ALIBABA CLOUD

Alibaba Cloud

Elasticsearch

Beats Data Shippers

Document Version: 20201030

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.Install a shipper ---- 05
2.Modify shipper configuration ---- 08
3.Prepare the YML configuration files for a shipper ---- 10

1.Install a shipper

Install a shipper

This topic uses Filebeat as an example to describe how to install and manage a Beats data shipper. A shipper collects the data of your Elastic Compute Service (ECS) instance, which includes log files, network data, and server metrics. Then, the shipper sends the data to your Alibaba Cloud Elasticsearch or Logstash cluster for further operations, such as monitoring and analytics.

install shipper

Prerequisites

You have completed the following operations:

- An Alibaba Cloud Elasticsearch cluster is created.

  For more information, see Create an Elasticsearch cluster.

- The Auto Indexing feature is enabled for the Alibaba Cloud Elasticsearch cluster.

  For security purposes, Alibaba Cloud Elasticsearch disables **Auto Indexing** by default. However, Beats depends on this feature. If you select **Elasticsearch** for **Output** when you create a shipper, you must enable the **Auto Indexing** feature. For more information, see Enable auto indexing.

- An Alibaba Cloud ECS instance is created, and it is in the same Virtual Private Cloud (VPC) as the Alibaba Cloud Elasticsearch or Logstash cluster.

  For more information, see Create an instance by using the provided wizard.

  > **Notice**
  > - The default installation directory of Beats is */opt/aliyunbeats/*. After you install Beats, the *conf*, *logs*, and *data* directories are generated on the ECS instance. The conf directory contains the configuration file. The logs directory contains the Beats log file. The data directory contains the Beats data file. We recommend that you do not delete or modify the content of these files. Otherwise, errors may occur or data may be altered. If an error occurs, you can view Beats logs in the *logs* directory to locate the error.
  > - Beats now supports only Aliyun Linux, Red Hat, and CentOS.

- Cloud Assistant and Docker are installed on the ECS instance.

  For more information, see Install the Cloud Assistant client and Deploy and use Docker.

Procedure

1. Log on to the Alibaba Cloud Elasticsearch console. In the left-side navigation pane, click **Beats Data Shippers**.
2. In the **Create Shipper** section, click **Filebeat**.

   

3. In the **Configure Shipper** step, specify parameters as required.

   

   ### Parameters used to create a Filebeat shipper

Parameter	Description
Shipper Name	Enter a name for the shipper.
Version	Select **6.8.5**, which is the only version supported by Filebeat.
Output	Select a destination for the data collected by Filebeat. The system provides Elasticsearch and Logstash for you to select. The access protocol must be consistent with that of the Alibaba Cloud Elasticsearch cluster.
Username/Password	If you select **Elasticsearch** for **Output**, enter the username and password for Filebeat to write data to the Alibaba Cloud Elasticsearch cluster.
Enable Kibana Monitoring	Determine whether to monitor the metrics of Filebeat. If you select **Elasticsearch** for **Output**, the Kibana monitor uses the same Alibaba Cloud Elasticsearch cluster as **Output**. If you select **Logstash** for **Output**, you must separately configure a monitor in the configuration file.
Enable Kibana Dashboard	Determine whether to enable the default Kibana dashboard. Alibaba Cloud Kibana is configured in a VPC. You must enable the Kibana private network access feature on the Kibana configuration page. For more information, see Configure a whitelist for access to the Kibana console over the Internet or an internal network.
Filebeat Log File Path	This parameter is specific to Filebeat. Alibaba Cloud deploys Beats with Docker. You must map the directory from which logs are collected to Docker. We recommend that you enter a directory that is consistent with `input.path` in *filebeat.yml*.
Shipper YML Configuration	Prepare configuration files for the shipper. You can modify the YML configuration files based on your business requirements. For more information, see Prepare the YML configuration files for a shipper.

Notice If you already specify **Output**, you do not need to specify it again in **Shipper YML Configuration**. If you specify it again, the system prompts a shipper installation error.

4. Click **Next**.
5. In the **Install Shipper** step, select the target ECS instance.

Notice The instance list displays all ECS instances that are in the same VPC as the Alibaba Cloud Elasticsearch or Logstash cluster that you select for **Output** and have Cloud Assistant and Docker installed.

6. Click **Enable**.
7. In the **Enable Shipper** dialog box, click **Back to Beats Shippers** to view the created shipper. After the **Status** of the shipper changes to **Enabled**, the shipper is created. The two numbers following **Enabled** indicate the number of ECS instances where the shipper is installed and the

number of total target ECS instances. If the shipper is installed on all the ECS instances, the two numbers are equal.

8. View running ECS instances. After the shipper is created, you can view running ECS instances to check whether the shipper installation succeeds and handle exceptions as prompted.
   i. In the **Manage Shippers** section, find the target shipper and click **View Instances** in the **Actions** column.
   ii. In the **View Instances** pane, check whether the shipper installation on the ECS instance succeeds. The Installed Shippers column provides the values **Heartbeat Normal**, **Heartbeat Abnormal**, and **Installation Failed** to indicate whether the shipper installation on an ECS instance succeeds. If the value of Installed Shippers is **Heartbeat Abnormal** or **Installation Failed**, you can remove the instance or retry the installation on the instance. If the retry also fails, check whether the prerequisites are met.

   

   iii. Click **Add Instance** to add ECS instances where you want to install a shipper with the same configuration and type as the created shipper.
9. (Optional)View monitoring and dashboard information. If you select **Enable Kibana Monitoring** or **Enable Kibana Dashboard** when you create the shipper, you can view the monitoring and dashboard information in the Kibana console after Beats is started.
   i. In the **Manage Shippers** section, find the target shipper and choose **More > Dashboard** in the **Actions** column.
   ii. On the logon page of the Kibana console, enter the username and password, and click Log in.
   iii. In the left-side navigation pane, click **Dashboard** to view the dashboard tables and graphs.

   

   iv. In the left-side navigation pane, click **Monitoring** to view monitoring information.

   

2.Modify shipper configuration

Modify shipper configuration

After you install a shipper, you can modify its configuration on the Beats Data Shippers page of the Elasticsearch console.

configure shipper

Prerequisites

A shipper is installed. For more information, see Install a shipper.

Procedure

1. Log on to the Alibaba Cloud Elasticsearch console. In the left-side navigation pane, click **Beats Data Shippers**.
2. In the Manage Shippers section, find the target shipper and click **Configure** in the Actions column.
3. Modify the configuration of the shipper as required, and click **Modify**.

Parameter	Description
Shipper Name	Enter a name for the shipper. The name must be 1 to 30 characters in length and can contain letters, digits, underscores (_), and hyphens (-). The name must start with a letter.
Version	Set Version to **6.8.5**, which is the only version supported by Filebeat.
Output	Select a destination for the data collected by Filebeat. The destination is the Elasticsearch cluster you created. The protocol must be the same as that of the selected Elasticsearch cluster.
Username/Password	If you select **Elasticsearch** for **Output**, enter the username and password for Filebeat to write data to the Alibaba Cloud Elasticsearch cluster.
Enable Kibana Monitoring	Used to monitor the metrics of Filebeat. If you select **Elasticsearch** for **Output**, the Kibana monitor uses the same Alibaba Cloud Elasticsearch cluster as **Output**. If you select **Logstash** for **Output**, you must configure a monitor in the configuration file.
Enable Kibana Dashboard	Used to enable the default Kibana dashboard. Alibaba Cloud Kibana is configured in a VPC. You must enable private network access for Kibana on the Kibana configuration page. For more information, see Configure a whitelist for access to the Kibana console over the Internet or an internal network.
Filebeat Log File Path	This parameter is specific to Filebeat. Alibaba Cloud deploys Beats with Docker. You must map the directory from which logs are collected to Docker. We recommend that you specify a directory that is consistent with `input.path` in *filebeat.yml*.

Parameter	Description
Shipper YML Configuration	Prepare configuration files for the shipper. You can modify the YML configuration files based on your business requirements. For more information, see Prepare the YML configuration files for a shipper.

Notice If you already specify **Output**, you do not need to specify it again in **Shipper YML Configuration**. If you specify it again, the system prompts a shipper installation error.

After you save the changes, the shipper state changes to Enabling. After the state changes to Enabled, the shipper configuration modification is complete.

3.Prepare the YML configuration files for a shipper

Prepare the YML configuration files for a shipper

You can modify and enable the YML configuration of a shipper to complete specific data collection tasks. This topic discusses specific parameters in the YML configuration files and describes how to modify the YML configuration.

Prerequisites

An Alibaba Cloud Elasticsearch cluster is created, and the **Auto Indexing** feature is enabled for the cluster. For more information about how to create an Elasticsearch cluster, see Create an Elasticsearch cluster.

For security purposes, Alibaba Cloud Elasticsearch disables the **Auto Indexing** feature by default. However, Beats depends on this feature. If you select **Elasticsearch** for **Output** when you install a shipper, you must enable the **Auto Indexing** feature. For more information, see Enable auto indexing.

Note Open source Beats provides multiple modules, but Alibaba Cloud Beats does not provide separate configuration for these modules. If you want to use them, you must configure them in the configuration files of different shippers. For example, if you want to enable the system module in a Metricbeat shipper, add the following script to metricbeat.yml:

```
metricbeat.modules:
- module: system
metricsets: ["diskio","network"]
diskio.include_devices: []
period: 1s
```

Filebeat configuration

You can specify `filebeat.inputs` in **filebeat.yml** to determine how to search for or handle input data sources. The following figure shows an example of a simple input configuration.

```
filebeat.inputs:
- type: log
 enabled: true
 paths:
  - /opt/test/logs/t1.log
  - /opt/test/logs/t2/*
 fields:
  alilogtype: usercenter_serverlog
```

> **Notice**
>
> - If you specify **Output** when you install a shipper, you do not need to specify it again in **Shipper YML Configuration**. Otherwise, the system prompts a shipper installation error.
> - Each input data source starts with a hyphen (-). You can use multiple hyphens to specify multiple input data sources.

Parameter	Description
type	The input type. Examples of valid values: stdin , redis , tcp , and syslog . Default value: log .
paths	The paths of the logs you want to monitor. You can specify a file or a directory to map to Docker.
enabled	Specifies whether the configuration takes effect. The value true indicates that the configuration takes effect. The value false indicates that the configuration does not take effect.
fields	Optional. Below this parameter, you can indent with two spaces to add fields. For example, enter alilogtype: usercenter_serverlog to add this field to each output log to identify the type of the log source. If logs are shipped to Logstash, they can be classified and processed based on this field.

For more information, see Log input in the open source Filebeat documentation.

Metricbeat configuration

Metricbeat delivers system and service statistics in a lightweight manner. You can specify metricbeat.modules in **metricbeat.yml** to configure a module .

```
metricbeat.modules:
- module: system
  metricsets: ["diskio","network"]
  enabled: true
  hosts: ["http://XX.XX.XX.XX/"]
  period: 10s
  fields:
    dc: west
  tags: ["tag"]
```

> **Notice** If you specify **Output** when you install a shipper, you do not need to specify it again in **Shipper YML Configuration**. Otherwise, the system prompts a shipper installation error.

Parameter	Description
module	The name of the module you want to run. For more information about supported modules, see Modules.
metricsets	Specifies the metricsets you want to execute. For more information about metricsets, see Modules.
enabled	Specifies whether the configuration takes effect. The value true indicates that the configuration takes effect. The value false indicates that the configuration does not take effect.
period	Specifies how often the metricsets are executed. If the system is inaccessible, Metricbeat returns an error for each period.
hosts	Optional. This parameter specifies the hosts from which you want to obtain information.
fields	Optional. This parameter specifies the fields that are sent with the metricset event.
tags	Optional. This parameter specifies the tags that are sent with the metricset event.

For more information, see open source Metricbeat documentation.

Heartbeat configuration

Heartbeat can be installed on a remote server in a lightweight manner. You can use Heartbeat to periodically check the status of your services and determine whether they are available. Unlike Metricbeat, Heartbeat checks whether your services are available but Metricbeat checks whether your services are running.

You can specify heartbeat.monitors in **heartbeat.yml** to specify the services you want to monitor.

Note You can configure only the services that you want to monitor for Heartbeat. To ensure the availability of Heartbeat, we recommend that you deploy at least two Elastic Compute Service (ECS) instances.

```
heartbeat.monitors:
- type: http
  name: ecs_monitor
  enabled: true
  urls: ["http://localhost:9200"]
  schedule: '@every 5s'
  fields:
    dc: west
```

Notice If you specify **Output** when you install a shipper, you do not need to specify it again in **Shipper YML Configuration**. Otherwise, the system prompts a shipper installation error.

Parameter	Description
type	The monitor type. Valid values: icmp, tcp, and http.
name	The monitor name. This value appears in Exported fields of the monitor field and is used as the job name. The type field is used as the job type.
enabled	Specifies whether the configuration takes effect. The value true indicates that the configuration takes effect. The value false indicates that the configuration does not take effect.
urls	Optional. This parameter specifies the servers to which you want to connect.
schedule	The task schedule. If you set the value to @every 5s, the system runs the task every five seconds from the time Heartbeat is started. If you set the value to */5 * * * * *, the system runs the task every five seconds.
fields	Optional. You can add the fields to output as additional information.

For more information, see open source Heartbeat documentation.

Auditbeat configuration

Auditbeat is a lightweight shipper that collects audit logs from the Linux audit framework and monitors file integrity. Auditbeat combines relevant messages into an event to generate structured data for analytics. It can also be seamlessly integrated with Logstash, Elasticsearch, and Kibana.

Notice Auditbeat is based on the Linux audit framework and requires an OS kernel version of 3.14 or later. The Auditd service must be in the stop state. You can run the service auditd status command to query its status.

You can specify auditbeat.modules in **auditbeat.yml** to configure the Auditbeat shipper. **auditbeat.yml** consists of two parts: module and output. If you want to enable a module, you must add specific parameters to **auditbeat.yml**. The following configurations use the auditd and file_integrity modules as examples:

```
auditbeat.modules:
- module: auditd
 audit_rules: |
  -w /etc/passwd -p wa -k identity
  -a always,exit -F arch=b32 -S open,creat,truncate,ftruncate,openat,open_by_handle_at -F exit=-EPERM -k
access
- module: file_integrity
 paths:
 - /bin
 - /usr/bin
 - /sbin
 - /usr/sbin
 - /etc
```

Notice If you specify **Output** when you install a shipper, you do not need to specify it again in **Shipper YML Configuration**. Otherwise, the system prompts a shipper installation error.

For more information about **auditbeat.yml** configuration, see Step 2: Configure Auditbeat in the open source Auditbeat documentation. For more information about `module` configuration, see Modules.