


TencentCloud Managed Service for Prometheus FAQs Product Documentation



Copyright Notice

©2013-2024 Tencent Cloud. All rights reserved.

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

Trademark Notice



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

FAQs

Basic Questions

Integration with TKE Cluster

Product Consulting

Use and Technology

FAQs

Basic Questions

Last updated : 2024-08-15 17:58:46

How to use Prometheus's native default UI and other features?

TencentCloud Managed Service for Prometheus (TMP) is different from the open source stand-alone Prometheus. TMP is of structure with separated collection and storage and does not provide native default UI features. Use Tencent Cloud Managed Service for Grafana as an alternative for querying features.

When the alarm is recovered, the \$value in the notification template is incorrect. How to handle it?

The \$value during the alarm recovering is the value that meets the alarm expression condition for the last time. The value that does not meet the condition cannot be obtained. Regarding the design, the alarm expression is a whole. The alarm expression has no difference from the PromQL query in the normal scene. If the series in the query result meets the duration, the alarm will be triggered. When the next query result does not contain a series, the alarm corresponding to the series will be recovered. Prometheus cannot deconstruct and interpret the alarm expression by itself, because some expressions themselves do not contain comparison relationships such as thresholds, for example: `a and b` and `123456789`.

How to view the alarm records?

Prometheus itself does not have the concept of alarm records. Due to special reasons, it cannot fully support the feature of alarm records, but the related alarms and status can be viewed through the ALERTS or ALERTS_FOR_STATE metrics. In addition, all Prometheus alarms are currently delivered to the Tencent Cloud Observability Platform Alarm Service, and then sent to the user side through various notification channels of the Platform. The alarm records feature provided by the Tencent Cloud Observability Platform can make up for this defect to a certain extent. Due to conceptual differences in design, it cannot be used as a complete alternative but can be used for troubleshooting and other requirements. You can view records through the [Alarm Records Feature](#).

Is the alarm repeat interval inconsistent with the configured duration?

Prometheus alarm repeat interval (`repeat_interval`) is not set for a single alarm but for the entire group. When there is a new alarm or alarms are recovered in the group, the entire group will be notified, and then notifications will be sent at a repeated interval. Currently, [TencentCloud Managed Service for Prometheus](#) is grouped by a single alarm policy. For example, a single alarm policy may be for the restart of instance A/B/C. When A/B restarts, the alarm is triggered and then notifications will be sent at a repeated interval. After a period of time, C also restarts, and the interval will be interrupted and recovered. At this time, it will be notified to the user taking A/B/C as a Group. Currently, due to implementation limitations, the alarms delivered based on the Tencent Cloud Observability Platform cannot be

grouped and aggregated as a single message to notify users. If necessary, you can use a custom Alertmanager or Tencent Cloud Managed Service for Grafana.

Why doesn't the rate/irate function produce any data when the original metrics exist?

The rate/irate function requires at least two data points to calculate, so the time range calculated by rate/irate must cover at least two data points. Considering the possibility of data point loss due to network exception, the official recommendation for the time range is **four times the collection interval**.

Why does the rate/irate function calculate a very large outlier?

rate/irate functions can only be used for [Counter](#) type metrics, which are defined as strictly increasing numbers. Querying through Prometheus will process the problem of Counter being reset to 0, such as server restart. Normally, this will not affect the calculation results unless there is a problem of data point disorder. For example, the disorder of two second-level data points 9999 and 10000 results in an outlier of $(10000+9999)-10000 = 9999$ (normally it should be 1). The typical scene of this situation is as follows:

There are multiple collection components collecting the same metric and repeatedly reporting it to the same Prometheus instance, which may cause disorder problems, resulting in a large outlier calculated after the reset processing logic of Prometheus. Solution: If the collection component collects the same metric, it does not meet the expectations. It is recommended to troubleshoot and solve the problem of repeated collection. If the collection solution is designed to ensure the high availability of the collection component, different collection components need to add replica or other tags to distinguish all collected metrics.

Report directly to [TencentCloud Managed Service for Prometheus](#) through pushgateway and other methods and no timestamp information is included. [TencentCloud Managed Service for Prometheus](#) side will store the current time as the timestamp of the metric. If the network jitters, the arrival time of the data points will become disordered or the timestamp information granted by different processes/threads processing different data points will be disordered, resulting in calculation errors. The solution is to include the timestamp information when reporting. [TencentCloud Managed Service for Prometheus](#) built-in pushgateway reporting method is only a supplement to remotewrite and does not realize the complete feature. It is recommended to use Agent(remotewrite) for collection and reporting when it is not necessary.

Why is the interval of data points returned by the query different from that of capturing?

The interval of data points returned by querying through Prometheus is determined by the query parameter interval/step. Each data point is filled in and aligned strictly according to the interval/step, and there is no one-to-one correspondence with the capturing interval. If there is too much data loss or the capturing interval is too large, data will not be filled in, and the storage side will not store any information related to the capturing configuration. Users need to process their capturing configuration.

Why does the query return extra data for the last five minutes?

By default, Prometheus will fill in data for certain queries. Even if there is only one data point in the last five minutes, it may return multiple data points in the past five minutes (according to the step/interval parameters of the query). This is the default behavior of open source Prometheus and cannot be adjusted currently. Generally, it does not affect normal use.

How do time-related functions in PromQL handle local time zones?

All time in Prometheus is in UTC time. There are no exceptions, and the concept of time zone is not considered in the design, which may cause `day_of_*` series of functions to be inconvenient to use, and the official support for them can not be realized in the short term. The temporary solution for the domestic time zone can be as follows:

```
day_of_week(vector(time() + 8*3600)) .
```

Can Prometheus-related use limits be adjusted?

The use limits related to Prometheus are adjustable. In most cases, exceeding these use limits may lead to degradation of user experience and performance. Therefore, adjusting the limits upward cannot guarantee the query and write performance benchmarks before adjustments. The service level agreement may no longer apply. Users need to have certain psychological expectations for potential related risks.

From the chart, it seems that the alarm meets the duration but is not actually triggered?

The basic principle of Prometheus alarm detection is to use Instant Query to query data every minute, and trigger alarms if the alarm condition is (continuously) met. In some cases, Prometheus's Range Query will fill in the missing data, and the continuous timeline on the chart may be discontinuous under the execution logic of the alarm component. In addition, due to the characteristics of alarm scheduled inspection, the time when the alarm is detected may be slightly delayed. You can check the status of the alarm by querying the `ALERTS` or `ALERTS_FOR_STATE` metrics.

Integration with TKE Cluster

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

What should I do if I can't access the TKE cluster over the private network?

When installing an agent, you need to access TKE over the private network. If **private network access** is not enabled for the corresponding TKE cluster, the installation will fail. You can solve this problem by following the steps below:

1. Log in to the [TKE console](#) and select a container cluster in the corresponding region.
2. Enable **Private Network Access** under **Basic Information > Cluster API Server Information**.

What should I do if the status of all kube-proxy collection targets is "Down"?

In TKE, the launch parameter `--metrics-bind-address` is not specified for kube-proxy, and the default listening address of the metrics service is 127.0.0.1; therefore, the agent cannot pull metrics by Pod IP. You can solve this problem by following the steps below:

1. Log in to the [TKE console](#) and select a container cluster in the corresponding region.
2. Go to **Basic Information > Cluster API Server Information > Connecting to Kubernetes cluster through Kubectl** to configure Kubectl as instructed.
3. Run `kubectl edit ds kube-proxy -n kube-system` and add the launch parameter `--metrics-bind-address=0.0.0.0` in `spec.template.spec.containers.args`.

What should I do if the status of all component collection targets on the master node in a dedicated TKE cluster is "Down"?

The inbound rule of the default security group of the master node in a dedicated TKE cluster does not allow access to the metrics ports of some components. You can solve this problem by following the steps below:

1. Log in to the [Security Group console](#) and select the corresponding region.
2. Enter `tke-master-security-for-<tke cluster id>` in the security group search box. For example, if the cluster ID is `cls-xxx`, then enter `tke-master-security-for-cls-xxx`.
3. Click the returned security group ID to enter the **Edit Inbound Rule** window.
4. The protocol port column of the rule to be edited should include `TCP:60001,60002`. Select the rules one by one and add ports 10249, 10252, 10251, 9100, and 9153 for the following purposes respectively:

10249: kube-proxy metrics port

10252: kube-controller-manager metrics port

10251: kube-scheduler metrics port

9100: node-exporter metrics port

9153: core-dns metrics port

Product Consulting

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

Does TMP support customizing the reported data?

Yes. TMP supports customizing the reported metric monitoring data in multiple languages and displaying such data in the integrated Grafana dashboards.

How does TMP collect monitoring data?

TMP can pull data through the Prometheus agent or write data through Pushgateway. It is fully compatible with open-source Prometheus in terms of collection method.

Is there a separate exporter for each instance?

Each MySQL or Kafka instance has one exporter, while multiple Redis instances can share the same exporter.

What is the difference between TMP and a self-built Prometheus service?

TMP is fully compatible with the open-source Prometheus ecosystem and is integrated with CM to help you quickly set up a monitoring system (for custom monitoring, component monitoring, basic monitoring, etc.). It supports Grafana, with preset common monitoring dashboards. It also supports diverse exporters, with preset common alarm templates. It well solves various problems with open-source Prometheus, such as complicated HA setup, poor performance, and labor-intensive Ops.

Which Tencent Cloud products does TMP support?

TMP supports CVM, TencentDB for MongoDB, TencentDB for MySQL, TencentDB for PostgreSQL, TencentDB for Redis, ES, and TKE. For more information, see [Integration Center](#).

Can I integrate other data sources on the Grafana page of TMP?

No. We recommend you integrate them by using [Grafana](#).

Can I pull monitoring data in Prometheus pull mode?

No, you can only pull data through APIs. For detailed directions, see [Querying Monitoring Data](#).

What is the series storage limit on pay-as-you-go TMP instances?

Each instance can have up to 4.5 million series. If you need to adjust it, [submit a ticket](#) for application.

Does TMP support integration with EKS?

Yes. Pay-as-you-go TMP instances support general clusters, EKS elastic clusters, edge clusters, and registered clusters.

Use and Technology

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

How do I migrate from self-built Prometheus to TMP?

In the configuration file of self-built Prometheus, add a remote write configuration pointing to TMP for migration. For more information, see [Migration from Self-Built Prometheus](#).

Can I batch import/export dashboards into/from Grafana?

You can do so through APIs as instructed in [HTTP API reference](#).

What should I do if one exporter has too much data?

You can use Prometheus in this scenario, but we recommend you not expose too many metrics in the exporter, as the exporter itself and the Prometheus agent consume a lot of memory. Only expose key metrics, and do not use parameters such as user ID and URL in metric labels.

How do I use TMP to monitor clusters in two different VPCs?

1. You can interconnect the VPCs of the two clusters through [CCN](#) and then integrate the clusters into Prometheus .
2. Install the [agent](#) in one of the clusters, and then use the Nginx proxy to remotely write the target address to the TMP instance.

How do I integrate native container services with TMP?

You can't integrate them directly. However, you can create data and add a remote write configuration pointing to TMP in the configuration file of your self-built Prometheus. For detailed directions, see [Migration from Self-Built Prometheus](#).

How do I set the parameters of the alarm cycle when creating an alarm rule through Prometheus APIs?

Add `key=_interval_ value=1m/5m/10m/15m/30m/60m` to the new `Labels` parameter.

Will data get lost during instance rebooting?

Data will not get lost during instance rebooting, as it is stored in TencentDB for CTSDB. However, because data cannot be reported normally during rebooting, data breakpoints may occur.

Is it normal to see a rise in generated data after the TMP instance is rebooted?

Retries will be performed after instance rebooting, so it is normal if the data volume seems to fluctuate greatly in a short period of time.

Can I batch add MySQL instances when integrating MySQL with Prometheus in the integration center?

No. You can only add instances one by one.

Where can I view the security group of the EKS cluster created by Prometheus?

Go to **Security > Security Group** in the [VPC console](#) and search by the current TMP instance ID.

Can I configure multiple scrape tasks for a PodMonitor?

Yes, but you should ensure that the tasks have the same port name and label.

(TMP Pushgateway usage) If the clients of multiple instances push the same job with different label metrics, metrics will be overwritten. What should I do?

You can use `groupingkey` to solve this problem. Below is the sample code:



```
if err := push.New("http://$IP:$PORT", "db_backup").  
    BasicAuth("$APPID", "$TOKEN").  
    Collector(completionTime).  
    Grouping("instance", "$INSTANCE").
```