ALIBABA CLOUD

Alibaba Cloud

Tracing Analysis

Use the console

Document Version: 20201222

Alibaba Cloud

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.

2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.

3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.

4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	Danger: Resetting will result in the loss of user configuration data.
Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	Notice: If the weight is set to 0, the server no longer receives new requests.
Note	A note indicates supplemental instructions, best practices, tips, and other content.	Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click **Settings**> **Network**> **Set network type**.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click **OK**.
`Courier font`	Courier font is used for commands	Run the `cd /d C:/window` command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`
{} or {a|b}	This format is used for a required value, where only one item can be selected.	`switch {active|stand}`

Table of Contents

1.Application management

1.1. Overview

On the overview page, you can view the overall metrics of your application, including the number of API request entries, average response time, and related metrics. In addition, you can view the access procedure and access point information.

Procedure

1.
2. In the left-side navigation pane, click **Overview**. On the top of the page, click the **Overview analysis** tab.

Overview

The **Overview analysis** tab shows these key metrics:

- The number of inbound requests, average response time, number of spans, and number of Span exceptions within the selected time, as well as the year-on-year increase or decrease compared with that of the previous day and the previous day.
- Displays the name, region, health check result, requests, errors, and response time of each application.
- The name, application, number of requests, response time, and constant of the entry.
- The name, alert trigger status, alert time, alert content, alert level, and rules to which the alert belongs.

Procedure

The **process** tab shows the process and process. In addition, you can view access point information.

Related information

- Activate related services and authorize roles

1.2. View applications

On the Applications page of the Tracing Analysis console, you can check key metrics of all monitored applications, including the request count and error count in the current day, and the health status. You can also set custom tags for your app and use tags to filter.

Context

The applications page displays multiple key metrics of the monitored Application. The scores of these metrics are calculated based on the APDEX Performance Index (Application Performance Index). It is an international standard for evaluating the performance of applications. The user experience of an application includes the following aspects:

- - Satisfied (0 to T)

- - Tolerating (T to 4T)
- - Frustrated (greater than 4T)

The following formula is used to calculate the APDEX score:

```
APDEX = (satisfactory number + tolerable number /2) /Total sample size
```

Tracing analysis uses the average response time of an application as a calculation metric and defines T as 500 milliseconds.

Procedure

Follow these steps to enter the **application list** page.

1. Log on Tracing Analysis console.
2. In the left-side navigation pane, choose **application list** , and select the target region at the top of the **application list** page.

Applications page

Sorting application

You can sort all apps in ascending or descending order by clicking the arrow next to the column headers:

- **Health Score**
- **Number of requests today**
- **Number of errors today**
- **Response time**

Set application tags

After you set custom tags for an application, you can use these tags to filter the application.

1. Place the pointer over the **label** column and click the pencil icon.
2. In the **manage account labels** dialog box, enter custom labels in the **add labels** field and click **add** . Click one or more labels at the top, and click **OK** .

> **Notice** If you delete an existing tag in the **manage account tags** dialog box, the app that previously added the tag will lose it.

Filter applications by tag

In the **select tags** section, click one or more tags to filter all applications that have at least one tag.

Related information

- Before you begin
- Manage apps and tags

1.3. View key application performance metrics and topology

The application overview page displays the key performance metrics and topology of your application.

Context

After the application data is reported to Tracing Analysis, Tracing Analysis monitors your application. On the **Application Overview** page, you can quickly view key metrics for application performance and view the upstream and downstream dependencies of an application using the topology.

View key application performance metrics

You can view the key metrics of application performance on the **Overview analysis** tab.

1. Log on Tracing Analysis console.
2. In the left-side navigation pane, click **application list** . On the top of the **application list** page, select a region, and then click the application name.
3. On the **Application Overview** page, view the following key metrics on the **Overview analysis** tab.
   - The total number of requests, average response time, number of abnormal calls, and Span in the selected time range, and how these indicators change from the previous week to the previous day or previous week.
   - The line chart of how many times your application is called by upstream components and the response time, and how many times your application calls downstream services and the response time.
   - The top 10 APIs with the lowest call speed and their average response time sequence curves.

     Overview analysis

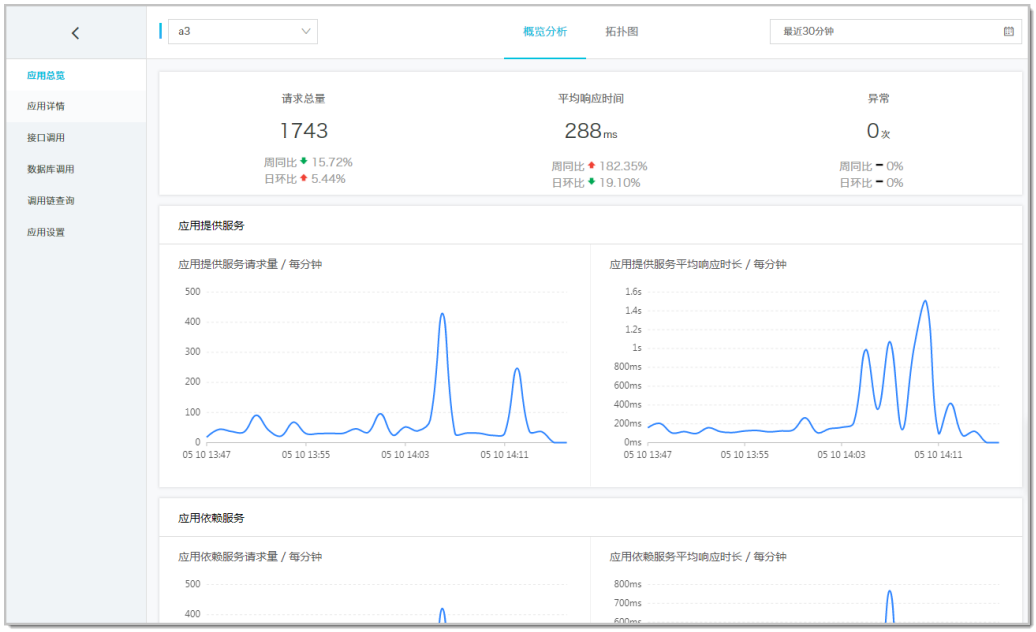

View the application topology

On the **topology** tab, you can have a better view of the upstream and downstream components of your application and their call relations, allowing you to quickly identify the performance bottlenecks of your application.

1. In the left-side navigation pane, click **application list** . On the top of the **application list** page, select a region, and then click the application name.
2. On the **Application Overview** page, click the **Topology Graph** tab. On the **Topology Graph** tab, you can view the following information:
   - The call topology of the application within the selected time range.
   - The number of calls, average response time, and error rate of the client, provider, and internal calls within the selected time.
   - The sequence time chart of the number of requests, response time, and error rate per minute within the specified period.

   Topology tab

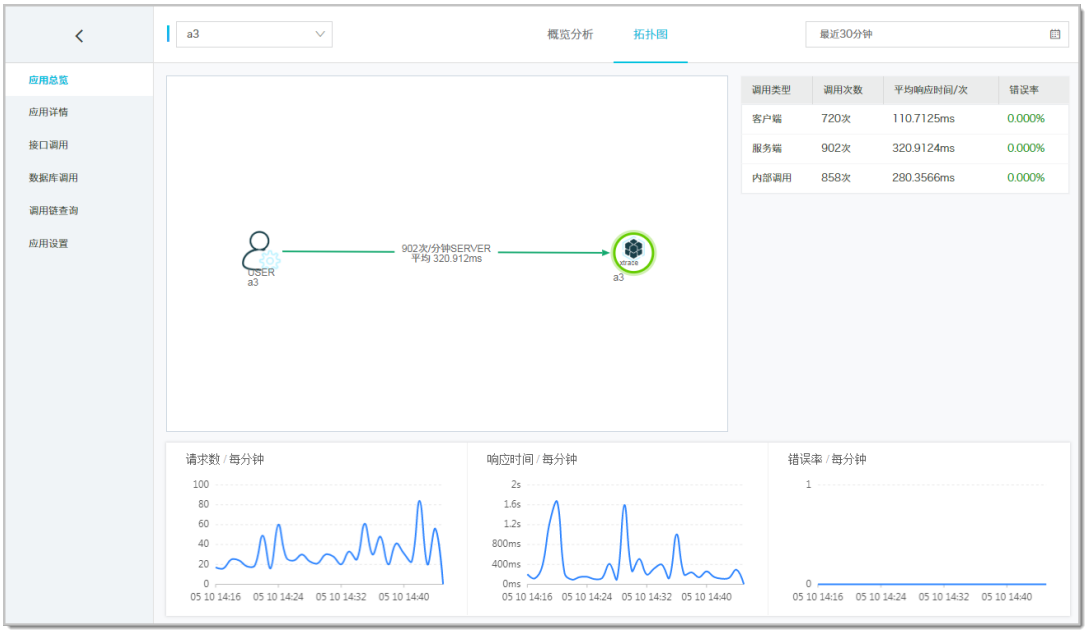

Set the query time range

You can select a preset time range or enter a custom time range.

- Click the time option in the upper-right corner of the page, and then click a preset time range, such as **Last 30 minutes**,**This week**,**Last 30 days**.
- If no preset time range meets your requirements, click **Custom**. Select the start time and end time from the calendar, or enter them manually in the text box. Then, click **OK**.

> **Note** The date format is `YYYY-MM-DD`, The time format is `HH:MM`.

Query time range selector

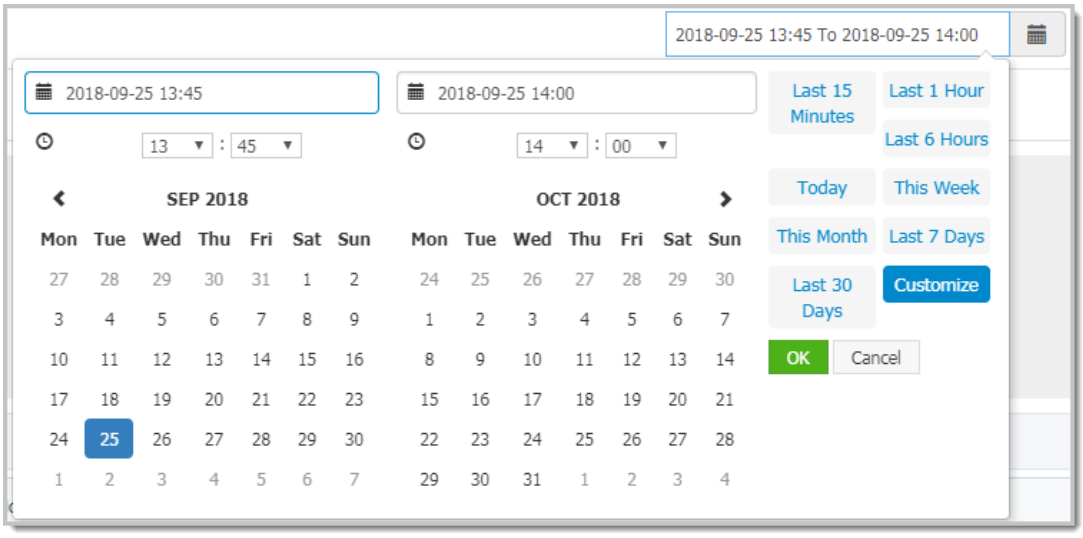

Related information

- Before you begin
- Manage apps and tags

1.4. View application details

On the details page of an application in the Tracing Analysis console, you can view the key metrics, call topology, and traces of each server.

View key performance metrics and topology

Application details page: the **overview** tab displays all instances deployed with the application. You can sort the interface calls by response time, request count, and error count. Select a single machine from the machine List. On the **overview** tab, you can view the detailed call topology of an application and the time sequence curves of the number of requests, response time, and number of errors.

1. Log on Tracing Analysis console.
2. In the left-side navigation pane, click **application list** . On the top of the **application list** page, select a region, and then click the application name.
3. In the left-side navigation pane, click **application details** . In the left-side server list, click **all** or a server. On the **overview** tab, view the call topology and key performance metrics.

   > **Note** Click the **response time** , **requests** , **errors** tab, and then click the arrow next to the tab. You can sort the instances by the required criteria. Enter a keyword in the search box to dynamically filter machines that meet the keyword.

> **Note** To switch to another application in the same region, click the application name drop-down list in the upper-left corner and select another application.

View traces

The **Call link**

Trace tab

> **Note** **Status** The green icon in the column indicates that the time consumed is less than 500 milliseconds, the yellow icon indicates that the time consumed is between 500 milliseconds and 1000 milliseconds, the red icon indicates that the time consumed is greater than 1000 milliseconds, or the Tag Key is `Error` .

On the **Call link** tab page, you can perform the following operations as needed:

- In **Time-consuming** enter a time value (in milliseconds) in the adjustment box, and click **Query** to filter traces whose time consumption is greater than the specified value.
- Select **Exception** and click **Query** to filter out abnormal call links.
- Click **Time consumed** Or **Status** By using the up and down arrows on the right, you can sort the query results in ascending or descending order.
- Click the TraceID to open it in a new window. **Call link** Page, and view the waterfall chart of the call link.

View trace waterfall chart

On the **Call link** page, you can check the log generation time, status, IP address/machine name, service name, and Timeline of the call link are displayed on the page.

> **Note** **IP address**Whether the IP address or machine name is displayed depends on **Application Settings** the display configuration on the page. For more information, see manage applications and tags.

Trace page

Place the cursor over a service name to view the service duration, start time, Tag, and log event information.

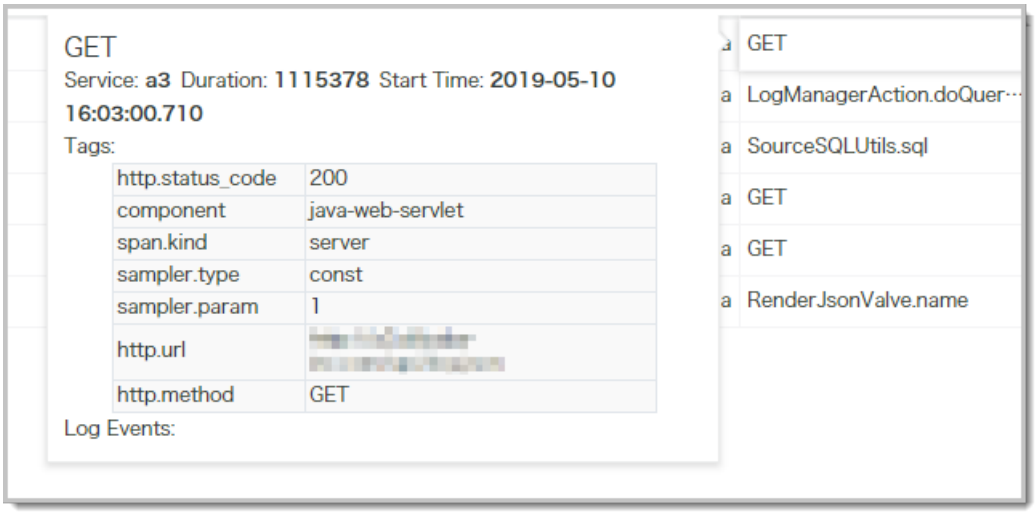

Set the query time range

You can select a preset time range or enter a custom time range.

- Click the time option in the upper-right corner of the page, and then click a preset time range, such as **Last 30 minutes**,**This week**,**Last 30 days**.
- If no preset time range meets your requirements, click **Custom**. Select the start time and end time from the calendar, or enter them manually in the text box. Then, click **OK**.

> **Note** The date format is `YYYY-MM-DD` , The time format is `HH:MM` .

Query time range selector

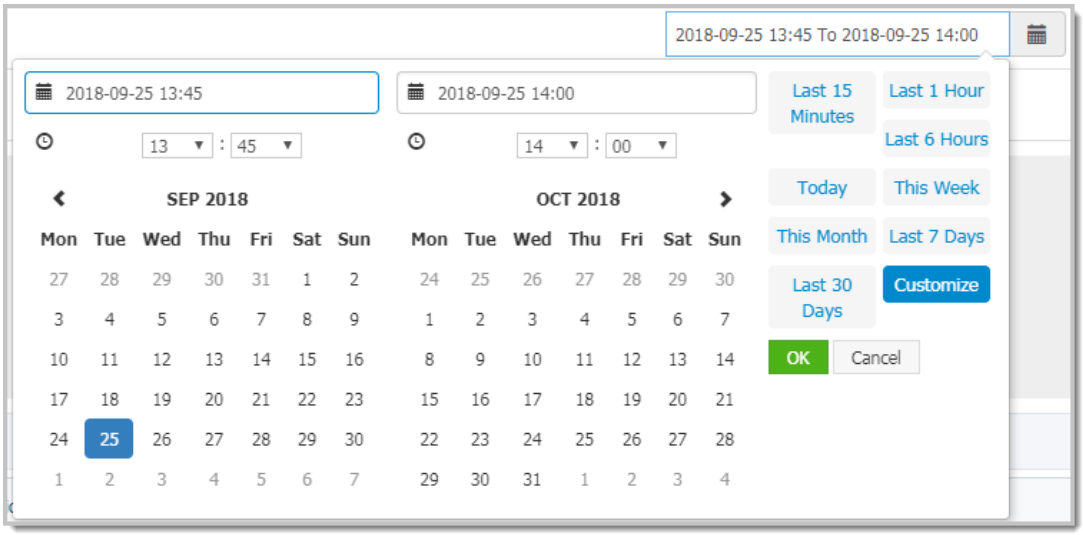

1.5. View API usage

The API invocation page displays the performance indicators of API calls in client calls, server calls, and local calls, as well as API calls of upstream and downstream links.

View API call performance metrics

The **interface invocation** page lists all the interfaces (spans) involved in the application call. You can sort the list by response time, number of requests, or number of exceptions. Select an interface from the interface list. On the **overview** tab, you can view the topology of an application and the time sequence curves of the API call performance metrics, including the number of requests, response time, and variation constants.

1. Log on Tracing Analysis console.
2. In the left-side navigation pane, click **application list** . On the top of the **application list** page, select a region, and then click the application name.
3. In the left-side navigation pane, click **interface invocation** . Click an interface in the left-side navigation pane, and then click the **overview** tab to view the topology and performance metrics of the interface.
   - Click the **response time** , **requests** , **exceptions** tab, and then click the upward or downward arrow next to the tab. Then, you can sort all APIs in ascending or descending order by the specified condition.
   - In the **call type** section, click **all** , **client** , **server** , or **local call** to filter the types of interfaces that need to be called.
   - You can enter a keyword in the search box to dynamically filter API operations that meet the keyword.

> **Note** To switch to another application in the same region, click the application name drop-down list in the upper-left corner and select another application.

View upstream and downstream services

The **upstream** interface and **downstream** interface tabs display the interfaces of the upstream application that calls the selected application and downstream application that calls the specified application and their performance metrics, including the number of requests, response time, and latency.

Downstream link tab

On the **upstream** and **downstream** tabs, you can perform the following operations as needed:

- At the top of the tab, click **show /hide** all.
- On the tabs, enter an application name or an interface (span) name in the search box, and click the magnifier icon to filter out the interfaces that meet corresponding conditions.
- Click the collapse panel where the interface call information is located, or click the up or down arrow at the end of the row to expand or collapse the performance metric information of this interface call.

View traces

The **Call link**

Trace tab

> **Note** **Status** The green icon in the column indicates that the time consumed is less than 500 milliseconds, the yellow icon indicates that the time consumed is between 500 milliseconds and 1000 milliseconds, the red icon indicates that the time consumed is greater than 1000 milliseconds, or the Tag Key is `Error` .

On the **Call link** tab page, you can perform the following operations as needed:

- In **Time-consuming** enter a time value (in milliseconds) in the adjustment box, and click **Query** to filter traces whose time consumption is greater than the specified value.
- Select **Exception** and click **Query** to filter out abnormal call links.
- Click **Time consumed** Or **Status** By using the up and down arrows on the right, you can sort the query results in ascending or descending order.
- Click the TraceID to open it in a new window. **Call link** Page, and view the waterfall chart of the call link.

View trace waterfall chart

On the **Call link** page, you can check the log generation time, status, IP address/machine name, service name, and Timeline of the call link are displayed on the page.

> **Note** **IP address** Whether the IP address or machine name is displayed depends on **Application Settings** the display configuration on the page. For more information, see manage applications and tags.

Trace page

Place the cursor over a service name to view the service duration, start time, Tag, and log event information.

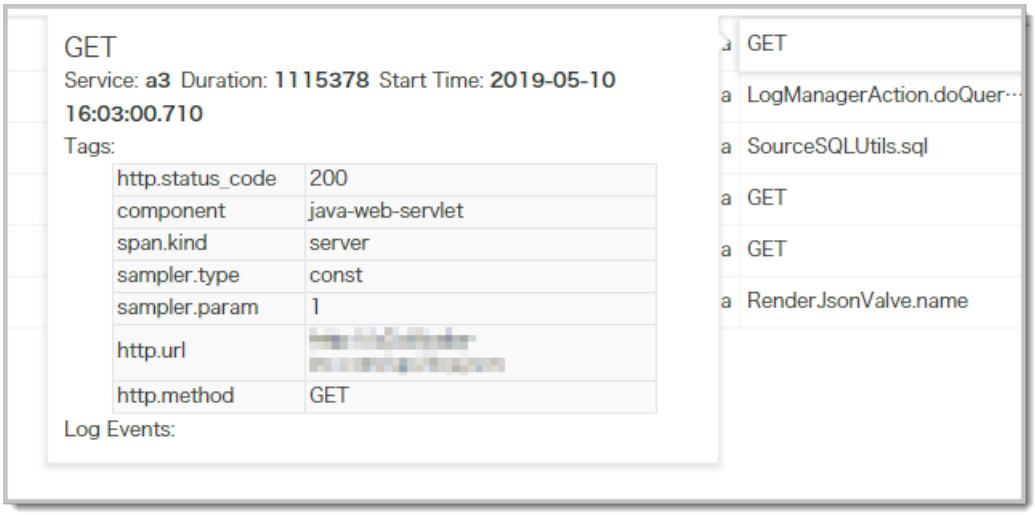

Set the query time range

You can select a preset time range or enter a custom time range.

- Click the time option in the upper-right corner of the page, and then click a preset time range, such as **Last 30 minutes**,**This week**,**Last 30 days**.
- If no preset time range meets your requirements, click **Custom**. Select the start time and end time from the calendar, or enter them manually in the text box. Then, click **OK**.

> **Note** The date format is `YYYY-MM-DD`, The time format is `HH:MM`.

Query time range selector

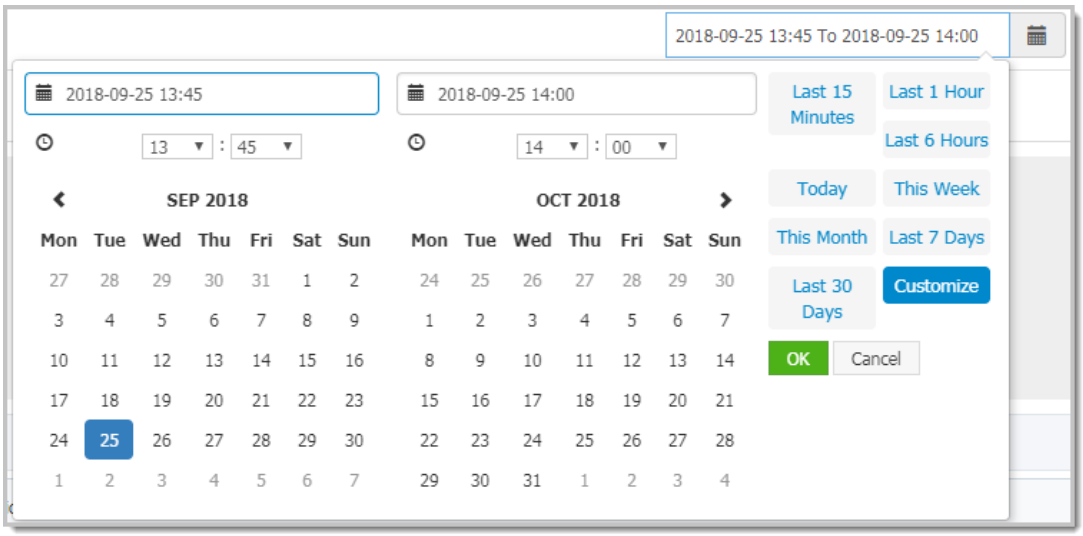

Related information

- Before you begin
- Manage apps and tags

1.6. View SQL performance analysis

The database call page displays the call times, average time consumption, and call links of each SQL statement to help you locate SQL performance problems.

View SQL analysis

Follow these steps to view the SQL statistics and analysis of the application.

1. Log on Tracing Analysis console.
2. In the left-side navigation pane, click**Application List**, And in**Application List**Select a region and click the application name.
3. In the left-side navigation pane, click**Database call**, And then choose**SQL analysis**On the Tab Page, view the following metrics:
   - Displays the number of SQL calls per minute and average elapsed time within a specified period.
   - The number of times an SQL statement is called and the average time consumed during the selected period.
4. In**SQL analysis**On the Tab Page, perform the following operations as needed:
   - In**Operation**Column, click**Call statistics**To view the chart of the number of calls per minute and the average time consumed by a specific SQL statement within the selected time range.
   - In**Operation**Column, click**Link query**In the**Call link**Tab page to view all call links related to the corresponding SQL statement.

> **Note** To switch to another application in the same region, click the application name drop-down list in the upper-left corner and select another application.

Set the query time range

You can select a preset time range or enter a custom time range.

- Click the time option in the upper-right corner of the page, and then click a preset time range, such as **Last 30 minutes**,**This week**,**Last 30 days**.
- If no preset time range meets your requirements, click **Custom**. Select the start time and end time from the calendar, or enter them manually in the text box. Then, click **OK**.

> **Note** The date format is `YYYY-MM-DD`, The time format is `HH:MM`.

Query time range selector

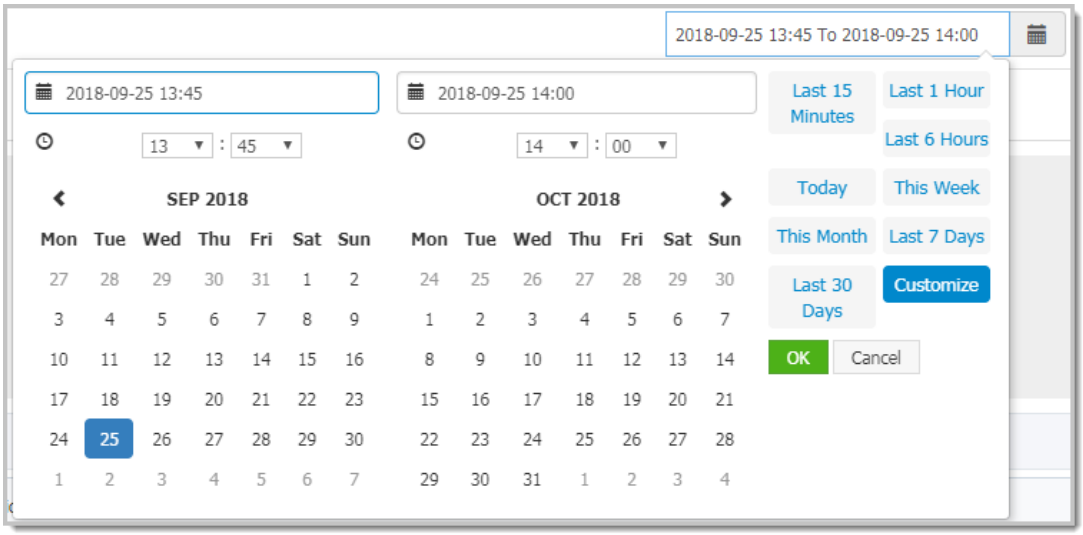

Related information

- Before you begin
- Manage apps and tags

1.7. Query trace

This paper introduces how to use multi-dimensional query function to query call chain.

Procedure

1. Log on Tracing Analysis console.
2. In the left-side navigation pane, click **multidimensional query** .

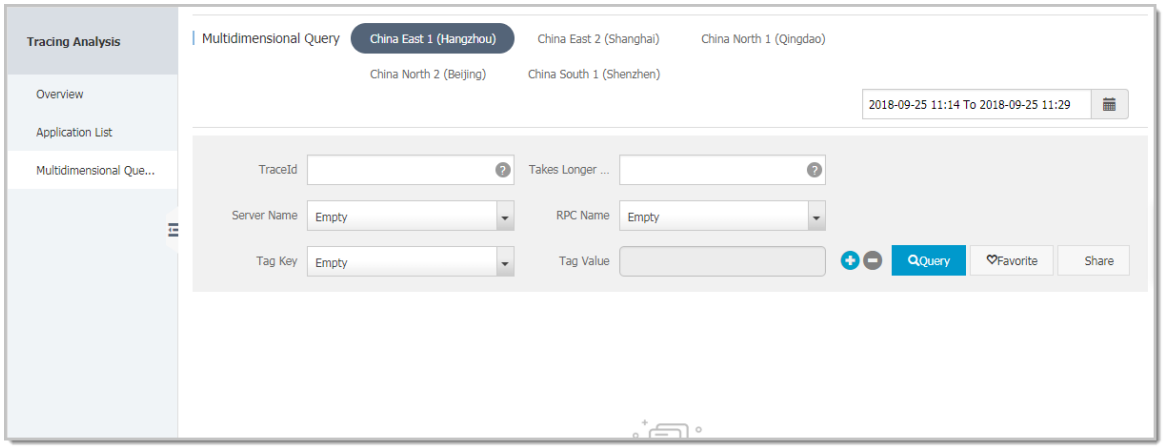

3. On the **multidimensional query** page, enter values for the following parameters as needed, and click **query** .
   - Fixed query parameters
     - TraceID
     - Time consumed: greater than (milliseconds)
     - Server Name
     - RPC name

- Optional query parameters: Tag key

  Select the key of the Tag from the **Tag key** drop-down list and fill in the **Tag value** field with the value of the Tag. To add a Tag key or Tag value as a query condition, click the blue plus sign icon to the right of the field.

4. In the search results, click the TraceID to go to the **invocation trace** tab and view the trace details.

What's next

- To save the current parameter configurations, click **favorites** . The parameter configurations of favorite queries are displayed in **the favorite queries** section.
- To delete all the query parameter configurations from your favorite, click **the favorite query** on the right side of **clear** .

Related information

- Before you begin
- View applications

1.8. Real-time diagnosis

After the data is reported to the Tracing Analysis component, the Tracing Analysis component performs real-time aggregate computing on the data. Generally, there is a certain delay. If you want to see the real-time result when locating a problem, you can use the real-time diagnosis function to quickly display the diagnostic result.

Context

- Real-time diagnostics enables temporary storage. After data is reported, you can quickly display diagnostic results without real-time statistics.
- Real-time diagnosis does not affect normal statistics, and is automatically disabled in five minutes. After the function is disabled, you can enable it again.
- By default, the real-time diagnosis page is refreshed regularly every 10 seconds. You can also disable the timed refresh function.

View trace information

1. Log on Tracing Analysis console.
2. In the left-side navigation pane, click **Application List** , And in **Application List** Select a region and click the name of the target application.
3. In the left-side navigation pane, click **Real-time diagnosis**, In **Real-time diagnosis** Page. Click the + icon on the top to add **Span name**,**IP** And **Tag** Three filter criteria.

   > **Note** When you add a filter, you can add one **Span name** Or **IP** And multiple **Tag**.

4. Click **Query** To view the filtered trace information, including:
   - The scatter chart of real-time request response time distribution.

> **Note** You can select an area in a scatter chart to obtain the real-time diagnostic results for this area.

- Request count/time consumption distribution chart.
- The list of trace information.

View the waterfall chart of call traces

1. Click a trace ID in the trace list. The waterfall chart appears.

   On the **trace** page, you can view the trace information such as the Span name, application name, status, IP address /machine name, log generation time, and timeline.

   > **Note** The **IP** address field displays the IP address or machine name, depending on the display configuration on the **app settings** page. For more information, see Manage apps and tags.

   Link page

   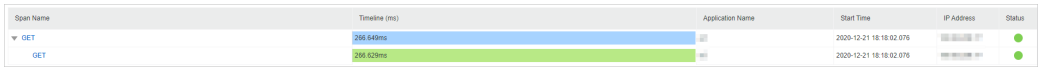

2. Move the pointer over a Span to view the time, start time, Tag, and event log of the Span.

   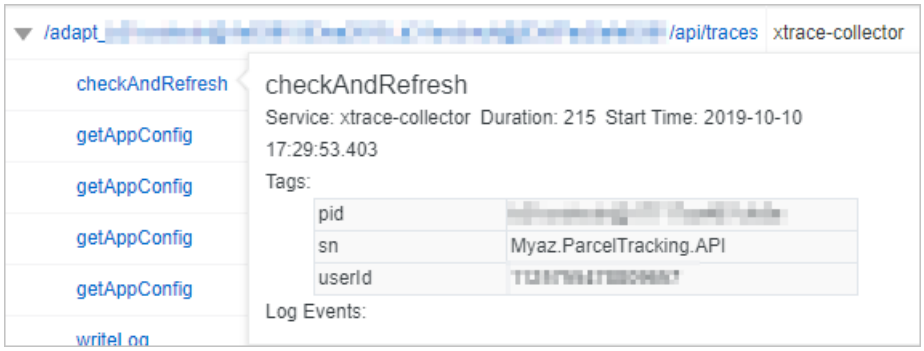

View the interface aggregation list

1. Click**Interface aggregation**Tab to view the list of API aggregates that aggregate call chains by Span name.

1.9. Analyze traces

After analyzing the call link information of an application on the call link analysis page, you can filter the call link information by condition. you can also view the link topology, real-time aggregate link table, and call link waterfall chart.

View trace information

1. Log on Tracing Analysis console.
2. In the left-side navigation pane, click**Application List**, And in**Application List** Select a region and click the name of the target application.
3. In the left-side navigation pane, click**Trace Analysis**, In**Trace Analysis**Page to filter trace

information, as shown in the following figure.

- Click **Comprehensive condition**input box, which can be added simultaneously or separately **Span name**, **IP** and**Label**. Three filter criteria. When you add a filter, you can add one **Span name** or **IP**, And multiple **Label**.
- In**Time consumed (> ms)**Enter a specific response time in the input box to query trace information that is later than this time consumption.
- Check**Exception**Check box to query trace information with exceptions.

4. Click**Search**To view the filtered trace information, including:
   - Time series curves of time consumption and number of spans.
   - The distribution chart of the number of spans and the time consumed.
   - You can filter call link information by Span, IP address, or Tag.

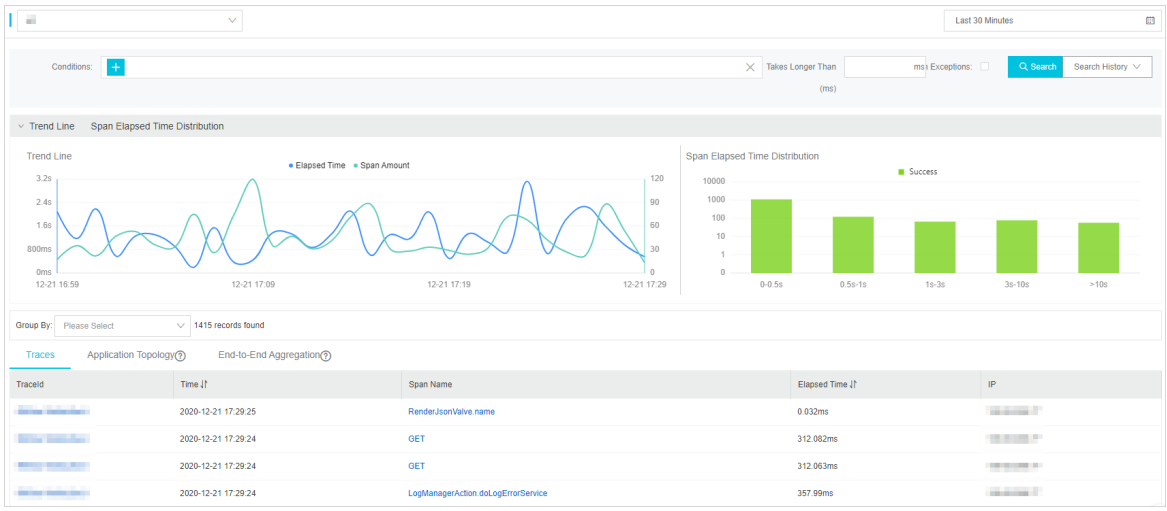

5. From**Group by**List, select Span, IP, or Tag to filter groups. For example, select Tag**IP**.

6. Click a IP. A list of trace information related to this IP is displayed.

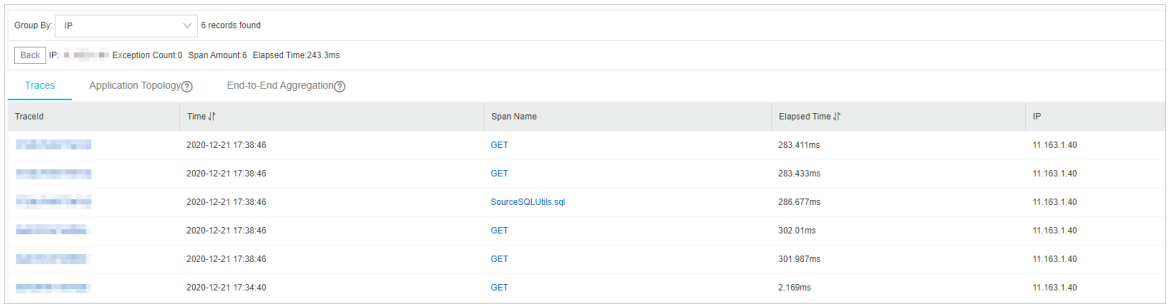

View the link topology

The link topology mainly displays the topology of the dependency between applications after conditional filtering, as well as the request ratio, call speed, and time consumption ratio between applications. In consideration of the performance experience, the link topology can pull up to 5000 link requests for aggregation.

1. Click **Link topology** Tab to view the link topology.

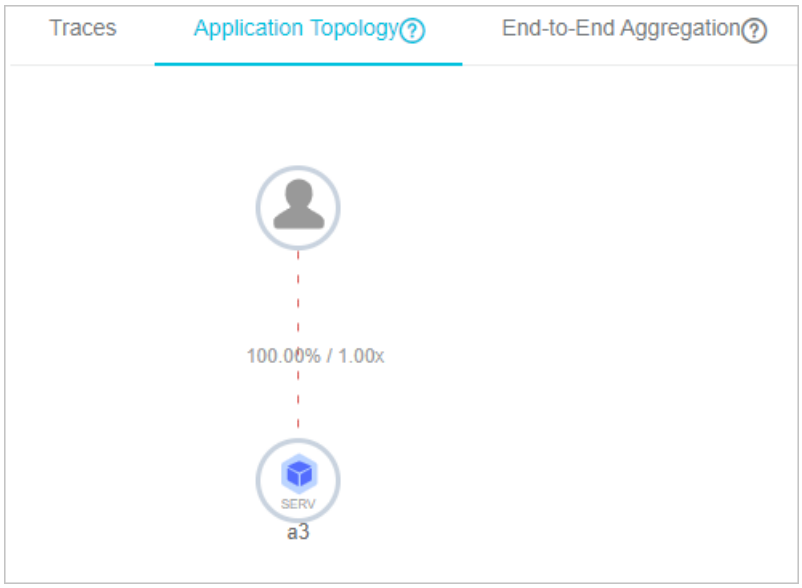

> Note
>
> - Request ratio = number of requests requested by an application for external calls/total number of requests of an application. For example, if 100 requests enter upper-layer application A and only 90 requests from A call lower-layer application B, the request ratio from A to B is 90%. (In application A, some requests may not go into application B because some requests may be filtered by if judgment.)
> - Call multiple = number of spans called by the application/total number of spans of the application. For example, if 100 spans enter upper-layer application A and 300 spans call lower-layer application B from A, the call multiple from A to B is 3. For example, the values of A to B are 90%/3x, indicating that 90% of requests from application A call Application B, and application A Calls application B three times on average.

View the real-time aggregation link table

Real-time aggregation is a call link table that aggregates call links that have been filtered by conditions based on Span names and application names. In consideration of performance experience, real-time aggregation supports pulling up to 5000 link requests for aggregation.

1. Click **Real-time aggregation** Tab to view the real-time aggregation link table.

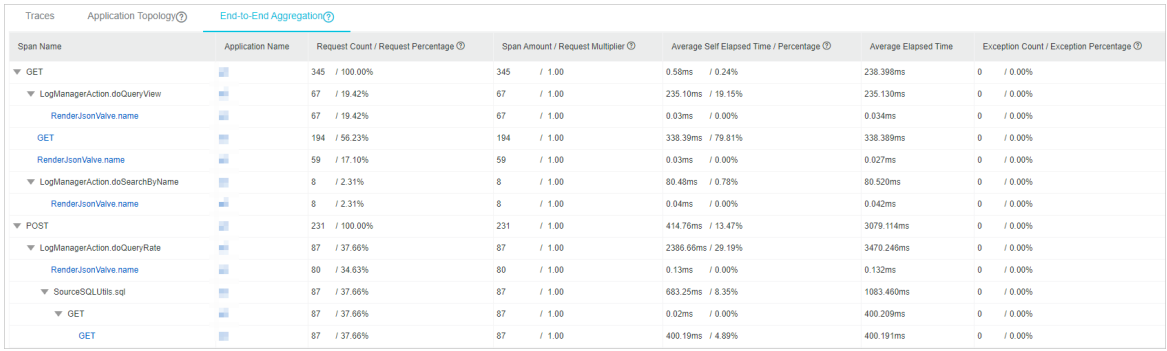

> **Note**
>
> - Requests/request ratio: The request ratio indicates the percentage of requests that call the current Span node. For example, if the total number of requests is 100 and the request ratio is 10%, 10 requests call the current Span. Calculation formula = number of requests for the current Span/total number of requests X 100%.
> - Span/request multiple: The request multiple indicates the average number of times each request calls the current Span. For example, 1.5x indicates that each request calls the current Span 1.5 times on average. Calculation formula = number of spans/number of requests per Span.
> - Average self-elapsed time/ratio: the average self-elapsed time does not include the average elapsed time of subspans. For example, if Span A to B, the elapsed time of A is 10 milliseconds, and that of B is 8 milliseconds. the time consumed by A is 2 milliseconds. The formula is as follows: Span elapsed time-Sum (subspan elapsed time). If an Asynchronous call is performed, the child time consumed is not subtracted. The formula is as follows: Span time consumed.
> - Exception count/exception percentage: The exception percentage indicates the percentage of requests with exceptions. For example, 3% indicates that 3% of requests have exceptions. Calculation formula = number of abnormal requests/total number of requests. The number of exception requests is not equal to the number of exceptions. If the request multiple is greater than 1, an exception request may correspond to multiple exceptions.

2. (Optional)Hover the mouse over the blue Span name to display**Recommended call link**Prompt information, you can view the call chains associated with this Span. Click a trace id. The trace waterfall chart is displayed. For more information, seeView trace waterfall chart.

View trace waterfall chart

1. Click a trace ID in the trace list. The trace waterfall chart is displayed.

   In**Call link**You can view the Span name, application name, status, IP address/machine name, log generation time, and Timeline of a call chain.

> **Note** **IP address**Whether the IP address or machine name is displayed depends on**Application Settings**The display configuration on the page. For more information, seeManage apps and tags.

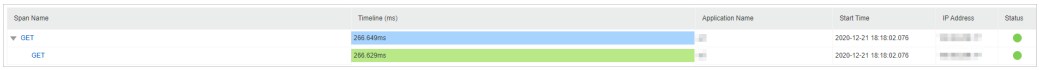

2. Place the cursor over a Span name to view the length, start time, Tag, and log events of the Span.

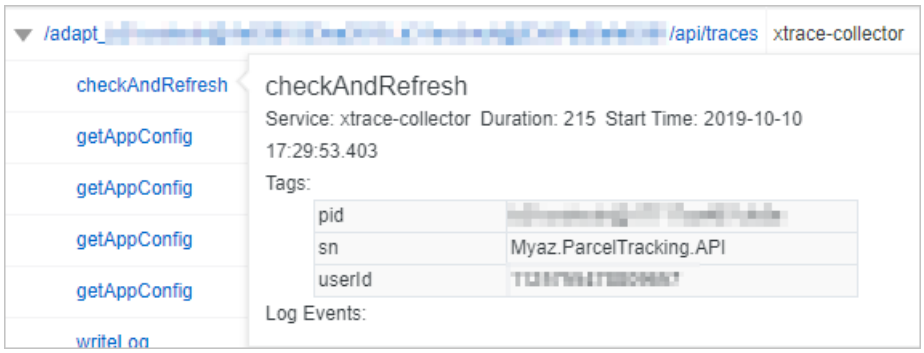

Set the query time range

You can select a preset time range or enter a custom time range.

- Click the time option in the upper-right corner of the page, and then click a preset time range, such as **Last 30 minutes**,**This week**,**Last 30 days**.
- If no preset time range meets your requirements, click **Custom**. Select the start time and end time from the calendar, or enter them manually in the text box. Then, click **OK**.

> **Note** The date format is `YYYY-MM-DD`, The time format is `HH:MM`.

Query time range selector

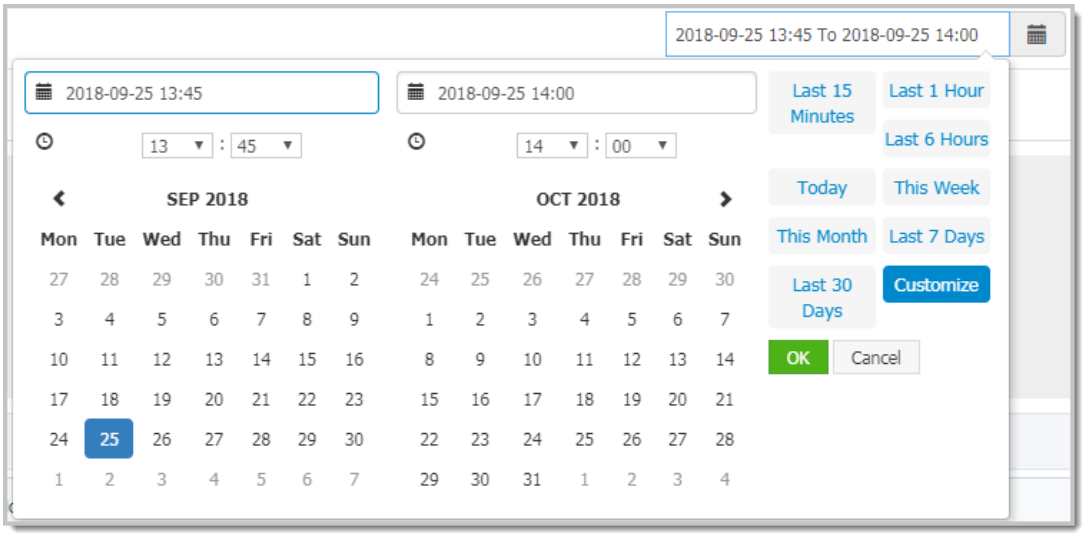

Related information

- Before you begin
- Manage apps and tags

1.10. Manage apps and tags

On the Application Settings page, you can specify whether to display the machine name and collect application data. You can also manage the custom tags of applications and delete applications.

Background

In the link Tracing Analysis console, when you need to display the IP address of the host that is deployed by application, the IP address of the host is displayed by default. However, you can select the host name on the **application settings** page.

> **Note** Only the data that is generated after the setting takes effect is affected. For example, if the **display machine name** switch is turned on and saved, the machine name will only be displayed in the new data generated thereafter, and the IP address will still be displayed in the previously generated data.

Example: Display the IP address of the machine

Example: Display Machine name

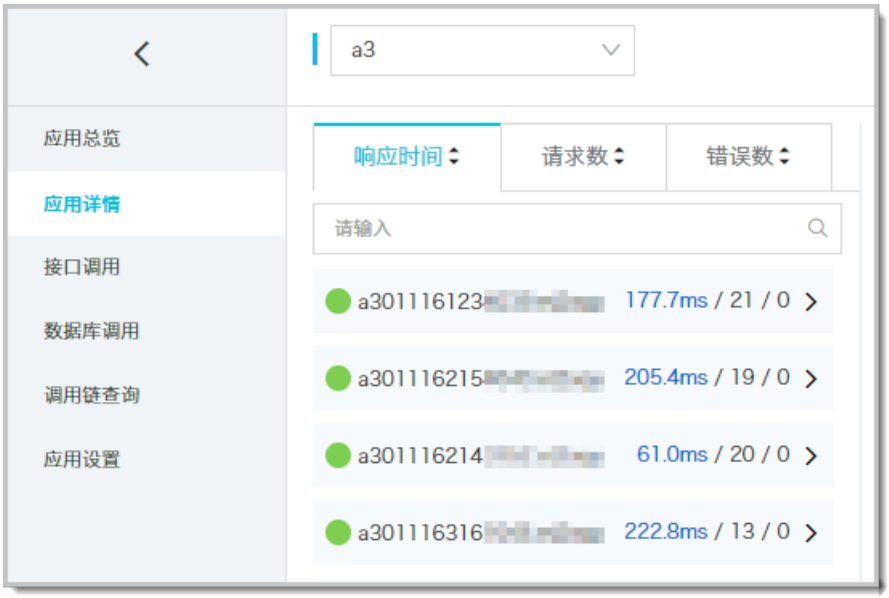

To stop billing for an application, disable data collection.

On the **application settings** page, you can also configure whether to use all tags for the current application and manage all tags under your account. You can delete an application when you no longer need it.

Open Application Settings

Follow these steps to open the application settings page.

1. Log on Tracing Analysis console.
2. In the left-side navigation pane, choose **application list** , and select a region at the top of the **application list** page.
3. On the **applications** page, click **actions** in the **settings** column of an application.

Display Machine name

Perform the following steps to display the name of the host where the application is deployed:

1. On the **application settings** page, click the **custom configuration** tab (displayed by default).
2. On the **custom configuration** tab, turn on **display configuration** switch in the **display configuration** area.
3. Click **save** at the bottom of the page.

> **Note** Only the data that is generated after the setting takes effect is affected. For example, if the **display machine name** switch is turned on and saved, the machine name will only be displayed in the new data generated thereafter, and the IP address will still be displayed in the previously generated data.

Stop Collecting application data

To stop application billing, perform the following steps to stop application data collection.

1. On the **application settings** page, click the **custom configuration** tab (displayed by default).
2. In the **collection configuration** section, turn on **disable data collection** .
3. Click **save** at the bottom of the page.

Enable or disable labels for an application

Follow these steps to enable or disable an existing label for your app.

1. On the **application settings** page, click the **labels** tab.
2. In the **apply tags** section, select the tags to be enabled and clear the corresponding tags.
3. Click **save** at the bottom of the page.

Manage all tags of an account

To manage all tags under an account, such as adding tags for all applications or deleting all existing tags, perform the following steps:

1. On the **application settings** page, click the **labels** tab.
2. Under **apply tags** , click **manage account tags** .
3. In the **manage account tags** dialog box, perform the following operations as needed.

Manage account labels dialog box

- To create a new tag, click the plus icon and enter a tag in the text box.
- To delete an existing tag, move the pointer over the tag and click the X icon on the left.

> **Notice** If you delete an existing tag in the **manage account tags** dialog box, all applications that have enabled the tag will lose it.

4. Click **confirm** at the bottom of the dialog box.

Delete an application

Follow these steps to delete unnecessary applications.

1. On the **application settings** page, click **delete** .
2. In the **delete application** section, click **delete** . In the **note** dialog box that appears, click **OK** .

2.Alerting

2.1. Manage alerts

You can manage all alert rules under your account and query alert events and alert notification records through the alert management module in the tracing analysis console.

Manage an alert rule

The created alarm rule is displayed on the **alarm rules and history** page. You can start, stop, edit, and delete the alert rules. You can also view alert details.

1. Log on Tracing Analysis console.
2. In the left-side navigation pane, choose **Alarm management > Alarm rules** .
3. On the **alarm rules and history** tab page, enter an alarm name in the search box and then click **search** .

   > **Note** You can enter part of an alert name in the search box to perform a fuzzy search.

4. In the search results list of **operation** column, on-demand target alarm rule to take the following actions:

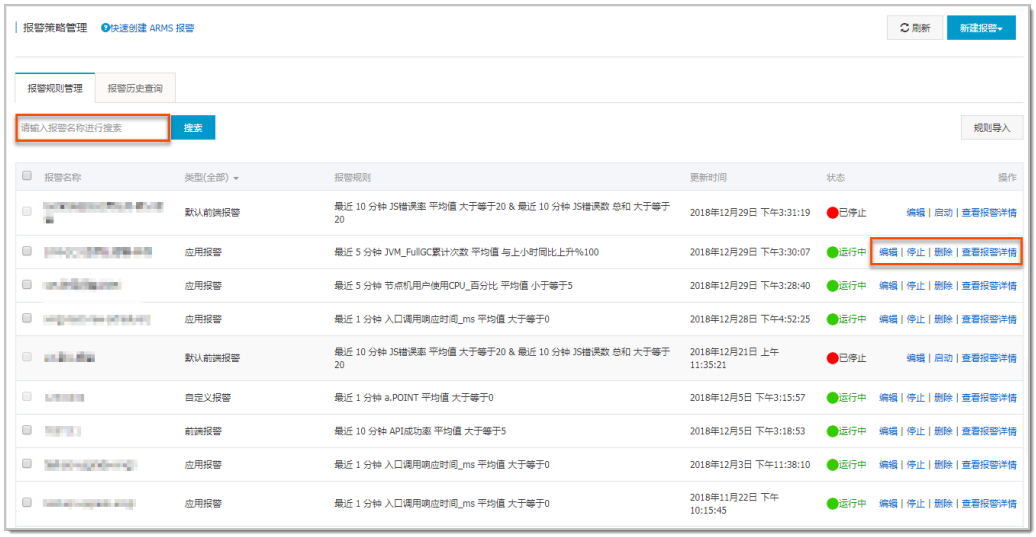

   - To edit an alarm rule, click **edit** . In the **edit alarm** dialog box, edit the alarm rule and click **save** .
   - To delete an alarm rule, click **delete** . In the **delete** dialog box, click **delete** .
   - To start a stopped alert rule, click **start** . In the **start** dialog box, click **OK** .
   - To stop the started alarm rule, click **stop** . In the **stop** dialog box, click **OK** .
   - To view the alert event history and alert notification history, click **view alert details** and view the **alert history** records on the alert history query tab.

Query the alert history

You can **click alarm history** to search for historical records about when an alarm rule is triggered and the alarm notifications sent to specific alarm contacts after the rule is triggered.

1. Log on Tracing Analysis console.

1. Log on Tracing Analysis console.
2. In the left-side navigation pane, choose **Alarm management > Alarm rules** .
3. On the **alarm rules and history** page, click the **alarm history** tab.
4. On the **alarm history** tab, select or enter the alarm **type** , **event trigger status** , and **alarm name** . Then click **search** . The line charts and bar charts on the tab show the relationship between alert data and alert trigger events, also the alert trigger details. The line chart represents the alert data and the bar chart represents the alert events.

5. Click the **alarm event history** tab at the bottom to view the history of alarm events.

> **Note** An alarm notification is sent only when the **triggering status** is **triggered** (red dot in the triggering status column).

6. Click the **alarm sending History** tab to view the records of alarm notifications (by SMS, email, or other means) that have triggered alarms.

Related information

- Create an alert contact

2.2. Create an alert contact

When the tracing alarm rule is triggered, the system sends a notification to the specified contact group. Before you create a contact group, you must create a contact first. When you create a contact, you can specify a mobile phone number and email address for the contact to receive notifications. You can also specify the webhook URL of a DingTalk chatbot to automatically receive alert notifications.

Prerequisites

Enable DingTalk chatbot alert: to add a DingTalk robot as a contact, you need to obtain the address of the DingTalk robot first.

Procedure

1. Log on Tracing Analysis console.
2. In the left-side navigation pane, choose **Alarm management > Alert contact** .
3. On the **contacts** tab, click **create contact** in the upper-right corner.
4. Edit the contact information in the **new contact** dialog box.
   - To add a contact, edit the contact **name** , **mobile number** , and **email** .

     > **Note** The mobile phone number and email address cannot be left blank at the same time. Each mobile phone number or email address must be used for only one contact. You can create a maximum of 100 contacts.

   - To add a DingTalk chatbot, enter the name and the webhook URL of the chatbot.

     > **Note** For more information about how to obtain DingTalk address of a robot, see Enable DingTalk chatbot alert.

What's next

- To search for a contact, on the **contact** tab, select the **name** , **mobile number** , or **Email** from the search drop-down box. Then, enter all or part of the contact's name, mobile number, or Email in the search box and click **search** .
- To edit a contact, click **actions** in the **edit** column corresponding to the contact. In the **update contact** dialog box, edit the information and click **OK** .
- To delete a single contact, click **actions** in the **delete** column for the right contact. In the **delete** dialog box, click **delete** .
- To delete multiple contacts, select the contacts and click **delete contacts** . In the **note** dialog box that appears, click **OK** .

Related information

- Manage alerts

2.3. Create an alert

By creating alerts, you can set alert rules for specific monitored objects. When a rule is triggered, the system will send an alert message to the specified contact group in the specified alerting mode. This reminds you to take necessary actions to solve the problem.

Prerequisites

- You have created contacts. You can only set a contact group as the notification receiver of an alert.

Context

Default behaviors of alert notifications:

- To prevent you from receiving a large number of alert notifications in a short period of time, the system only sends one message for repeated alerts within 24 hours.
- If no duplicate alerts are generated within five minutes, Application Real-Time Monitoring Service (ARMS) sends a recovery email to notify you that the alert has been cleared.
- After a recovery email is sent, the alert status is reset. If this alert arises again, it is deemed as a new one.

An alert widget is essentially a data display method of datasets. When you create an alert widget, a dataset is created to store the underlying data of the alert widget.

> **Note** New alerts take effect within 10 minutes. The alert check may have a delay of 1 to 3 minutes.

Create an alert

To create an alert for an application monitoring job on Java Virtual Machine-Garbage Collection (JVM-GC) times in corresponding-period comparison, perform the following steps:

1. Log on to the console. Click the target application in **Applications**. In the left-side navigation pane, choose **Alerts > Alert Policies**.
2. On the **Alert Policies** page, click **Create Alert** in the upper-right corner.
3. In the **Create Alert** dialog box, enter all required information and click **Save**.
   i. Enter **Alert Name**, for example, alert on JVM-GC times in corresponding-period comparison.
   ii. Select an application for **Application Site** and an application group for **Application Group**.
   iii. In the **Type** drop-down list, select the type of the monitoring metrics, for example, **JVM_Monitoring**.
   iv. Set Dimension to **Traverse**.
   v. Set alert rules.
      a. Select **Meet All of the Following Criteria**.
      b. Edit the alert rule. For example, an alert is triggered when the value of N is 5 and the average value of JVM_FullGC increases by 100% compared with that in the previous hour.

      > **Note** To add another alert rule, click + on the right of **Alert Rules**.

   vi. Set Notification Mode. For example, select Email.
   vii. Set Notification Receiver. In the **Contact Groups** box, click the name of a contact group. If the contact group appears in the **Selected Groups** box, the setting is successful.

Description of basic fields

The following table describes the basic fields of the **Create Alert** dialog box.

Field	Description	Remarks
Application Site	The monitoring job that has been created.	Select a value from the drop-down list.
Type	The type of the metric.	The types for the three alerts are different: • Application monitoring alerts: This displays application entry calls, the statistics of application call types, database metrics, JVM monitoring, host monitoring, and abnormal interface calls. • Browser monitoring alerts: This shows page metrics, interface metrics, custom metrics, and page interface metrics. • Custom monitoring alerts: This creates alerts based on existing drilled-down datasets and existing general datasets.
Dimension	The dimensions for alert metrics (datasets). You can select None,"=", or Traverse.	• When it is set to None, the alert content shows the sum of all values of this dimension. • When it is set to "=", you need to enter the specific content. • When it is set to Traverse, the alert content shows the dimension content that actually triggers the alert.
Last N Minutes	The system checks whether the data results in the last N minutes meet the trigger condition.	Range of N: 3 to 3600 minutes.
Notification Mode	Email, SMS, and DingTalk chatbot are supported.	You can select multiple modes. If you want to set DingTalk chatbot alert, see Enable DingTalk chatbot alert.
Alert Quiet Period	You can enable or disable Alert Quiet Period. By default, it is enabled.	• When it is enabled: if data remains in the triggered state, the second alert message will only be sent 24 hours after the first alert is triggered. When data recovers, you will receive a data recovery notification and the alert will be cleared. If the data triggers the alert one more time, the alert message is sent again. • When it is disabled: if the alert is continually triggered, the system sends the alert message every minute.

Field	Description	Remarks
Alert Severity	Valid values include Warn, Error and Fatal.	None
Notification Time	The time when the alert was sent. No alert notification is sent out of this time period, but alert events are recorded.	For more information about alert event history, see Alert management.
Notification Content	The custom content of the alert.	You can edit the default template. In the template, the four variables, $AlertName, $AlertFilter, $AlertTime, and $AlertContent, are preset. (Other preset variables are not supported currently.) The rest of the content can be customized.

Description of complex general fields: Period-over-period comparison

- N-minute-on-N-minute comparison: Assume that β is the data (optionally average, sum, maximum, or minimum) in the last N minutes, and α is the N-minute data starting from 2N minutes ago. The N-minute-on-N-minute comparison is the percentage increase or decrease of β as compared to α.
- N-minute-on-N-minute hourly comparison: Assume that β is the data (optionally average, sum, maximum or minimum) in the last N minutes, and α is the N-minute data from an hour ago. The N-minute-on-N-minute hourly comparison is the percentage increase or decrease of β as compared to α.
- N-minute-on-N-minute daily comparison: Assume that β is the data (optionally average, sum, maximum or minimum) in the last N minutes, and α is the N-minute data a day ago. The N-minute-on-N-minute daily comparison is the percentage increase or decrease of β as compared to α.

Description of complex general fields: Alert data revision strategy

You can select "Zero fill", "One fill", or "Zero fill null" (default). This feature is generally used to fix anomalies in data, including no data, abnormal composite metrics, or abnormal period-on-period comparison.

- Zero fill: fixes the value checked to 0.
- One fill: fixes the value checked to 1.
- Zero fill null: does not trigger the alert.

Scenarios:

- Anomaly 1: No data

  User A wants to use the alert feature to monitor the page views. When creating the alert, user A selects Browser Monitoring Alert. User A sets the alert rule as follows: N is 5 and the sum of the page views is at most 10. If the page is not hit, no data is reported and no alert is sent. To solve this problem, you can select "Zero fill" as the alert data revision policy. If you do not receive any data, it considered that zero data is received. This meets the alert rule and an alert is sent.

- Anomaly 2: abnormal composite metrics

User B wants to use the alert feature to monitor the real-time unit price of a product. When creating the alert, user B selects Custom Monitoring Alert. User B sets the dataset of variable a to the current total price, and the dataset of variable b to the current total items. User B also sets the alert rule as follows: N is 3 and the minimum value of current total price divided by current total items is at most 10. If the current total of items is 0, the value of the composite metric, current total price divided by current total items, does not exist. No alert will be sent. To solve this problem, you can select "Zero fill" as the alert data revision policy. The value of the composite metric, current total price divided by current total items, is now considered as 0. This meets the alert rule and an alert will be sent.

- Anomaly 3: Abnormal period-on-period comparisons

  User C wants to use the alert function to monitor the CPU utilization of the node machine. When user C creates the alert, C selects Application Monitoring Alert, and sets the alert rule as follows: N is 3 and the average user CPU utilization of the node machine decreases by 100% compared with the previous monitoring period. If the user's CPU fails to work in the last N minutes, the α cannot be obtained. This means the period-on-period result does not exist. No alert is sent. To solve this problem, you can select the alert data revision strategy as "One fill", and consider the period-on-period comparison result as a decrease of 100%. This meets the alert rule and an alert will be sent.

What's next

You can query and delete alert records in alert management.

2.4. Manage alerts

On the alert Policies page, you can manage all the alert rules under your account and query the history of alert events and alert posts.

Manage alert rules

1. Log on to the console. Click the target application in **Applications**. In the left-side navigation pane, choose **Alerts > Alert Policies**.
2. (Optional)On the **Alert Rules** tab, enter the alert name in the search box, then click **Search**.

   > **Note** You can enter part of an alert name in the search box to perform a fuzzy search.

3. In the **Actions** column, you may take actions on the filtered alert rules, as needed:
   - To edit an alert rule, click **Edit**. In the **Edit Alert** dialog box, edit the alert rule, and click **Save**.
   - To delete an alert rule, click **Delete**. In the **Delete** dialog box, click **Delete**.
   - To start a stopped alert rule, click **Start**, and in **OK** dialog box, click **Start**.
   - To stop a running alert rule, click **Stop**, and then click **OK** in the **Stop** dialog box.
   - To view the alert event history and alert post history, click **View Alert Detail**, and view related records on the **Alert Event History** tab.

Query alert history

You can view historical records about when and why an alert rule was triggered, and about the alert notification records sent to specified alert contacts on the **Alert History** tab.

1. On the **Alert Policies** page, click the **Alert History** tab.
2. On the **Alert History** tab, select or enter the **Alert Type**, **Trigger State** and **Alert Name**, and

then click **Search**.

The line charts and bar charts on the tab show the relationship between alert data and alert trigger events, also the alert trigger details. The line chart represents the alert data and the bar chart represents the alert events.

3. Scroll down to the bottom, view the history of alert events on the **Alert Event History** tab.

> **Note** Alert notifications are sent only when the trigger state is **triggered**. (**Trigger** column contains a red dot.)

4. Click the **Alert Post History** tab to view the records of alert notifications (such as SMS and email) that were sent for triggered alerts.

2.5. Create contacts

When an alert rule is triggered, notifications are sent to the contact group that you specified. Before you create a contact group, you must create contacts. When creating a contact, you can specify the mobile phone number and email address of the contact to receive notifications. You can also provide a DingTalk chatbot webhook URL used to automatically send alert notifications.

Prerequisites

To add a DingTalk chatbot as a contact, you must obtain its webhook URL first. For more information, see Enable DingTalk chatbot alert.

Procedure

1. Log on to the console. Click the target application in **Applications**. In the left-side navigation pane, choose **Alerts > Alert Policies**.
2. On the **Alert Policies** page, click **Create Alert** in the upper-right corner.
3. On the **Contacts** tab, click **Create Contact** in the upper-right corner.
4. In **Create Contact** dialog box, edit contact information.
   - To add a contact, enter the **Name**, **Phone Number** and **Email**.

     > **Note** The phone number and email address cannot be blank at the same time. Each phone number or email address must be used for only one contact. You can create a maximum of 100 contacts.

   - To add a DingTalk chatbot, enter the name and the webhook URL of the chatbot.

     > **Note** For more information about how to obtain the webhook URL of the DingTalk chatbot, see Enable DingTalk chatbot alert.

What to do next

- To search for contacts, on the **Contacts** tab, select **Name**, **Phone Number**, or **Email** in the drop-down list, then enter the entire or a part of the selected name, phone number or email in the search box, and click **Search**.
- To edit a contact, click **Edit** in the **Actions** column of the contact, edit the information in the **Update Contact** dialog box, then click **OK**.

- To delete a single contact, click **Delete** in the **Actions** column of the contact, then click **Delete** in the **Delete** dialog box.
- To delete multiple contacts, select the target contacts, click **Batch Delete Contacts**, then click **OK** in the **Note** dialog box.

2.6. Create contact groups

When creating an alert rule, you can specify a contact group as the alert notifications receiver. When the alert rule is triggered, Application Real-Time Monitoring Service (ARMS) sends alert notifications to the contacts in this contact group. This topic describes how to create contact groups.

Prerequisites

You have created contacts.

Procedure

1. Log on to the console. Click the target application in **Applications**. In the left-side navigation pane, choose **Alerts > Contact Management**.
2. On the **Contact Groups** tab page, click **New Contact Group** in the upper-right corner.
3. In **Create Contact Group** dialog box, enter **Group Name**, select **Contact Members**, and click **OK**.

> **Note** If there are no options in the **Contact Members** list, you need to first Create a contact.

What to do next

- To search for a contact group, go to the **Contact Groups** tab, enter all or some characters of the contact group name in the search box, then click **Search**.

> **Notice** English keywords are case-sensitive.

- To edit a contact group, click the pencil icon on the right side of the contact group, and edit the information in the **Edit Contact Group** dialog box.
- To show the contacts under a contact group, click the downward arrow on the right side of a contact group to expand the group.

> **Note** You can remove one or more contacts from an expanded contact group. To remove a contact, click **Delete** in the **Actions** column of the target contact.

- To delete a contact group, click the X icon on the right side of a contact group.

> **Notice** Before deleting a contact group, make sure that no monitoring job is running. Otherwise, alerting and other functions may be ineffective.

2.7. Enable DingTalk robot alerts

ARMS allows you to receive alert notifications from DingTalk groups. After enabling the DingTalk robot alert function, you can receive alert notifications from DingTalk groups. This topic describes how to enable the DingTalk robot alert function.

1. Obtain the address of the DingTalk robot.
   i. Run the DingTalk client on a PC, click to enter the DingTalk group to which you want to add an alert robot, and click the Group Settings icon in the upper-right corner.
   ii. In the Group Settings dialog box that appears, choose **ChatBot**.
   iii. On the ChatBot page that appears, click **+** in the **Add Robot** section, and then click **Custom**.
   iv. In the Add Robot dialog box that appears, edit the robot avatar and name, and click **Finish**.
   v. In the **Add Robot** dialog box that appears, copy the address that the system generates for the robot.
2. In the ARMS console, add the DingTalk robot as the contact. For more information about how to add a contact, see 创建联系人.
3. Create a contact group, and add the contact that you created in the previous step as the alert contact. For more information about how to create a contact group, see Create a contact group.
4. Set alert rules.
   - If you have not created an alert job yet, create an alert first, set the notification mode to **DingTalk Robot**, and set the notification receiver to the contact group that you created in Step 3.
   - If you have created an alert job, click Modify Alert Rules, set the notification mode to DingTalk Robot, and set the notification receiver to the contact group that you created in Step 3.

Now, you have enabled the DingTalk robot alert function. When an alert is triggered, you can receive the alert notification in the specified DingTalk group, for example: