Application Performance Management Operation Guide

Product Documentation

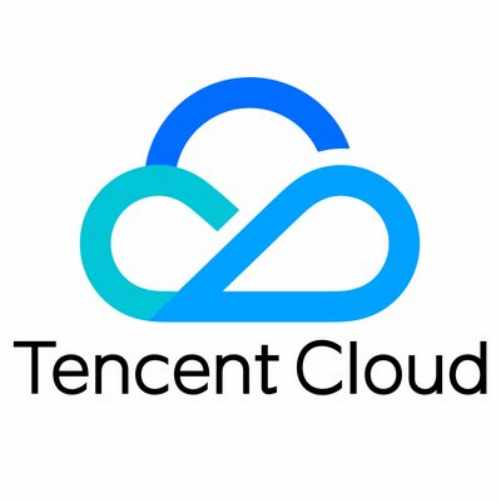

Contents

Operation Guide
Resource Management
Creating Business System

Last updated：2023-12-25 15:47:08

This document describes how to create a business system to manage your applications by category. You can set different retention periods and report quotas for different business systems.

Directions

1. Go to the **Resource management** page in the APM console.
2. Click **Create** to pop up the **Create business system** window.
3. Configure the system information as follows.

Configuration Item	Description
Business system name	Enter a custom business system name.
Billing mode	Currently, only pay-as-you-go is supported.
Reporting region	Guangzhou, Shanghai, Beijing, Hong Kong (China), Beijing Finance, Shanghai Finance, Singapore, and Silicon Valley are supported currently.
Retention period	It can be **1 day**, **3 days**, **7 days**, **15 days**, or **30 days**. During the trial, the default value is **1 day**. The longer the time, the higher the fees. Traces stored longer than the retention period will not be displayed in the APM console.
Business system overview	Describe the purpose of the business system.
Tag	APM can be used with the Tencent Cloud resource tag feature to perform tag-based sub-account authorization and cost allocation. For more information, see Access Management.

4. After completing the configuration, click **OK**.

Application Monitoring
Application List

Last updated：2023-12-25 15:47:27

Prerequisites

Go to the **Application list** page in the APM console.

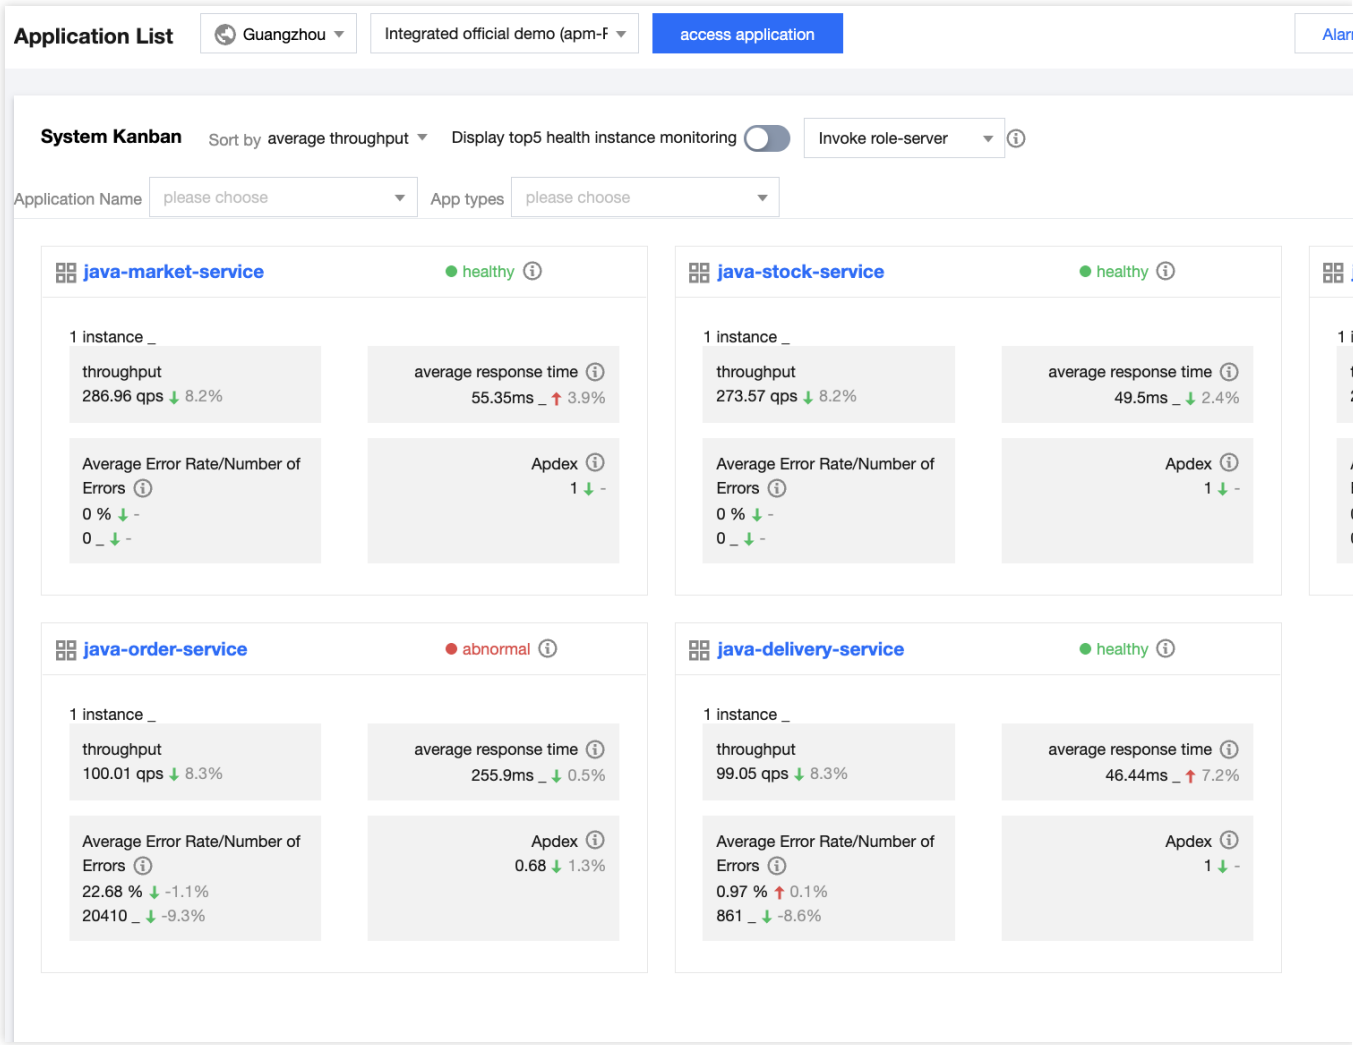

System Overview

The system overview section displays the core performance metrics and call dependencies of all services and components in the selected region as a matrix, list, or topology.

Application access

You can access a new service or add more services as follows.

1. At the top of the overview page, click **Access application**.

2. On the access guide page, select the programming language, access method, and reporting method and complete reporting as instructed in the console.

Note:

APM supports multiple programming languages with different access methods. For more information, see Access Guide.

System dashboard

The system dashboard displays the health, number of instances, throughput, average response time, average error rate, average number of errors, Apdex, and change rate of each metric of the accessed application as a list (default) or topology.

Below each metric, day-over-day comparison data is available to help you better monitor the fluctuations of the application performance. You can also click the target application to drill down on the **Application monitoring** page, where you can view more detailed monitoring data of the application.

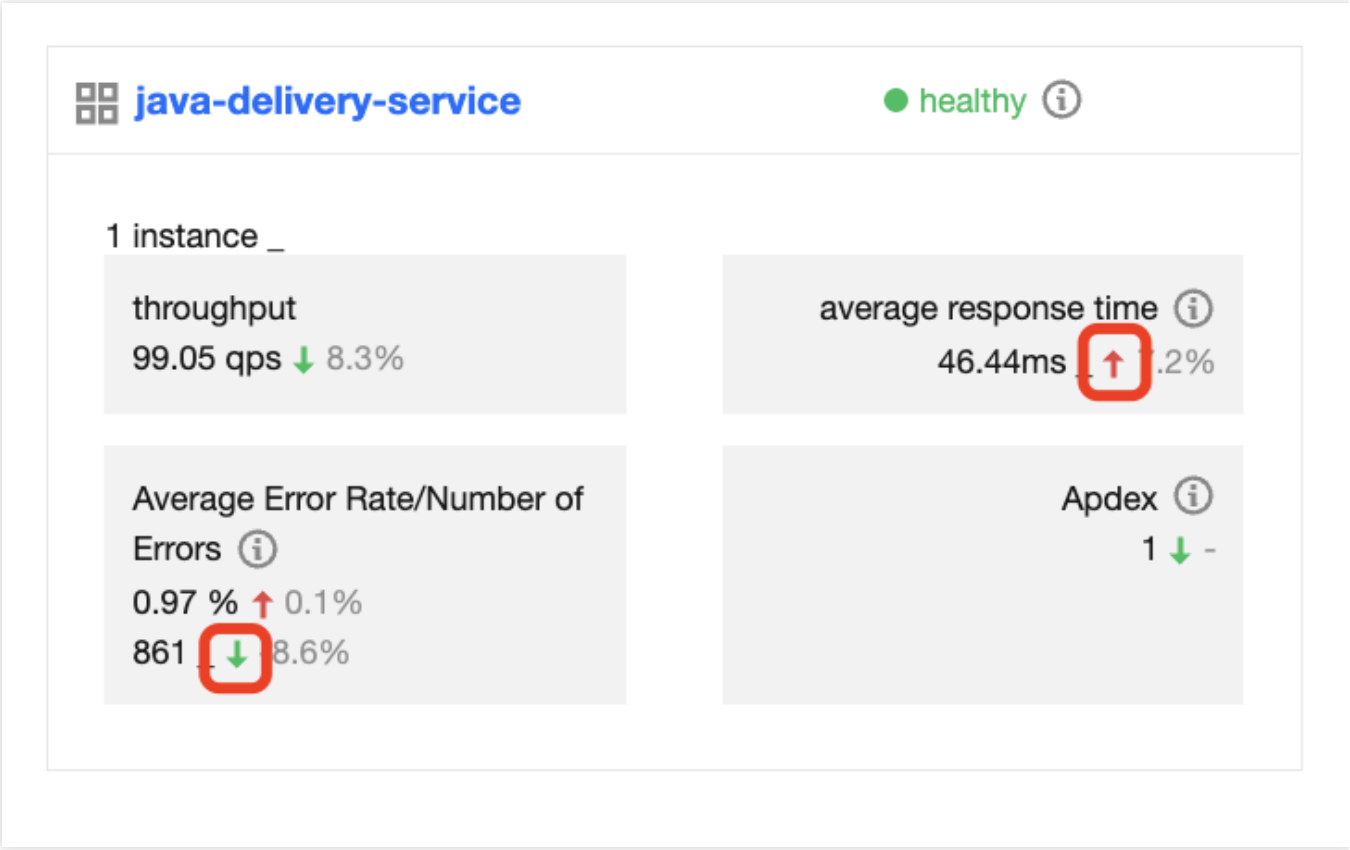

Analysis order

By default, services are sorted by health so that you can focus on those with a high error rate and slow response. In the top-left corner of the dashboard, you can switch between sorting dimensions, such as average response time, average throughput, average error rate, Apdex, and health status. If you select **Avg response time**, the slower the response, the higher the order of the service in the system dashboard.

Displaying top 5 healthy instances monitoring

After this option is enabled, sorting will be based on the health status, and top 5 instances will be displayed at the bottom of each module.

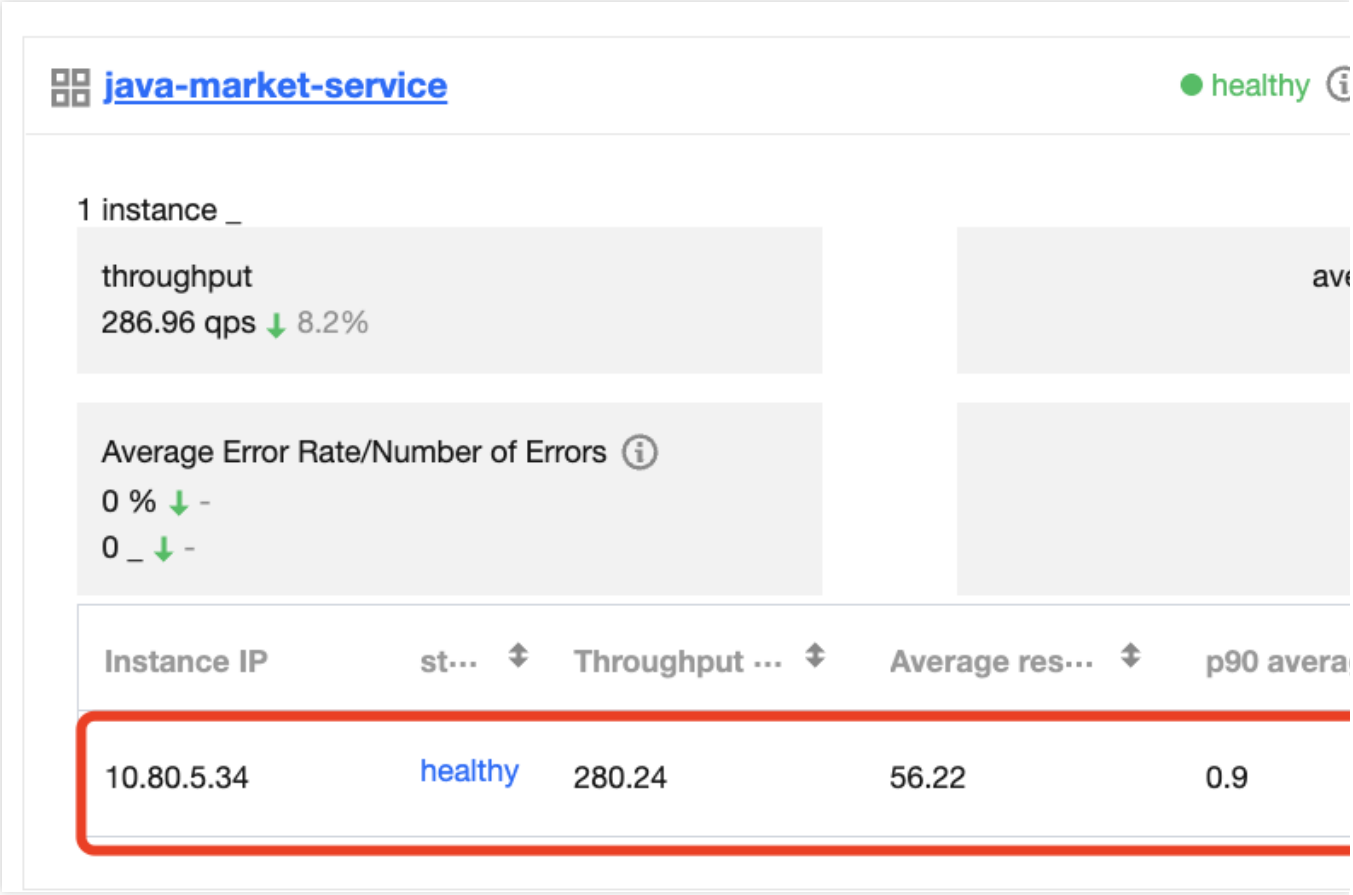

Switching reporting roles

You can switch between the application data on the client and server for isolated data analysis.

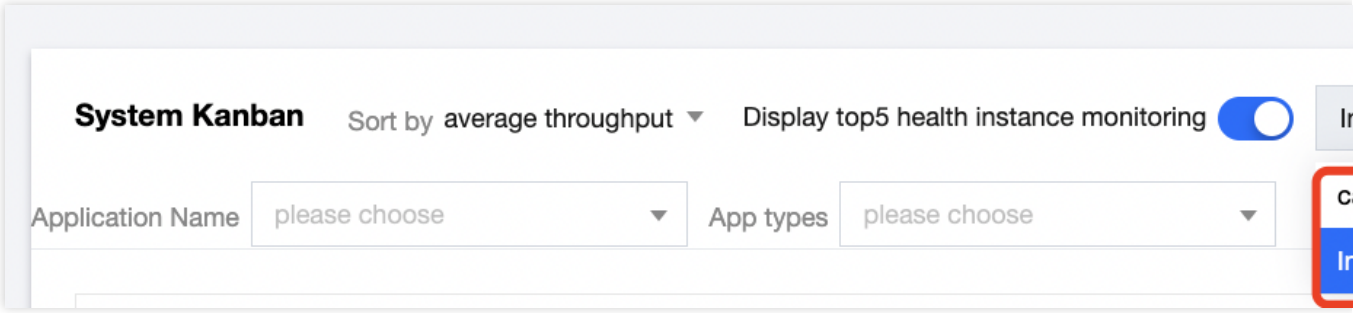

Global topology dependency

In addition to the list mode, you can also use the topology mode by clicking **Topology** in the top-right corner on the **Application list** page, which helps you organize the dependencies and call relationships between services.
Hover over the target node to view the total number of requests, average response time, and average error rate of the current application. You can also double-click the node to drill down on the **Application details** page.

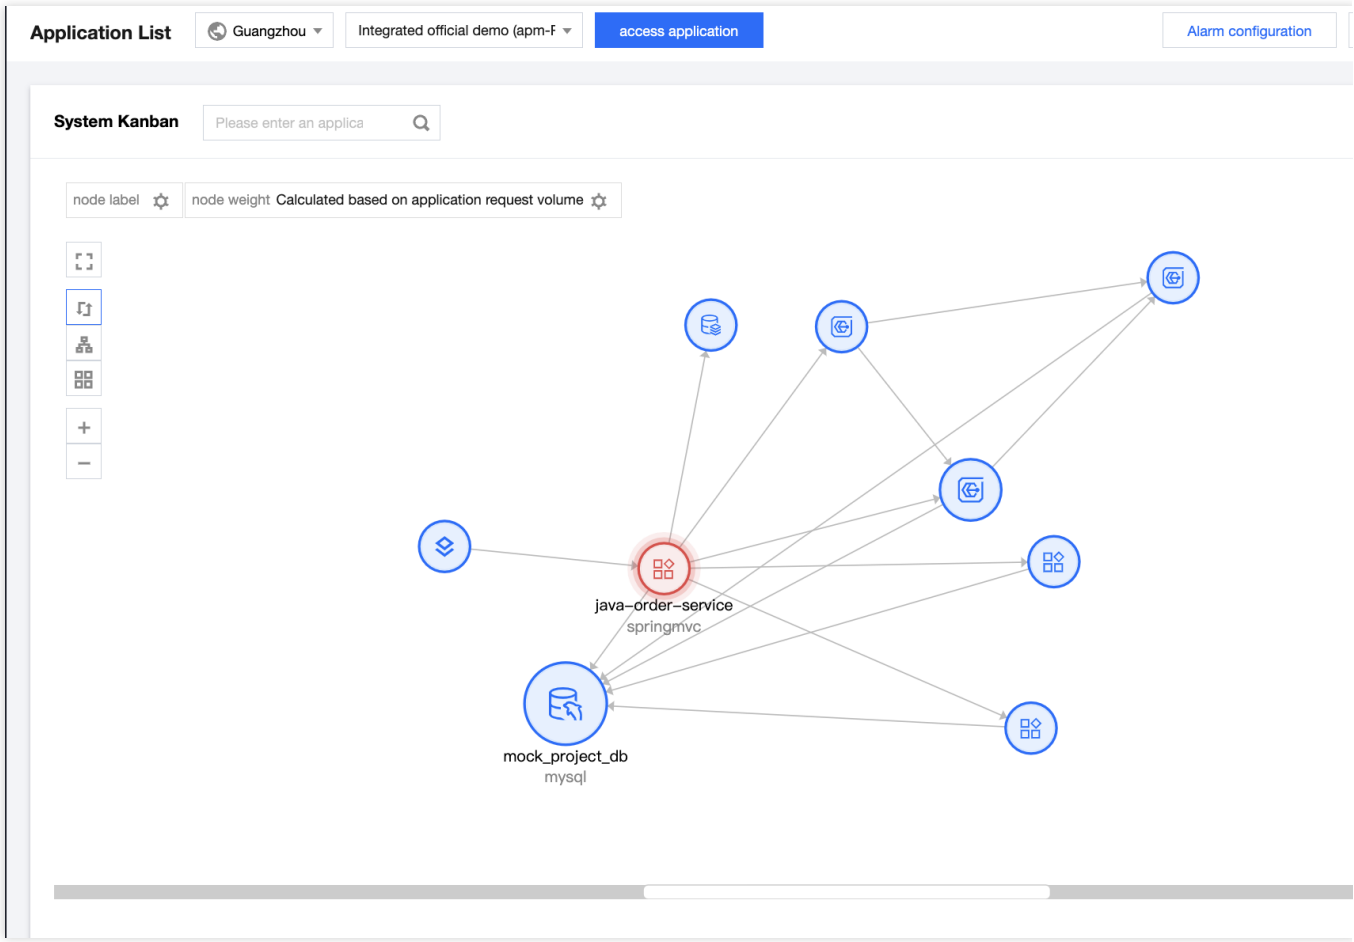

Metric description

Metric	Description

Health status	Assesses whether your application is healthy based on the response time and error rate. Healthy: Your application is assessed as healthy based on the response time and error rate. Warning: The current average response time of your application exceeds the satisfaction threshold, but the average error rate does not exceed the alarm threshold. Abnormal: The average error rate of your application exceeds the alarm threshold. By default, the satisfaction threshold for the average response time is 500 ms, and the alarm threshold for the average error rate is 5%.
Response time	The average response time of all service APIs of the application. The percentage below shows the day-over-day change.
Throughput	The average number of requests at the 1-minute granularity.
Error rate	The average error rate of all APIs at the 1-minute granularity among the call relationships of the selected service. You can click a data point on the curve to drill down to typical requests in the selected time range.
Apdex	The user satisfaction metric calculated based on the response time of all APIs at the 1-minute granularity among the call relationships of the selected service.
Change rate	The rate of change between the current value and the value yesterday.

Health status description

Healthy: Your application is assessed as healthy based on the response time and error rate.

Warning: The current average response time of your application exceeds the satisfaction threshold, but the average error rate does not exceed the alarm threshold.

Abnormal: The average error rate of your application exceeds the alarm threshold.

The satisfaction threshold for the average response time is 500 ms by default.

The alarm threshold for the average error rate is 5% by default.

Application Details

Last updated：2023-12-25 15:47:43

This document describes how to view application details, such as the topology, number of requests, response time, number of errors, and throughput.

Application Details

1. Go to the **Application details** page in the APM console.
2. On the **Application details** page, you can view the total number of requests, average response time, average error rate, slow call, slow SQL, number of exceptions, and full GC of the service in the current time range.

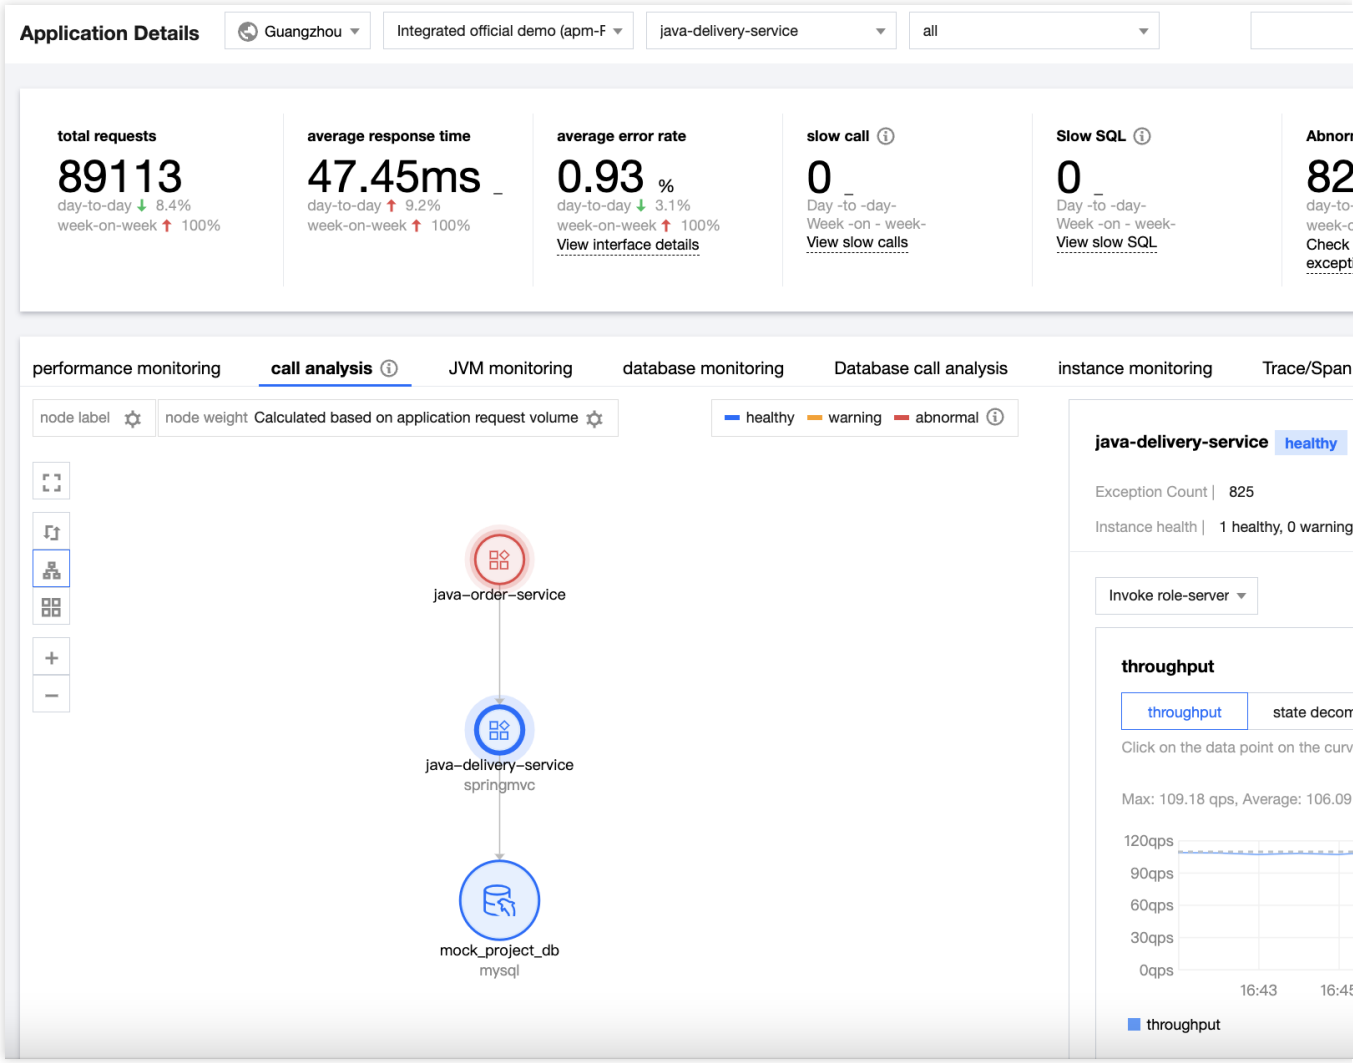

Performance monitoring

This module displays the trend of the average response time, average throughput, average error rate, and Apdex of the application on the selected server or client within the specified time range. You can click the clock icon in the top-right corner of each chart to compare the data of any day in the last 30 days with the current data.

Call analysis

This module displays the upstream and downstream local topology centered on the current service. Hover over the target node to view the average throughput, response time, and error rate of the corresponding application.
APM uses topology icons in different colors for identification. Green indicates that the application is healthy, orange delayed, and red abnormal.

JVM monitoring

This module displays the trend of critical JVM metrics, including average/maximum garbage collection (GC) count, CPU utilization, heap space, NoHeap space, heap space refinement, and number of JVM threads.

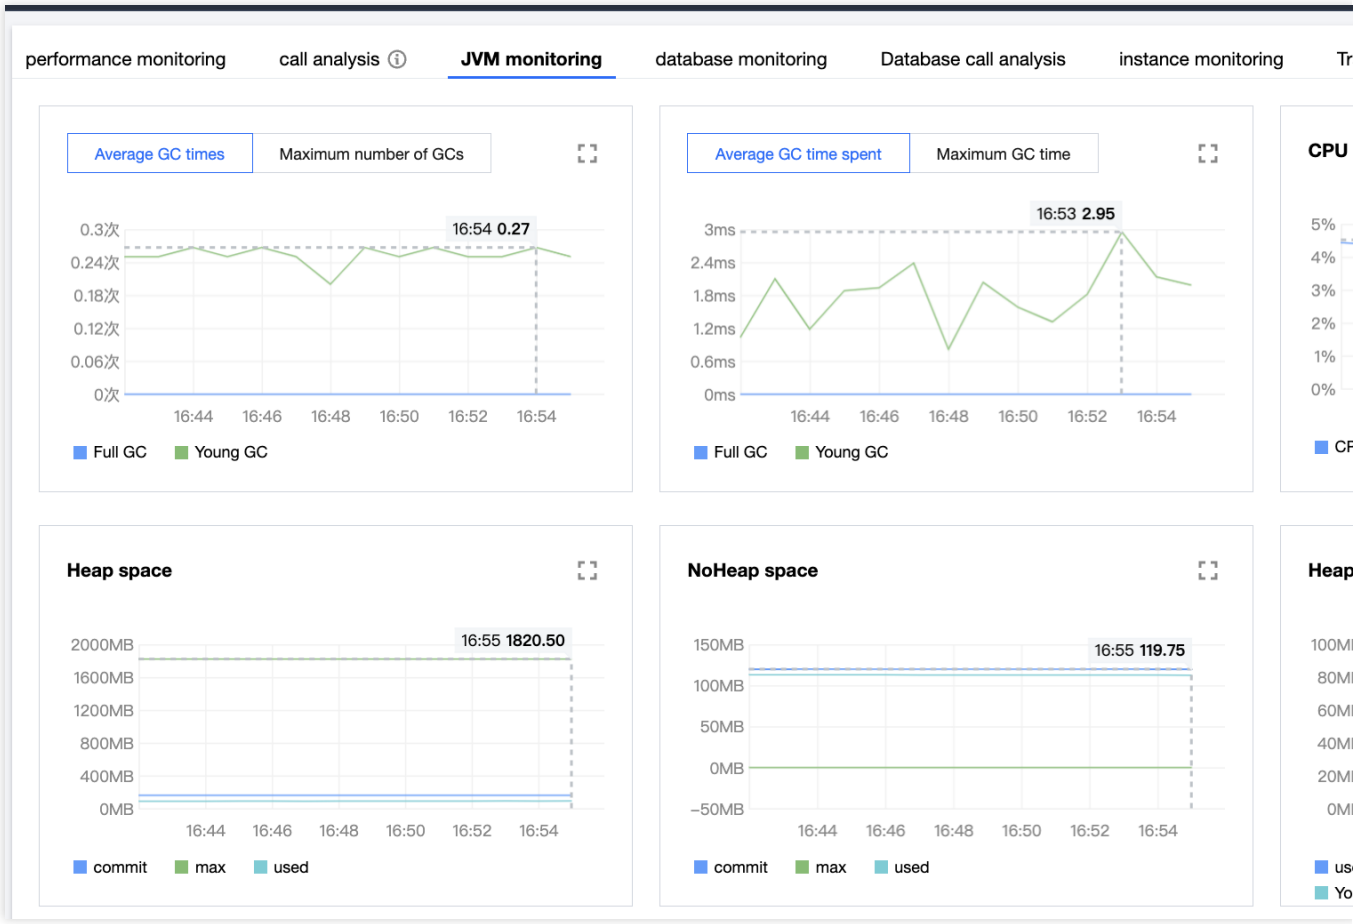

Module	Monitoring Metric
GC count	Full GC count
	Young GC count

GC time	Full GC time
	Young GC time
Heap/NoHeap space	Committed
	Max
	Used
Heap space refinement	Used young generation
	Eden space
	Survivor space
	Old generation
JVM threads/minute	Number of TIMED_WAITING threads
	Number of WAITING threads
	Number of RUNNABLE threads
	Number of created threads
	Number of terminated threads
	Number of blocked threads

Database monitoring

This module displays the information of database calls, including the average response time, throughput, and top 5 slow calls. You can also click **Overview** or **View SQL** for more details.

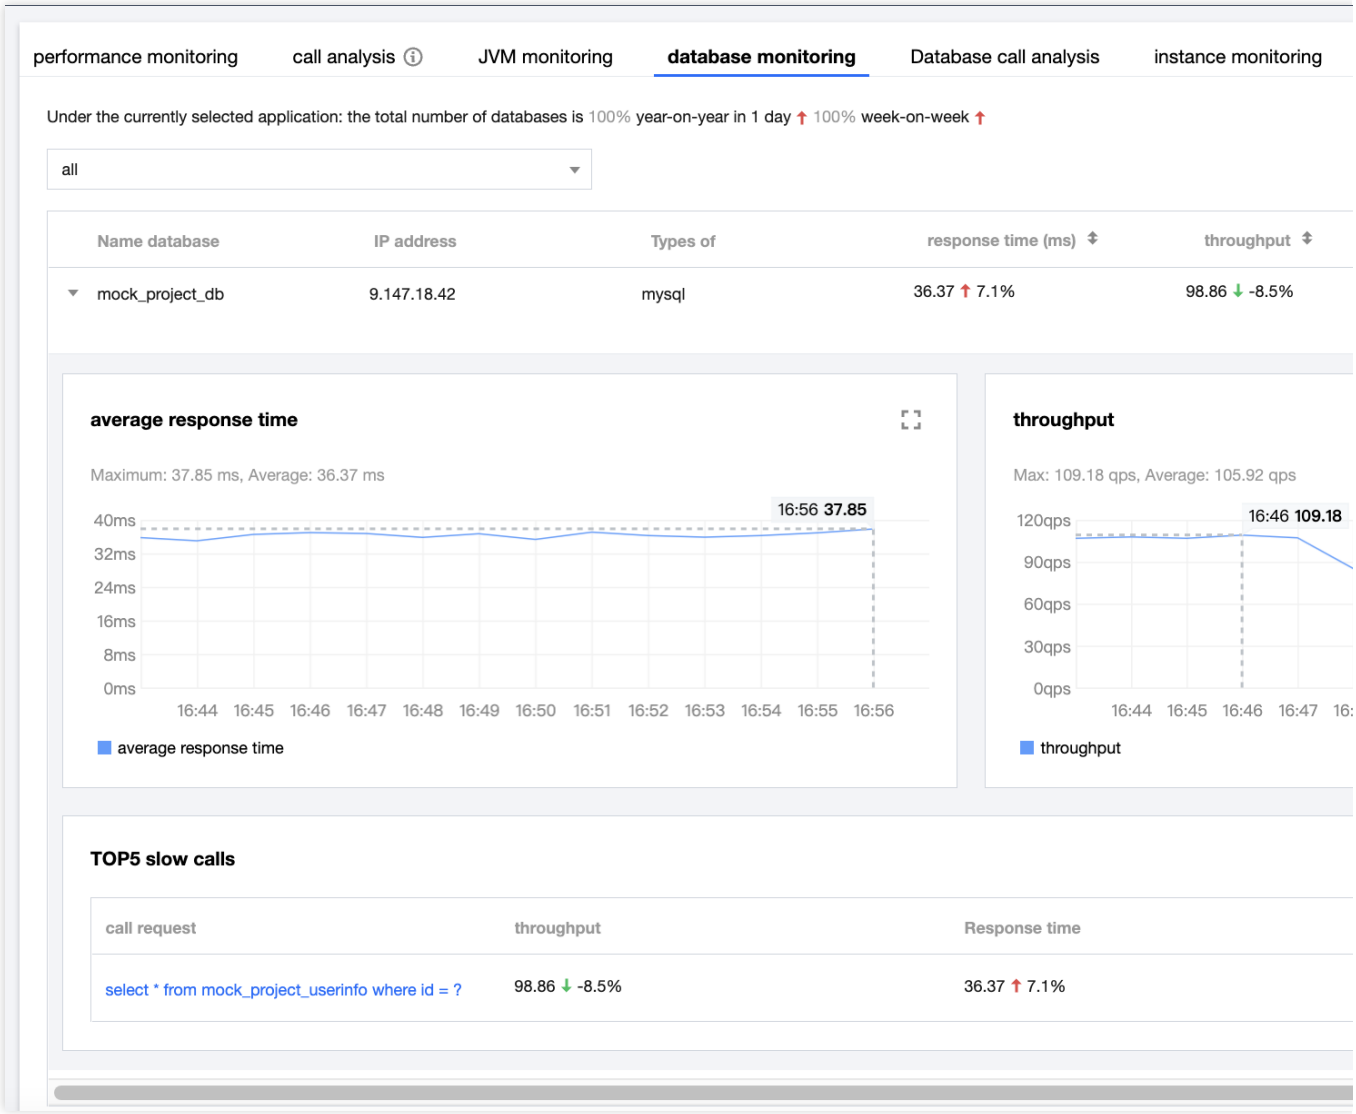

Metric description

Metric	Description
Throughput	The average throughput of the current application
Throughput status breakdown	The proportions of successful and failed requests
Slow call	The API calls that took more than 500 ms to execute
Slow SQL	The SQL queries that took more than 2,000 ms to execute
Full GC	The number of full GCs performed by JVM
Avg response time	The average response time of all APIs at the 1-minute granularity among the call relationships of the selected service

Avg error rate	The average error rate of all APIs at the 1-minute granularity among the call relationships of the selected service
Top 5 slow APIs	The top 5 slow APIs at the 1-minute granularity among the call relationships of the selected service
Top 5 erroneous APIs	The top 5 erroneous APIs at the 1-minute granularity among the call relationships of the selected service
Avg GC times	The average number of GCs executed by all JVM instances per minute
Max GC times	The maximum number of GCs executed by all JVM instances per minute
Avg GC time	The average duration of GC executions by all JVM instances per minute
Max GC time	The maximum duration of GC executions by all JVM instances per minute
CPU utilization	The CPU resource utilization of running programs per minute
Heap space	The status of the heap space per minute (committed, max, or used)
NoHeap space	The status of the NoHeap space per minute (committed, max, or used)
Thread pool	The number of active threads in the thread pool per minute
Throughput (database)	The average throughput of the current database
Average response time	The average response time of all calls of the selected database and instance at the 1-minute granularity
Top 5 callers	The top 5 upstream applications/components calling the selected database most frequently
Top 5 slow calls	The top 5 slow statements

API Monitoring

Last updated：2023-12-25 15:48:04

The API monitoring module displays the APIs in client, server, and local calls as well as the call traces in the upstream and downstream services, allowing you to check the key metrics like the number of requests, average response time, error rate, and throughput.

Prerequisites

Go to the **API monitoring** page in the APM console.

API overview

On the **API monitoring** page, select the target API on the left to display **API analysis** and **Exception analysis** on the right. You can select the target caller to view its details.

API analysis

API analysis allows you to select the target API service and view its total throughput, top 5 callers, average response time (99, 90, or 50 pct), error rate, and error code breakdown.

Metric	Description
Top 5 callers	The top 5 upstream applications/components calling the selected application API most frequently
Error code breakdown	The distribution trend of error codes returned by the current API

Note:
99 pct: Data in the 99th percentile in ascending order.
90 pct: Data in the 90th percentile in ascending order.
50 pct: Data in the 50th percentile in ascending order.

Exception statistics

Besides basic metrics, you can also view top 5 APIs in the current application in terms of average response time and error rate, which are intelligently filtered by the exception analysis module. You can click **View details** above the curve to view the call trace of the target API in the last 15 minutes, so you can drill down to the issue to troubleshoot it quickly.

Upstream and downstream analysis

You can switch between the **Upstream analysis** and **Downstream analysis** tabs to analyze upstream and downstream calls and quickly identify performance bottlenecks.

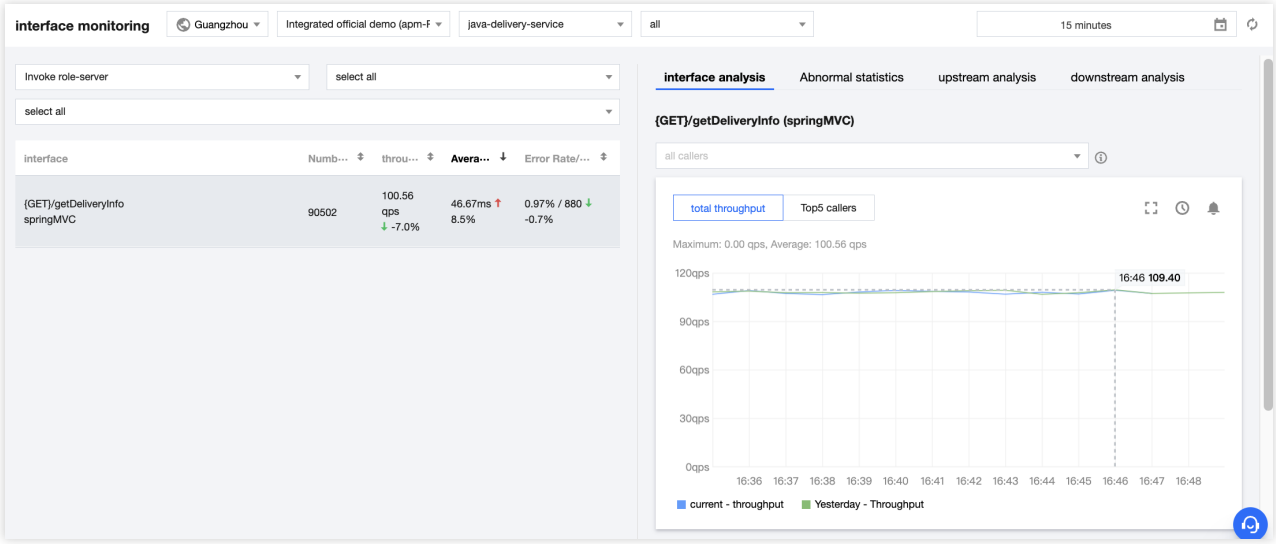

Exception Monitoring

Last updated：2023-12-25 15:48:36

This document describes how to view and analyze an application exception.

Prerequisites

Go to the **Exception monitoring** page in the APM console.

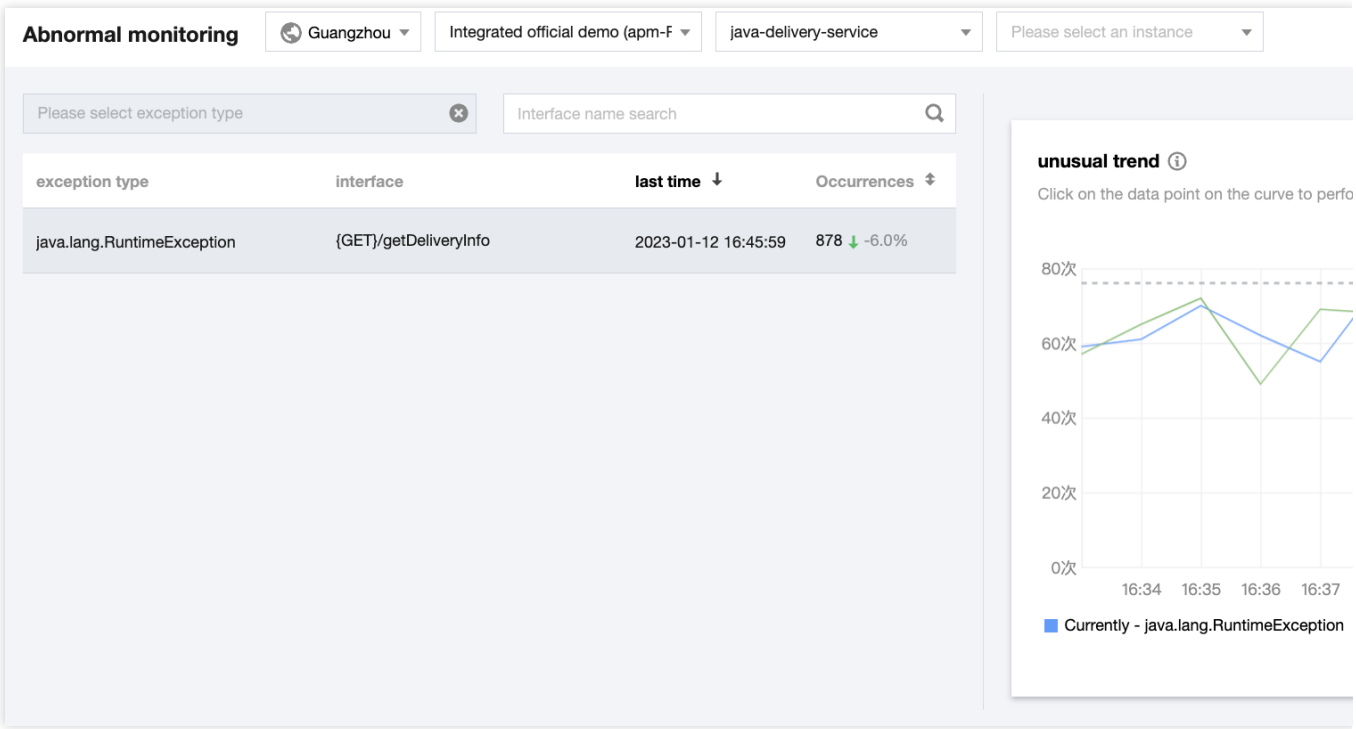

Exception list

It displays all exceptions within the specified time range on the left of the **Exception monitoring** page, including the service exception type, API, and number of occurrences.

Exception trend

The chart on the right displays the exception trend, i.e., number of exception occurrences per minute.

Exception analysis

Click the target application. Then, the number of occurrences of the exception and day-over-day comparison will be displayed as curves on the right of the **Exception monitoring** page.

Note:

You can click **Select comparison curve** above the chart to add a month-over-month comparison curve. You can also click a data point in the curve to view the list of related requests. In the list, click **Request details** to drill down to call trace details for exception analysis.

Tracing
Call Query

Last updated：2023-12-25 15:49:14

This document describes the capabilities and directions of API call search and analysis. With this feature, you can query the trace information, trace IDs, and trace details of servers and clients. In addition, the method stack information can be displayed for applications that use the proprietary probe to report data.

Prerequisites

Go to the **Call query** page in the APM console.

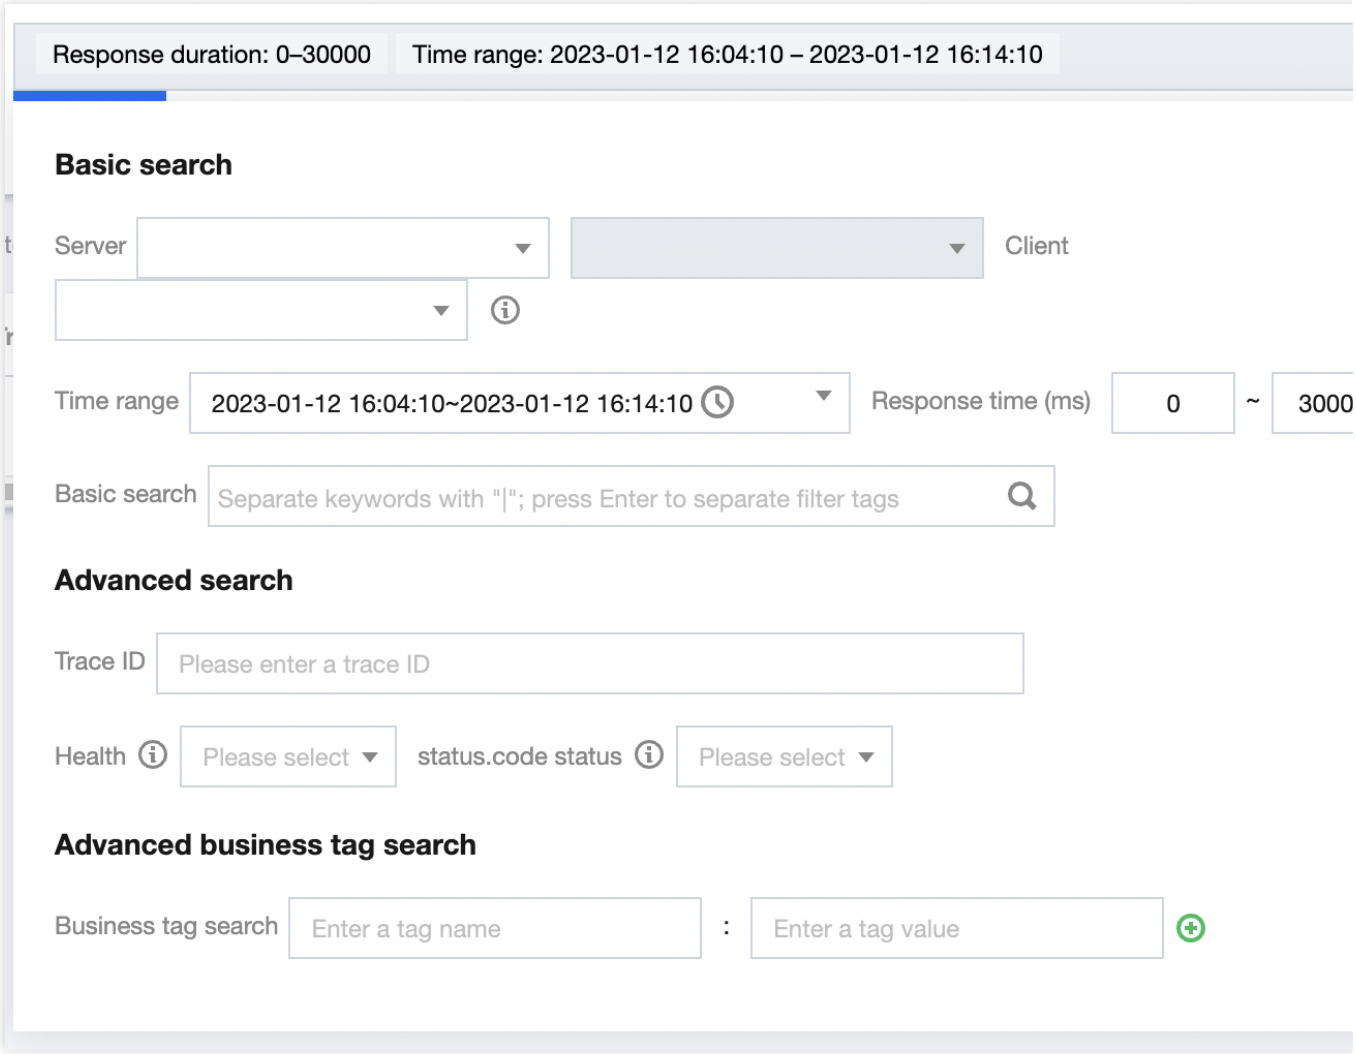

Basic query

You can query API calls by server name, server IP, client name, client IP, target API, trace ID, or business tag in the basic query module on the **Call query** page.

Note:

If you don't select the server or client, all traces in the application will be displayed in the list.

Business tag

You can customize tags when reporting data based on the business type, such as order and cart tags. If there is an exception in a user's order, you can quickly get the order call conditions, including response time, execution result, and service status.

For example, when reporting the data of a PHP application, you can customize a tag and then search for the business bound to the tag by the tag key and value (key:value).

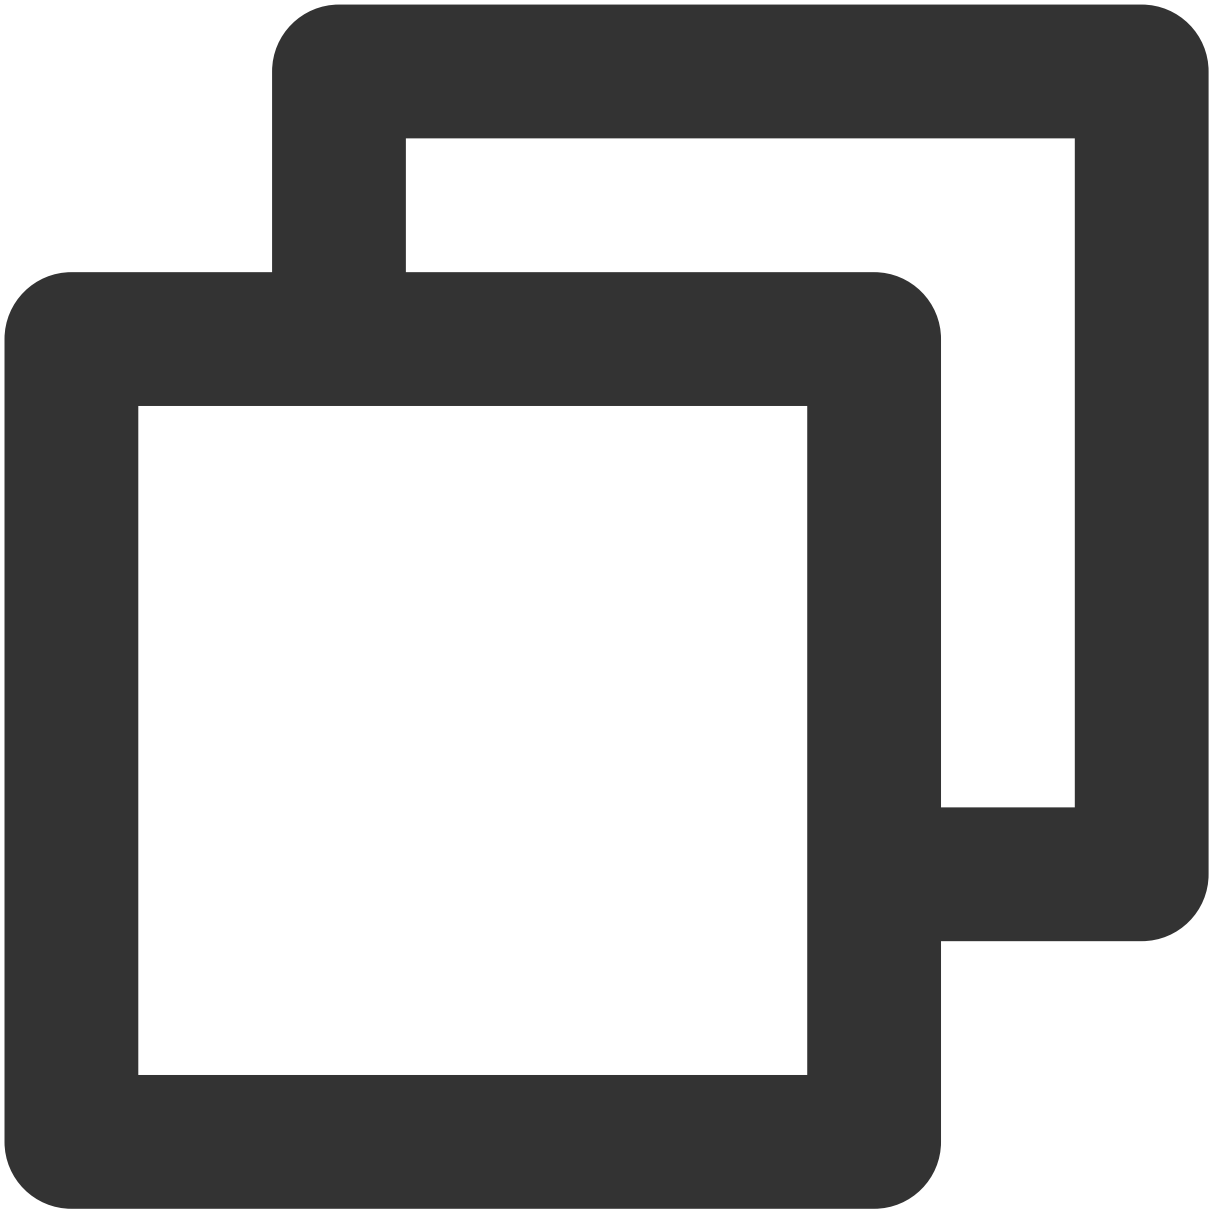

```
$span->tag('key', $value);
```

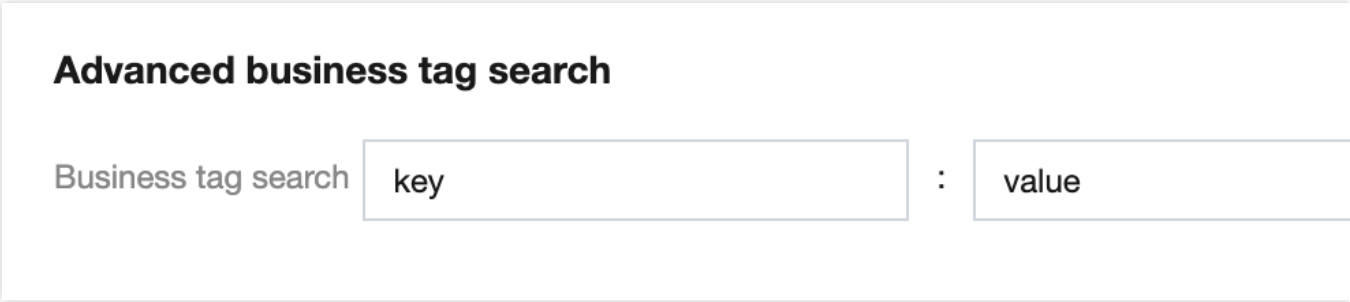

Displaying the trace entry only

You can quickly search for the trace entry (span) in an application to locate abnormal traces.
Assume a trace is as shown below. If you select **Display the trace entry only**, the list will only display the call status of the trace entry **Service A**.

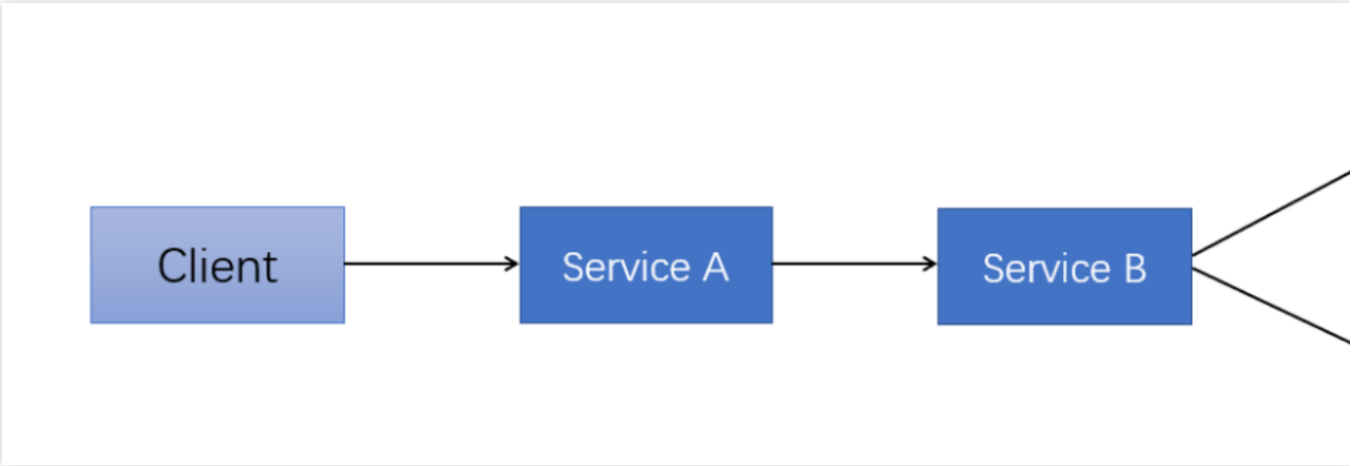

Call trace details

Click **Request details** or **TraceID/SpanID** in the **Operation** column of the target API to enter the trace analysis page, where you can view the duration of each stage of the trace and the reported information of the entire trace, including the trace health and duration.
The method stack information can be displayed for applications that use the proprietary Java probe as described in Reporting via TAPM to report data. You can view the line numbers of method stacks in the trace list to quickly troubleshoot slow calls and abnormal method stacks.

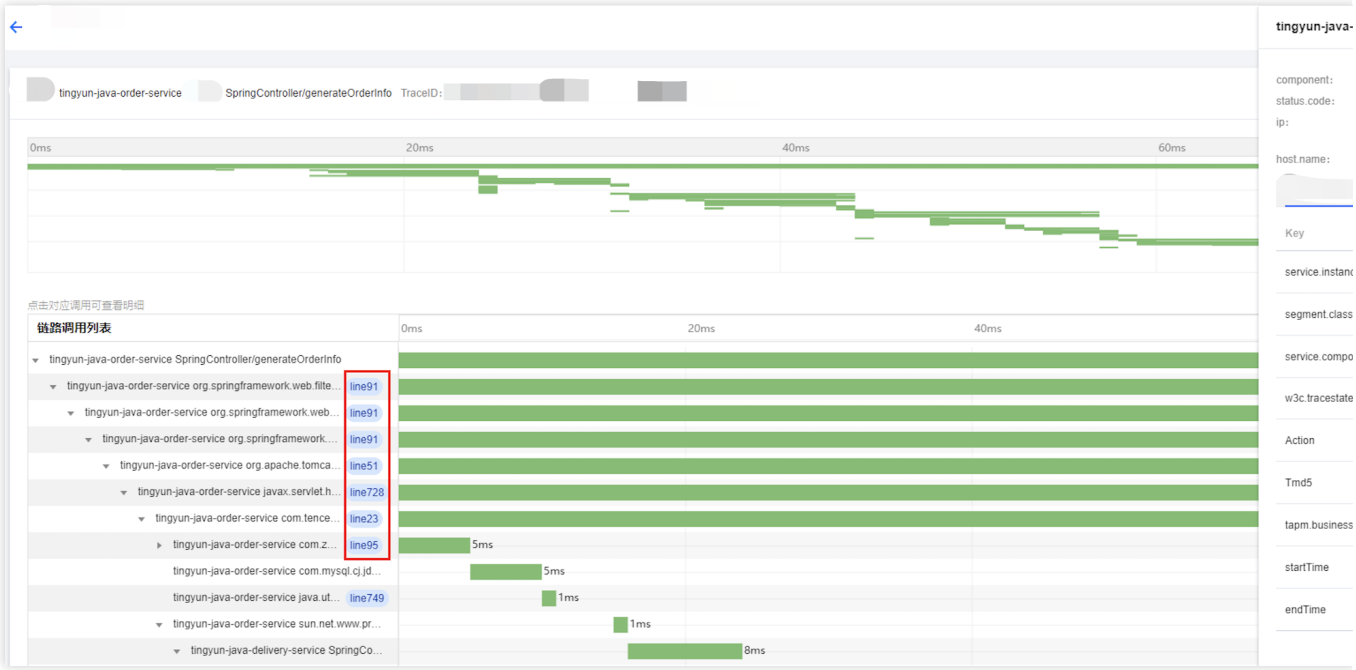

Database Call Monitoring
Database Overview

Last updated：2024-07-23 14:50:51

The database overview page displays the overall call conditions of the current database, including call topology, database response time, throughput, top 5 slow calls, top 5 callers, and exception list.

Prerequisites

Go to the Database query > Database overview page in the APM console.

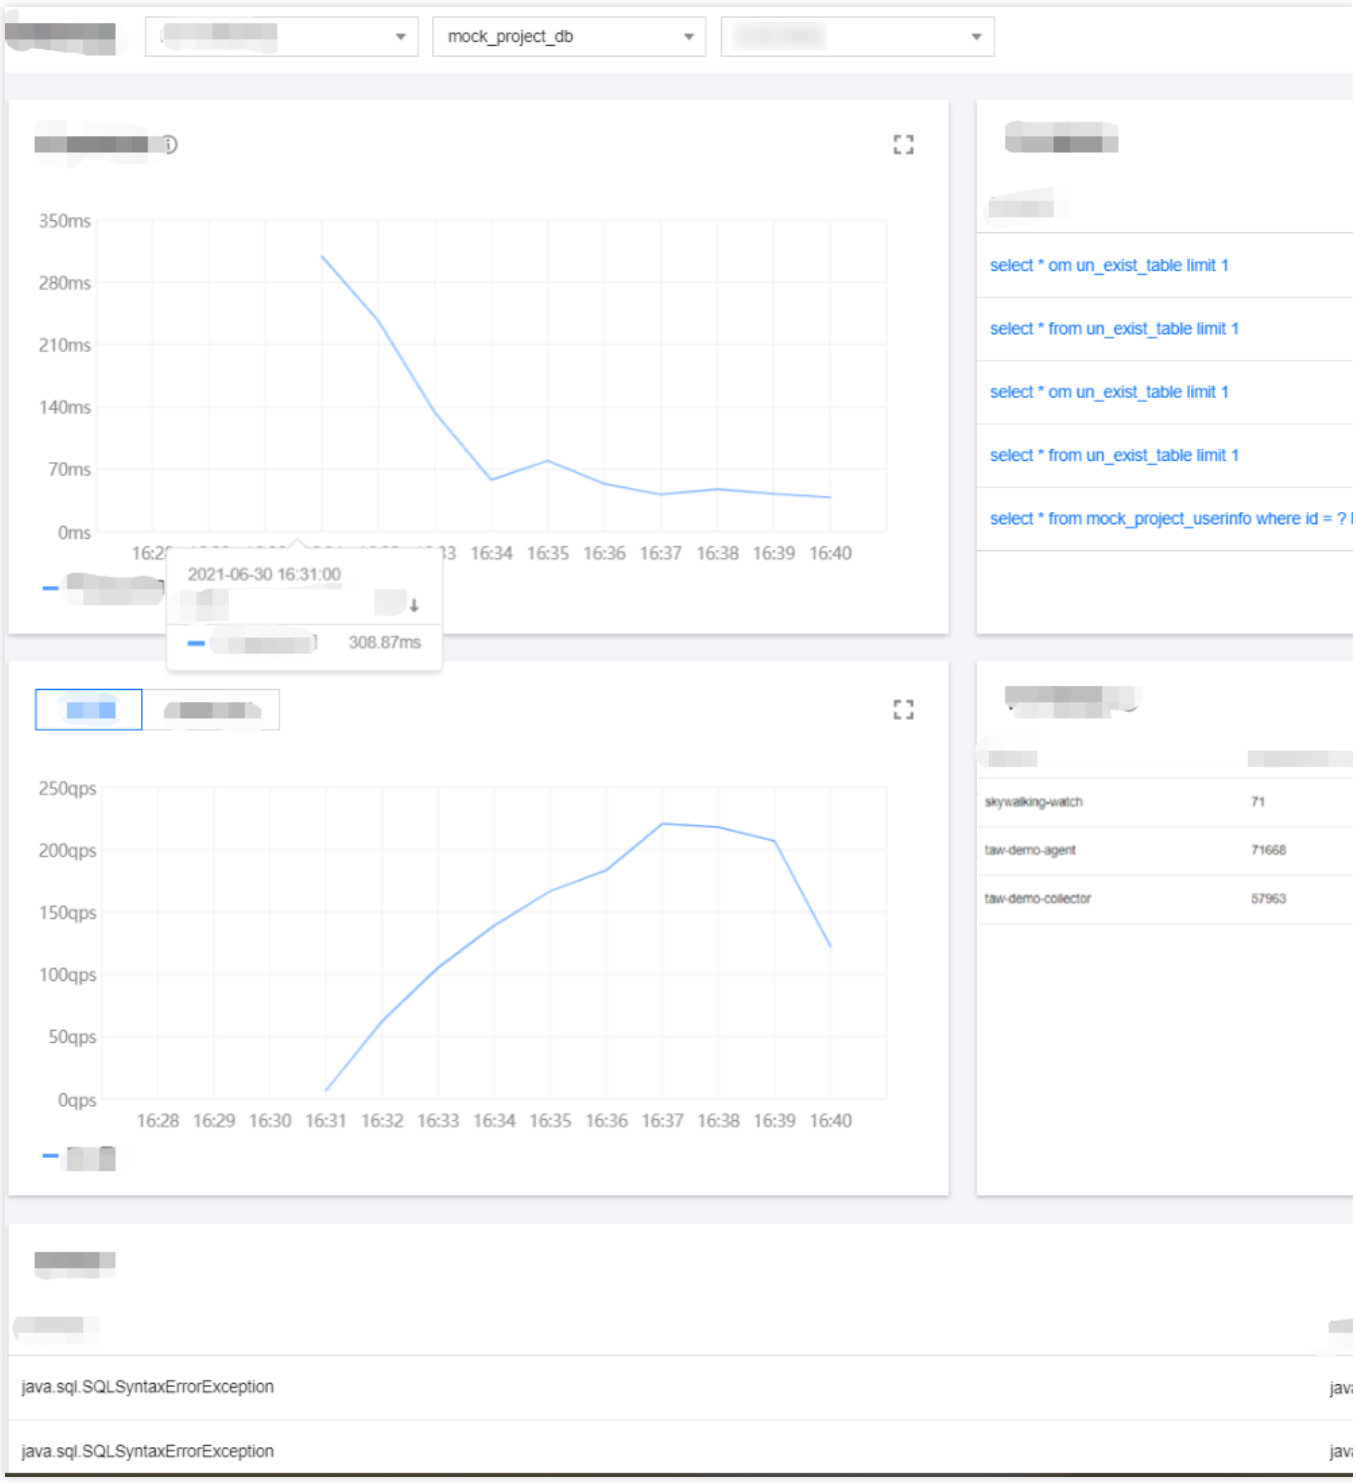

Database call topology

Database call relationships can be displayed in a topology, making it easier for you to view the details of calls between the application and database.

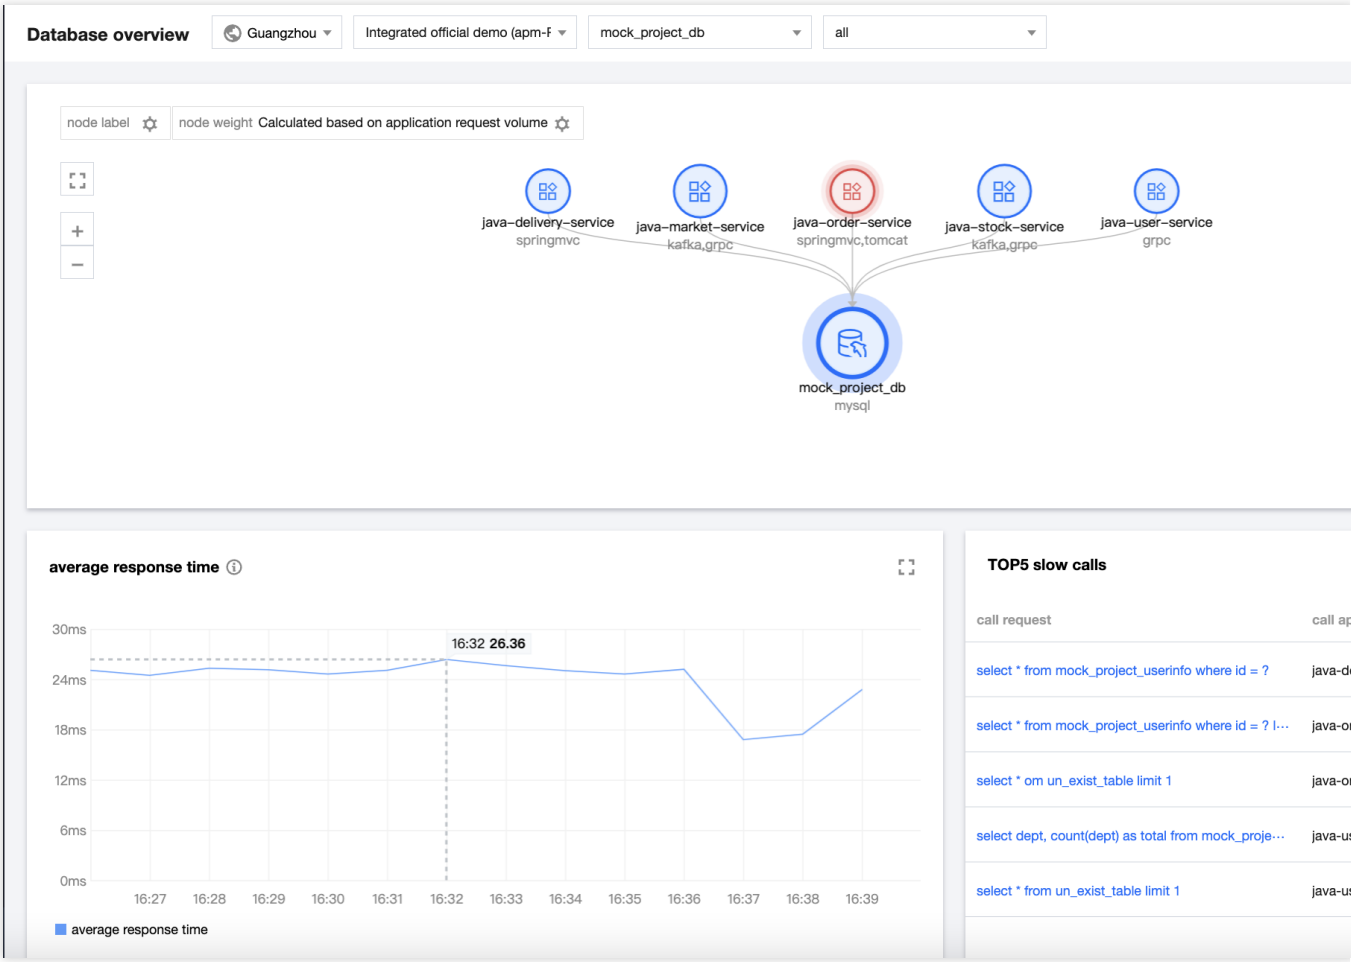

Health:

Healthy: Your application is assessed as healthy based on the response time and error rate.

Warning: The current average response time of your application exceeds the satisfaction threshold, but the average error rate does not exceed the alarm threshold.

Abnormal: The average error rate of your application exceeds the alarm threshold.

The satisfaction threshold for the average response time is 500 ms by default.

The alarm threshold for the average error rate is 5% by default.

Exception statistics

The list of database exceptions displays the exception type, caller, number of exceptions, first occurrence time, and last occurrence time. You can click **View details** to view specific database exceptions or click **Trace** to view trace details.

Metric description

Metric	Description
Throughput	The average throughput of the current database

Avg response time	The average response time of all calls of the selected database and instance at the 1-minute granularity
Top 5 callers	The top 5 upstream applications calling the selected database most frequently

Database Call Analysis

Last updated：2023-12-25 15:49:47

This document describes the capabilities and directions of database call search and analysis.

Prerequisites

Go to the **Database call analysis** page in the APM console.

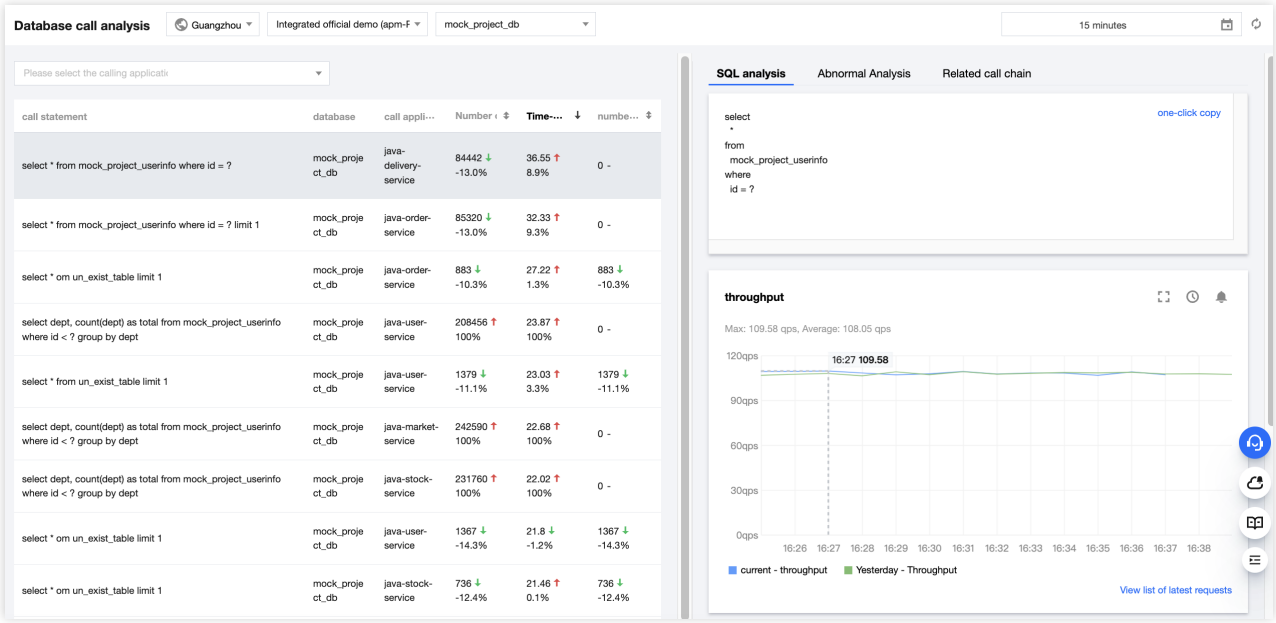

SQL monitoring list

The SQL list on the left displays all SQL calls between the application and database in the selected time range, as well as the number of calls, duration, error rate, and DoD change of each SQL statement. You can directly locate the target SQL statement through the search bar above the list.

SQL analysis

The SQL analysis module on the right displays the change trends of calls and response time of the currently selected SQL statement in the specified time range. The list of exceptions in the current database is displayed below the trend

chart, including exception type, caller, number of exceptions, first occurrence time, and last occurrence time. You can click **View details** to view specific database exceptions or click **Trace** to view trace details.

Exception analysis

You can switch the sub-window menu on the right to enter the exception analysis page. This page displays the list of SQL exceptions, including service exception types, APIs, and the number of exception occurrences.

Access Management

Overview

Last updated：2023-12-25 15:50:04

If you have multiple users managing the APM service, and they all share your Tencent Cloud account access key, you may face the following problems:
Your key will be easily compromised because it is shared by several users.
You cannot restrict the access from other users and your service will be vulnerable to the security risks caused by their maloperations.
You can avoid the above problems by allowing different users to manage different services through sub-accounts. By default, sub-accounts have no permissions to use APM. Therefore, you need to create a policy to grant different permissions to sub-accounts.

Overview

Cloud Access Management (CAM) is a Tencent Cloud web service that helps you securely manage and control access to your Tencent Cloud resources. CAM allows you to create, manage, or terminate users (groups), and control who have access to which Tencent Cloud resources based on identity and policy management.
When using CAM, you can associate a policy with a user or user group to allow or forbid them to use specified resources to complete specified tasks. For more information on CAM policies, see Syntax Logic. For more information on how to use CAM policies, see Concepts.

Authorization method

APM supports two authorization methods: resource-level authorization and authorization by tag.
Resource-level authorization: You can use policy syntax or the default policy to grant sub-accounts permissions to manage individual resources. For more information, see Policy Syntax and Granting Policy.
Authorization by tag: You can tag resources and grant sub-accounts permissions to manage resources with particular tags. For more information, see Resource Tag.
You can skip this section if you don't need to manage permissions of APM resources for sub-accounts. This won't affect your understanding and use of the other sections of the document.

Policy Syntax

Last updated：2023-12-25 15:50:24

Overview

An access policy that employs the JSON-based access policy language is used to grant access to APM resources. You can authorize a specified principal to perform actions on a specified CM resource through the access policy language.
The access policy syntax describes the basic elements and usage of the policy. For the description of the policy syntax, see Concepts.

Policy Syntax

CAM policy:

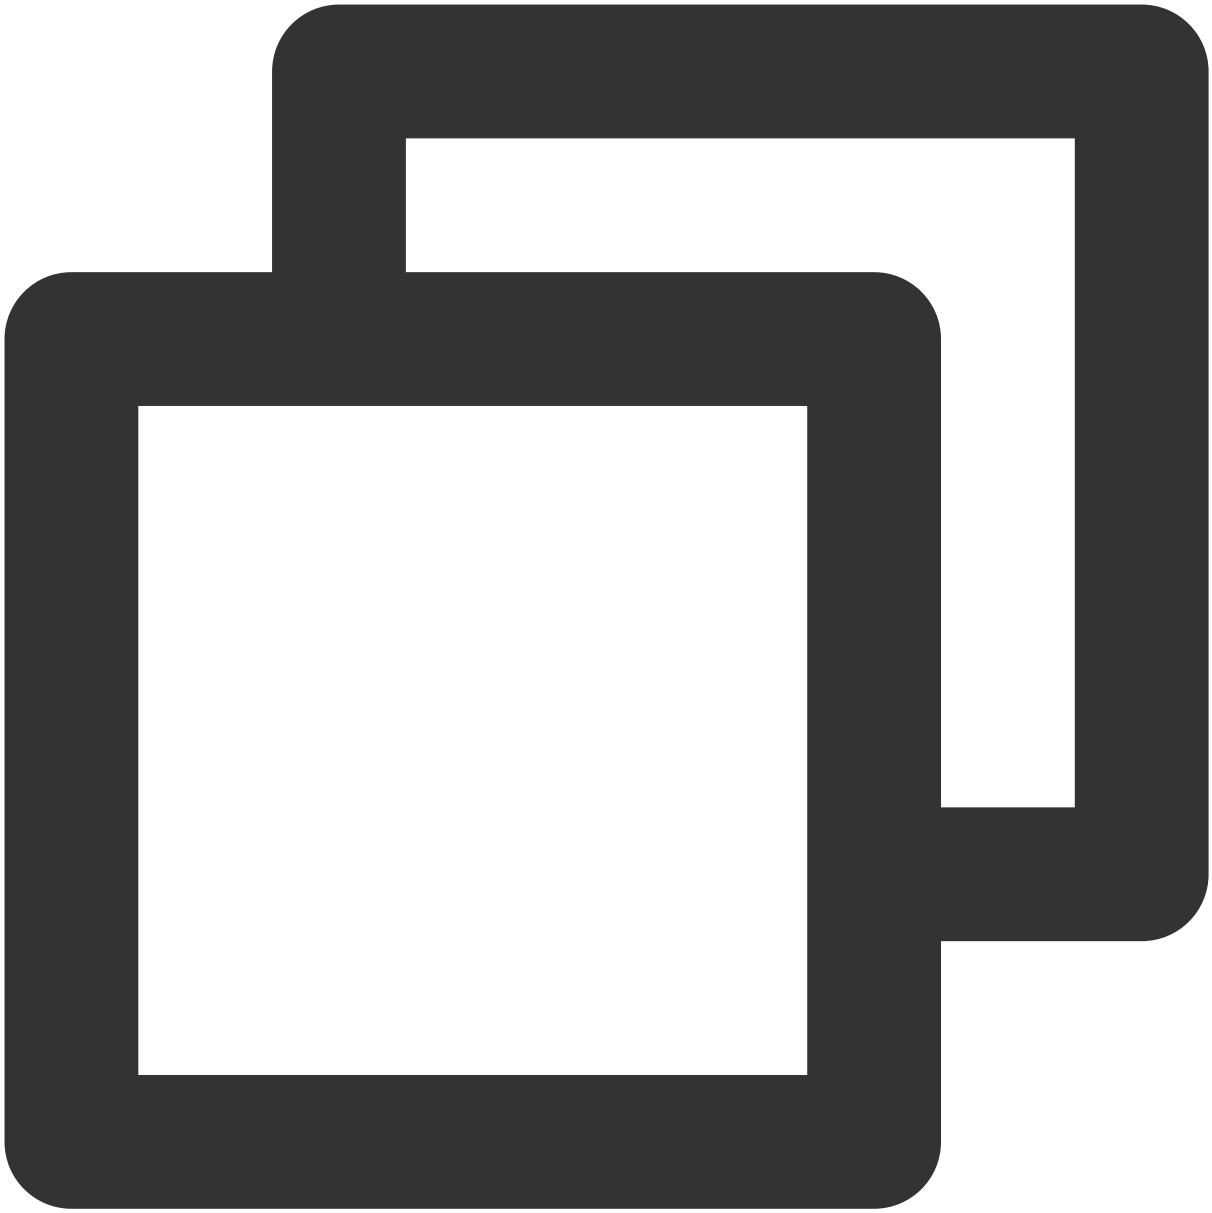

```
{
        "version":"2.0",
        "statement":
        [
            {
                "effect":"effect",
                "action":["action"],
                "resource":["resource"],
                 "condition": {"key":{"value"}}
            }
        ]
```

```
}
```

Element usage

version is required. Currently, only the value "2.0" is allowed.
statement describes the details of one or more permissions. This element contains a permission or permission set of other elements such as `effect` , `action` , `resource` , and `condition` . One policy has only one statement.
effect is required. It describes whether the declaration result is `allow` or explicit `deny` .
action is required. It specifies whether to allow or deny the operation. The operation can be an API (prefixed with `name`) or a feature set (a group of APIs, prefixed with `permid`).
resource is required. It describes the details of authorization. A resource is described in a six-segment format. Detailed resource definitions vary by product. For more information on how to specify resources, see the product documentation corresponding to the resource statement you are writing.
condition is optional. It describes the condition for the policy to take effect. A condition consists of operator, action key, and action value. A condition value may contain information such as time and IP address. Some services allow you to specify additional values in a condition.

Specifying an effect

If you don't explicitly grant access to (`allow`) a resource, access is implicitly denied. You can also explicitly `deny` access to a resource to ensure that a user cannot access it, even if another policy has granted access to it. The following example specifies an `allow` effect.

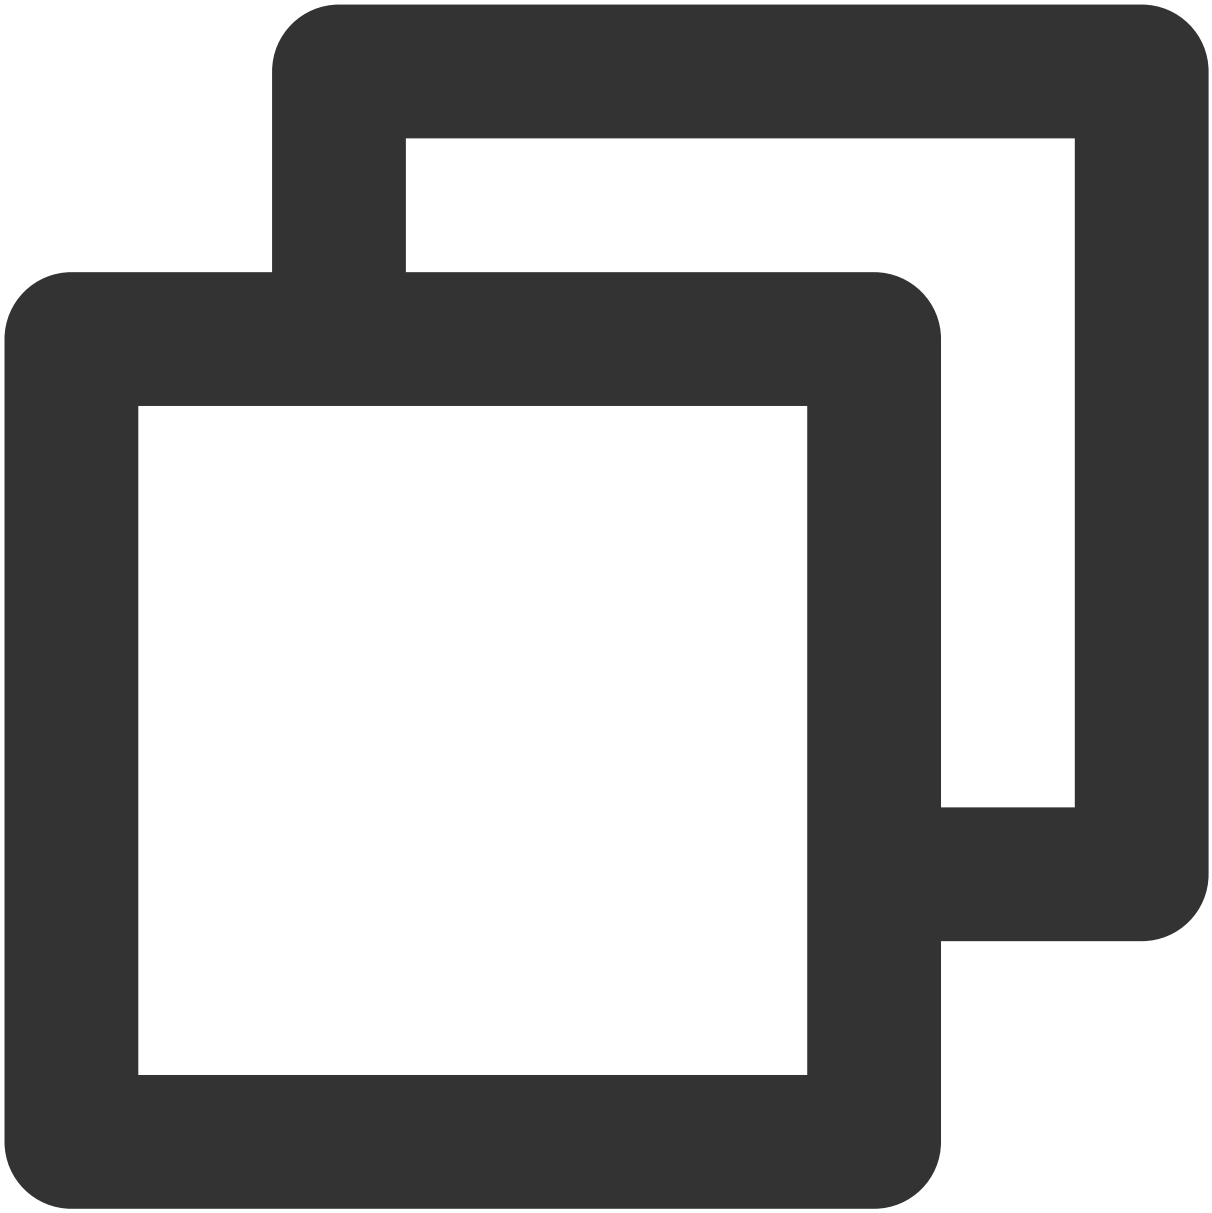

```
"effect" : "allow"
```

Specifying an action

APM defines console operations that can be specified in a policy. The specified operations are divided into reading part of APIs (`apm:Describe\\*`) and all APIs (`apm:\\*`) based on the operation nature.
Below is an example of specifying the allowed operations:

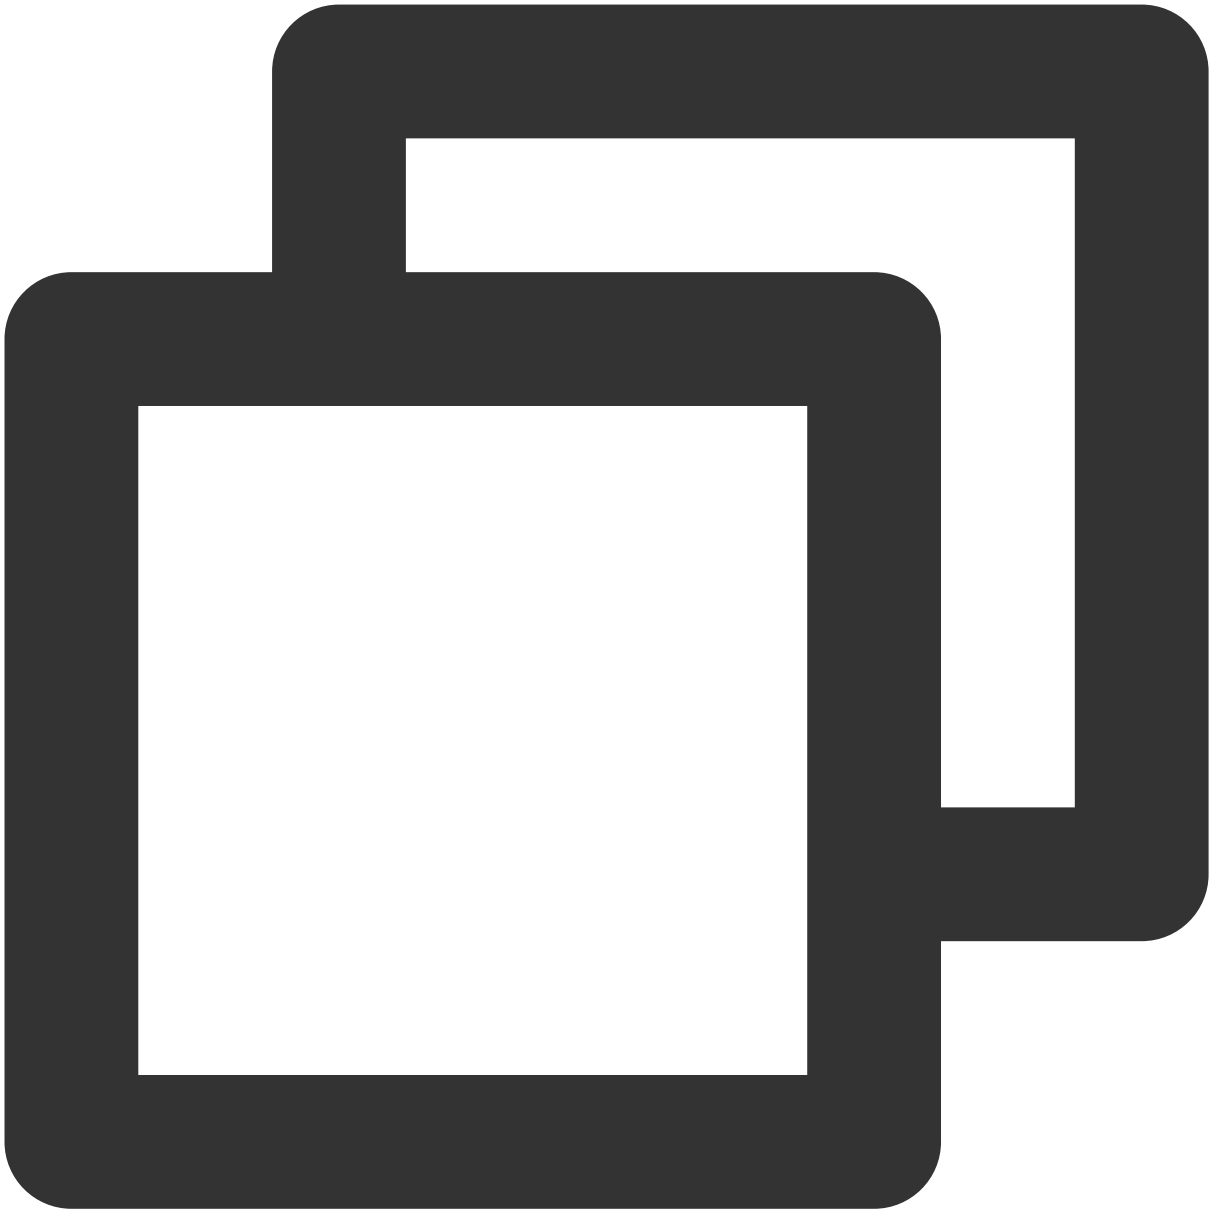

```
"action": [
  "name/apm:Describe*"
]
```

Specifying a resource

The `resource` element describes one or multiple operation objects, such as APM resources. All the resources can be described in the following 6-segment format.

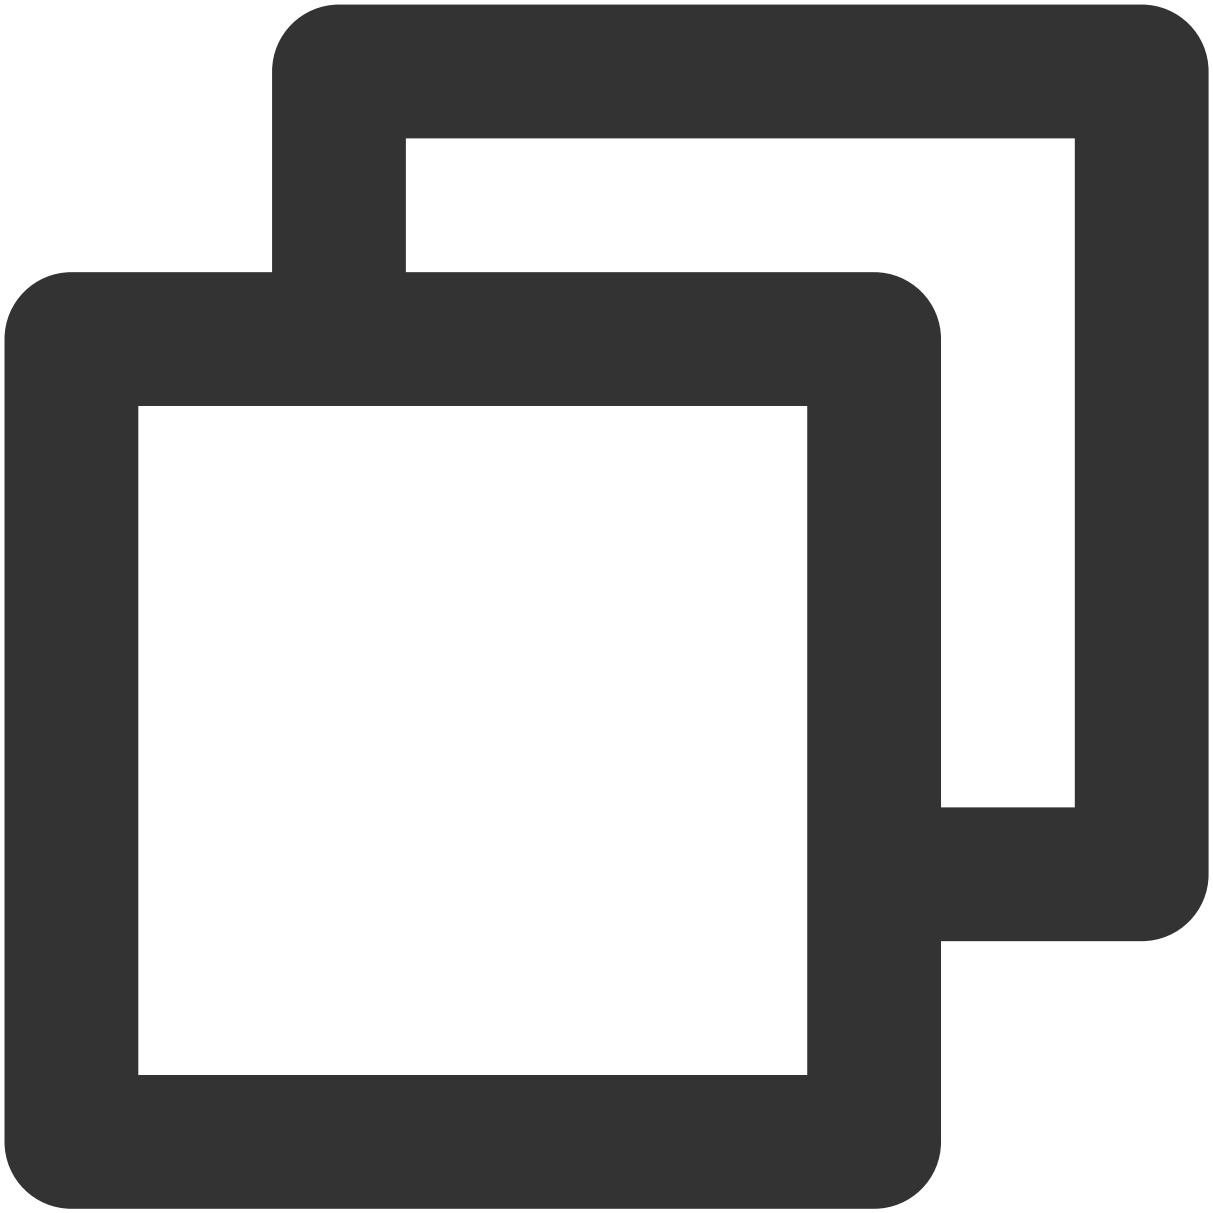

```
qcs:project_id:service_type:region:account:resource
```

The parameters are as described below:

Parameter	Description	Required
qcs	Abbreviation for "qcloud service", which indicates a Tencent Cloud service	Yes
project_id	Project information, which is only used to enable compatibility with legacy CAM logic and generally can be left empty	No

service_type	Product name abbreviation, which is `apm` here	Yes
region	Region information	Yes
account	Root account information of the resource owner, which is the root account ID in the format of `uin/${OwnerUin}` , such as `uin/100000000001`	Yes
resource	Resource details prefixed with `instance`	Yes

Below is a sample six-segment description of an APM resource:

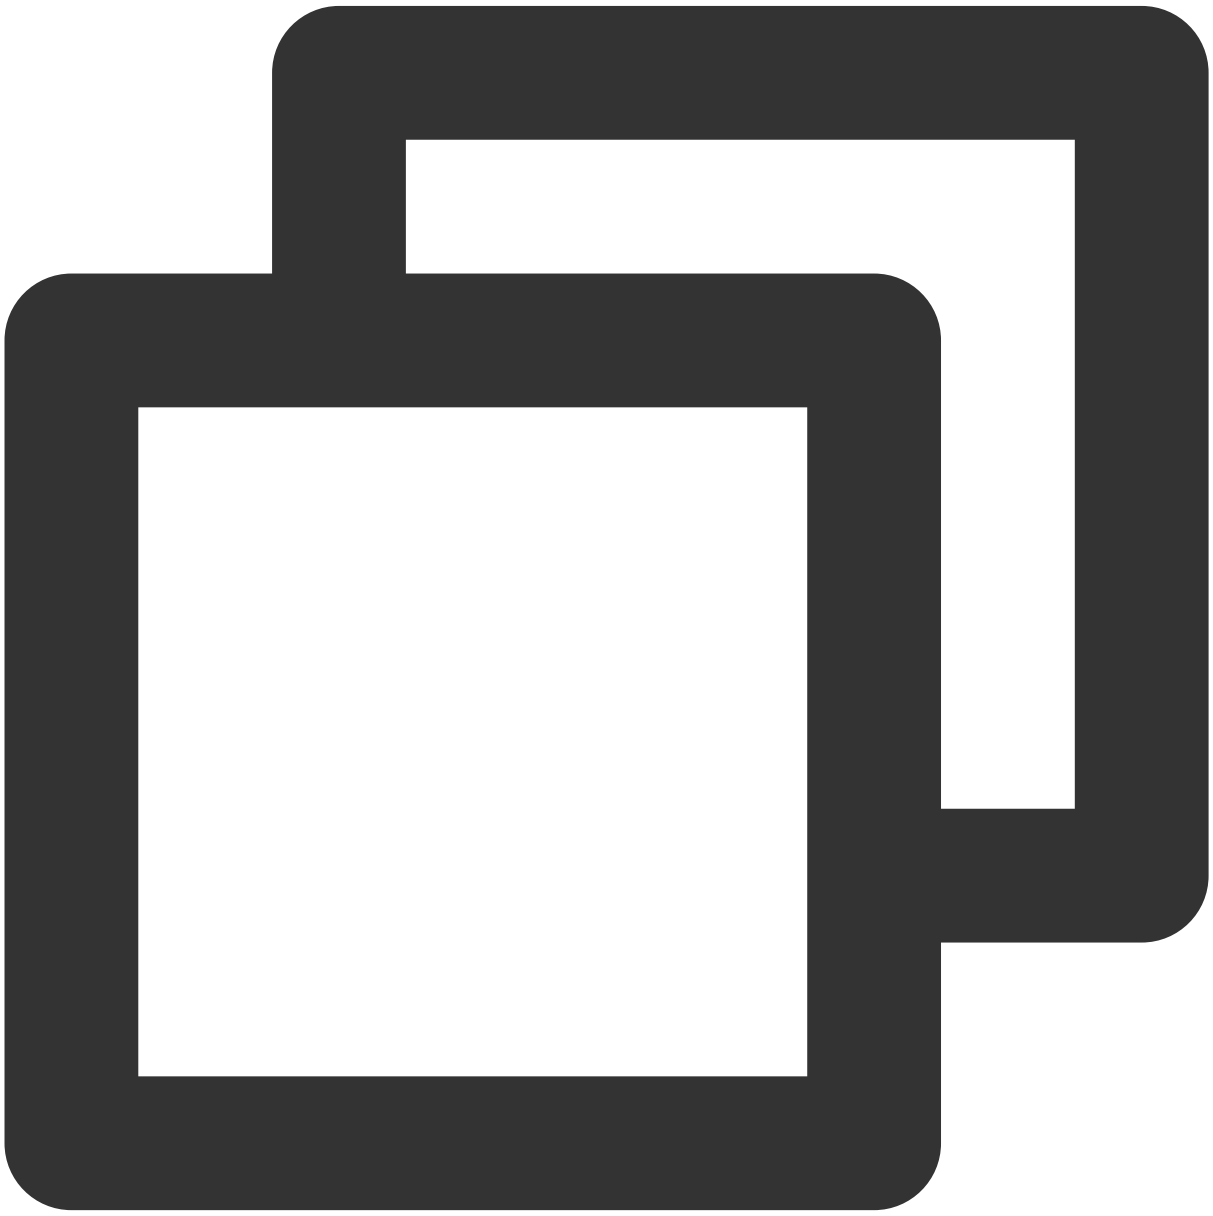

```
"resource":["qcs::apm:ap-guangzhou:uin/1250000000:apm/apm-btzsrI123"]
```

Samples

Grant the read and write permissions of specified resources based on resource ID. The root account ID is `1250000000` :

Sample: Granting the sub-user the permission to modify the business system (ID: apm-btzsrI123)

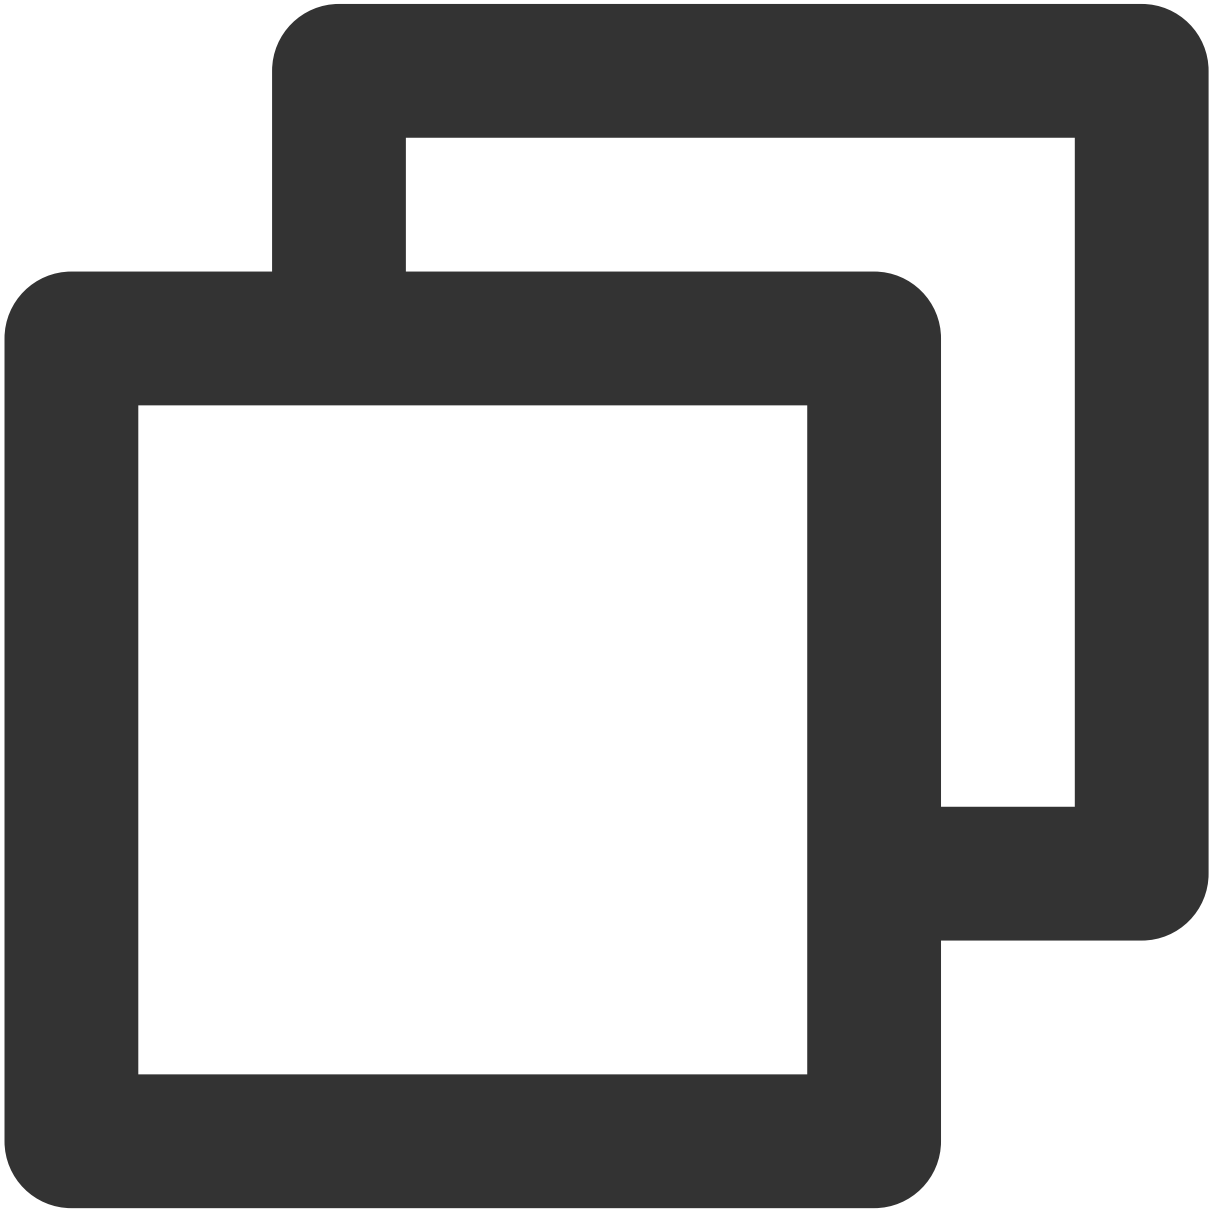

```
{
    "version": "2.0",
    "statement": [
        {
            "effect": "allow",
            "action": [
                "apm:ModifyApmInstance"
            ],
            "resource": [
                "qcs::apm:ap-guangzhou:uin/1250000000:apm-instance/apm-btzsrI123"
            ]
        }
    ]
}
```

List of APIs supporting resource-level authorization

API	Description
DescribeApmAgent	Gets the APM agent
DescribeApmInstances	Queries APM business systems
DescribeApmRegions	Gets APM regions
DescribeGeneralSpanList	Queries spans
DescribeInstanceBriefs	Queries the business system overview
DescribeMetricLineData	Pulls metric curve data
DescribeMetricRecords	Queries the list of metrics
DescribePAASGeneralSpanList	Queries spans
DescribePAASMetricLineData	Queries the metric curve data
DescribePAASMetricPointData	Queries the metric point data
DescribePAASTagValues	Queries the dimension information
DescribePAASTopology	Queries the topology data
DescribeServiceNodes	Gets the list of services
DescribeServiceOverview	Gets the APM system overview

CreateApmInstance	Creates an APM business system
CreatePAASInstance	Creates a PaaS APM business system
DeletePAASInstance	Deletes an APM business system
ModifyApmInstance	Modifies an APM business system
TerminateApmInstance	Terminates an APM business system

Granting Policy

Last updated：2024-07-23 14:50:51

A sub-account has no APM permissions by default and can access APM resources only after being granted relevant permissions by the root account.

Prerequisites

Log in to the Tencent Cloud console with the root account or a sub-account with the `QcloudCamFullAccess` permission and create a sub-account as instructed in Creating Sub-User.

Custom policy

1. Use the root account or a sub-account with the `QcloudCamFullAccess` permission to log in to the **CAM** console and go to the **Policies** page.
2. Click **Create Custom Policy** > **Create by Policy Syntax** and select **Blank Template**. Edit the policy as instructed in Policy Syntax Description.

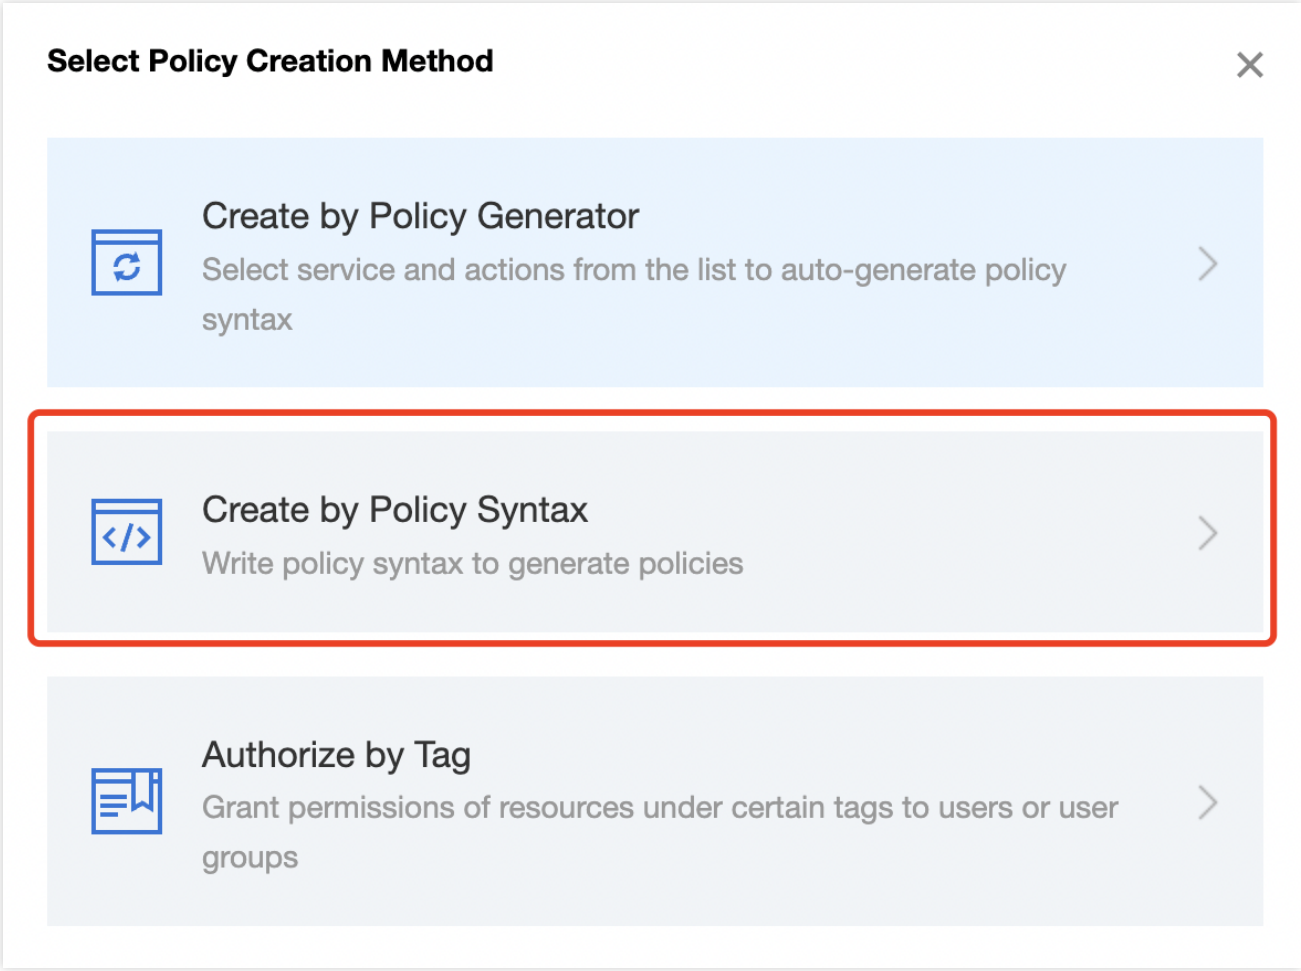

Policy authorization

Note:

APM creates default permission policies `QcloudAPMFullAccess` (full access to APM) and `QcloudAPMReadOnlyFullAccess` (read-only access to APM) for you. You can search for a default policy for quick authorization. You can also use a custom policy for authorization. Then, the sub-account can access the relevant resources.

1. Use the root account or a sub-account with the `QcloudCamFullAccess` permission to log in to the **CAM** console and go to the **Policies** page.
2. Go to the policy management page and enter a policy name in the policy search box.
3. Select `QcloudAPMFullAccess` or `QcloudAPMReadOnlyFullAccess` and click **Associate User/Group/Role** in the **Operation** column.

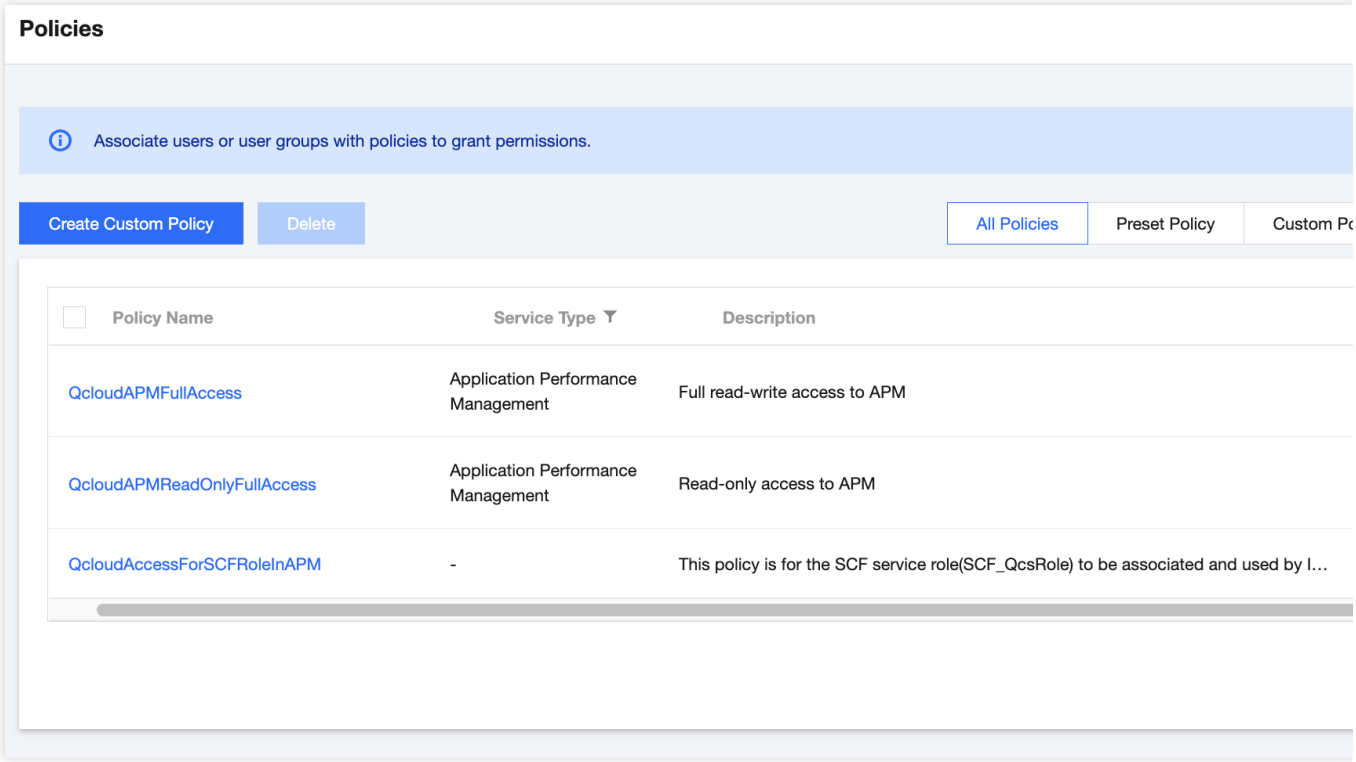

4. In the pop-up window, select the target user and click **OK**.

Resource Tag

Last updated：2023-12-25 15:51:02

APM can be used with the Tencent Cloud resource tag feature to perform tag-based sub-account authorization and cost allocation.
Resource tag is a resource management tool provided by Tencent Cloud, which has two parts: tag key and tag value. One tag key can correspond to multiple tag values. You can use tags for cost allocation and authorization in the following steps.

Use Cases

A company has multiple business systems connected to APM, which are developed and operated separately by departments A and B. You want to create tags, bind resources, and grant permissions to the two departments as follows:
Create tag A and bind it to all business systems of department A.
Create tag B and bind it all business systems of department B.

Authorization by tag

User A is a developer in department A and is responsible for the development of all business systems in the department. You want to grant tag A's permission to user A.

Cost allocation by tag

User B is a company accountant responsible for the separate accounting of the financial expenditures of departments A and B. You want to grant user B the permissions of tags A and B to allocate costs by tag.

Preparations

Step 1. Create a tag

Create tags A and B respectively in the following steps:
1. Go to the Tag List page.
2. Click **Create** to enter the tag creation page and enter the tag key and the corresponding tag value.

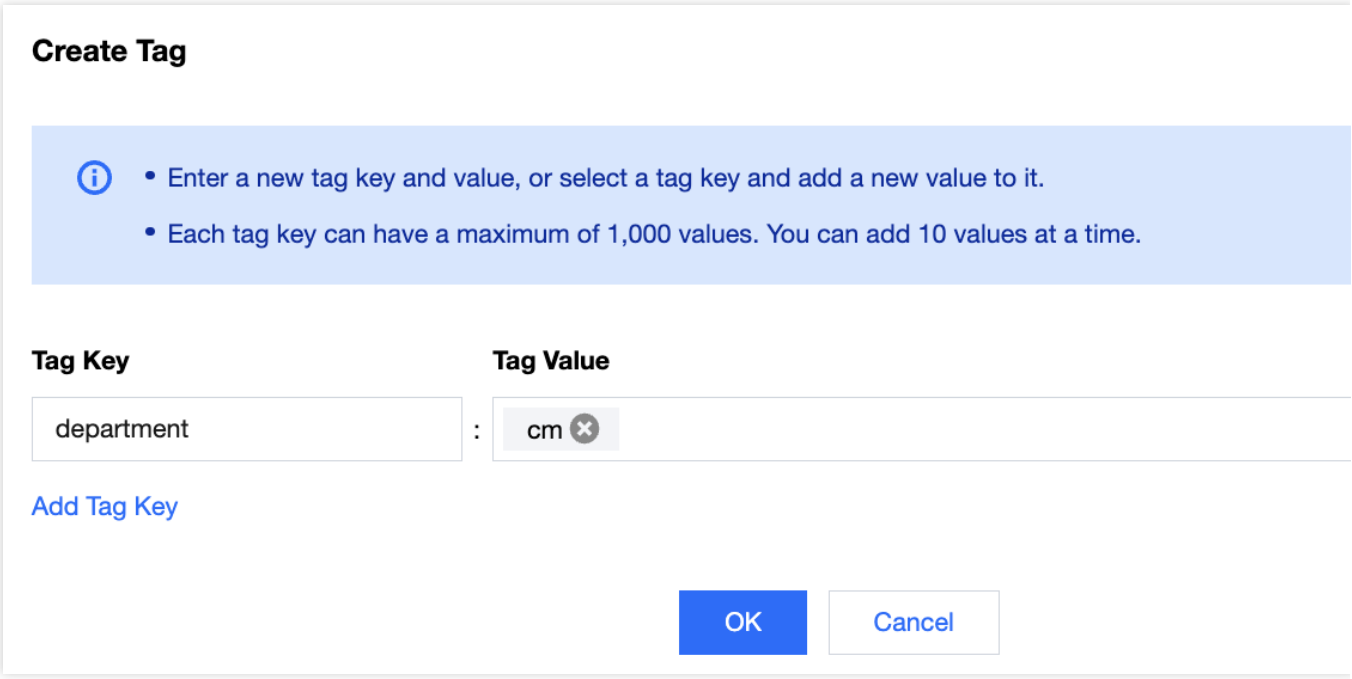

3. Click **OK**.

Step 2. Assign tags to resources

Bind tag A to all business systems in department A and tag B to all business systems in department B in the following steps:

1. Go to the **Resource management** page in the APM console.
2. Click **Create**. In the pop-up window, enter the information and bind a tag. You can also find an existing business system in the list, click **Modify configuration** in the **Operation** column, and select a tag.

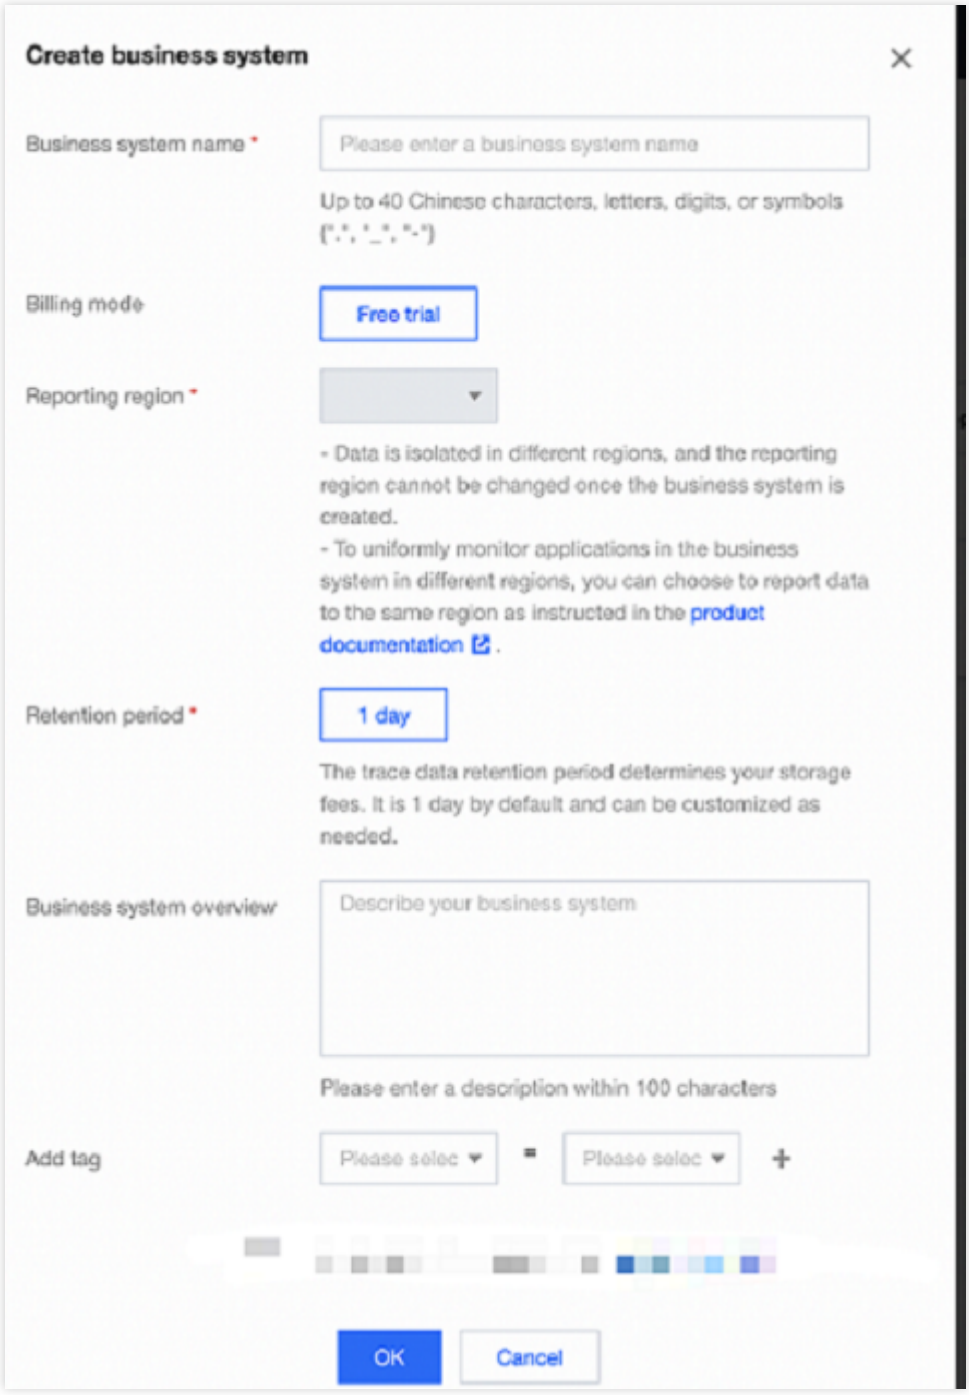

Authorization by Tag

Grant user A the permission of tag A and user B the permission of tags A and B according to the tag authorization policy in the following steps:

1. Go to the Policies page and click **Create Custom Policy** in the top-left corner.
2. In the creation method selection window that pops up, click **Authorize by Tag** to enter the **Authorize by Tag** page.

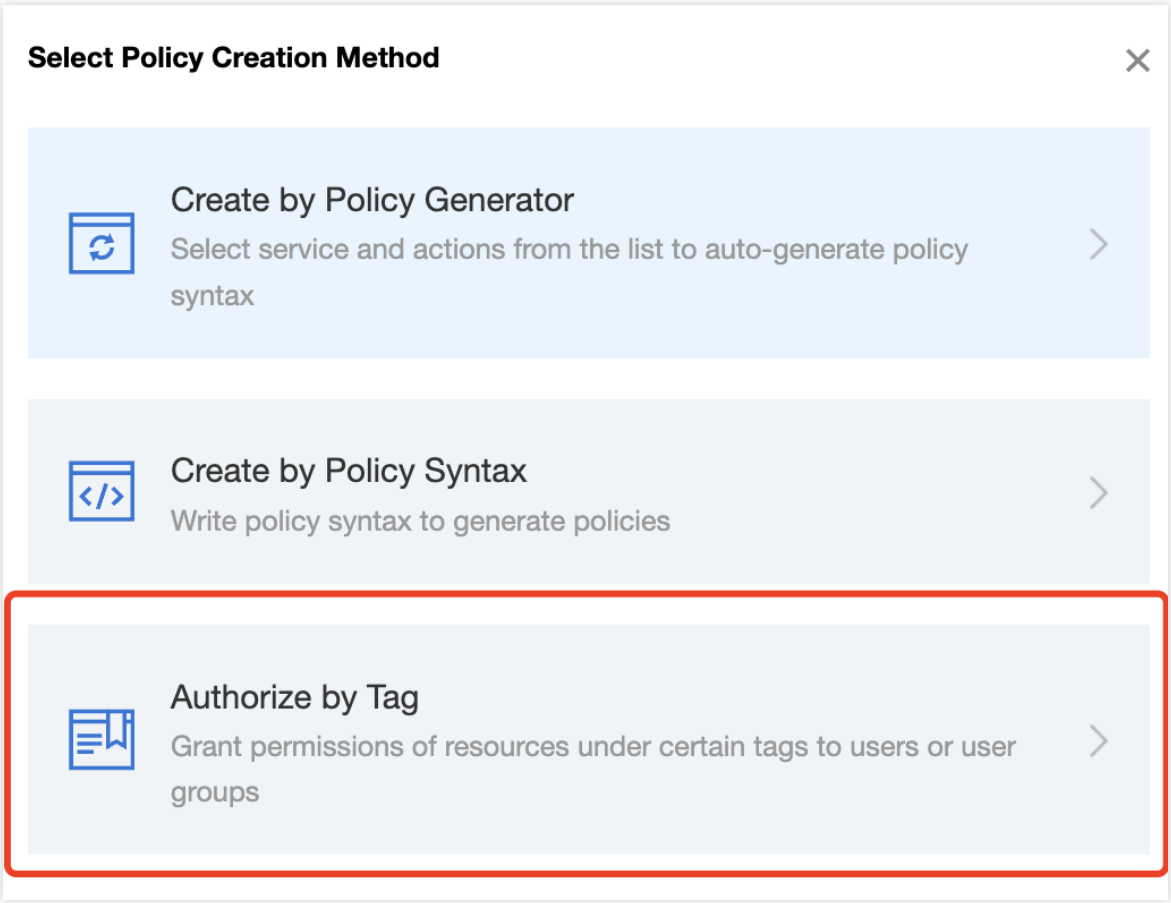

3. On the **Authorize by Tag** page, select the following information and click **Next** to enter the check page.

Authorized Users/User Groups: Select the user to be authorized (user A or B).

Bound Tag Key: Select the tag key to be authorized (tag key of tag A or B).

Bound Tag Value: Select the tag value to be authorized (tag value of tag A or B).

Service Resource: Select `apm` and select all operations.

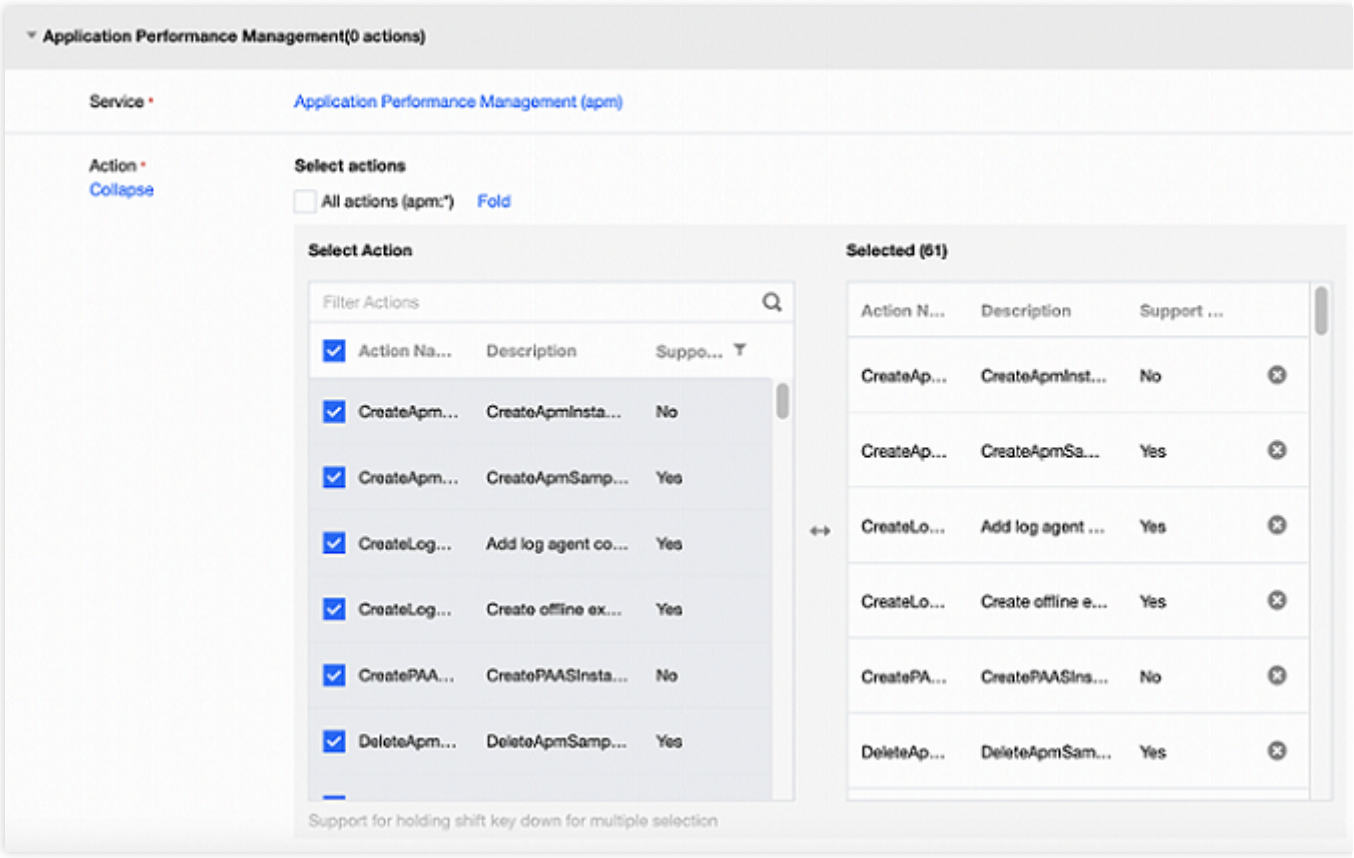

4. Click **Next**, check the policy (which can be renamed), and click **Complete**.

Cost Allocation by Tag

Step 1. Set a cost allocation tag

Set tags A and B as cost allocation tags in the following steps:

1. To use the tag feature for bills, you need to go to the Billing Center and select **Bills** > **Cost Allocation Tags** on the left sidebar. The tag key set as a cost allocation tag will be displayed as a separate column of the bill. You can filter and categorize bills based on this tag key.

2. On this page, you can view the list of created tag keys. Select the tag key to be displayed and click **Set as Cost Allocation Tag** to set the tag key as a cost allocation tag in the bill.

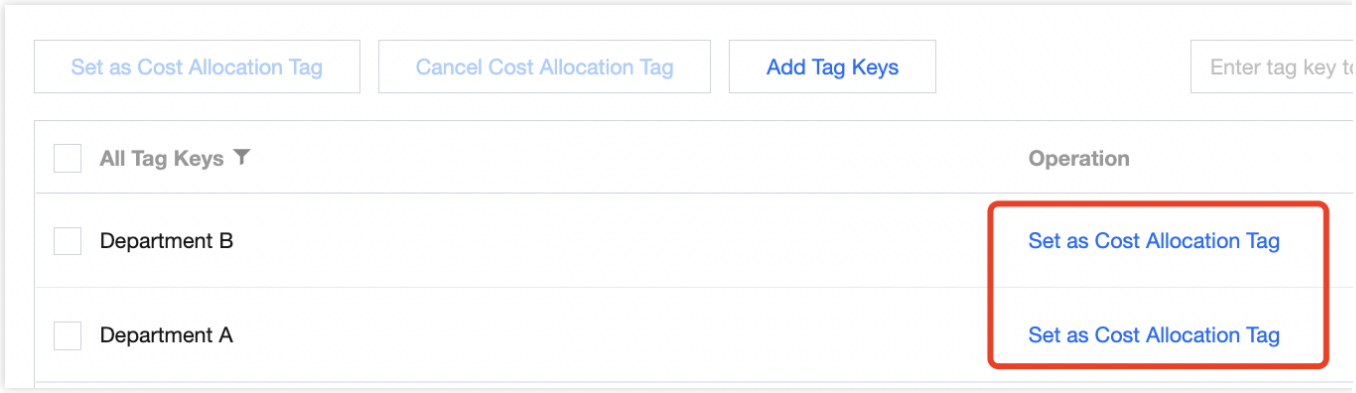

Note:

You can set 5 cost allocation tags at most. A small number of such tags makes it easier for you to manage your costs.

Step 2. Display bills by tag

You can view and click the new option **By Tag** on the **Bill Overview** page. Then, you can select a specific **tag key** to view the column chart and list of relevant resources aggregated by the tag key.

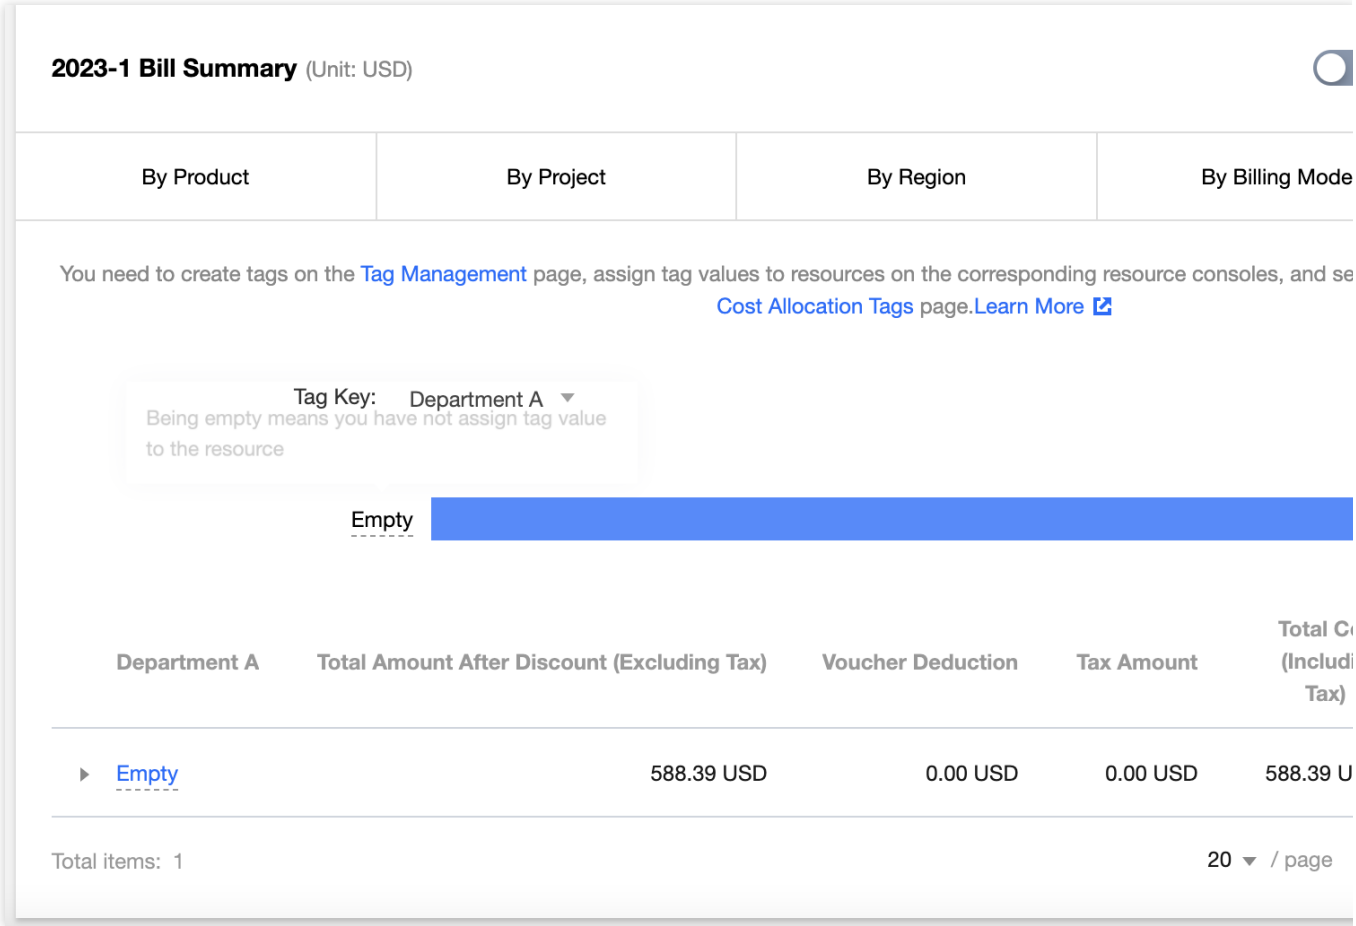

Alarm Service
Creating Alarm

Last updated：2023-12-25 15:51:16

This document describes how to set an alarm for a key application performance metric, so that you can receive timely notifications when the metric is abnormal.

Directions

1. Log in to the APM console.
2. On the left sidebar, click **Alarm configuration** to enter the alarm policy configuration page.
3. Click **Create** and configure a new alarm policy as follows:

Configuration Type	Configuration Item	Description
Basic information	Policy name	The custom policy name.
	Remarks	The custom policy remarks.
	Monitoring type	Select the application performance monitoring type.
	Policy type	Performance metrics are selected by default.
Alarm policy	Filter condition (AND)	You can filter objects that meet filter conditions for alarming, and the relationship between these conditions is AND. Only objects with reported data will be displayed after filtering. Business system (required): You can set alarms by business system. Currently, you can select only a single business system. Application (required): You can filter the performance data of a specific application in a specific business system for alarm detection. Call role (required): You can filter the performance data of a specific call role in a specific application for alarm detection. Instance: You can filter the performance data of a specific instance in a specific business system for alarm detection. API: You can filter the performance data of a specific API in a specific application for alarm detection.
	Alarm object dimension	You can customize the alarm objects displayed in the alarm notification. Suppose you select business system, application, and

		call role, the following alarm objects will be displayed: business system=xxx, application=xxx, call role=xxx.
	Trigger condition	You can specify whether to trigger an alarm when any condition or all conditions are met.
Alarm notification	Notification template	The system provides default notification templates. You can also create a custom one as instructed in Creating Notification Template.
Advanced configuration	Auto scaling	If enabled, the auto scaling policy can be triggered when the alarm condition is met.

4. After configuring the above information, click **Save**. When a metric is abnormal, alarm notifications will be sent through the alarm channels you configured.

Viewing Alarm

Last updated：2024-07-24 16:05:46

This document describes how to view APM alarm records.

Viewing alarm records

1. Log in to the TCOP console and go to the **Alarm Records** page.
2. Select **APM** as the **Monitoring Type**, and click **OK**.
3. You can also set a time range in the top-left corner to filter alarm records during this period.

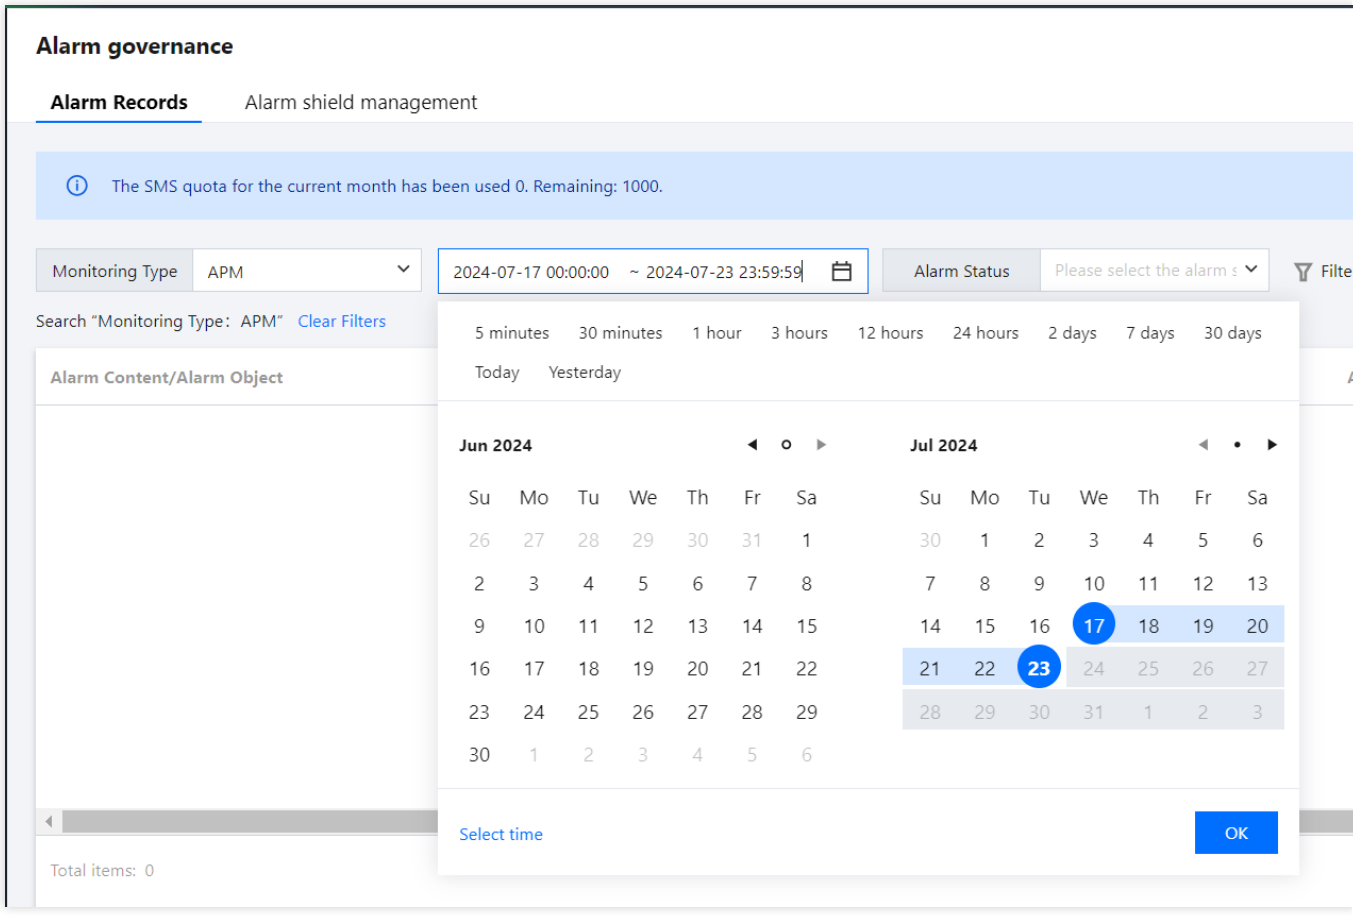

System Configuration
Business System Configuration

Last updated：2023-12-25 15:52:00

You can modify the basic information of the target business system on the **System configuration** > **Business system configuration** page.

Prerequisites

1. Go to the **System management** page in the APM console.
2. Go to **Business system configuration**.

Modifying basic information

At the top of the page, select the target business system and click **Edit** on the right of **Basic info**. You can modify the retention period, business system name, and tag.

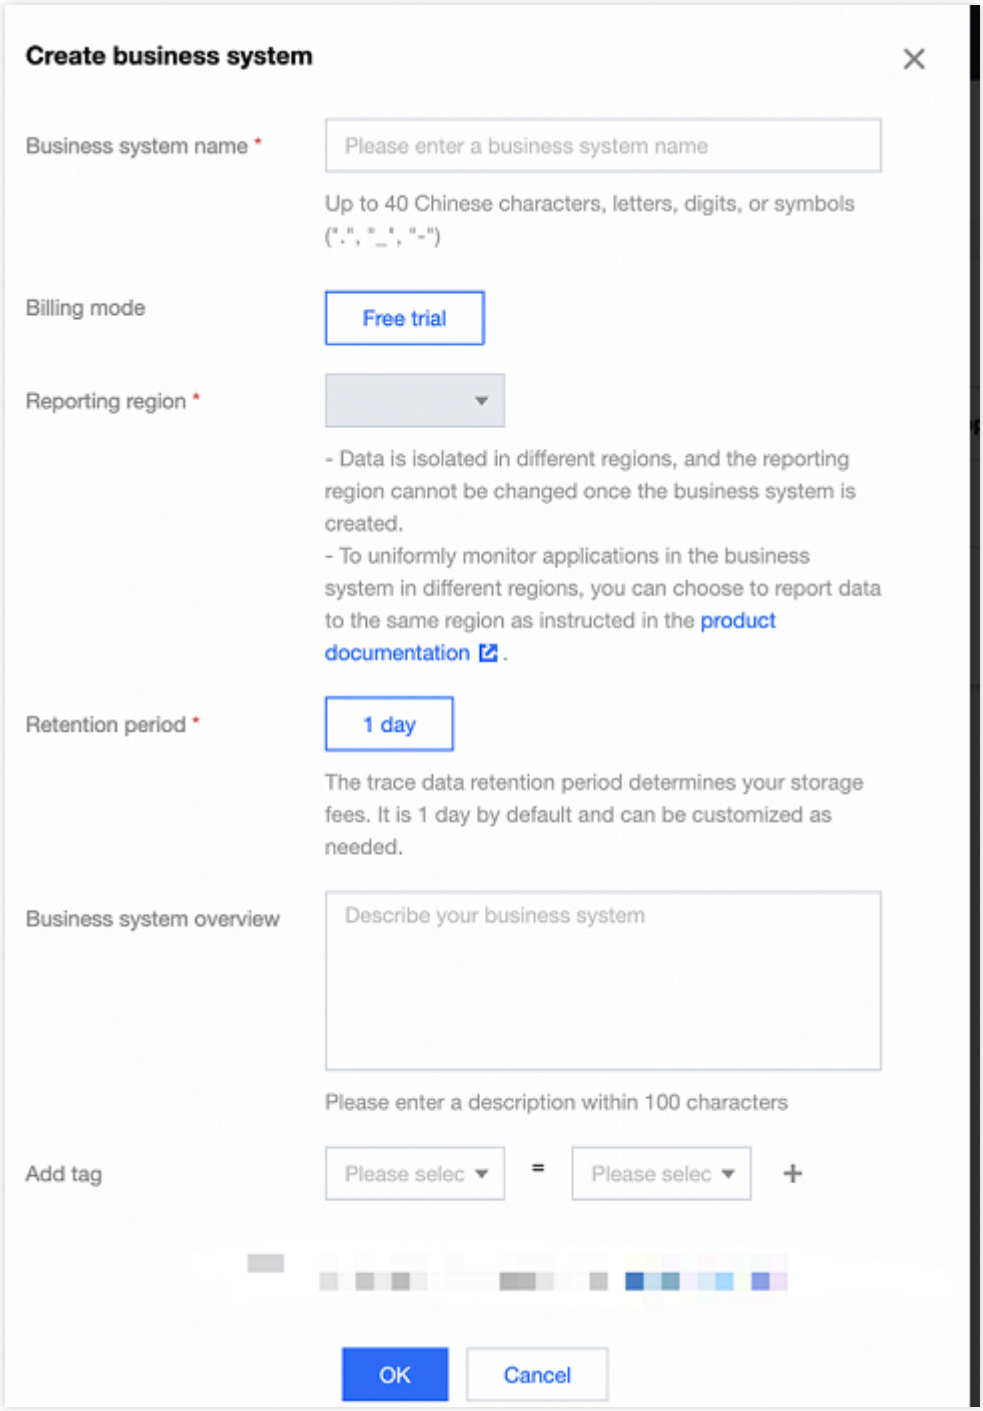

Application Configuration

Last updated：2023-12-25 15:52:16

On the **Application configuration** page, you can customize the exception types in the application details and URL convergence rules and configure whether to include errors with specified status codes in the number of errors.

Prerequisites

1. Go to the **System management** page in the APM console.
2. Go to the **Application configuration** page.

Directions

Exception filtering

This parameter filters the exception type in the exception analysis chart on the application details page. You can use the regex to match the full name of the exception class and separate multiple items by comma, such as `java.lang.inter,java.lang.index`. The exception types entered here will neither be displayed on the **Application details** page nor included in the number of exceptions and exception analysis.

Error code filtering

This parameter filters the status code of errors included in the number of errors. You can set the error codes that need to be ignored here and separate them by comma, such as "429,512". Errors with the status codes entered here will not be included in the number of errors.

URL convergence

This parameter converges multiple similar URLs into one, such as URLs starting with `/service/demo/`.

Convergence threshold

This parameter specifies the minimum quantity for URL convergence. For example, if the threshold is `100`, convergence will be performed when the number of URLs in line with the regex reaches 100.

Convergence rule regex

This parameter specifies the URL convergence rule in regex. For example, to converge URLs starting with `/service/demo/` when their number exceeds 100, the convergence threshold should be `100`, and the convergence rule regex should be `/service/demo/(.*?)`.

Exclusion rule regex

This parameter excludes specified URLs from convergence. For example, to exclude URLs starting with `service/demo/example` from convergence, the regex should be `/service/demo/example/(.*?)`.