Alibaba Cloud Elasticsearch

Monitoring Alarms

Issue: 20200216

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 due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall 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 the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. 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 bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequent**

ial, exemplary, incidental, 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 contact Alibaba Cloud directly if you discover any errors in this document.**

Document conventions

Style	Description	Example
	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.
	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.
	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.
	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]`

Style	Description	Example
{} or {a|b}	**This format is used for a required value, where only one item can be selected.**	`switch {active|stand}`

Contents

1 Configure Elasticsearch alerts in CloudMonitor

Elasticsearch supports instance monitoring and sends SMS messages to alert users. You can customize the thresholds for triggering alerts. This topic describes how to configure alerts for your Alibaba Cloud Elasticsearch instance in CloudMonitor. With this feature, you can use CloudMonitor to monitor your instance in real time and trigger alerts.

Monitor metrics

> **Notice:**
> **We recommend that you enable the alerting feature.**

The following monitor metrics are essential to the health of your instance. We recommend that you enable alerting for the following monitor metrics:

- **Cluster status**

  Monitors whether the status of the cluster is yellow or red.
- **Node Disk Space Usage (%)**

  We recommend that you set an alerting threshold lower than 75%. The threshold must not be higher than 80%.
- **Node Heap Memory Usage (%)**

  We recommend that you set an alerting threshold lower than 85%. The threshold must not be higher than 90%.

We recommend that you also enable alerting for the following monitor metrics:

- **Node CPU Usage (%)**

  We recommend that you set an alerting threshold lower than 95%. The threshold must not be higher than 95%.
- **load_1m**

  Set a threshold according to the number of CPU cores per node multiplied by 80 %.
- **Cluster Index Queries (QPS)**

  Set a threshold based on the actual test result.

- **Cluster Write Queries (QPS)**

  Set a threshold based on the actual test result.

Navigate to the Elasticsearch monitor page in CloudMonitor

You can use one of the following methods to navigate to the Elasticsearch monitor page in CloudMonitor.

- **From the Elasticsearch console.**

  Log on to the *Alibaba Cloud Elasticsearch console*, and click the ID of your Elasticsearch instance. On the Basic Information page of your instance, click Cluster Monitoring in the upper-right corner. You are then navigated to the Elasticsearch monitor page in the CloudMonitor console.

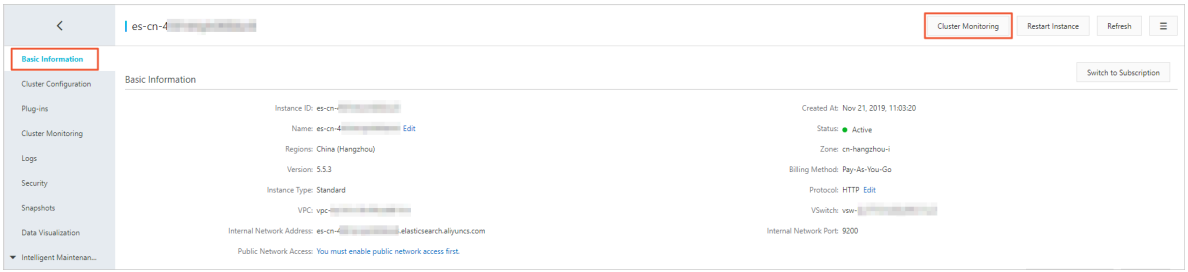

- **From the CloudMonitor console.**

  1. **Log on to the *Alibaba Cloud console*, choose CloudMonitor from the products list.**
  2. **In the left-side navigation pane, choose Cloud Service Monitoring > Elasticsearch.**
  3. **Select the region where your instance is deployed, and then click the ID of the instance to open the Elasticsearch monitor page.**

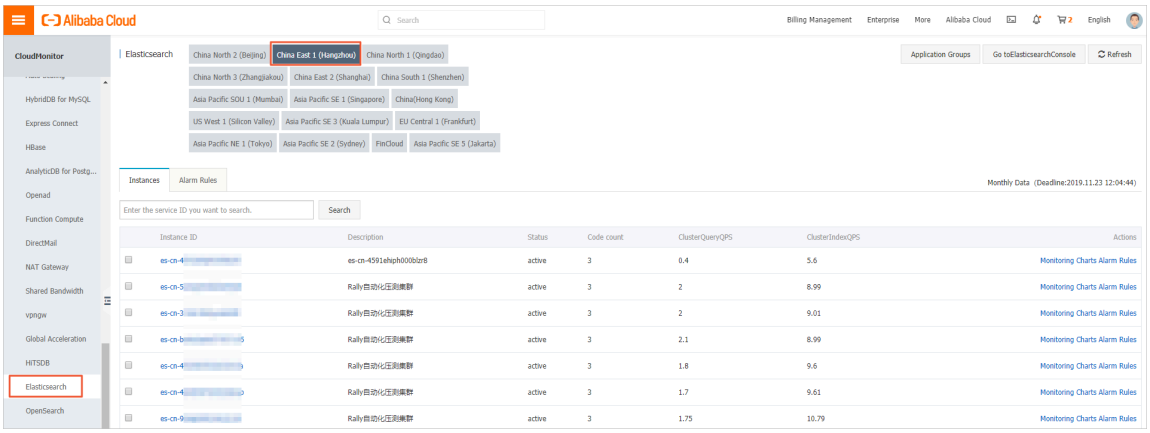

Create alert rules

1. **Go to the *Elasticsearch monitor* page of your Elasticsearch instance in the CloudMonitor console.**

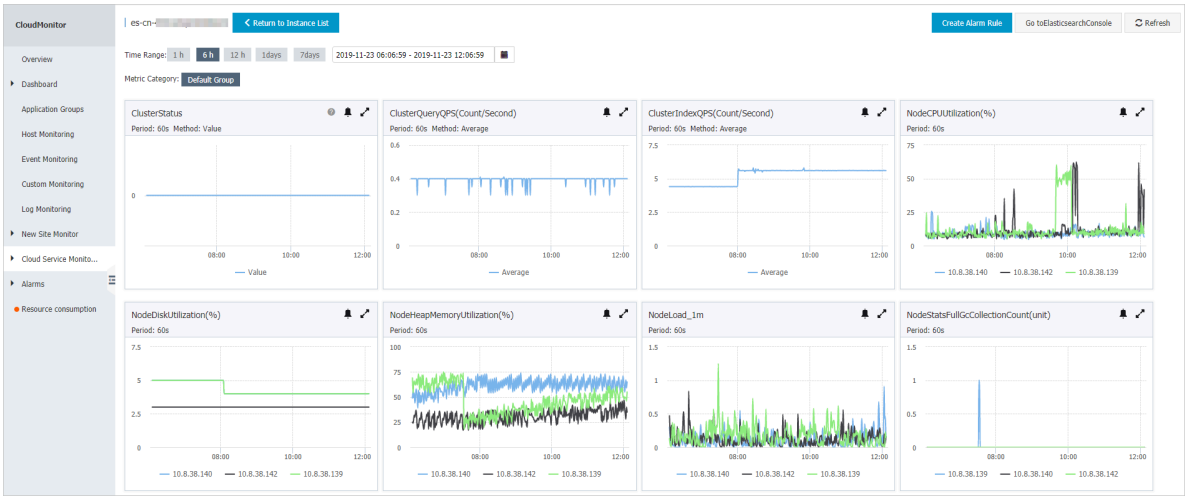

On this page, you can view the monitoring data of your instance. Currently, only the monitoring data collected in the last month is retained. You can enable alerting for monitor metrics by creating alert rules.

2. **On the Elasticsearch monitor page, click Create Alarm Rule in the upper-right corner.**

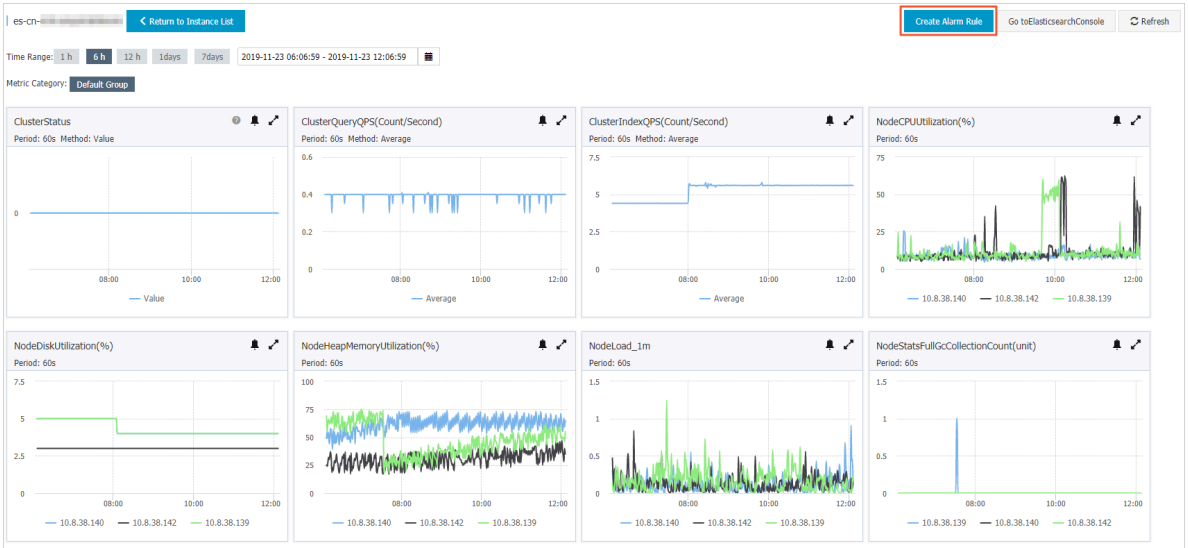

3. **On the Create Alarm Rule page, set the parameters in the Set Alarm Rules section.**

   As shown in the following figure, a rule is added for the NodeDiskUtilization, ClusterStatus, and NodeHeapMemoryUtilization monitor metrics.

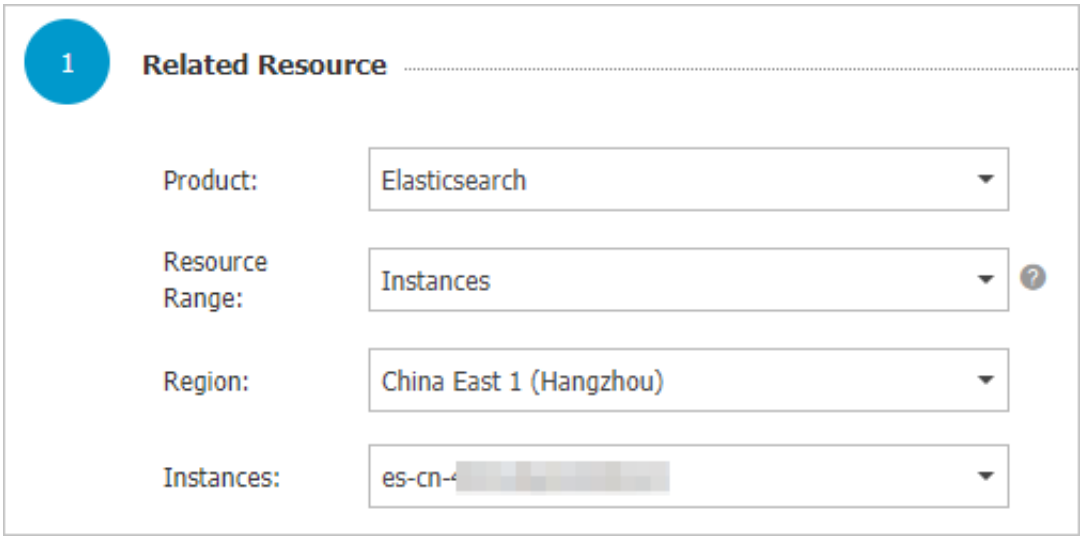

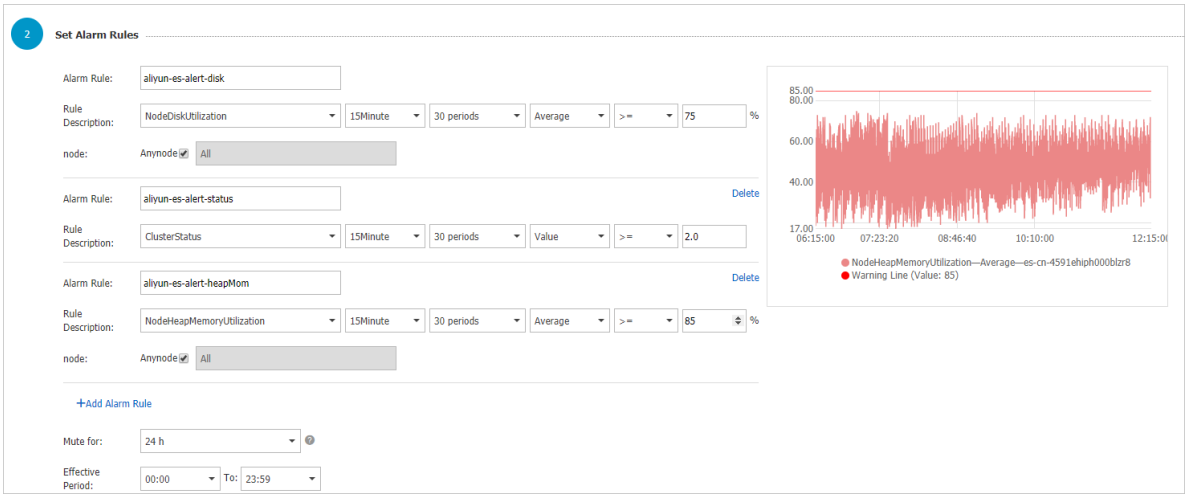

- **The values for the cluster states Green, Yellow, and Red are 0.0, 1.0, and 2.0. Reference these values and set an appropriate threshold for the ClusterStatus metric.**
- **The Mute for parameter specifies the interval for re-sending an alert when a threshold is triggered.**

> **Note:**
> **For more information about other parameters, see** *#unique_4*.

4. In the Notification Method section, select Default Contact Group.

   If you do not have a contact group, click Quickly create a contact group.

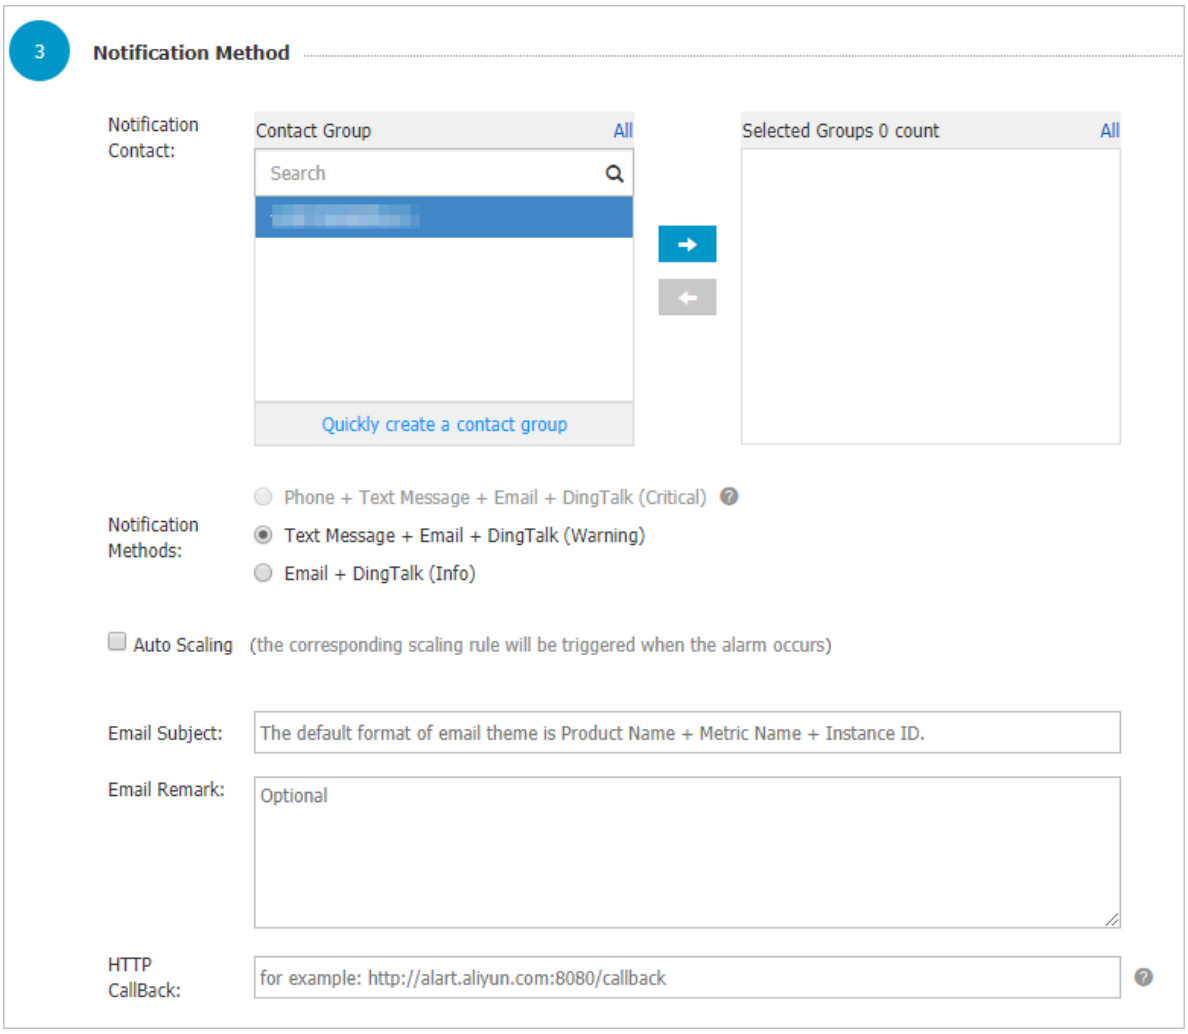

5. Click Confirm to save the rules.

   The system starts to collect the monitoring data of the Elasticsearch instance five minutes after the instance runs normally. You can then view the monitoring data on the monitor page.

2 X-Pack Watcher

This topic describes how to configure X-Pack Watcher for your Alibaba Cloud Elasticsearch instance. With X-Pack Watcher, you can use a watch to trigger specific actions. For example, you can create a watch to search the `logs` index for `errors` and then send alerts through emails or DingTalk messages. You can consider X-Pack Watcher as an alerting service based on Elasticsearch.

Notice:

- **X-Pack Watcher only supports Alibaba Cloud Elasticsearch instances deployed in one zone. Elasticsearch instances deployed across zones are not supported.**
- **By default, X-Pack Watcher is disabled. Before you use X-Pack Watcher, you must go to the Cluster Configuration > YML Configuration page of your Elasticsearch instance, and enable X-Pack Watcher. For more information, see *FAQ*.**

Features

X-Pack Watcher allows you to create watches. A watch consists of the `Trigger`, `Input`, `Condition`, and `Actions`.

- `Trigger`

  Determines when the watch is executed. You must configure a trigger for each watch. X-Pack Watcher allows you to create various types of triggers. For more information, see *Schedule Trigger*.

- `Input`

  Loads data into the payload of a watch. Inputs are used as filters to match the specified type of index data. For more information, see *Inputs*.

- `Condition`

  Controls whether the actions of a watch are executed.

- `Actions`

  Determines the actions to be executed when the specified conditions are met.

Procedure

You cannot directly connect to X-Pack Watcher from the Internet. You must connect to the internal network address of your Elasticsearch instance through a VPC network. To use X-Pack Watcher, you must also purchase an Alibaba Cloud Elastic Compute Service (ECS) instance that has access to both the Internet and Alibaba Cloud Elasticsearch instance. The ECS instance runs as a proxy to execute actions.

The following example shows how to use a webhook action to connect the DingTalk Chatbot to X-Pack Watcher.

1. *Purchase an Alibaba Cloud ECS instance*.

   The purchased ECS instance must meet the following requirements:

   

   Notice:

   - **The Alibaba Cloud ECS instance and Elasticsearch instance must be connected to the same VPC network.**
   - **The ECS instance must have access to the Internet.**

2. **Configure a security group rule for the ECS instance.**
   a. **On the Instances page of the Alibaba Cloud ECS console, choose More > Network and Security Group > Configure Security Group.**
   b. **Find your security group in the Security Groups list, and click Add Rules in the Actions column.**
   c. **On the Security Group Rules page, click Add Security Group Rule.**
   d. **Set the parameters, and click OK to add the rule.**

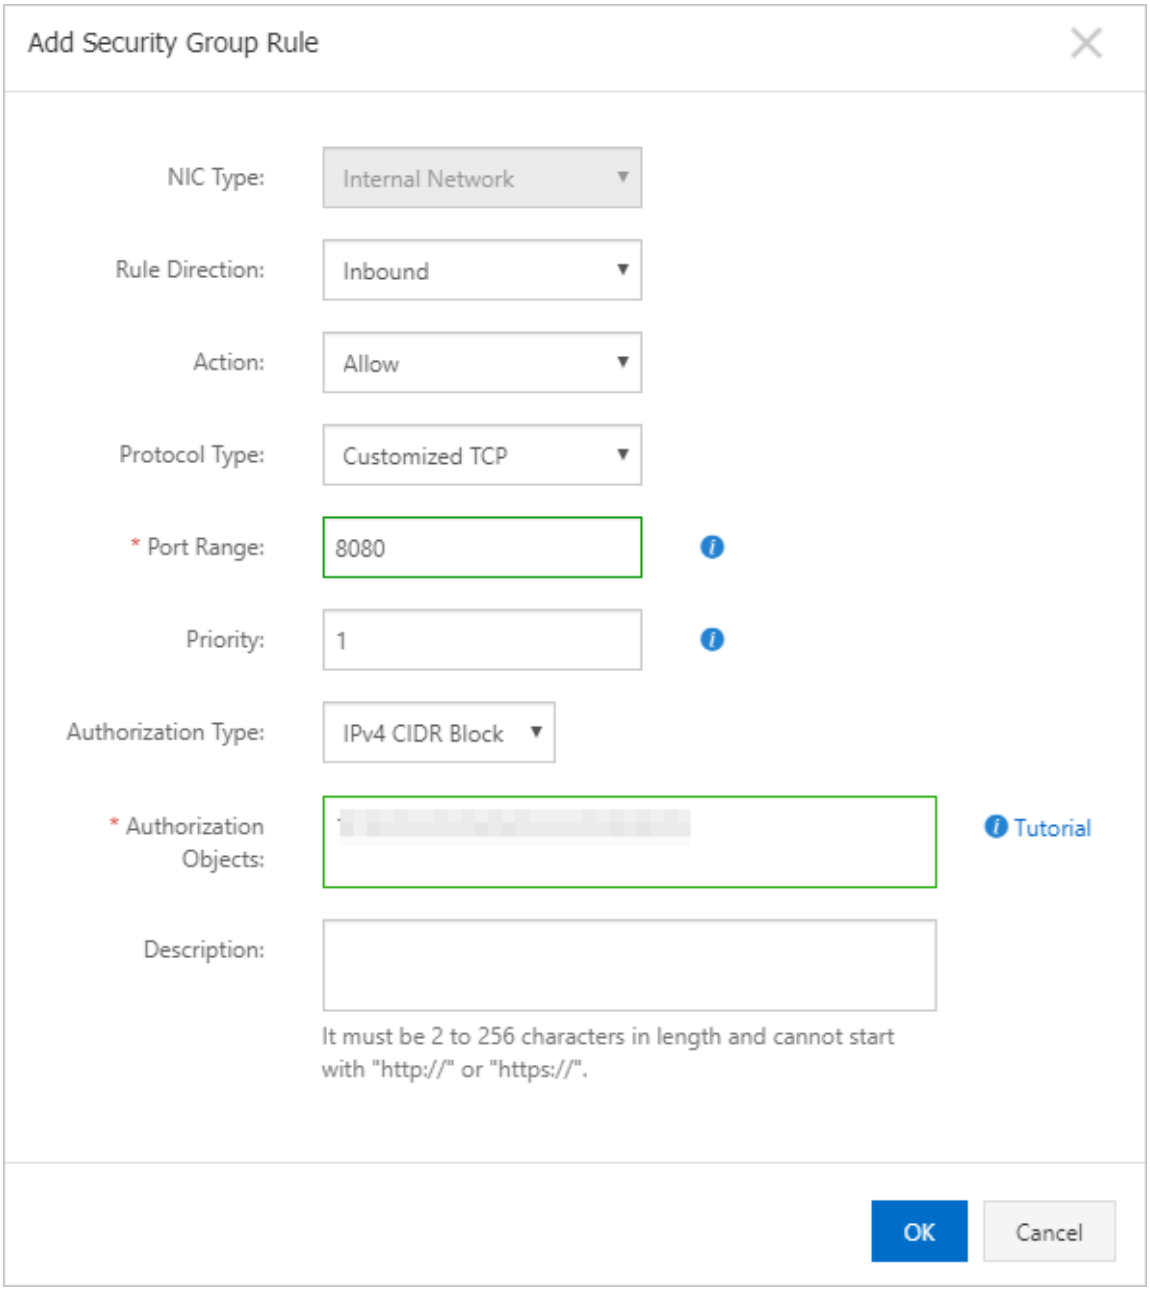

Parameter	Description
Rule Direction	Select Inbound.

Parameter	Description
Action	Select Allow.
Protocol Type	Select Customized TCP.
Priority	Use the default setting.
Port Range	Set the port to your frequently used port. This port is required for NGINX configuration. In this example, port 8080 is specified.
Authorization Type	Select IPv4 CIDR Block.
Authorization Objects	Add the IP addresses of all the nodes in your Elasticsearch instance. **Note:** Use the following method to query the IP addresses of the nodes. Log on to the Kibana console of the Alibaba Cloud Elasticsearch instance, choose Monitoring in the left-side navigation pane, and then click Nodes.

3. Configure an NGINX proxy.

   a. Modify the NGINX configuration file. Reference the following configuration and then replace the `server` configuration in Install and configure NGINX with this configuration.

```
server
  {
    listen 8080;#Listening port
    server_name localhost;#Domain name
    index index.html index.htm index.php;
    root /usr/local/webserver/nginx/html;#Website directory
      location ~ . *\.(php|php5)? $
    {
      #fastcgi_pass unix:/tmp/php-cgi.sock;
      fastcgi_pass 127.0.0.1:9000;
      fastcgi_index index.php;
      include fastcgi.conf;
    }
    location ~ . *\.(gif|jpg|jpeg|png|bmp|swf|ico)$
    {
      expires 30d;
  # access_log off;
    }
    location / {
      proxy_pass <Webhook address of the DingTalk Chatbot>;
    }
    location ~ . *\.(js|css)? $
    {
      expires 15d;
```

```
    # access_log off;
     }
     access_log off;
   }
}
```

<Webhook address of the DingTalk Chatbot>: **replace it with the webhook address of the DingTalk Chatbot that is used to receive alerts.**

Note:

You can use the following method to query the webhook address of the DingTalk Chatbot.

Create an alert contact group in DingTalk. Click the DingTalk group, click the More icon in the upper-right corner, click ChatBot, and select Custom to add a webhook ChatBot. You can then view the webhook address of the ChatBot.

b. **After you complete the configuration, reload the NGINX configuration file and then restart NGINX.**

```
/usr/local/webserver/nginx/sbin/nginx -s reload            #
Reload the NGINX configuration file
/usr/local/webserver/nginx/sbin/nginx -s reopen            #
Restart NGINX
```

4. **Create a watch.**

   Log on to the Kibana console of your Alibaba Cloud Elasticsearch instance. In the left-side navigation pane, click Dev Tools. On the Console tab, call the corresponding API operation to create a watch.

   The following example shows how to create a watch named `log_error_watch` to search the `logs` index for `errors` at intervals of `10s`. If more than `0` errors are found, an alert is triggered.

```
PUT _xpack/watcher/watch/log_error_watch
{
  "trigger": {
    "schedule": {
      "interval": "10s"
    }
  },
  "inputs": [
    "search": {
      "request": {
        "indices": ["logs"],
        "body": {
          "query": {
            "match": {
              "message": "error"
            }
```

```
          }
        }
      }
    }
  },
  "condition": {
    "compare": {
      "ctx.payload.hits.total": {
        "gt": 0
      }
    }
  },
  "actions" : {
  "test_issue" : {
    "webhook" : {
      "method" : "POST",
      "url" : "http://Private IP address of your ECS instance:8080",
      "body" : "{\"msgtype\": \"text\", \"text\": { \"content\": \"
An error has been found. Handle the issue immediately.\"}}"
    }
  }
}
}
```

Note:

The `url` specified in `actions` must contain the private IP address of the purchased Alibaba Cloud ECS instance that is deployed in the same region and VPC network as your Elasticsearch instance. You must also make sure that you have followed the preceding procedure to create a security group rule for the ECS instance. Otherwise, you cannot connect to X-Pack Watcher.

If you no longer need this watch, run the following command to delete the watch.

```
DELETE _xpack/watcher/watch/log_error_watch
```

FAQ

Q: How do I resolve the error `No handler found for uri [/_xpack/watcher/watch/log_error_watch_2] and method [PUT]` that occurs when I configure the alerting settings?

A: This error indicates that X-Pack Watcher is disabled for your Alibaba Cloud Elasticsearch instance. Follow these steps to enable X-Pack Watcher for your instance.

1. *Log on to the Alibaba Cloud Elasticsearch console*, **and choose Instance ID > Cluster Configuration.**
2. **On the Cluster Configuration page, click Modify Configuration on the right side of YML Configuration.**

3. On the YML Configuration page, select Enable for Watcher.

> **Notice:**
> After you enable X-Pack Watcher, the Elasticsearch instance is restarted. Make sure that your workloads are not adversely affected before you confirm the operation.

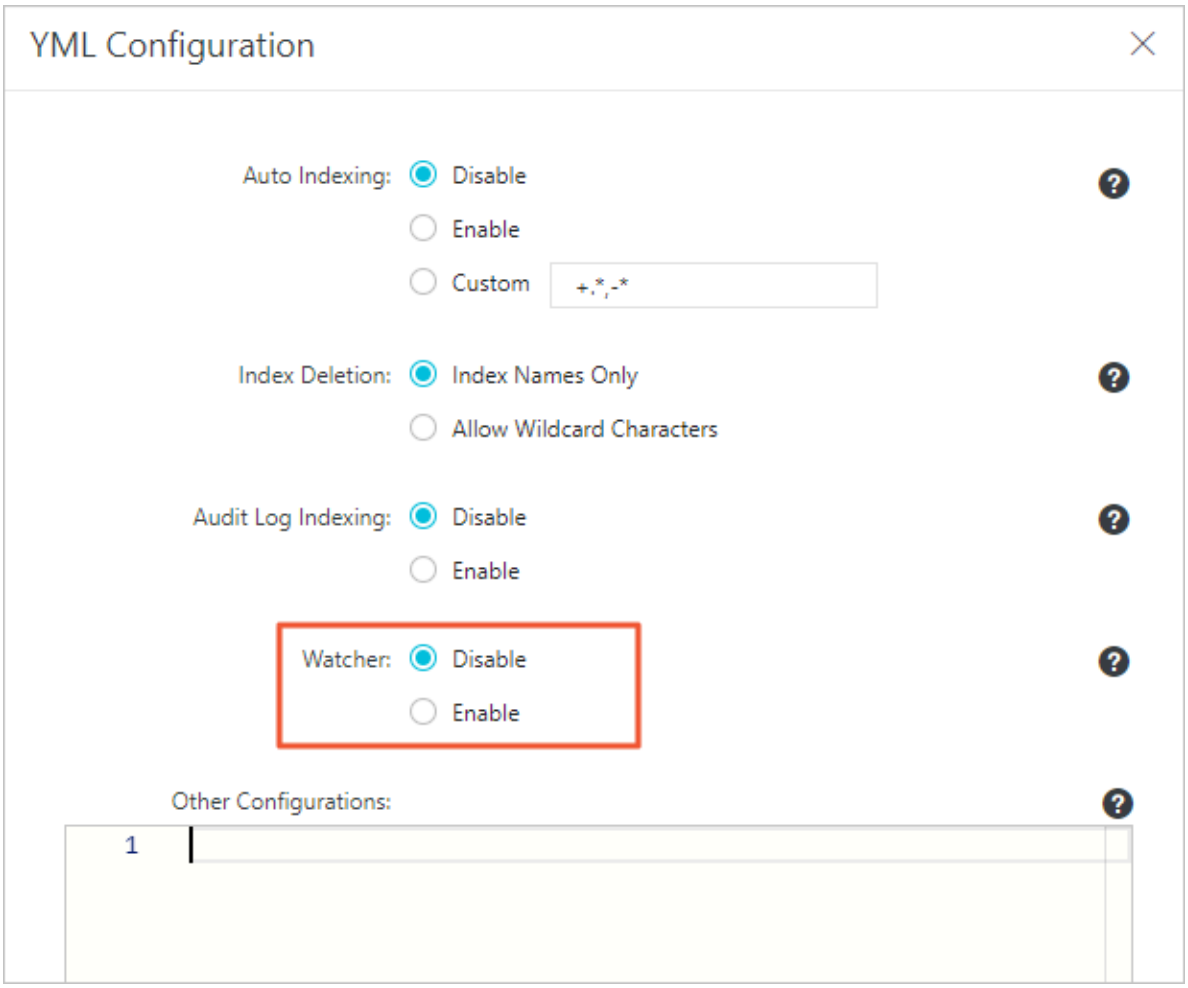

4. Select the This operation will restart the instance. Continue? check box, and then click OK.

   During the restart process, you can check the progress on the *Tasks* page. After the Elasticsearch instance is restarted, X-Pack Watcher is enabled.

3 Monitoring settings

This topic describes how to configure monitoring settings for your Alibaba Cloud Elasticsearch instance. You can view log data on the Monitoring page of the Kibana console. You can also specify a retention period for the indexes that store the log data. This ensures that your storage space is not exhausted by the log data and the workloads on your instance are not affected by storage space shortage.

Log data collection

By default, the X-Pack monitoring component collects the monitoring data at intervals of 10 seconds, and saves the data to the indexes starting with `.monitoring-*` on your Elasticsearch instance.

Currently, the `.monitoring-es-6-*` and `.monitoring-kibana-6-*` indexes are used to store the monitoring data. The instance rolls over to a new index each day. The name of a `.monitoring-es-6-` index ends with date of the day when the monitoring data is saved.

The `.monitoring-es-6-*` index stores information about the cluster status, cluster statistics, node statistics, and index statistics, which consumes a large amount of disk space.

View index information

Log on to the Kibana console, click Monitoring > Indices to enter the Indices page, and open the System indices to show the disk space consumed by the indexes.

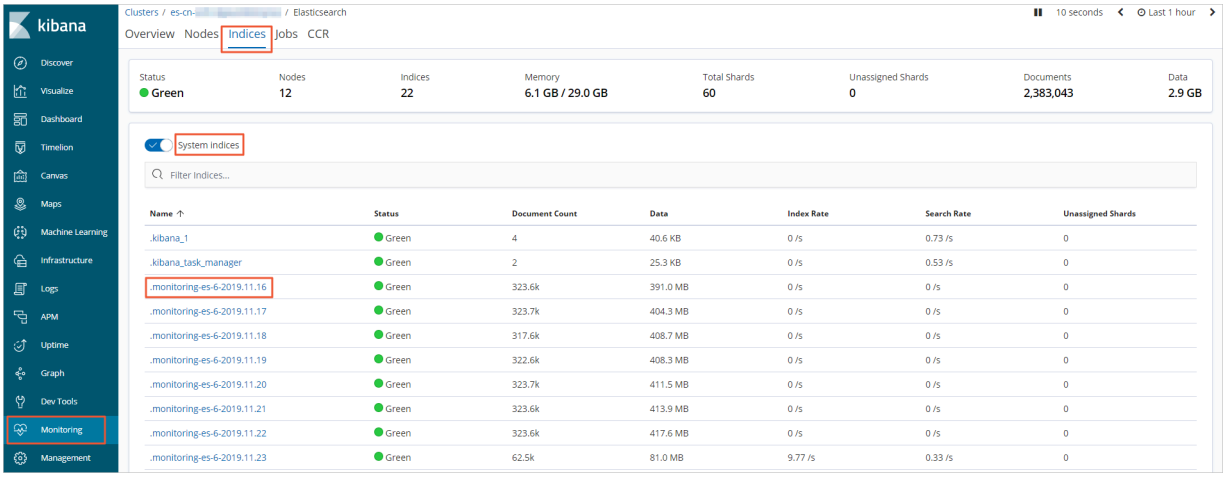

Set the index retention period

By default, the system retains the indexes created in the last seven days. Indexes that store the monitoring data, such as `.monitoring-es-6-*`, consume the storage space of your Elasticsearch instance. The size of each index depends on the number of indexes (including system indexes) on your cluster and the number of nodes in the cluster. You can use one of the following methods or use both methods if needed to prevent the indexes that store the monitoring data from consuming too much disk space of your instance.

- **Call the following API operation to specify the index retention period.**

```
PUT _cluster/settings
{"persistent": {"xpack.monitoring.history.duration":"1d"}}
```

Note:

Specify the index retention period as needed. The minimum retention period is one day.

- **Specify the indexes to be collected.**

  Call the API operation to control how indexes are collected by specifying an inclusion or exclusion list. This helps you reduce the size of the `.monitoring-es-6-*` indexes. The following example shows how to create an exclusion list for index collection.

```
PUT _cluster/settings
{"persistent": {"xpack.monitoring.collection.indices": "*,-. *"}}
```

Note:

The Monitoring page in the Kibana console displays the indexes and relevant monitoring data. The excluded indexes are not collected or displayed on this page. However, these indexes are listed in the index list retrieved by calling the `GET _cat/indices` operation. The status of the indexes, open or close, is also shown.