster="demo"} 1.9926654e+07
elasticsearch_breakers_estimated_size_bytes{breaker="accounting",cluster="demo"} 1.9685163e+07
elasticsearch_breakers_estimated_size_bytes{breaker="fielddata",cluster="demo"} 1.9926654e+07
elasticsearch_breakers_estimated_size_bytes{breaker="fielddata",cluster="demo"} 1.9685163e+07
elasticsearch_breakers_estimated_size_bytes{breaker="fielddata",cluster="demo"} 1.9685163e+07
elasticsearch_breakers_estimated_size_bytes{breaker="in_flight_requests",cluster="demo"} 0
elasticsearch_breakers_estimated_size_bytes{breaker="in_flight_requests",cluster="demo"} 1167
elasticsearch_breakers_estimated_size_bytes{breaker="in_flight_requests",cluster="demo"} 1167
elasticsearch_breakers_estimated_size_bytes{breaker="parent",cluster="demo"} 2.0102643e+07
elasticsearch_breakers_estimated_size_bytes{breaker="parent",cluster="demo"} 2.0102643e+07
```

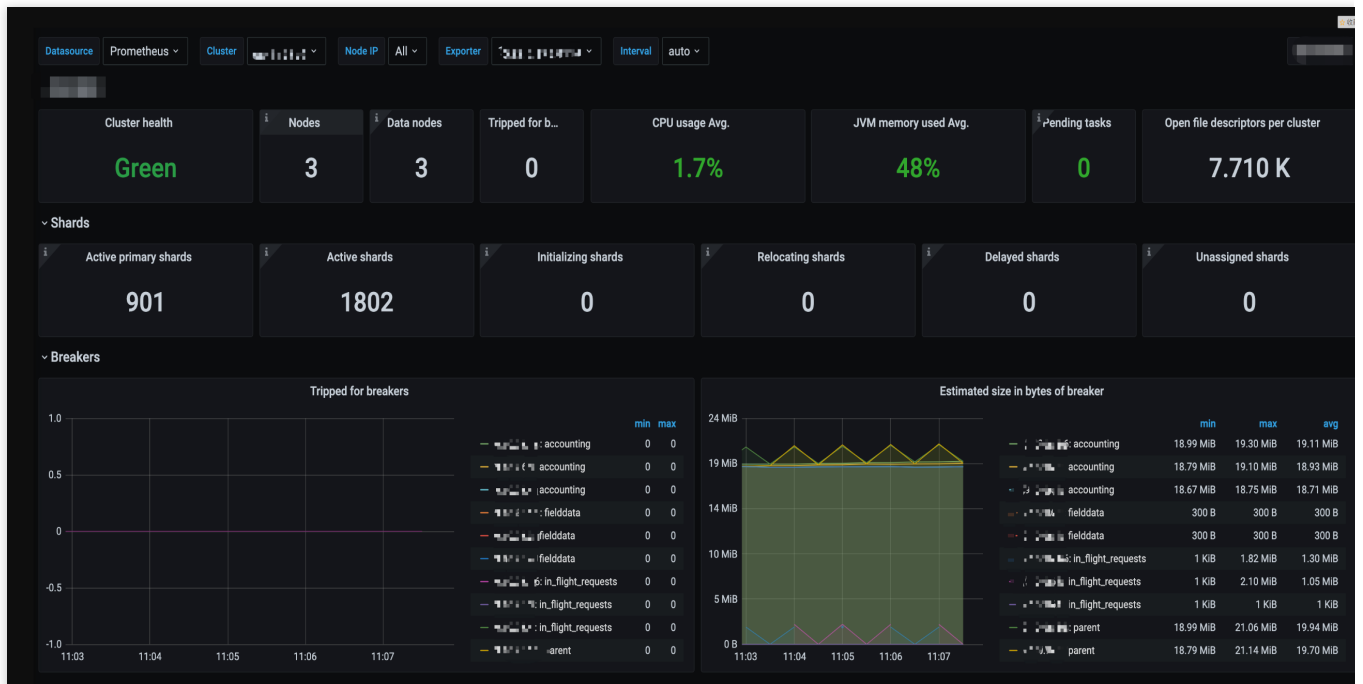
Adding scrape task

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 : 2024-01-29 15:55:07

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: danielqsj/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

```

Note:

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 Prometheus

```

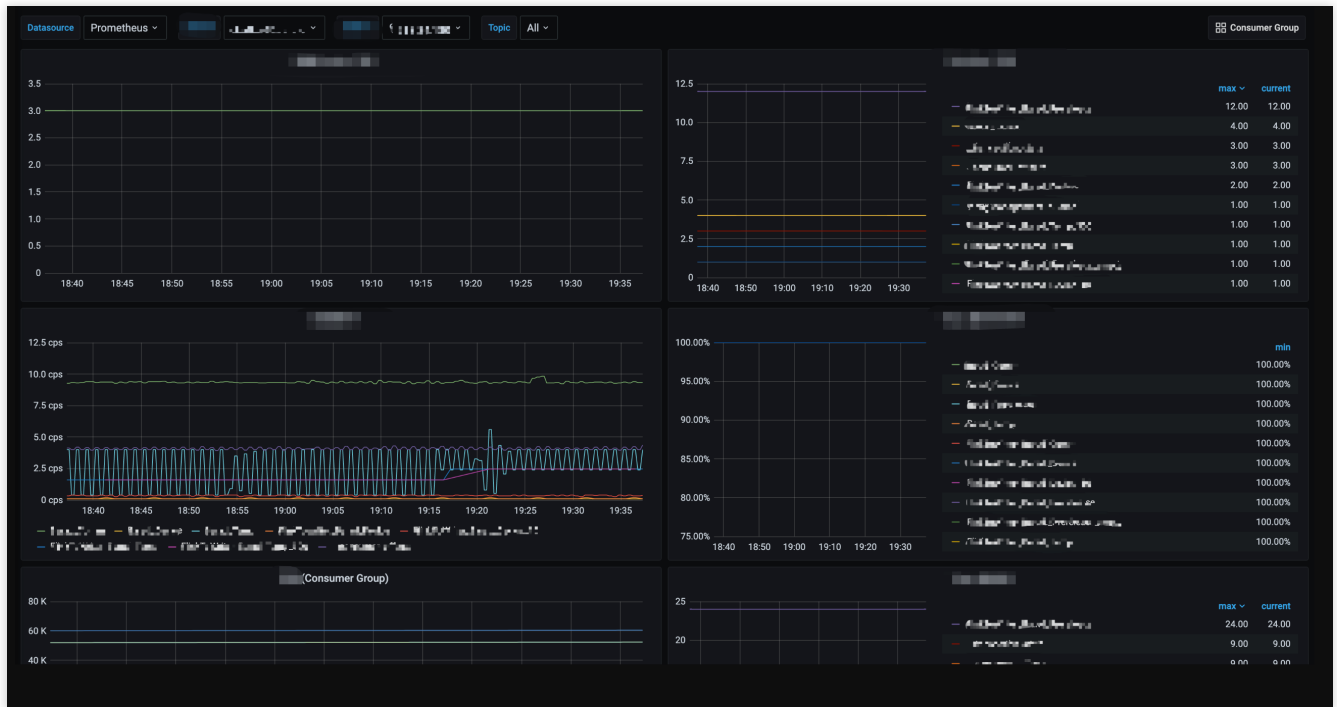
```
path: /metrics # Enter the value of the corresponding path of the Prometheus
relabelings:
- action: replace
  sourceLabels:
  - instance
  regex: (.*)
  targetLabel: instance
  replacement: 'ckafka-xxxxxx' # Change it to the corresponding Kafka instance
- action: replace
  sourceLabels:
  - instance
  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 target
matchLabels:
  k8s-app: kafka-exporter
```

Note:

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:



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 : 2024-01-29 15:55:08

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.

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 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

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:
          - --collect.database      # Enable the collection of `Database` metric
          - --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

```

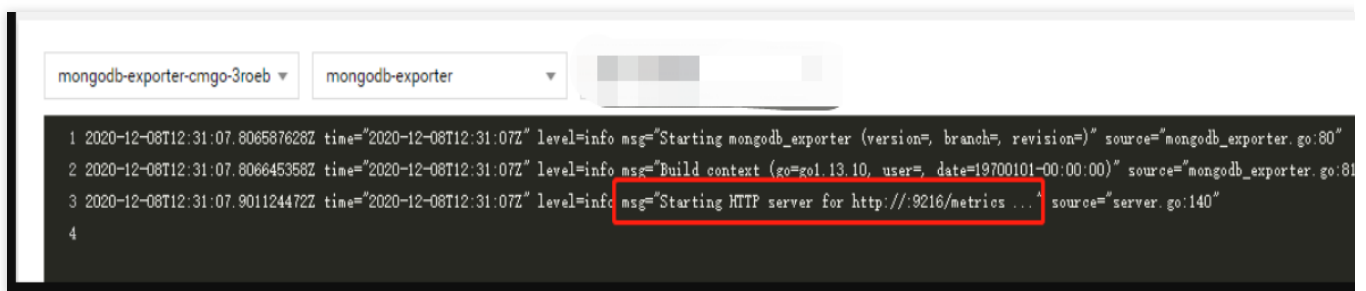
```
terminationMessagePath: /dev/termination-log
terminationMessagePolicy: File
dnsPolicy: ClusterFirst
imagePullSecrets:
  - name: qcloudregistrykey
restartPolicy: Always
schedulerName: default-scheduler
securityContext: { }
terminationGracePeriodSeconds: 30
```

Note:

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:

```
# TYPE mongodb_connections gauge
mongodb_connections{state="available"} 9971
mongodb_connections{state="current"} 29
# HELP mongodb_connections_metrics_created_total totalCreated provides a count of all incoming connections created to the server. This number inc
# TYPE mongodb_connections_metrics_created_total counter
mongodb_connections_metrics_created_total 1.543107e+06
# HELP mongodb_connpoolstats_connection_sync Corresponds to the total number of client connections to mongo.
# TYPE mongodb_connpoolstats_connection_sync gauge
mongodb_connpoolstats_connection_sync 6
# HELP mongodb_connpoolstats_connections_available Corresponds to the total number of client connections to mongo that are currently available.
# TYPE mongodb_connpoolstats_connections_available gauge
mongodb_connpoolstats_connections_available 13
# HELP mongodb_connpoolstats_connections_created_total Corresponds to the total number of client connections to mongo created since instance star
# TYPE mongodb_connpoolstats_connections_created_total counter
mongodb_connpoolstats_connections_created_total 17
# HELP mongodb_connpoolstats_connections_in_use Corresponds to the total number of client connections to mongo currently in use.
# TYPE mongodb_connpoolstats_connections_in_use gauge
mongodb_connpoolstats_connections_in_use 0
# HELP mongodb_connpoolstats_connections_scoped_sync Corresponds to the number of active and stored outgoing scoped synchronous connections from
lica set.
# TYPE mongodb_connpoolstats_connections_scoped_sync gauge
mongodb_connpoolstats_connections_scoped_sync 0
# HELP mongodb_exporter_build_info A metric with a constant '1' value labeled by version, revision, branch, and goversion from which mongodb_expo
# TYPE mongodb_exporter_build_info gauge
mongodb_exporter_build_info(branch="",goversion="go1.13.10",revision="",version="") 1
# HELP mongodb_exporter_last_scrape_duration_seconds Duration of the last scrape of metrics from MongoDB.
# TYPE mongodb_exporter_last_scrape_duration_seconds gauge
mongodb_exporter_last_scrape_duration_seconds 0.026315909
# HELP mongodb_exporter_last_scrape_error Whether the last scrape of metrics from MongoDB resulted in an error (1 for error, 0 for success).
# TYPE mongodb_exporter_last_scrape_error gauge
mongodb_exporter_last_scrape_error 0
```

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:
        - action: replace
          sourceLabels:
            - instance
          regex: (.*)
          targetLabel: instance
          replacement: 'cmgo-xxxxxxx' # Change it to the corresponding MongoDB insta
      namespaceSelector: # Select the namespace where the Pod to be monitored reside
      matchNames:
        - mongodb-test
      selector: # Enter the label value of the Pod to be monitored to locate the targ
```

```
matchLabels:
  k8s-app: mongodb-exporter
```

Note:

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

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 monitoring data as shown below:

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.

**Note:**

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

Integrating with alert feature

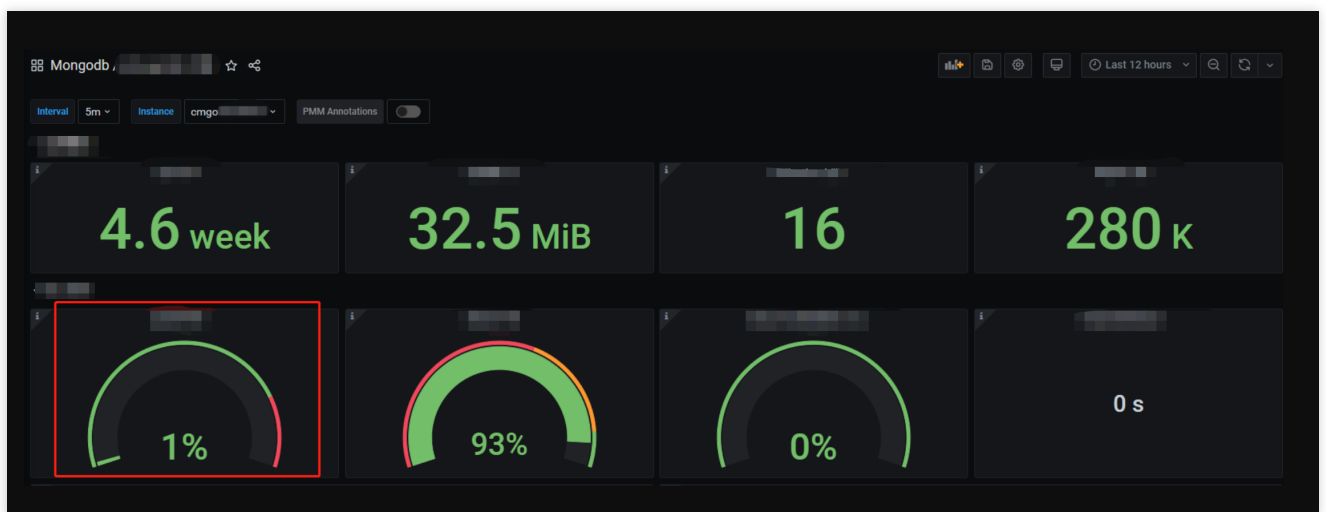
1. Log in to the [TMP console](#) and select the target TMP instance to enter the management page.

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 : 2024-01-29 15:55:07

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](#) > [Deploying PostgreSQL exporter](#).

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

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
```

Note:

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

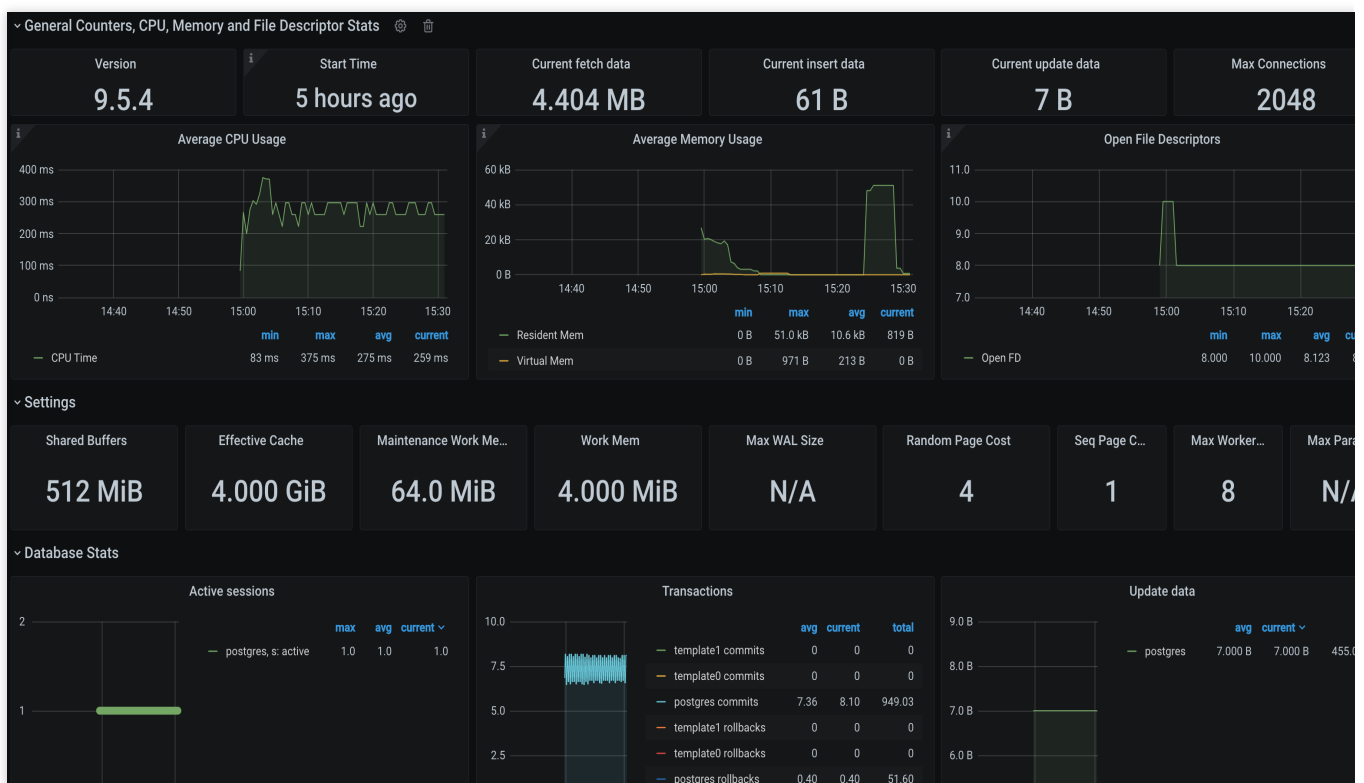


on the right of the instance ID to open your Grafana page, and enter your account and password to access the Grafana visual dashboard operation section.

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.



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 : 2024-10-30 16:02:09

Overview

Nginx exposes some monitoring metrics through the `stub_status` page. Nginx Prometheus Exporter collects metrics from a single Nginx instance, converts them into Prometheus-compatible monitoring data, and exposes such data to the Prometheus service through the HTTP protocol for collection. Through Exporter, key monitoring metrics can be reported for exception alarming and dashboard display.

Prerequisites

Enabling the NGINX `stub_status` Feature

Note:

1. The following example is for Nginx deployed in Tencent Kubernetes Engine (TKE). For other deployment methods, adjust the login and configuration modification methods accordingly.
2. For TKE-related operations, see the [TKE](#) documentation.

Because Nginx Prometheus Exporter monitors Nginx through the `stub_status` module of Nginx, you need to ensure that the `stub_status` module has been enabled for the Nginx service. The steps for enabling this module are as follows:

1. Log in to the [TKE console](#).
2. Click **Cluster** in the left sidebar, find the cluster where the Nginx server is located, enter the cluster, and find the Nginx server.
3. Log in to the Nginx server and execute the following command to check whether this module has been enabled for Nginx:

```
nginx -V 2>&1 | grep -o with-http_stub_status_module
```

If `with-http_stub_status_module` is output in the terminal, the `stub_status` module has been enabled for Nginx.

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. The example is as follows:

```
./configure \\  
... \\  
## Command required to compile nginx previously.  
--with-http_stub_status_module  
make  
sudo make install
```

4. If the Nginx service-related ConfigMap is not added, you can log in to the Nginx server, copy the default.conf configuration information in the configuration directory (`/etc/nginx/conf.d` for the official image), create a ConfigMap, and add the configuration information to the ConfigMap. For the ConfigMap operation guide, see [ConfigMap Management](#).

5. After you confirm that the stub_status module is enabled, add the following configuration to default.conf of ConfigMap. The example is as follows:

```
server {  
    listen 8080; # Adjust the configuration based on the business situation.  
    listen [::]:8080; # Adjust the configuration based on the business situation.  
    server_name localhost; # Adjust the configuration based on the business  
situation.  
    location = /stub_status { # Adjust the specific path based on the business  
situation.  
        stub_status;  
    }  
}
```

The configuration example in ConfigMap is as follows:

Search for the required CAM policy as needed, and click to complete policy association.

Manually add Import via file

```
nginx -t
nginx -s reload
```

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

Page 71 of 165

Directions

1. Log in to [TMP Console](#).
2. Select the corresponding Prometheus instance from the instance list.
3. Go to the instance details page, select **Data Collection > Integration Center**.
4. Search for **Nginx** in the integration center, and click it to pop up an installation window.
5. On the Installation tab of the pop-up window, fill in the metric name, address, path, and other information, and click **Save**.

Search for the required CAM policy as needed, and click to complete policy association.

Nginx (nginx-exporter)

InstallDashboardIntegrated

Current subnet [\[lucy-subnet-4\]](#) Remaining IP count: 189

Installation method

One-click installation

[Installation instruction document](#)

nginx metric collection

name *

Global unique name

Nginx Instance

address *

http://host:port

path *

/stub_status

user name

password

tag ⓘ

+ Add

Collector estimated resource occupancy ⓘ: CPU-0.25 cores Memory-0.5GiB

Configuration cost:**0.0018 Dollar/Hour** Original price:0.0065Dollar/Hour No charge for collect only free metrics[Billing explanation](#)

Save

Cancel

Configuration Instructions

Parameters	Description
name	Exporter name, which should meet the following requirements: The name should be unique. The name should conform to the following regular expression: '^([a-z0-9]([-a-z0-9]*[a-z0-9])?(\.[a-z0-9]([-a-z0-9]*[a-z0-9])?)*\$'.
address	Nginx service connection address.
path	Nginx service status path, which is specified in the configuration.
user name	Username for HTTP authentication of the Nginx service.
password	Password for HTTP authentication of the Nginx service.
tag	Custom labels for metrics.

Method 2: Custom Installation

Note:

[TKE](#) is recommended for convenient installation and management of the Exporter.

Prerequisites

A [TKE cluster](#) has been created in the region and VPC of the corresponding Prometheus instance, and a [namespace](#) has been created for the cluster.

In the [TMP Console](#) > **select the corresponding Prometheus instance** > **Data Collection** > **Integrate with TKE** to find the corresponding container cluster and complete the cluster association operation. See the guide [Associate Cluster](#) for reference.

Directions

Step 1: Deploying the Exporter

1. Log in to the [TKE console](#).
2. Click **Cluster** in the left sidebar.
3. Click the ID/name of the cluster whose access credential is required to go to the management page of the cluster.
4. Follow the steps below [to deploy Nginx Exporter](#) and [verify](#) the deployment status.

Step 2: Deploying the Nginx Exporter

1. Choose **Workload** > **Deployment** in the left sidebar to enter the Deployment page.
2. Click **Create via YAML** in the upper right corner of the page to create a YAML file, and select the corresponding namespace for server deployment. The following part shows how to deploy the Exporter by using a YAML file. Sample configurations are as follows:

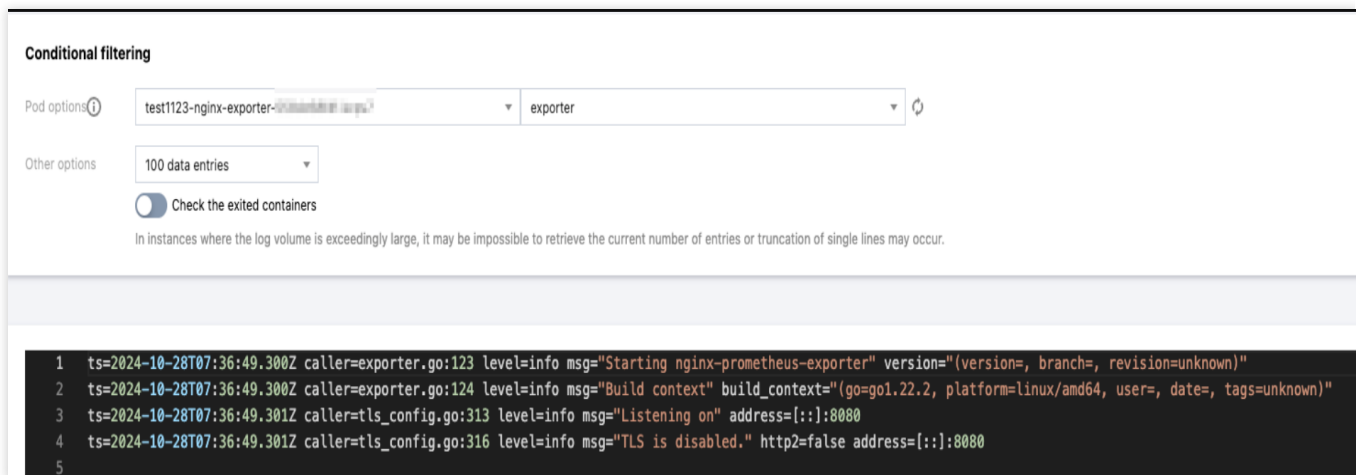
```

apiVersion: apps/v1
kind: Deployment
metadata:
  labels:
    k8s-app: nginx-exporter # Use the actual name based on business needs. It
is recommended to include the information on the corresponding Nginx instance.
    name: nginx-exporter # Use the actual name based on business needs. It is
recommended to include the information on the corresponding Nginx instance.
    namespace: nginx-demo # Use the actual namespace based on business needs.
spec:
  replicas: 1
  selector:
    matchLabels:
      k8s-app: nginx-exporter # Use the actual name based on business needs.
It is recommended to include the information on the corresponding Nginx
instance.
  template:
    metadata:
      labels:
        k8s-app: nginx-exporter # Use the actual name based on business needs.
It is recommended to include the information on the corresponding Nginx
instance.
    spec:
      containers:
        - args:
            - --web.listen-address=:8080
            - --nginx.scrape-uri=http://127.0.0.1:8080/stub_status # Use the
actual address corresponding to the Nginx instance based on business needs.
          image: ccr.ccs.tencentyun.com/rig-agent/common-image:nginx-exporter-
v1.1.0
          name: nginx-exporter
          ports:
            - containerPort: 9113
              name: metric-port
          terminationMessagePath: /dev/termination-log
          terminationMessagePolicy: File
        dnsPolicy: ClusterFirst
        imagePullSecrets:
          - name: qcloudregistrykey
        restartPolicy: Always
        schedulerName: default-scheduler
        securityContext: {}
        terminationGracePeriodSeconds: 30

```

Validation

1. Click the Deployment created in the previous step on the Deployment page to go to the Deployment management page.
2. Click the **Log** tab. The Exporter is started, and the corresponding access address is exposed, as shown below:
Search for the required CAM policy as needed, and click to complete policy association.



3. Click the **Pod** tab to enter the Pod page.
4. Click **Remote login to** in the operation bar to log in to the Pod. Execute the following wget command on the command line interface to access the exposed Exporter address. In this way, data of corresponding Nginx metrics can be collected. If no data is collected, check whether the **connection string** is correct. The command is as follows:

```
wget -qO- http://localhost:8080/metrics
```

The successful outcome is shown in the following figure:

Search for the required CAM policy as needed, and click to complete policy association.

```
# HELP nginx_connections_active Active client connections
# TYPE nginx_connections_active gauge
nginx_connections_active 2
# HELP nginx_connections_handled Handled client connections
# TYPE nginx_connections_handled counter
nginx_connections_handled 128
# HELP nginx_connections_reading Connections where NGINX is reading the request header
# TYPE nginx_connections_reading gauge
nginx_connections_reading 0
# HELP nginx_connections_waiting Idle client connections
# TYPE nginx_connections_waiting gauge
nginx_connections_waiting 1
# HELP nginx_connections_writing Connections where NGINX is writing the response back to the client
# TYPE nginx_connections_writing gauge
nginx_connections_writing 1
# HELP nginx_exporter_build_info A metric with a constant '1' value labeled by version, revision, branch, goversion from which nginx_exporter was built, and the goos and goarch for the build
# TYPE nginx_exporter_build_info gauge
nginx_exporter_build_info{branch="",goarch="amd64",goos="linux",goversion="go1.22.2",revision="c68dd0b5518795457adf1cce7c2fe791f04a0250-modified",tags="unknown",version="1.1.0"} 1
# HELP nginx_http_requests_total Total http requests
# TYPE nginx_http_requests_total counter
nginx_http_requests_total 29564
# HELP nginx_up Status of the last metric scrape
# TYPE nginx_up gauge
nginx_up 1
```

Step 4: Adding a Collection Task

1. Log in to the [TMP console](#) and select the corresponding Prometheus instance to go to the management page.
2. Choose **Data Collection > Integrate with TKE**, select the associated cluster, and choose **Data Collection Configuration > Customize Monitoring Configuration > Via YAML** to add a collection task.
3. Add a `PodMonitor` via service discovery to define the collection task. The YAML example is as follows:

```
apiVersion: monitoring.coreos.com/v1
kind: PodMonitor
metadata:
  name: nginx-exporter # Enter a unique name.
  namespace: cm-prometheus # Pay-as-you-go instance: Use the namespace of
the cluster. Monthly subscription instance (no longer available): The namespace
is fixed. Do not change it.
spec:
  podMetricsEndpoints:
  - interval: 30s
    port: metric-port # Enter the port of the Prometheus Exporter in the Pod
YAML file.
    path: /metrics # Enter the path of the Prometheus Exporter. Default
value: /metrics.
    relabelings:
    - action: replace
      sourceLabels:
      - instance
      regex: (.*)
      targetLabel: instance
```



```
replacement: 'crs-xxxxxx' # Enter the information on the corresponding
Nginx instance.
namespaceSelector: # Select the namespace where the Pod to be monitored is
located.
  matchNames:
  - nginx-demo
selector: # Enter the labels of the Pod to be monitored to locate the
target Pod.
  matchLabels:
    k8s-app: nginx-exporter
```


Viewing Monitoring Information

Prerequisites

The Prometheus instance has been bound to a Grafana instance.

Directions

1. Log in to the [TMP console](#) and select the corresponding Prometheus instance to go to the management page.
2. On the **Basic Info** page of the instance, find the bound Grafana address, open it, and log in to Grafana. Then, find the Nginx instance monitoring panel in the middleware folder to view relevant monitoring data of the instance, as shown below:



Configure Alarm

TMP supports configuring alerting rules based on the actual business situation. For details, see [Creating Alerting Rules](#).

Appendix: Data Collection Parameters of Nginx Exporter

Global Configuration Parameters

Name	Description
web.telemetry-path	Path for exposing metrics. Default value <code>/metrics</code> .
nginx.scrape-uri	URL for Nginx metric collection. Default value: <code>http://127.0.0.1:8080/stub_status</code> .
[no-]nginx.plus	Whether to enable Nginx Plus. Default value: enabled.
[no-]nginx.ssl-verify	Whether to verify the SSL certificate.
nginx.ssl-ca-cert	SSL certificate path.

nginx.ssl-client-cert	SSL certificate path.
nginx.ssl-client-key	SSL certificate path.
nginx.timeout	Nginx metric collection timeout interval.
prometheus.const-label	Tag to be used for each metric, which is in the format of label=value. One tag can be used multiple times.
[no-]web.systemd-socket	Use a systemd socket listener instead of a port listener (Linux only).
web.listen-address	Listening address. Default value: 9113.
web.config.file	Configuration file path. TLS or authentication can be enabled. (This parameter is used for testing.)
log.level	Log level. Default value: info.
log.format	Log message output format. Valid values: logfmt and json. Default value: logfmt.
version	Printed Apache version information.

Redis Exporter Integration

Last updated : 2024-01-29 15:55:07

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
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: 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:



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 `REDIS_ADDR` and `REDIS_PASSWORD` 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 keyspace_misses_total metric
# TYPE redis_keyspace_misses_total counter
redis_keyspace_misses_total 29
# HELP redis_last_slow_execution_duration_seconds The amount of time needed for last slow execution, in
# TYPE redis_last_slow_execution_duration_seconds gauge
redis_last_slow_execution_duration_seconds 0.011276
# HELP redis_latency_spike_duration_seconds Length of the last latency spike in seconds
# TYPE redis_latency_spike_duration_seconds gauge
redis_latency_spike_duration_seconds{event_name="command"} 0.011
redis_latency_spike_duration_seconds{event_name="fast-command"} 0.022
# HELP redis_latency_spike_last When the latency spike last occurred
# TYPE redis_latency_spike_last gauge
redis_latency_spike_last{event_name="command"} 1.604752448e+09
redis_latency_spike_last{event_name="fast-command"} 1.604738646e+09
# HELP redis_latest_fork_seconds latest_fork_seconds metric
# TYPE redis_latest_fork_seconds gauge
redis_latest_fork_seconds 0
# HELP redis_lazyfree_pending_objects lazyfree_pending_objects metric
# TYPE redis_lazyfree_pending_objects gauge
redis_lazyfree_pending_objects 0
# HELP redis_loading_dump_file loading_dump_file metric
# TYPE redis_loading_dump_file gauge
redis_loading_dump_file 0
# HELP redis_master_repl_offset master_repl_offset metric
# TYPE redis_master_repl_offset gauge
redis_master_repl_offset 2.37644710082e+11
# HELP redis_mem_fragmentation_ratio mem_fragmentation_ratio metric
# TYPE redis_mem_fragmentation_ratio gauge
redis_mem_fragmentation_ratio 1.43
# HELP redis_memory_max_bytes memory_max_bytes metric
# TYPE redis_memory_max_bytes gauge
redis_memory_max_bytes 1.2884901888e+10
# HELP redis_memory_used_bytes memory_used_bytes metric
# TYPE redis_memory_used_bytes gauge
redis_memory_used_bytes 1.1479248e+07
```

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 Prometheus
    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 : 2024-01-29 15:55:08

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

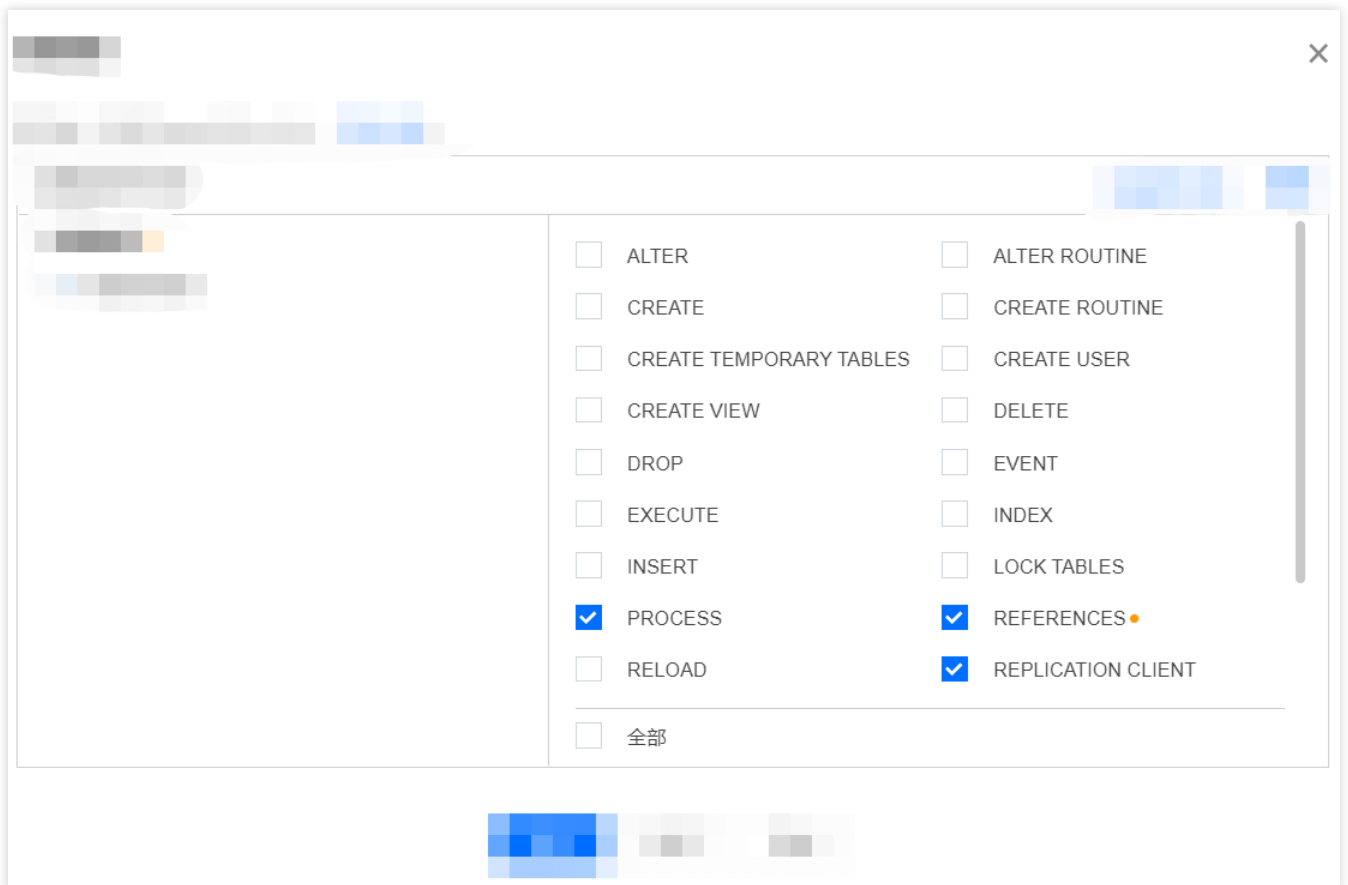
You have created a [TKE cluster](#) in the region and VPC of your TMP instance and created a [namespace](#) for the cluster. You have located and integrated the target TKE cluster in the **Integrate with TKE** section of the **target TMP instance** in the [TMP console](#). For more information, please 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:



You can run the following command for authorization:

```
CREATE USER 'exporter'@'ip' IDENTIFIED BY 'XXXXXXXX' 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 string
```

Deploying MySQL exporter

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  mysql-exporter
1 2020-12-08T09:55:18.315462103Z time="2020-12-08T09:55:18Z" level=info msg="Starting mysql_exporter (version=0.12.1, branch=HEAD, revision=48667bf7c3b438b5e93b259f3d17b70a7c9aff96)" source="mysql_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="mysql_exporter.go:258"
3 2020-12-08T09:55:18.315537718Z time="2020-12-08T09:55:18Z" level=info msg="Enabled scrapers:" source="mysql_exporter.go:269"
4 2020-12-08T09:55:18.315541954Z time="2020-12-08T09:55:18Z" level=info msg="--collect.global_status" source="mysql_exporter.go:273"
5 2020-12-08T09:55:18.315546174Z time="2020-12-08T09:55:18Z" level=info msg="--collect.global_variables" source="mysql_exporter.go:273"
6 2020-12-08T09:55:18.315549924Z time="2020-12-08T09:55:18Z" level=info msg="--collect.slave_status" source="mysql_exporter.go:273"
7 2020-12-08T09:55:18.315748537Z time="2020-12-08T09:55:18Z" level=info msg="--collect.info_schema.innodb_cmp" source="mysql_exporter.go:273"
8 2020-12-08T09:55:18.315765268Z time="2020-12-08T09:55:18Z" level=info msg="--collect.info_schema.innodb_cmpmem" source="mysql_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="mysql_exporter.go:273"
10 2020-12-08T09:55:18.315774561Z time="2020-12-08T09:55:18Z" level=info msg="Listening on :9104" source="mysql_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:

```
mysql_info_schema_innodb_cmpmem_pages_used_total{buffer_pool="0",page_size="4096"} 0
mysql_info_schema_innodb_cmpmem_pages_used_total{buffer_pool="0",page_size="8192"} 0
# HELP mysql_info_schema_innodb_cmpmem_relocation_ops_total Number of times a block of the size PAGE_SIZE has b
# TYPE mysql_info_schema_innodb_cmpmem_relocation_ops_total counter
mysql_info_schema_innodb_cmpmem_relocation_ops_total{buffer_pool="0",page_size="1024"} 0
mysql_info_schema_innodb_cmpmem_relocation_ops_total{buffer_pool="0",page_size="16384"} 0
mysql_info_schema_innodb_cmpmem_relocation_ops_total{buffer_pool="0",page_size="2048"} 0
mysql_info_schema_innodb_cmpmem_relocation_ops_total{buffer_pool="0",page_size="4096"} 0
mysql_info_schema_innodb_cmpmem_relocation_ops_total{buffer_pool="0",page_size="8192"} 0
# HELP mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total Total time in seconds spent in relocating
# TYPE mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total counter
mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total{buffer_pool="0",page_size="1024"} 0
mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total{buffer_pool="0",page_size="16384"} 0
mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total{buffer_pool="0",page_size="2048"} 0
mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total{buffer_pool="0",page_size="4096"} 0
mysql_info_schema_innodb_cmpmem_relocation_time_seconds_total{buffer_pool="0",page_size="8192"} 0
# 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 resides
matchNames:
  - mysql-demo
selector: # Enter the label value of the Pod to be monitored to locate the target
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

Strategy name *

MySQL Shut down

Rules PromQL *

mysql_up != 1

[Click to preview rules](#)

duration

1

minu

Alarm notification period ⓘ

please choose

Alarm Object (Summary) *

MySQL Not running

Alarm message (Description) *

MySQL Not running, Instance: {{labels.instance}} . |

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, 3 more can be selected

Notification template name

Contains operations

The current notification template list is empty, you can select the corresponding notification template

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 <code>auto_increment</code> columns
collect.binlog_size	5.1	Collects the current size of all registered
collect.engine_innodb_status	5.1	Collects the status data from <code>SHOW EN</code>

collect.engine_tokudb_status	5.6	Collects the status data from <code>SHOW EN</code>
collect.global_status	5.1	Collects the status data from <code>SHOW GL</code>
collect.global_variables	5.1	Collects the status data from <code>SHOW GL</code>
collect.info_schema.clientstats	5.5	If <code>userstat=1</code> is set, this parameter collection.
collect.info_schema.innodb_metrics	5.6	Collects the monitoring data from <code>info</code>
collect.info_schema.innodb_tablespace	5.7	Collects the monitoring data from <code>information_schema.innodb_sy</code>
collect.info_schema.innodb_cmp	5.5	Collects the monitoring data of compress <code>information_schema.innodb_cm</code>
collect.info_schema.innodb_cmpmem	5.5	Collects the monitoring data of InnoDB b <code>information_schema.innodb_cm</code>
collect.info_schema.processlist	5.1	Collects the monitoring data of the threa <code>information_schema.processli</code>
collect.info_schema.processlist.min_time	5.1	Minimum time a thread must be in each s
collect.info_schema.query_response_time	5.5	Collects query response time distribution to <code>ON</code> .
collect.info_schema.replica_host	5.6	Collects the status data from <code>informa</code>
collect.info_schema.tables	5.1	Collects the status data from <code>informa</code>
collect.info_schema.tables.databases	5.1	Sets the list of databases to collect table
collect.info_schema.tablestats	5.1	If <code>userstat=1</code> is set, this parameter statistics.
collect.info_schema.schemastats	5.1	If <code>userstat=1</code> is set, this parameter statistics.
collect.info_schema.userstats	5.1	If <code>userstat=1</code> is set, this parameter statistics.
collect.perf_schema.eventsstatements	5.6	Collects the monitoring data from <code>performance_schema.events_st</code>
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 <code>performance_schema.events_st summed</code> .
collect.perf_schema.eventswaits	5.5	Collects the monitoring data from <code>performance_schema.events_wa</code>
collect.perf_schema.file_events	5.6	Collects the monitoring data from <code>performance_schema.file_summ</code>
collect.perf_schema.file_instances	5.5	Collects the monitoring data from <code>performance_schema.file_summ</code>
collect.perf_schema.indexiowaits	5.6	Collects the monitoring data from <code>performance_schema.table_io_</code>
collect.perf_schema.tableiowaits	5.6	Collects the monitoring data from <code>performance_schema.table_io_</code>
collect.perf_schema.tablelocks	5.6	Collects the monitoring data from <code>performance_schema.table_loc</code>
collect.perf_schema.replication_group_members	5.7	Collects the monitoring data from <code>performance_schema.replicati</code>
collect.perf_schema.replication_group_member_stats	5.7	Collects the monitoring data from <code>performance_schema.replicati</code>
collect.perf_schema.replication_applier_status_by_worker	5.7	Collects the monitoring data from <code>performance_schema.replicati</code>
collect.slave_status	5.1	Collects the monitoring data from <code>SHOW</code>
collect.slave_hosts	5.1	Collects the monitoring data from <code>SHOW</code>
collect.heartbeat	5.1	Collects the monitoring data from heartbe
collect.heartbeat.database	5.1	Database from where to collect heartbea
collect.heartbeat.table	5.1	Table from where to collect heartbeat da
collect.heartbeat.utc	5.1	Uses UTC for timestamps of the current <code>utc</code>). Default value: false.

Global configuration parameters

Item	Description
config.my.cnf	Path of <code>.my.cnf</code> file to read MySQL credentials from. Default value: <code>~/.my.cnf</code> .
log.level	Log level. Default value: info.
exporter.lock_wait_timeout	Sets a <code>lock_wait_timeout</code> (in seconds) on the connection to avoid long metadata locking. Default value: 2.
exporter.log_slow_filter	Adds a <code>log_slow_filter</code> 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 : 2024-01-29 15:55:08

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

Consul indicator collection

name *

example

Consul instance

address *

192.1.1.1

Label ⓘ

+ Add to

save

Cancel

Will incur additional costs , [billing overview](#)

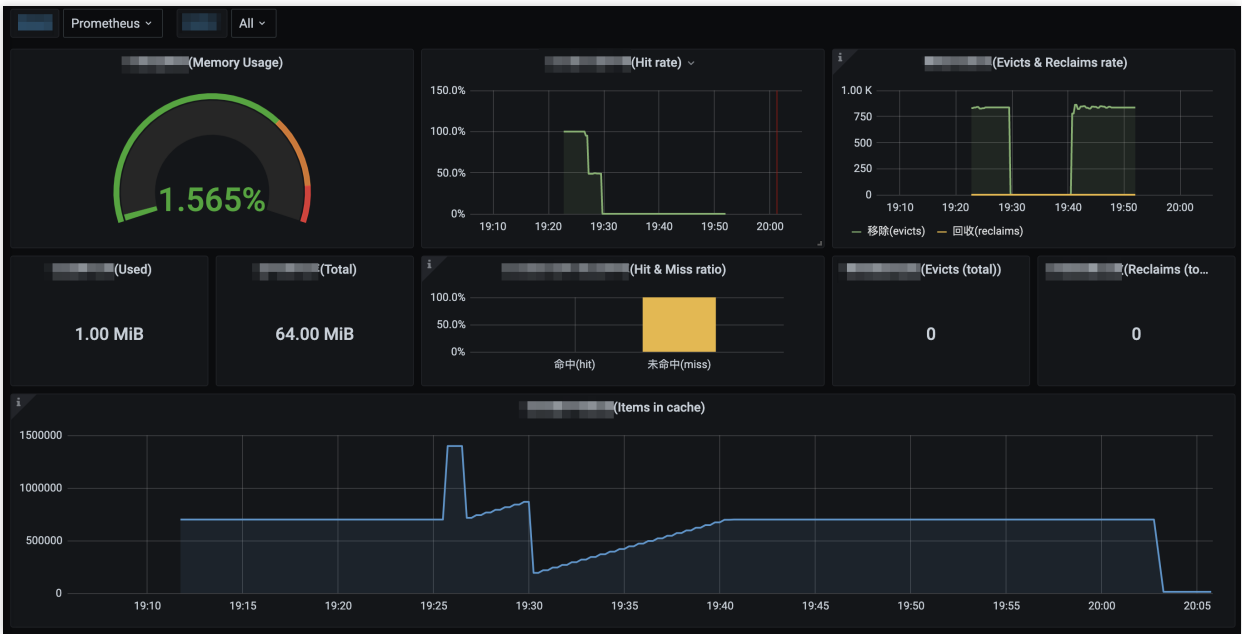
Item	Description
Name	Unique integration name

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 : 2024-01-29 15:55:08

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

Memcached metrics collection

name *

example

Memcached instance

address *

192.168.1.1:3600

Label ⓘ

+ Add to

save

Cancel

Will incur additional costs , [billing overview](#)

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 : 2024-01-29 15:55:07

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 : 2024-01-29 15:55:08

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_exporter-1.3.1.linux-amd64.tar.gz && tar -xvf node_exporter-1.3.1.linux-amd64.tar.gz
```

The file directory is as follows:



```
[root@VM-0-7-centos node_exporter-1.2.2.linux-amd64]# ll
total 18080
-rw-r--r-- 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
[root@VM-0-7-centos node_exporter-1.2.2.linux-amd64]# ./node_exporter
```

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.

```
FW-F--F-- 1 3434 3434 463 Aug 6 2021 NOTICE
[root@VM-0-7-centos node_exporter-1.2.2.linux-amd64]# ./node_exporter
level=info ts=2022-02-11T07:15:26.555Z caller=node_exporter.go:182 msg="Starting node_exporter" version="(version=1.2.2, branch=HEAD,
bn=26645363b486e12be40af7ce4fc91e731a33104e)"
level=info ts=2022-02-11T07:15:26.555Z caller=node_exporter.go:183 msg="Build context" build_context="(go=go1.16.7, user=root@b9cb4aa2
late=20210806-13:44:18)"
level=warn ts=2022-02-11T07:15:26.555Z caller=node_exporter.go:185 msg="Node Exporter is running as root user. This exporter is design
run as unprivileged user, root is not required."
level=info ts=2022-02-11T07:15:26.555Z caller=filesystem_common.go:110 collector=filesystem msg="Parsed flag --collector.filesystem.mo
nts-exclude" flag=~((dev|proc|sys|var|lib|docker|.+)($|/))
level=info ts=2022-02-11T07:15:26.555Z caller=filesystem_common.go:112 collector=filesystem msg="Parsed flag --collector.filesystem.fs
exclude" flag=~(autofs|binfmt_misc|bpf|cgroup2?|configfs|debugfs|devpts|devtmpfs|fusectl|hugetlbfs|iso9660|mqueue|nsfs|overlay|proc|p
store|rpc_pipefs|securityfs|selinuxfs|squashfs|sysfs|tracefs)$
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:108 msg="Enabled collectors"
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=arp
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=bcache
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=bonding
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=btrfs
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=conntrack
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=cpu
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=cputime
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=diskstats
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=edac
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=entropy
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=fibrechannel
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=filefd
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=filesystem
level=info ts=2022-02-11T07:15:26.556Z caller=node_exporter.go:115 collector=hwmon
level=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.

```
[root@VM-0-7-centos node_exporter-1.2.2.linux-amd64]# clear
[root@VM-0-7-centos node_exporter-1.2.2.linux-amd64]# curl 127.0.0.1:9100/metrics
# HELP go_gc_duration_seconds A summary of the pause duration of garbage collection cycles.
# TYPE go_gc_duration_seconds summary
go_gc_duration_seconds{quantile="0"} 0
go_gc_duration_seconds{quantile="0.25"} 0
go_gc_duration_seconds{quantile="0.5"} 0
go_gc_duration_seconds{quantile="0.75"} 0
go_gc_duration_seconds{quantile="1"} 0
go_gc_duration_seconds_sum 0
go_gc_duration_seconds_count 0
# HELP go_goroutines Number of goroutines that currently exist.
# TYPE go_goroutines gauge
go_goroutines 7
# HELP go_info Information about the Go environment.
# TYPE go_info gauge
go_info{version="go1.16.7"} 1
# HELP go_memstats_alloc_bytes Number of bytes allocated and still in use.
# TYPE go_memstats_alloc_bytes gauge
go_memstats_alloc_bytes 2.344136e+06
# HELP go_memstats_alloc_bytes_total Total number of bytes allocated, even if freed.
# TYPE go_memstats_alloc_bytes_total counter
go_memstats_alloc_bytes_total 2.344136e+06
# HELP go_memstats_buck_hash_sys_bytes Number of bytes used by the profiling bucket hash table.
# TYPE go_memstats_buck_hash_sys_bytes gauge
go_memstats_buck_hash_sys_bytes 4562
# HELP go_memstats_frees_total Total number of frees.
# TYPE go_memstats_frees_total counter
go_memstats_frees_total 1362
# HELP go_memstats_gc_cpu_fraction The fraction of this program's available CPU time used by the GC since the program s
# TYPE go_memstats_gc_cpu_fraction gauge
```

Step 3. Configure the collection

Log in to the [TMP console](#), select **Integration Center > CVM**, and configure the information in **Task Configuration** as prompted.

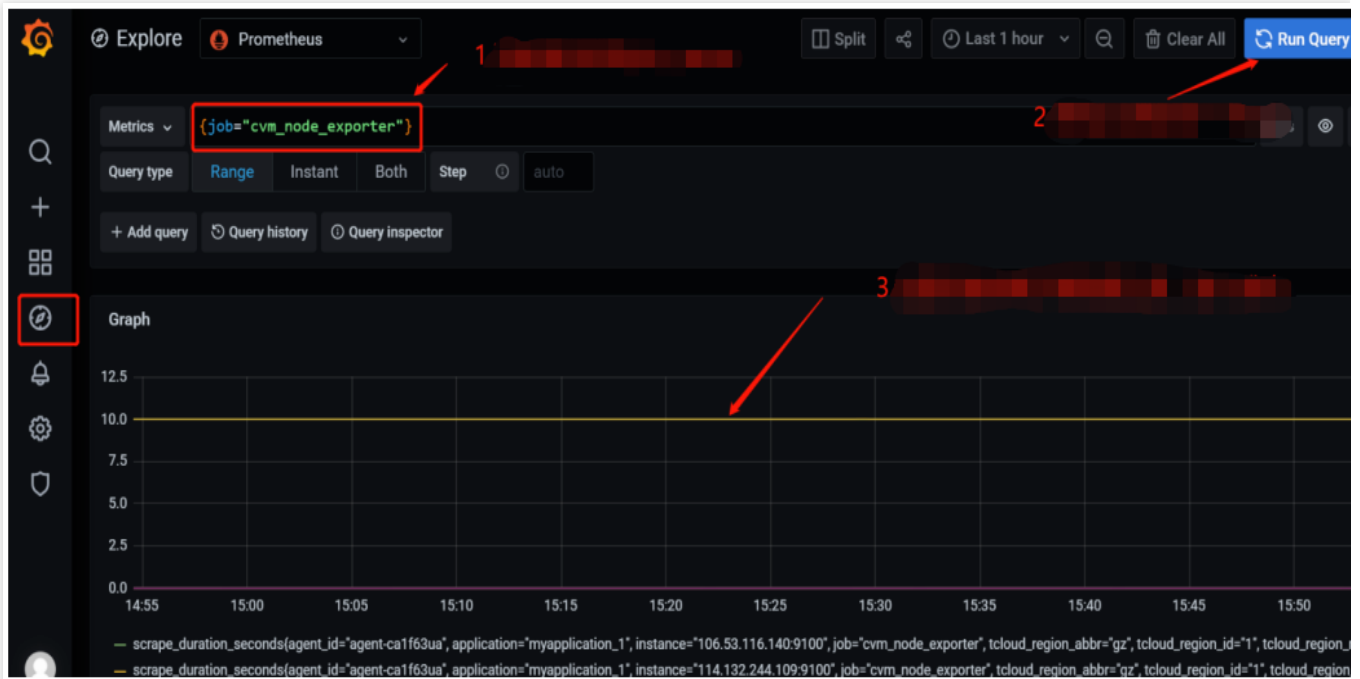
Below is a sample configuration of a scrape task:

```
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.

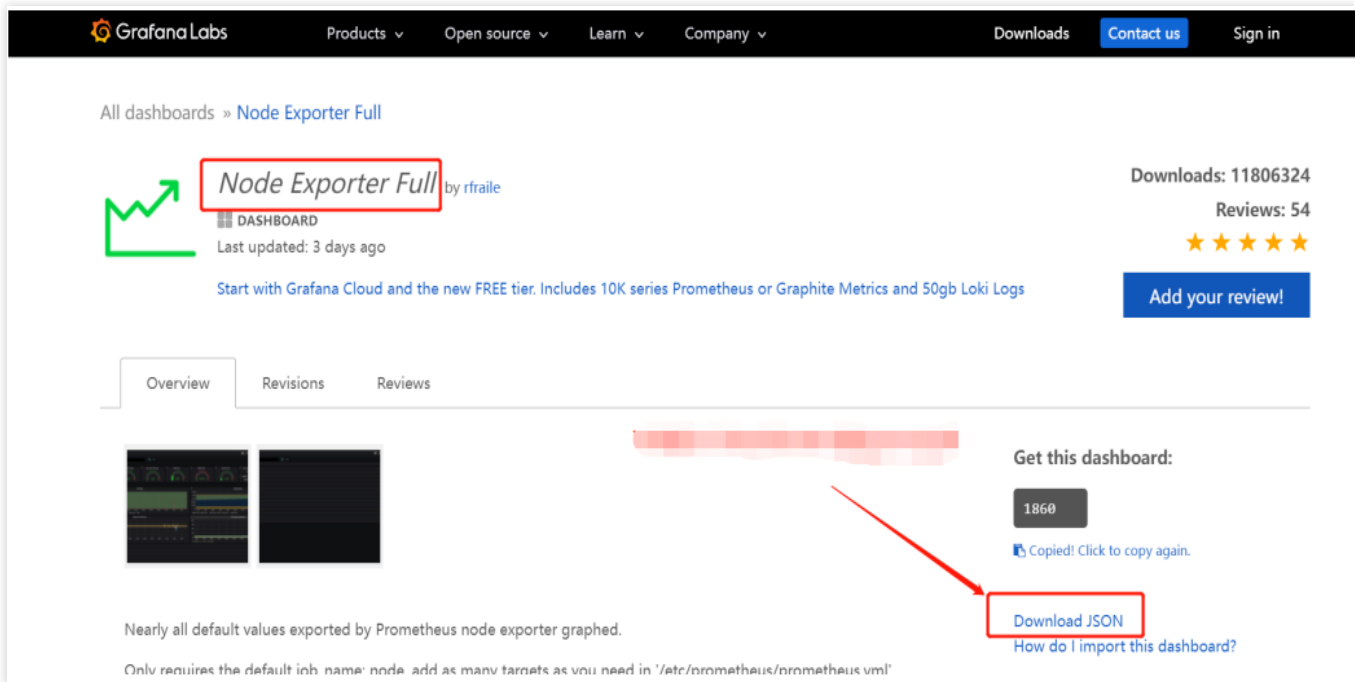
Search for `{job="cvm_node_exporter"}` in **Explore** to see whether there is data, and if so, data is reported successfully.



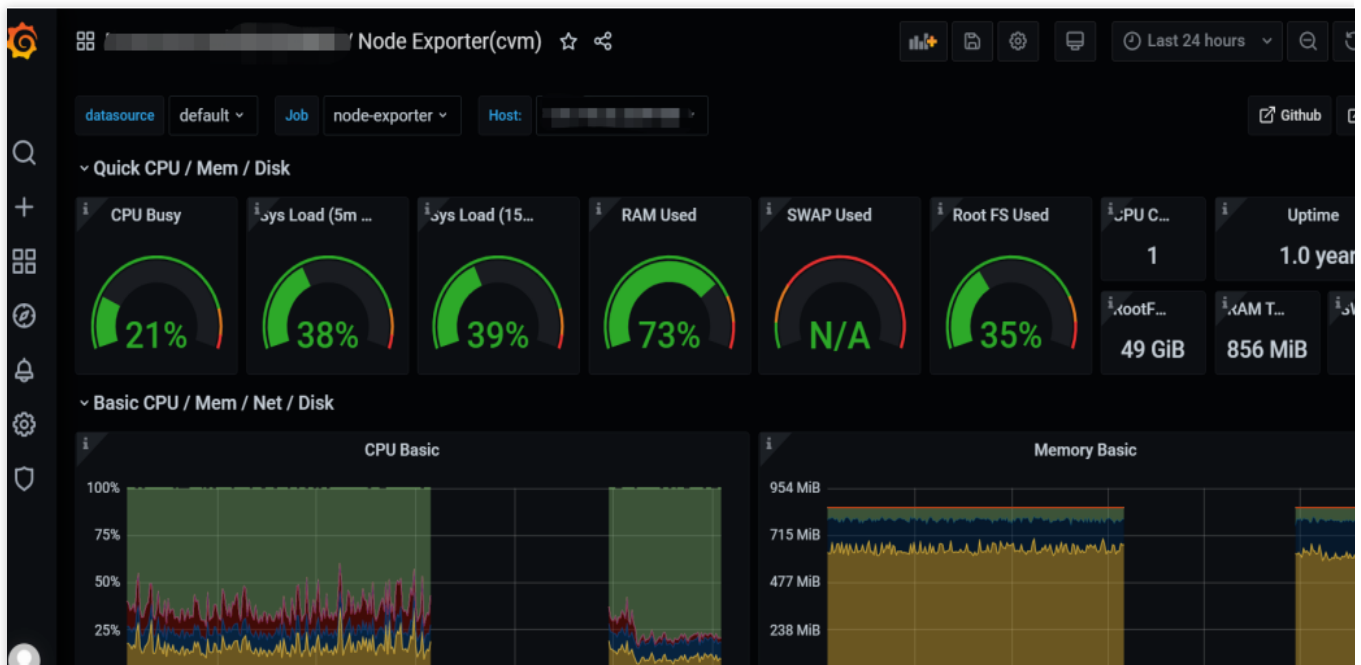
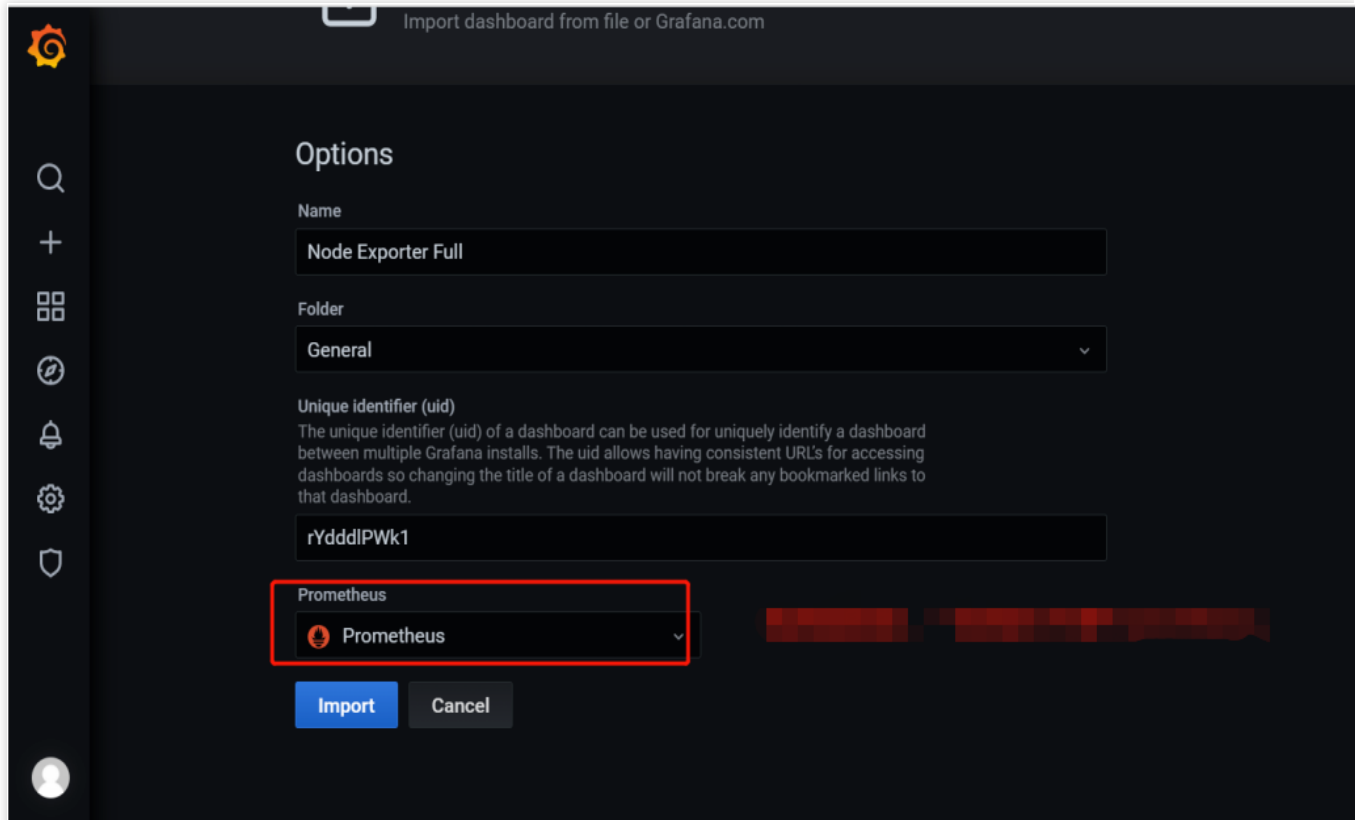
Step 5. Configure the dashboard

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 `node_exporter`, and select the latest dashboard for 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**.



Apache Exporter Integration

Last updated : 2024-10-24 19:11:24

Overview

Apache Exporter is a tool used for collecting data on Apache HTTP server metrics. The data of core metrics reported by the Exporter can be used to trigger alarms and is displayed on the monitoring dashboard. TMP on Tencent Cloud Observability Platform (TCOP) provides the Apache Exporter connection feature and an out-of-the-box Grafana monitoring dashboard.

Note:

To ensure the Exporter can collect data, make sure the Apache HTTP server is running. For details, see [this document](#).

Connection Method

Method 1: One-Click Installation (Recommended)

Directions

1. Log in to [TMP Console](#).
2. Select the corresponding Prometheus instance from the instance list.
3. Go to the instance details page, select **Data Collection > Integration Center**.
4. Search for **Apache** in the integration center, and click it to pop up the installation window.
5. On the Installation tab of the pop-up window, fill in the metric name, address, path, and other information, and click **Save**.

Search for the required CAM policy as needed, and click to complete policy association.

Apache (apache-exporter)

Install

Dashboard

Integrated

Current subnet

lucy-subnet-4

Remaining IP count: 191

Installation method

One-click installation

Installation instruction document

Apache metric collection

name *

Global unique name

Apache HTTP Service

address *

http://host:port

path *

/server-status

user name

password

tag

+ Add

Collector estimated resource occupancy

CPU-0.25 cores Memory-0.5GiB

Configuration cost:

0.00Dollar/Hour

Original price:0.00Dollar/Hour

No charge for collecting only metrics

Billing explanation

Save

Cancel

Configuration Instructions

Parameters	Note
name	Exporter name, which should meet the following requirements: The name should be unique. The name should conform to the following regular expression: '[a-z0-9]([-a-z0-9]*[a-z0-9])?(\\.[a-z0-9]([-a-z0-9]*[a-z0-9])?)*\$'.
address	Address of the connected Apache HTTP server.
path	Path for viewing the Apache HTTP Server Status page. Default value: /server-status .

user name	Username for accessing the Apache HTTP server.
password	Password for accessing the Apache HTTP server.
tag	Custom labels for metrics.

Method 2: Custom Installation

Note:

[TKE](#) is recommended for convenient installation and management of the Exporter.

Prerequisites

A [TKE cluster](#) has been created in the region and VPC of the corresponding Prometheus instance, and a [namespace](#) has been created for the cluster.

In the [TMP Console](#) > **select the corresponding Prometheus instance** > **Data Collection** > **Integrate with TKE** to find the corresponding container cluster and complete the cluster association operation. See the guide [Associate Cluster](#) for reference.

Directions

Step 1: Enabling the mod_status Module of the Apache Server

Note:

For TKE-related operations, see the [TKE](#) documentation.

The Apache Exporter collects data via the mod_status module of the Apache server. Therefore, you need to ensure that the mod_status module is enabled for the Apache server. The specific steps are as follows:

1. Log in to the [TKE console](#).
2. Click **Cluster** in the left sidebar, find the cluster where the Apache server is located, enter the cluster, and find the Apache server.
3. If no ConfigMap is configured in the Apache server, log in to the Apache server, copy the configuration files such as httpd.conf, mime.types, and extra/httpd-info.conf in the configuration directory, create a ConfigMap, and add the configuration files to the ConfigMap. For ConfigMap-related operations, see [ConfigMap Management](#).
4. In httpd.conf, delete the comment for the line LoadModule status_module modules/mod_status.so (remove that part starting with #). If extra-related configurations exist for the server, enable them in httpd.conf by deleting corresponding comments. Example:

Search for the required CAM policy as needed, and click to complete policy association.

5. Modify `httpd-info.conf` as required and enable `ExtendedStatus`. If no extra-related configuration exists, modify `httpd.conf` directly. Example:

Search for the required CAM policy as needed, and click to complete policy association.

Manually add Import via file

Search for the required CAM policy as needed, and click to complete policy association.

```
# curl -s http://10.10.10.10/server-status?auto

ServerVersion: Apache/2.4.53 (Unix) OpenSSL/1.1.1o
ServerMPM: event
Server Built: May 24 2022 19:22:43
CurrentTime: Friday, 29-Mar-2024 03:27:56
RestartTime: Thursday, 28-Mar-2024 07:16:19
ParentServerConfigGeneration: 1
ParentServerMPMGeneration: 0
ServerUptimeSeconds: 72696
ServerUptime: 20 hours 11 minutes 36 seconds
Load1: 6.14
Load5: 3.92
Load15: 3.31
Total Accesses: 4891
Total kBytes: 7363
Total Duration: 6526
```

Step 2: Deploying the Exporter

1. Log in to the [TKE console](#).
2. Click **Cluster** in the left sidebar.
3. Click the ID/name of the cluster whose access credential is required to go to the management page of the cluster.
4. Following the steps below to [deploy the Apache Exporter](#) and [verify](#) the deployment status.

Step 3: Deploying the Apache Exporter

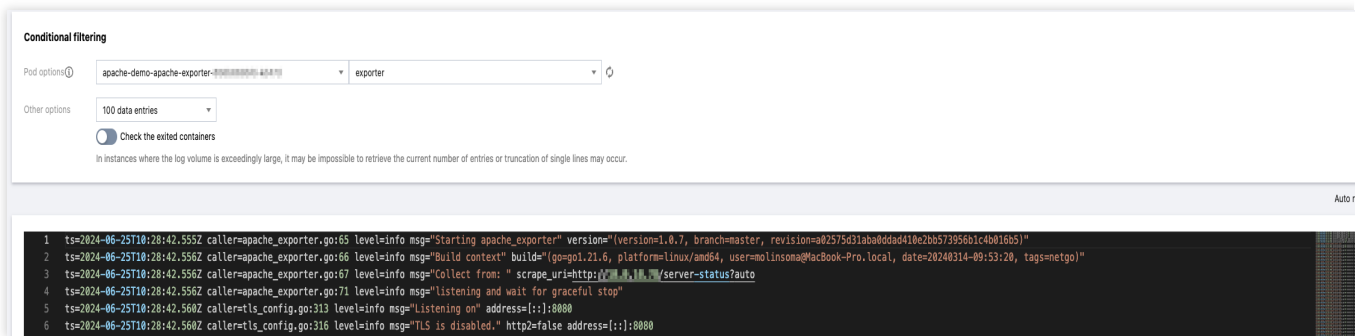
1. Choose **Workload > Deployment** in the left sidebar to enter the Deployment page.
2. Click **Create via YAML** in the upper right corner of the page to create a YAML file, and select the corresponding namespace for server deployment. The following part shows how to deploy the Exporter by using a YAML file. Sample configurations are as follows:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  labels:
    k8s-app: apache-exporter # Use the actual name based on the business
                             needs. It is recommended to include the information on the corresponding
                             Prometheus instance.
    name: apache-exporter # Use the actual name based on the business needs. It
                           is recommended to include the information on the corresponding Prometheus
                           instance.
  namespace: apache-demo # Use the actual name based on the business needs.
```

```
spec:
  replicas: 1
  selector:
    matchLabels:
      k8s-app: apache-exporter # Use the actual name based on the business
needs. It is recommended to include the information on the corresponding
Prometheus instance.
  template:
    metadata:
      labels:
        k8s-app: apache-exporter # Use the actual name based on the business
needs. It is recommended to include the information on the corresponding
Prometheus instance.
    spec:
      containers:
      - args:
        - --web.listen-address=:9117
        - --scrape_uri=http://192.1.1.2:8080/server-status?auto # Use the
address of the corresponding Prometheus instance based on business needs.
        image: ccr.ccs.tencentyun.com/rig-agent/common-image:apache-exporter-
v1.0.7
        name: apache-exporter
        ports:
        - containerPort: 9117
          name: metric-port
        terminationMessagePath: /dev/termination-log
        terminationMessagePolicy: File
      dnsPolicy: ClusterFirst
      imagePullSecrets:
      - name: qcloudregistrykey
      restartPolicy: Always
      schedulerName: default-scheduler
      securityContext: {}
      terminationGracePeriodSeconds: 30
```

Validation

1. Click the Deployment created in the previous step on the Deployment page to go to the Deployment management page.
2. Click the **Log** tab. The Exporter is started, and the corresponding access address is exposed, as shown below: Search for the required CAM policy as needed, and click to complete policy association.



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

4. Click **Remote Login** in the operation bar to log in to the Pod. Execute the following curl command in the command line window to access the exposed address. In this way, data of corresponding Apache server metrics can be collected. If no data is collected, check if the **connection string** is correct. The command is as follows:

```
curl localhost:9117/metrics
```

The successful outcome is shown in the following figure:

Search for the required CAM policy as needed, and click to complete policy association.

```
# HELP apache_accesses_total Current total apache accesses (*)
# TYPE apache_accesses_total counter
apache_accesses_total 6031
# HELP apache_connections Apache connection statuses
# TYPE apache_connections gauge
apache_connections{state="closing"} 1
apache_connections{state="keepalive"} 0
apache_connections{state="total"} 1
apache_connections{state="writing"} 0
# HELP apache_cpu_time_ms_total Apache CPU time
# TYPE apache_cpu_time_ms_total counter
apache_cpu_time_ms_total{type="system"} 3670
apache_cpu_time_ms_total{type="user"} 4019.9999999999999
# HELP apache_cpuload The current percentage CPU used by each worker and in total by all workers combined (*)
# TYPE apache_cpuload gauge
apache_cpuload 0.00858278
# HELP apache_duration_ms_total Total duration of all registered requests in ms
# TYPE apache_duration_ms_total counter
apache_duration_ms_total 8188
# HELP apache_exporter_build_info A metric with a constant '1' value labeled by version, revision, branch, goversion from which apache_exporter was built, and the goos and goarch for the
# TYPE apache_exporter_build_info gauge
apache_exporter_build_info{branch="master",goarch="amd64",goos="linux",goversion="go1.21.6",revision="a02575d31aba0ddad410e2bb573956b1c4b016b5",tags="netgo",version="1.0.7"} 1
# HELP apache_generation Apache restart generation
# TYPE apache_generation gauge
apache_generation{type="config"} 1
apache_generation{type="mpm"} 0
# HELP apache_info Apache version information
# TYPE apache_info gauge
apache_info{mpm="event",version="Apache/2.4.53 (Unix) OpenSSL/1.1.1o"} 1
```

Step 4: Adding a Collection Task

1. Log in to the [TMP console](#) and select the corresponding Prometheus instance to go to the management page.
2. Choose **Data Collection > Integrate with TKE**, select the associated cluster, and choose **Data Collection Configuration > Customize Monitoring Configuration > Via YAML** to add a collection task.

3. Add a `PodMonitor` via service discovery to define the collection task. The YAML example is as follows:

```
apiVersion: monitoring.coreos.com/v1
kind: PodMonitor
metadata:
  name: apache-exporter # Enter a unique name.
  namespace: cm-prometheus # Pay-as-you-go instance: Use the namespace of
the cluster. Monthly subscription instance (no longer available): The namespace
is fixed. Do not change it.
spec:
  podMetricsEndpoints:
  - interval: 30s
    port: metric-port # Enter the port of the Prometheus Exporter in the Pod
YAML file.
    path: /metrics # Enter the path of the Prometheus Exporter. Default
value: /metrics.
    relabelings:
    - action: replace
      sourceLabels:
      - instance
      regex: (.*)
      targetLabel: instance
      replacement: 'crs-xxxxxx' # Enter the information on the corresponding
Prometheus instance.
    namespaceSelector: # Select the namespace where the Pod to be monitored is
located.
      matchNames:
      - apache-demo
    selector: # Enter the labels of the Pod to be monitored to locate the
target Pod.
      matchLabels:
        k8s-app: apache-exporter
```

Viewing Monitoring Information

Prerequisites

The Prometheus instance has been bound to a Grafana instance.

Directions

1. Log in to [the TMP console](#) and select the corresponding Prometheus instance to go to the management page.

2. Choose **Data Collection > Integration Center**, find the Apache Exporter, and install the corresponding Grafana dashboard to display related monitoring data, as shown below:

Search for the required CAM policy as needed, and click to complete policy association.



Configure Alarm

TMP supports configuring alert rules based on the actual business situation. For details, see [Creating Alerting Rules](#).

Appendix: Data Collection Parameters of Apache Exporter

Global Configuration Parameters

Name	Description
telemetry.endpoint	Path for exposing metrics. Default value <code>/metrics</code> .
scrape_uri	URL of the Apache Server Status page. Default value: <code>http://localhost/server-status/?auto</code> .
host_override	String for overriding the HTTP Host request header. A null string indicates that the header is not overridden.
[no-]insecure	Ignore the server certificate if HTTPS is used.
custom_headers	Add custom headers to the Exporter.
[no-]web.systemd-socket	Use a systemd socket listener instead of a port listener (Linux only).
web.listen-address	Listening address. Default value: 9117.
web.config.file	Configuration file path. TLS or authentication can be enabled. (This parameter is used for testing.)
log.level	Log level. Default value: info.
log.format	Log message output format. Valid values: logfmt and json. Default value: logfmt.
version	Printed version information.

Health Check

Last updated : 2024-01-29 15:55:08

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

Integration list / **New**

The number of remaining IPs in the current subnet [2221]: 238

Detect

name *

ping-pp

Probe configuration

Detection method *

http_get

Detection target *

https://console.cloud.tencent.com

X

+ Add to

Label

i

+ Add to

save

Cancel

Will incur additional costs , [billing overview](#)

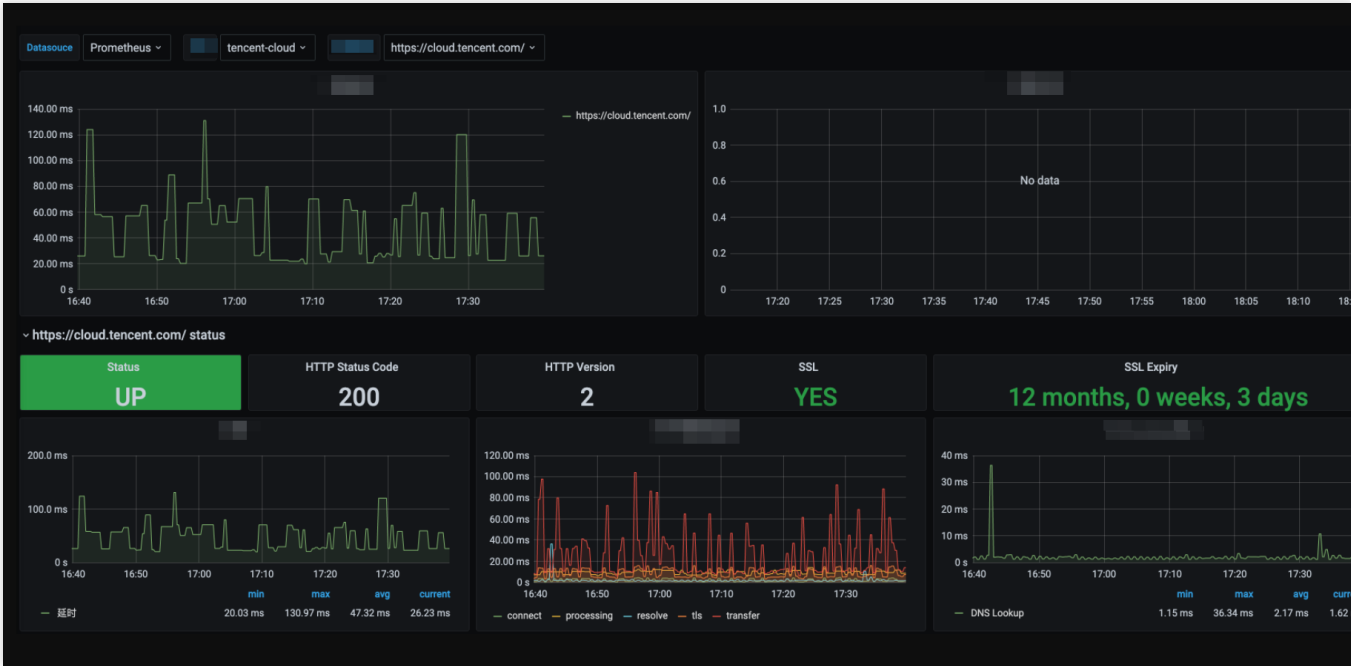
🔗

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.



Instructions for Installing Components in the TKE Cluster

Last updated : 2024-07-23 17:53:35

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	Type	Resource Amount	Description
<Prometheus instance ID>	proxy-agent	Deployment	0.25C256Mi*2	Collection proxy
<Prometheus instance ID>	<Prometheus instance ID>	ServiceAccount	-	Permission carrier
-	<Prometheus instance ID>	ClusterRole	-	Collection permissions related
-	<Prometheus instance ID>-crb	ClusterRoleBinding	-	Collection permissions related

Component Permission Description

Permission Scenarios

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

		Operati Permiss
Collection configuration management	scrapeconfigs,servicemonitors,podmonitors,probes,configmaps,secrets,namespaces	get/list/
Service discovery	services,endpoints,nodes,pods,ingresses	get/list/
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	Type	Resource Amount	Description
kube-system	tke-kube-state-metrics	Statefulset	0.5C512Mi	Collection program
kube-	tke-kube-state-	ServiceAccount	-	Permission carrier

system	metrics			
-	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

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: tke-kube-state-metrics
rules:
  - apiGroups:
    - ""
    resources:
      - configmaps
      - secrets
      - nodes
```



```
- pods
- services
- serviceaccounts
- resourcequotas
- replicationcontrollers
- limitranges
- persistentvolumeclaims
- persistentvolumes
- namespaces
- endpoints
verbs:
- list
- watch
- apiGroups:
  - apps
resources:
  - statefulsets
  - daemonsets
  - deployments
  - replicaset
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 Object Name	Type	Resource Amount	Description

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.

Cloud Monitoring

Last updated : 2025-02-17 15:47:31

Overview

The Cloud Monitor module of TencentCloud Managed Service for Prometheus (TMP) integrates the basic monitoring data of Tencent Cloud products, and implements unified collection, storage, and visualization through TMP.

Note:

Data collection interval: 1 minute. Currently, smaller collection intervals are not supported.

Monitoring data granularity: 1 minute. If a metric does not support the 1-minute granularity, you can select the 5-minute granularity.

The integrated monitoring data includes tag data (not supported by some cloud products) of cloud products. The tag key should conform to the regular expression `[a-zA-Z_][a-zA-Z0-9_]*`. Otherwise, it will be filtered out.

Multi-region is not supported. If cloud products are distributed in multiple regions, multiple integration modules need to be installed.

Operation Steps

1. Log in to the [TMP console](#).
2. Select and enter the corresponding Prometheus instance from the instance list.
3. On the instance details page, select **Data Collection > Integration Center**.
4. On the Integration Center page, click **Cloud Monitor** to enter the installation tab by default. Define the integration name, configure Exporter, and select the corresponding cloud product.

Cloud Monitor (qcloud-exporter)

Install

Metric

Dashboard

Alarm

Integrated

Cloud Monitor

name *

Global unique name

Exporter Configuration

Region *

Instance Region

Data Collection Latency/s ⓘ *

0

Instance Refresh Interval/min ⓘ

10

Instance ID Filtering ⓘ

+ Add

Cloud Tag Key Filtering ⓘ

+ Add

Cloud Tag Key Replacement ⓘ

+ Add

Cloud Tag Key Operations ⓘ

ToUnderLineAndLower

Dimension Whitelist ⓘ

+ Add

Label ⓘ

+ Add

Authentication

Authentication type *

This account collection

Service role

CM_QCSLinkedRoleInTMP

Tencent Cloud Products

☐ CVM

☐ Lighthouse

☐ Load Balancer(public)

☐ Load Balancer(internal)

Configuration Instructions

Parameter	Description
name	Exporter name, which should meet the following requirements: The name should be unique.

	The name should conform to the following regular expression: <code>^[a-z0-9]([-a-z0-9]*[a-z0-9])?(\.[a-z0-9]([-a-z0-9]*[a-z0-9])?)*\$</code> .
Region	Required. Region where the cloud product is located. If the cloud product does not distinguish regions, enter any region.
Data Collection Latency	Unit: second. If it is set to 0, the timestamp of the original data will be ignored. If it is set to a value greater than 0, the timestamp of the original data will be reported. Since there is a certain delay in reporting cloud product monitoring data to basic monitoring, this delay will be reflected in the latest data. Data pulling range: (Current time - Data collection delay - Fixed interval, Current time - Data collection delay).
Instance Refresh Interval	Unit: minute. The minimum value is 10. At each instance refresh interval, the integration module will re-pull cloud product instance information. If the instance name or cloud tag is modified or an instance is added or deleted, the monitoring data will be updated within one instance refresh interval.
Instance ID Filtering	Optional. If it is left blank, data will be collected from all instances under the root account by default. Enter a value in the form of the key-value pair, where the key is the unique ID of the cloud product defined by the integration module, and the value is a list of comma-separated cloud product instance IDs. Only data of the instances of the cloud products specified in the key-value pair form will be collected.
Cloud Tag Key Filtering	Optional. Enter a value in the form of the key-value pair, where one tag key can correspond to multiple tag values that are separated by . Take the intersection of different tag keys and take the union of multiple tag values under the same tag key. For products that support cloud tag filtering, if instance ID filtering is also configured, cloud tag filtering for that product will not take effect.
Cloud Tag Key Replacement	Optional. Replace illegal cloud product tag keys with valid values. For example, convert Chinese names into custom English names.
Cloud Tag Key Operations	By default, the integration module converts uppercase letters in tag keys into underscores followed by lowercase letters. It supports the conversion of tag keys of cloud products. ToUnderLineAndLower: default operation. ToLower: full conversion to lowercase letters. NoOperation: no conversion.
Dimension Whitelist	Optional. Some cloud products have dimensions with the same indicator name and features that need to be whitelisted. By default, collection is not performed, but it can be enabled through this configuration. lb_public:listener: Cloud Load Balancer (public network) - Listener dimension. lb_public:target: Cloud Load Balancer (public network) - Real server dimension. lb_public:domain: Cloud Load Balancer (public network) - Forwarding rule domain name dimension.

	<p>lb_private:listener: Cloud Load Balancer (private network) - Listener dimension.</p> <p>lb_private:domain: Cloud Load Balancer (private network) - Forwarding rule domain name dimension.</p> <p>apigw_cloudnative:node: Cloud Native Gateway - Node dimension.</p> <p>vbc:qosid: Cloud Connect Network - Scheduling queue dimension.</p>
Label	Optional. You can add additional custom tags to the metrics collected by the integration module.
Authentication	<p>Authentication type: You can choose This account collection or Cross-account collection.</p> <p>Service role: Configure for collection within this account. Fixed as CM_QCSLinkedRoleInTMP.</p> <p>This account role: Configure for cross-account collection. Custom role used to obtain a temporary key for this account.</p> <p>Target account role: Configure for cross-account collection. Custom role used to obtain a temporary key for the target account.</p> <p>Target account uin: Configure for cross-account collection. Root account ID of the target account.</p>
Tencent Cloud Products	Select the cloud product you want to collect.
Metric Relabel	Optional. Native metricRelabelings configuration for Prometheus Operator. The configuration method is the same as metric_relabel_configs in Prometheus scraping configuration, with only some field naming conventions being different.

Metric Relabel Configuration Examples

The common metricRelabelings examples are as follows:

```
metricRelabelings:
- action: labeldrop # Remove the label named labelA. regex indicates a regular exp
  regex: labelA
- regex: ins-(.*) # Add a label named id, whose value is derived from the value of
  replacement: $1
  sourceLabels:
  - instance_id
  targetLabel: id
- targetLabel: region # Add a label with region being ap-guangzhou.
  replacement: ap-guangzhou
- action: drop # Remove the metric named metricA or metricB.
  sourceLabels:
  - __name__
  regex: metricA|metricB
```


Supported Cloud Products

Cloud Product/Metric Documentation	Whether to Support Collecting Cloud Tags	Unique ID	Additional Notes
CVM	Yes	cvm	Only metrics at the instance dimension are supported.
Cloud Block Storage	Yes	cbs	-
CLB (public network)	Yes	lb_public	<p>By default, metrics at the instance dimension are collected. If metrics at the listener, forwarding rule domain name, or backend server dimension are required, submit a ticket.</p> <p>The names of metrics at different dimensions are the same name, which can be distinguished by the monitor_view tag.</p> <p>Instance dimension: instance.</p> <p>Listener dimension: listener.</p> <p>Backend server dimension: target.</p> <p>Forwarding rule domain name dimension: domain.</p>
CLB (private network)	Yes	lb_private	<p>By default, metrics at the instance dimension are collected. If metrics at the listener or forwarding rule domain name dimension are required, submit a ticket. The names of metrics at different dimensions are the same name, which can be distinguished by the monitor_view tag.</p> <p>Instance dimension: instance.</p> <p>Listener dimension: listener.</p> <p>Forwarding rule domain name dimension: domain.</p>
TencentDB for MongoDB	Yes	cmongo	-
TencentDB for MySQL (CDB)	Yes	cdb	-
TencentDB for Redis (CKV)	Yes	redis	-

edition)			
TencentDB for Redis (memory edition)	Yes	redis_mem	Metrics at the instance and node dimensions are supported.
TencentDB for MariaDB	Yes	mariadb	Only metrics at the instance dimension are supported.
TencentDB for PostgreSQL	Yes	postgres	-
TDSQL for MySQL	Yes	tdmysql	Only metrics at the instance dimension are supported.
TDSQL-C for MySQL	Yes	cynosdb_mysql	Only metrics at the instance dimension are supported.
TencentDB for SQL Server	Yes	sqlserver	Only metrics at the instance dimension are supported.
NAT Gateway	Yes	nat_gateway	-
TDMQ for CKafka	Yes	ckafka	Metrics at the broker_ip dimension are not supported.
Elastic IP	Yes	lb	-
VPN gateway	Yes	vpngw	-
VPN tunnel	Yes	vpn	-
Network probing	Tags are not supported.	vpc_net_detect	-
CDN (domain name for the Chinese mainland)	Yes	cdn	It does not distinguish by region.
CDN (domain name for countries outside China)	Yes	ov_cdn	It does not distinguish by region.
COS	Yes	cos	Storage-related metrics have a high delay (about 2 hours), and the original timestamp of the data is not retained. Storage-related metrics

			do not support the 1-minute granularity. By default, 5-minute data is pulled.
DC - connection	Yes	dc	It does not distinguish by region.
DC - dedicated tunnel	Yes	dcx	It does not distinguish by region.
DC - DC gateway	Yes	dcg	They are the same as the VPC, network connection, and DC gateway.
Lighthouse	Yes	Lighthouse	-
Cloud-native API gateway	Yes	apigw_cloudnative	By default, metrics at the instance and public network CLB dimensions are collected. If metrics at the node dimension are required, submit a ticket . The names of metrics at the instance and node dimensions are the same name, which can be distinguished by the monitor_view tag. Instance dimension: gateway. Public network CLB dimension: loadbalancer. Node dimension: node.
Elasticsearch	Yes	ces	Only metrics at the instance dimension are supported.
Tencent Cloud TCHouse-D	Yes	cdwdrs	-
Data Transmission Service	Yes	dts	Kafka-related dimension metrics are not supported.
CCN	Yes	vbc	-
GAAP	Yes	gaap	-
EdgeOne (layer-7)	Yes	edgeone_l7	-
WAF	Yes	waf	-
CFS	Yes	cfs	Currently, no metadata-related metrics are collected.
BWP	Yes	bwp	-

SCF	Yes	scf_v2	By default, metrics at the alias dimension are collected. If metrics at the version dimension are required, submit a ticket . The names of metrics at the alias and version dimensions are the same name, which can be distinguished by the monitor_view tag. Alias dimension: alias. Version dimension: version.
CLS - log topic	Yes	cls	-
API Gateway	Yes	apigateway	Only metrics at the API dimension are supported.

Metric Description

To distinguish metrics of different cloud products, Cloud Monitor integrates and converts the metric names (metric English names in the metric documentation) of cloud products. The metric page provides information on metrics supported by Cloud Monitor integration, making it convenient for users to directly view and use.

Cloud Monitor (qcloud-exporter)

[Install](#) **[Metric](#)** [Dashboard](#) [Alarm](#) [Integrated](#)

① The list below supports all metric information that can be collected by this integration type. After installation, the actual metric collection situation is subject to the metric details on the integrated page.

Please select product type

Please enter the metric name

Total 1406 Metrics

Metric name	Metric Chinese Name	Product type	Metric Description	Unit
qce_cbs_diskawait_avg	DiskAwait	CBS	-	ms
qce_cbs_diskreadiops_avg	DiskReadIops	CBS	-	count
qce_cbs_diskreadtraffic_avg	DiskReadTraffic	CBS	-	KB/s
qce_cbs_disksvctm_avg	DiskSvctm	CBS	-	ms
qce_cbs_diskutil_avg	DiskUtil	CBS	-	%
qce_cbs_diskwriteiops_avg	DiskWriteIops	CBS	-	count
qce_cbs_diskwritetrafic_avg	DiskWriteTraffic	CBS	-	KB/s
qce_cdb_abortedclients_sum	AbortedClients	MySQL(CDB)	-	Count
qce_cdb_abortedconnects_sum	AbortedConnects	MySQL(CDB)	-	Count
qce_cdb_bytesreceived_max	BytesReceived	MySQL(CDB)	-	Bytes/s
qce_cdb_bytessent_max	BytesSent	MySQL(CDB)	-	Bytes/s

Cross-Account Collection

Note:

Cross-site collection is not supported (Chinese mainland site accounts and international site accounts cannot collect data from each other).

Scenario: Account A collects monitoring data from Account B through cross-account collection.

Configuration entries:

Create a Cloud Monitor integration in the Prometheus monitoring service instance under account A.

Select **Authentication type** as **Cross-account collection**.

Select **This account role** as the custom role created by Account A.

Enter the custom role created by Account B in **Target account role**.

Enter the root account ID of Account B in **Target account uin**.

Authentication

Authentication type *

Cross-account collection

This account role *

Select role

Target account role *

Please enter the target account role

Target account uin *

Please enter the target account uin

Brief Flowchart



Creating Custom Roles

Account A Users Creating Custom Roles

1. On the [Policies](#) page, create a [Custom Policy](#) using policy syntax, and add the sts:AssumeRole permission, which is used to assume the role of account B. The policy syntax is as follows:

```
{
  "version": "2.0",
  "statement": [
    {
      "effect": "allow",
      "action": ["sts:AssumeRole"],
      "resource": ["*"]
    }
  ]
}
```

Note:

If you need to limit permissions, such as only assuming a custom role of Account B, you can modify resource to "qcs::cam::uin/[Root account ID of Account B]:roleName/[Custom role of Account B]".

2. On the [Roles](#) page, click **Create Role**.
3. In the pop-up window for selecting the role entity, choose **Tencent Cloud Product Service** to enter the role information page.
4. Check **Cloud Virtual Machine (cvm)** as the role entity, select **Cloud Virtual Machine** as the use case, and click **Next**.
5. In the policy list, select the policy created in Step 1 as the role configuration policy, and click **Next**.
6. Tag the role with tag keys and tag values, which can be left blank, and click **Next**.
7. Enter your role name and click **Complete** to finish creating the custom role.

Account B Users Creating Custom Roles

1. On the role list page, click **Create Role**.
2. In the pop-up information window for selecting the role entity, choose **Tencent Cloud Account** as the role entity, and enter the role information page.
3. On the role entity information page, select other root account for **Tencent Cloud Account Type**, enter the Account A main account ID for **Account ID**, leave others blank, and click **Next**.
4. In the policy list, select the preset policy **ReadOnlyAccess** as the role configuration policy, and click **Next**.
5. Tag the role with tag keys and tag values, which can be left blank, and click **Next**.
6. Enter your role name and click **Complete** to finish creating the custom role.

FAQs

How to Configure "Data Pull Configuration"?

1. If the configuration is 0, Prometheus will use the current timestamp to overwrite the original timestamp of data.
Use case: Ensure the real-timeness of data timestamps to maximize the timely issuance of alarms by Prometheus.

2. If the configuration is a value x greater than 0:

As long as the value is greater than 0, Prometheus will retain the original timestamp of the data.

Use case: Keep the timestamps consistent with those on the console monitoring page.

Time window for delayed data pulls (latency equals x).

Background: To be compatible with the latency of monitoring data reporting links of cloud products, Prometheus pulls data within the time range of `(now-fixed latency, now)` by default.

Use case: If the reporting link latency of certain products is too high, set x here to change the time range for pulling data to `(now-fixed latency-x, now-x)`, to ensure that data can be retrieved to the greatest extent possible within this delayed window.

Are There Issues with Targets Display?

No collection objects: A newly-created integration needs to wait for a few minutes before displaying the correct targets.

(1/2)down: Because the integration uses rolling update, it will continue to collect from the old pod until the new pod runs successfully. During that period, two targets will be displayed.

Certain Cloud Product Failed to Collect Metrics

On the **Integrated** tab, you can check the following information:

Instance information: Check whether it contains the cloud product. If not, it means the cloud product was not selected.

Ensure that **Targets** are in **up** status.

Metric details: Check whether there are metrics for the cloud product. If there is, wait for a minute before querying again.

Cloud Monitor (qcloud-exporter)

Install

Metric







Dashboard

Alarm

Integrated

Create

Support search by name

Name	Type	Instance information	Running status	Charged met...	Targets	Operation
	Cloud Monitor	CBS,CL [redacted] ...	 Deployed	83.88 per second	(1/1) 	Metric details Delete Disable Log
	Cloud Monitor	CVM	 Deployed	16.98 per second	(1/1) 	Metric details Delete Disable Log

Ensure that there are cloud product instances in the selected region.

Check whether **Instance ID Filtering** or **Cloud Tag Key Filtering** is configured, and confirm that the corresponding configuration can help obtain the cloud product instance.

Check whether **Metric Relabel** is configured, and ensure that the corresponding configuration does not filter out the cloud product metrics.

Viewing Monitoring Information

Prerequisites

The Prometheus instance has been bound to a Grafana instance.

Operation Steps




1. Log in to the [TMP console](#) and select the corresponding Prometheus instance to enter its management page.
2. Click **Data Collection > Integration Center**, on the Integration Center page, find and click **Cloud Monitor**, select **Dashboard > Dashboard operation > Install/Upgrade** in the pop-up window to install the corresponding Grafana Dashboard.
3. Select **Integrated**. In the integrated list, click the Grafana icon to automatically open the list of Cloud Monitor integrated dashboards. Select the corresponding cloud product dashboard to view monitoring data related to the instance, as shown below.

Cloud Monitor (qcloud-exporter)

Install Dashboard Alarm **Integrated**

Create

Support search by name

Name	Type	Instance infor...	Running status	Charged metri...	Targets	Operation
test	 Cloud Monitor		 Deployed	0.05 per second	(1/1) 	Metric details Delete Disabl Log

Cloud Monitor

Starred

Sort (Default A~Z)

CBS

cloud_monitor

Cloud Mo

CDB

cloud_monitor

Cloud Mo

CKafka

cloud_monitor

Cloud Mo



Read Cloud-Hosted Prometheus Instance Data via Remote Read

Last updated : 2024-01-29 15:55:08

Overview

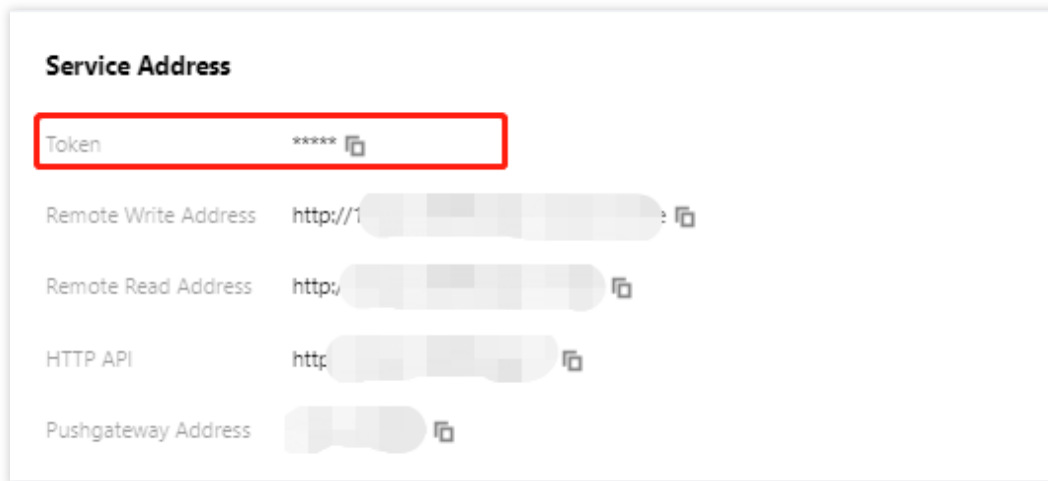
TMP provides the remote read API, which supports organizing a series of data sources of the Prometheus protocol into a single data source for query. This document describes how to use self-built Prometheus to read data from a cloud-managed TMP instance through the remote read API.

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...:

Remote Read Address http://...

HTTP API http://...

Pushgateway Address ...

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 read
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 the
headers:
[ <string>: <string> ... ]

# Whether to perform remote read query in the time range with complete local data series
[ read_recent: <boolean> | default = false ]

# Add Authorization header to each remote read request, and choose password or password
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 the
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 17:08:56

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.

Service Address

Token

***** 

Remote Write Address

Remote Read Address

HTTP API

Pushgateway Address

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/prometheus victoriametrics/vmagent:latest \
-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 `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 `OK` .

```
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](#).

```
global:
  scrape_interval: 30s
scrape_configs:
  - job_name: agent-monitor
    static_configs:
      - targets:
        - localhost:8429
```

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
```

Pushgateway Integration

Last updated : 2024-10-29 11:48:09

Application Scenario

Pushgateway is a crucial member of the Prometheus ecosystem. It allows any client to push custom monitoring metrics that comply with the standards, which are then collected and monitored by Prometheus. Prometheus Pushgateway is used to receive metric data from short-term tasks, which cannot be directly monitored through the service discovery monitoring system. Pushgateway allows temporary jobs (such as batch processing jobs) to push metrics to a central location, without directly exposing their metrics. Such data can be pulled and persistently stored by the Prometheus server.

One-Click Installation

1. Log in to [TCOP](#).
2. In the left sidebar, click **Managed Service for Prometheus**.
3. Select the corresponding Prometheus instance from the instance list.
4. On the instance details page, click **Data Collection > Integration Center**.
5. Search for **Pushgateway** in the integration center, and click it to pop up the installation window.
6. On the Installation tab of the pop-up window, fill in the relevant information as prompted and click **Save**.

Pushgateway (pushgateway)

Install Integrated

Current subnet [lucy-subnet-4](#) Remaining IP count: 190

Installation method [One-click installation](#) [Installation instruction document](#)

Pushgateway metric collection

name

Pushgateway Instance

scrape timeout

scrape interval

Pushgateway resource limits

CPU/core

Memory/Gi

Collector estimated resource occupancy ⓘ: CPU-0.25 cores Memory-0.5GiB

Configuration cost:**0.0018 Dollar/Hour** Original price:0.0065Dollar/Hour No charge for collect only free metrics[Billing explanation](#)

[Save](#) [Cancel](#)






Configuration Note

Parameters	Description
name	Exporter name, which should meet the following requirements: The name should be unique. The name should conform to the following regular expression: '[a-z0-9]([-a-z0-9]*[a-z0-9])?(\.[a-z0-9]([-a-z0-9]*[a-z0-9])?)*\$'.
scrape timeout	Pushgateway collection timeout, which is in time format and cannot be greater than the collection interval.

scrape interval	Pushgateway collection interval, which is in time format.
CPU/core	Number limit of Pushgateway CPU cores, which cannot be greater than 64.
Memory/Gi	Pushgateway memory limit. During the configuration, the value should include the unit Gi and cannot be greater than 512 Gi.

7. Obtain Pushgateway address information from the integrated list.

Search for the required CAM policy as needed, and click to complete policy association.

Pushgateway (pushgateway)						
<div>Install <u>Integrated</u></div>						
Name	Type	Instance information	Running status	Charged metrics	Targets	Operation
pushgateway	Pushgateway	    8080	 Deployed	0.00 per second	(-/-)	Metric details Delete Disable Log Monitoring ↗

Data Push

After Pushgateway is installed successfully, you can obtain the address for interaction and use this address to perform related operations on Pushgateway.

1. Obtain the component status:

```
curl -X GET http://10.*.*.*:8080/api/v1/status
```

2. Add a single data record to {job="some_job"}:

```
curl --location --request POST '10.*.*.*:8080/metrics/job/some_job' \\  
--header 'Content-Type: text/plain' \\  
--data 'some_metric 3.14'  
,
```

3. Add complex data to a specific instance:

```
curl --location --request PUT
'10.*.*.*:8080/metrics/job/some_job/instance/some_instance' \
--header 'Content-Type: text/plain' \
--data '# TYPE some_metric counter
some_metric{label="val1"} 42
# TYPE another_metric gauge
# HELP another_metric Just an example.
another_metric 2398.283
'
```

4. Delete all data under {job="some_job",instance="some_instance"}:

```
curl -X DELETE http://10.*.*.*:8080/metrics/job/some_job/instance/some_instance
```

5. Delete all data under {job="some_job"} (excluding data under {job="some_job",instance="some_instance"}):

```
curl -X DELETE http://10.*.*.*:8080/metrics/job/some_job
```

Viewing Monitoring Information

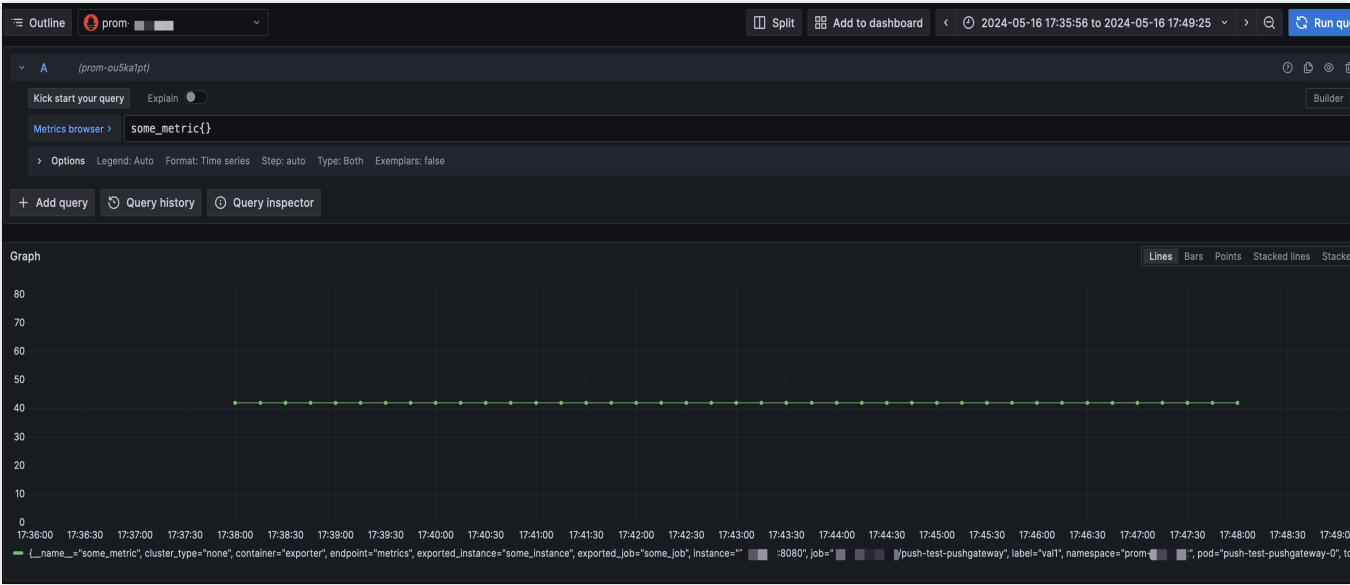
Prerequisites

The Prometheus instance has been bound to a Grafana instance.

Directions

1. Log in to the [TMP Console](#) and select the corresponding Prometheus instance to enter its management page.
2. On the **Basic Info** page of the instance, find the bound Grafana address, open it, and log in to Grafana. Then, you can view the pushed metrics in Explore or create a panel to view metrics:

Search for the required CAM policy as needed, and click to complete policy association.



Security Group Open Description

Last updated : 2024-08-15 17:08:56

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

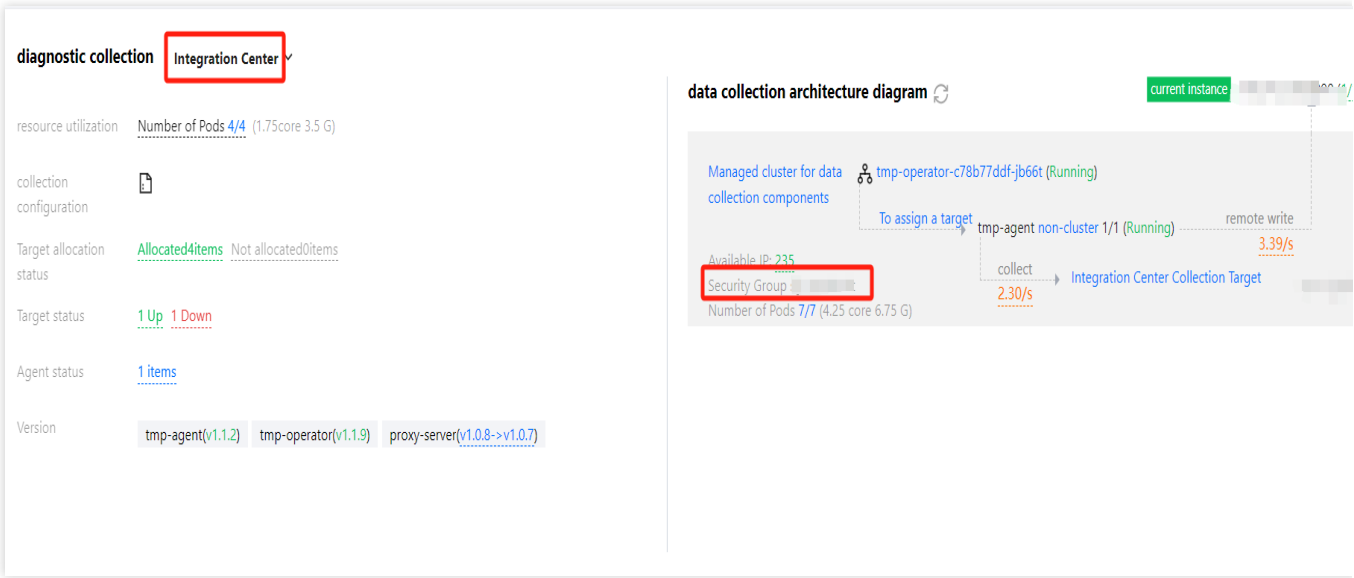
Port Description

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.



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

Details

Operation logs

Node pool information

Node pool name

np-emkhjrii(test)

Node pool status

Running...

Maintenance level

Medium

K8s version

1.26.1-tke.7

Number of nodes

Current number: 1, desired number: 1

Time created

2024-08-14 17:33:47

Deletion Protection

Enabled

Security reinforcement

Disabled

Tag

View

Node launch configuration info

Operating system

TencentOS Server 3.1

Billing mode

Pay-as-you-go

Supported subnets

Security group

Bind an SSH key

Model

SA2.MEDIUM2(Primary)

System disk

Premium cloud disk 50GB

Data disk

-

Node name

Auto-generated

Operation configuration

Node self-heal

Enabled

Self-healing rule

wudi

qGPU sharing

Disabled

Auto scaling

Enabled(min number of nodes: 0, max number of nodes: 1)

Scale-out policy

Preferred availability zone first

Removal Policy

Remove the latest instance

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**:

Cluster(Guangzhou) / (wudi2) Operation G

Details Scaling logs

Node pool information

Node pool name	(wudi2)	Scaling group name	
Node pool status	Running...	Launch configuration name	
Labels/Taints/Annotations	View	Number of nodes in the scaling group	Current number: 1, desired number: 1
Number of manually-added nodes	0	Retry policy	Retry with incremental intervals
Auto scaling	On(Min nodes:0,Max nodes:1)	Tag	View
Scaling mode	Release mode	Deletion Protection	Enabled
Instance creation policy	Preferred availability zone (subnet) first		
Removal Policy	Remove the latest instance		

Node configuration details

Operating system	TencentOS Server 3.2 (Final) Public image -Basic image	Runtime components	containerd 1.6.9
Model	SA2.MEDIUM2(Primary)	Subnet	
Data disk	-	Custom data	View
Custom Kubelet parameters	View	Placement group	

[Adjust quantity](#) [Add existing node](#) [Remove](#) [More](#)

Separate filters with carriage return

<input type="checkbox"/>	Node ID/name	S..	Availabili...	Configuration	Removal ...	IP address	How to...	Resource us...	Billing mode	Operation
<input type="checkbox"/>	View details									
<input type="checkbox"/>	as-tke-n...	Health	Guangzho...	SA2.MEDIUM2 2 core, 2 GB, - Mbps System disk: 50 GB	Disabled		Scaling ...	CPU: 0.11 / 1.90 core Memory: 0.21 / 1.07 Gi	Pay-as-you-go Created by 2024-1	Service Upgrade Remove More

Total items: 1

20 / page

1 / 1 page

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

The screenshot shows the Tencent Cloud console interface for an instance named 'as-tke-np-...'. The instance is in a 'Running' state. Below the instance name, there is a note about the initial login name and password, and a row of action buttons: 'Log in', 'Shutdown', 'Restart', 'Reset password', 'Terminate/Return', and 'More actions'. Below this, there is a tabbed interface with the following tabs: 'Basic information', 'ENI', 'Public IP', 'Monitoring', 'Security groups' (which is highlighted with a red box), 'Operation logs', 'Run commands', and 'Uploading a file'. The 'Security groups' tab contains two panels: 'Bound security groups' and 'Rule preview'. The 'Bound security groups' panel shows a table with one entry: a priority of 1, a security group ID/name of 'cm-eks-cls-iexxi79u-security-group', and an 'Unbind' operation. The 'Rule preview' panel shows 'Inbound rules' and 'Outbound rules' tabs, with one rule listed: 'cm-eks-cls-iexxi79u-security-group' with an 'Edit rule' link.

as-tke-np-... Running Expand

The initial login name is root. If you select "Random password" when purchasing the instance, check the password in [Message Center](#). You can [reset the password](#) if you forget it.

[Log in](#) [Shutdown](#) [Restart](#) [Reset password](#) [Terminate/Return](#) [More actions](#) ▼

Basic information ENI Public IP Monitoring **Security groups** Operation logs Run commands Uploading a file

Bound security groups [Sort](#) [Configuration](#)

Priority ⓘ	Security group ID/na...	Operation
1	cm-eks-cls-iexxi79u-security-group	Unbind

Rule preview

Inbound rules Outbound rules

▶ cm-eks-cls-iexxi79u-security-group [Edit rule](#)

Super Nodes

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

Node pool information

Node pool name	
Node pool status	Running...
Security group	sg-
Labels	View
Taints	View
Deletion Protection	Enabled
Node type	Linux

Create

Remove

Renew

Set to "Auto-renewal"

Set to "Manual renewal"

Cordon

Uncordon

Related Issues

Issue Description

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









		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 resourceInformer failed,Object:TMPAgent/pro /tke-cls-,Message:failed to sync cache ServiceMonitor informer}; context deadline exceeded	

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

Cluster name	
Cluster ID	
Deployment type	General cluster
Status	Running... 
Region	South China(Guangzhou)
Addition of Resource Allocated Project 	DEFAULT PROJECT 
Cluster specification	L5  <div>The application size does not exceed the recommended management size. Up to 5 nodes, 150 Pods, 128 ConfigMap and 150 CRDs are allowed under the current cluster specification. Please read Choosing Cluster Specification  carefully before you make the choice.</div> <div><input checked="" type="checkbox"/> Auto Cluster Upgrade </div> <div>Check specification adjustment history</div>
Kubernetes version	Master 1.26.1-tke.3(Updates available) Upgrade Node 1.26.1-tke.7、1.26.1-tke.3(Updates available) Upgrade

Node and Network Information

Number of nodes	3
	Check CPU and MEM usage on Node Map
Default OS	tlinux3.2x86_64 
System image source	Public image - Basic image
Node hostname naming pattern	Auto-generated 
Node network	vpc-  
Container network add-on	Global Router
Container network	<div>CIDR block</div> <div> Current VPC is not associated with any CCN instance</div> <div>Up to 1024 services per cluster, 64 Pods per node, 16 nodes per cluster</div>
Network mode	cni
VPC-CNI mode	<input type="checkbox"/> Disabled
Service CIDR block	172.16.4.0/22
Kube-proxy mode	iptables

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

Basic Info

Name		Region	Guangzhou	Network	
Instance ID		AZ	Guangzhou Zone 4	Subnet	
Status		Billing Mode	Pay as you go	IPv4 Address	
Tag		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

Help of VPC and Subr

Basic information

Classiclink

Monitoring

Basic information

ID

Name

test

IPv4 CIDR

many)

17

DNS

8

Domain Name

-

Tags

No tags found.

Associate with CCN

CCN provides multi-point intranet interconnection service between VPCs, or between VPCs and customer IDCs. [Learn more](#)

The current VPC is not associated with any CCN instance. [Associate now](#)

Select Create Public Network CLB Instance:

Associate Cluster

• Remaining IP(s) of the subnet [10.0.0.4]: 160

Cluster Type

Standard cluster

Cross-VPC Association

☒ Enable

After enabling this option, you can use one instance to monitor clusters in different regions and VPCs.

☒ Create Public Network CLB Instance

If the VPC where your instance resides is not interconnected with the cluster to be associated, you must create a public network CLB instance because data cannot be collected otherwise. You don't need to do so if the two are already interconnected.

Cluster Region

Chengdu

Tencent Cloud services in different regions cannot communicate with each other over the private network. We recommend that you select a region closest to your end users to minimize access latency and improve the download speed. You cannot modify the 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:

[illegible]

Check if the security group rules meet the requirements.