# Elasticsearch Service

# ES Serverless Guide

# Product Documentation

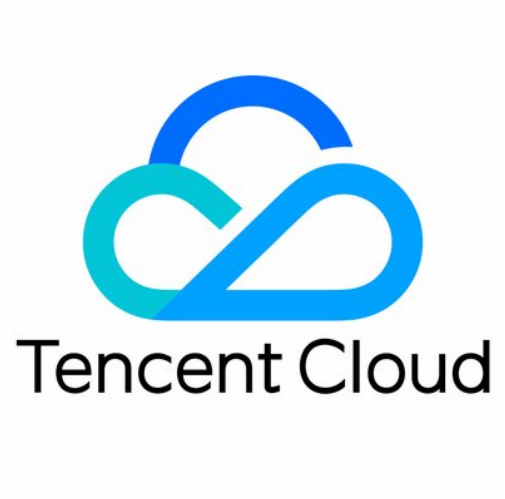

# Contents

# ES Serverless Guide
# Service Overview

Last updated：2024-08-20 16:57:46

## Industry Challenges

When using open-source Elasticsearch for log analysis, users often need to estimate cluster configuration based on write traffic, peak write, and storage days, including CPU, memory, and disk size, to ensure smooth business operation. However, as per extensive online operational experience, this method has the following problems:
Elastic capability is difficult to adapt to business development. In scenarios such as big promotions and holidays, log data presents obvious peak and trough effects, high write throughput, and high availability requirements, and it is impossible to predict sudden read-write traffic and scale out a cluster in advance, making it difficult to ensure the stability of the Elasticsearch cluster.
Resource costs are high. Insufficient resources affect traffic write during peak periods, and planning cluster capacity based on peak traffic results in resource redundancy and waste during off-peak periods, leading to high costs.
Operations and management costs are high. Enterprises need to plan and configure clusters and indices, and build monitoring and alert platforms. Moreover, enterprises have a strong demand for optimizing Ops and management costs with the focus on cost reduction and efficiency improvement, aiming to further reduce these expenses.

## Overview

Elasticsearch Serverless service is a one-stop, fully managed Elasticsearch service built by Tencent Cloud based on its proprietary cloud-native Serverless  technology architecture. It offers automatic scalability and a completely maintenance-free product capability, effectively addressing the problems of high resource costs caused by peaks and troughs in log analysis, metric monitoring, and other business scenarios. Meanwhile, it is fully compatible with the ELK ecosystem, featuring end-to-end data access, data management, and data visualization product features, providing an out-of-the-box product experience.
At the Enterprise Cloud Adoption and Cloud Computing Integration Industry Conference held on March 29, 2023, Tencent Cloud Elasticsearch Serverless  service was awarded the "2022 Trusted Computing Power Service · Leadership Plan" Excellent Case Award.

## Benefits and Features

Auto Scaling: It features automatic index-level AS to smoothly handle unexpected traffic growth, reducing high Ops and management costs during peaks and troughs in scenarios like log analysis and observability while ensuring business continuity.
Completely Ops-free: Built-in automatic sharding optimization, intelligent lifecycle management, and failures self-healing capabilities allow users to create and use indices as needed without worrying about underlying resource configuration, cluster scaling, and index settings, ensuring a completely Ops-free experience.
Cost-saving: Self-developed, low-cost, high-performance, and high-availability storage-compute separation architecture charges based on actual access and storage volumes, enabling pay-as-you-go in the scenario of dynamic matching of service load and resources. This reduces redundant cost expenditures due to idle resources, significantly lowering costs.
Flexible and easy to use: It provides end-to-end one-stop product capability featuring data access, data management, and data analysis and exploration, significantly lowering the barrier to business cloud adoption. Users can achieve minute-level business implementation.
Open integration: It is fully compatible with the ELK ecosystem and retains users' original usage habits, ensuring seamless migration and facilitating rapid cloud adoption. Meanwhile, it connects cloud data sources (such as CVM and TKE) to lower the data access threshold, achieving minute-level business implementation.
Stable and reliable: Cluster configuration and read-write performance are optimized by the backend, reducing fault issues caused by improper use, enhancing stability, and safeguarding business operations.

## Contact Us

Scan the code to join Tencent Cloud Big Data Elasticsearch Serverless community group, with occasional activities and exquisite gifts.

# Basic Concepts

Last updated：2024-12-04 15:51:12

This document introduces the basic concepts related to the project space and index in the ES Serverless service.

## Project Space

A project space is the basic resource unit in the ES Serverless service. You can create indexes for the same business within a single project space to facilitate index management. To read and write data, use the access address, username, and password for the project space to access the indexes within it.

## Index

In the ES Serverless service, an index is the smallest unit for data storage and management. It leverages Tencent Cloud ES's proprietary self-managing index capabilities, including built-in shard auto-optimization, intelligent lifecycle management, and fault self-recovery. Unlike traditional usage methods, you do not need to worry about index rollover or shard size; instead, you can focus solely on data writing, querying, and visual analysis.

**Upgrade Notes:**

The ES Serverless service has been upgraded to improve user experience, and currently supports unified access addresses and Kibana for managing and accessing multiple indexes, aligning better with traditional usage habits. The differences before and after the upgrade are as follows:

The project spaces created before January 23, 2024, lack independent access control and do not support simultaneous access to multiple indexes in Kibana. Data writing and querying are performed through each index's access address.

For the project spaces created after January 23, 2024, unified management and access for all indexes within the space can be achieved through the project space's access address and Kibana. Additionally, permissions can be configured through a visualized user management feature, allowing you to set permission types and scopes, aligning closely with traditional ES cluster usage to meet various scenarios. The upgrade requires no changes to business code -- simply migrate existing indexes to the new space. We strongly recommend migrating indexes to the new space.

# 5-Minute Quick Experience

Last updated：2024-12-04 15:56:05

## Overview

The ES Serverless service is a fully managed, cloud-native ES service by Tencent Cloud, built on a self-developed Serverless architecture with no cluster concept. Users can create and use indexes as needed, benefiting from **auto-scaling and completely maintenance-free** capabilities, which effectively solve the issue of high resource costs associated with peak and off-peak fluctuations in **log analysis and metric monitoring** scenarios. Fully compatible with the ELK ecosystem, it provides end-to-end data writing, data management, and data visualization features, providing a **plug-and-play log analysis experience**.

## Quick Start

The ES Serverless service supports writing data into indexes through methods such as **native ES APIs, Logstash, Flink, or Kafka**. If you require log collection for services such as CVM, TKE, or TCHouse-C, a one-stop visualized configuration option is also available. By simply setting up data sources and index information, logs can be collected into indexes for efficient retrieval and analysis. This document will guide you through the full process of **index creation > data writing > retrieval and analysis**, giving you a quick overview of using the ES Serverless service in log analysis scenarios.

## Basic Concepts

Before diving into the experience, let us review several relevant basic concepts:

| Name | Description |
| --- | --- |
| **Project space** | Project space is a fundamental resource unit in the ES Serverless service. You can create indexes related to the same business within a single project space, facilitating index management. |
| **Index** | Index is the smallest unit for data storage and management, providing log storage and near real-time query capabilities. Collected log data can be stored in indexes. |
| **Kibana** | Kibana is a data analysis and visualization platform integrated with ES, allowing for log writing, retrieval, and chart creation (such as maps and line charts). |
| **Logs** | Logs are records generated during the operation of application systems, including operation logs, access logs, and error logs. |

# Creating a Space

1. Log in to the ES Serverless console.
2. In the space list, click **Create Project** to enter the project creation page.
3. On the project creation page, configure the following settings:

**Project Name**: Use this name to identify the project. Follow the naming guidelines provided on the page.

**VPC** / **AZ and Subnet**: The project space is created within a VPC to ensure secure access. Select the appropriate VPC, availability zone, and subnet. if creation is needed, see Create New VPC and Create New Subnet.

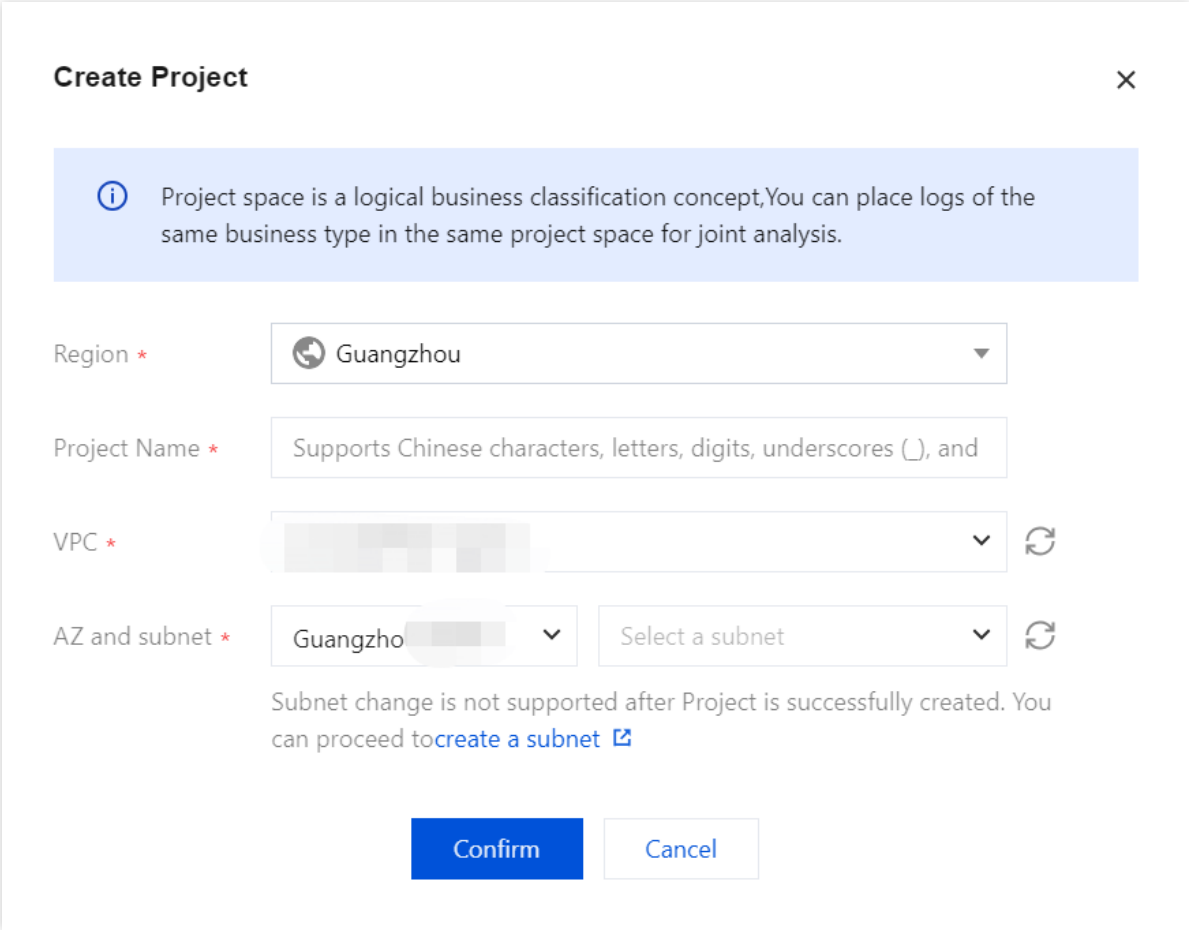


4. After completing the information, click **Confirm** to create the project.

# Creating an Index

There are two methods to create an index: directly from the Project list page or from the Project Basic page. The following example demonstrates the process on the Project list page.

1. On the Project list page, enter the Quick Access Data page and select your data source. Here, we will use API write as an example.

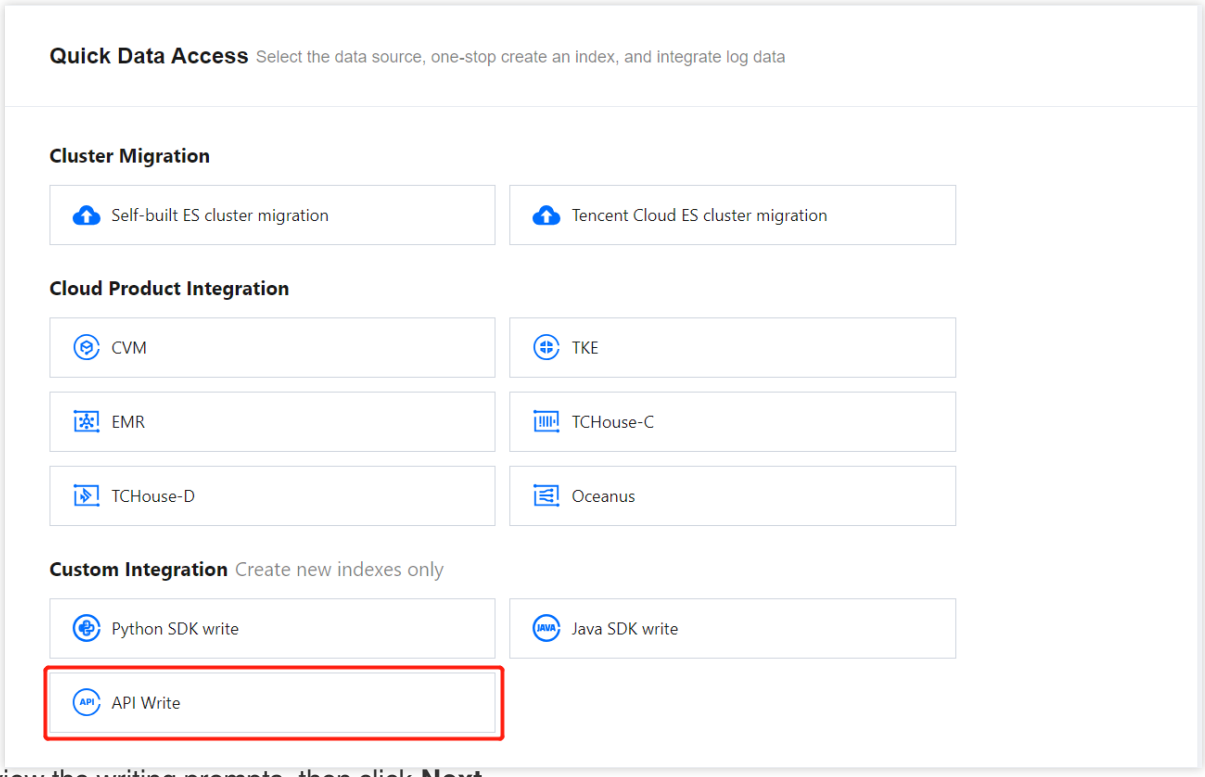


2. Review the writing prompts, then click **Next**.

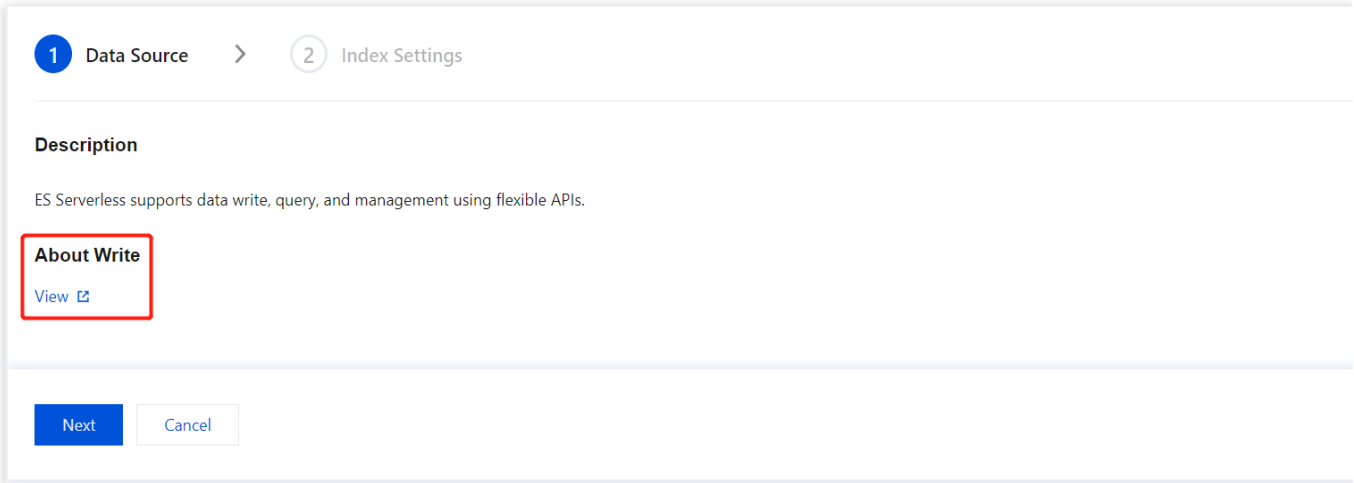


3. On the Index Settings page, enter the basic information and index configuration, then click **Create**.

**Region:** Select the regional information from the dropdown list.

**Project:** Choose the project space to organize the index for easier management. If no options are available in the dropdown, click **Create Project** and follow the previous instructions to create one.

**Index name:** This name will be used for subsequent data writing and querying. Follow the naming prompts shown on the page.

**Field mapping:** Used to set the field details of the data. You can select **Dynamic creation**, which will automatically generate field settings based on the data you write in, or choose to customize the field settings.

**Time field:** Select or input a field with a date type from your data. Once the index is created, this field cannot be modified.

**Data retention period:** The retention period of the data. For example, if it is set to **Limited 30 days**, data will be deleted on the 30th day after being written to the index.

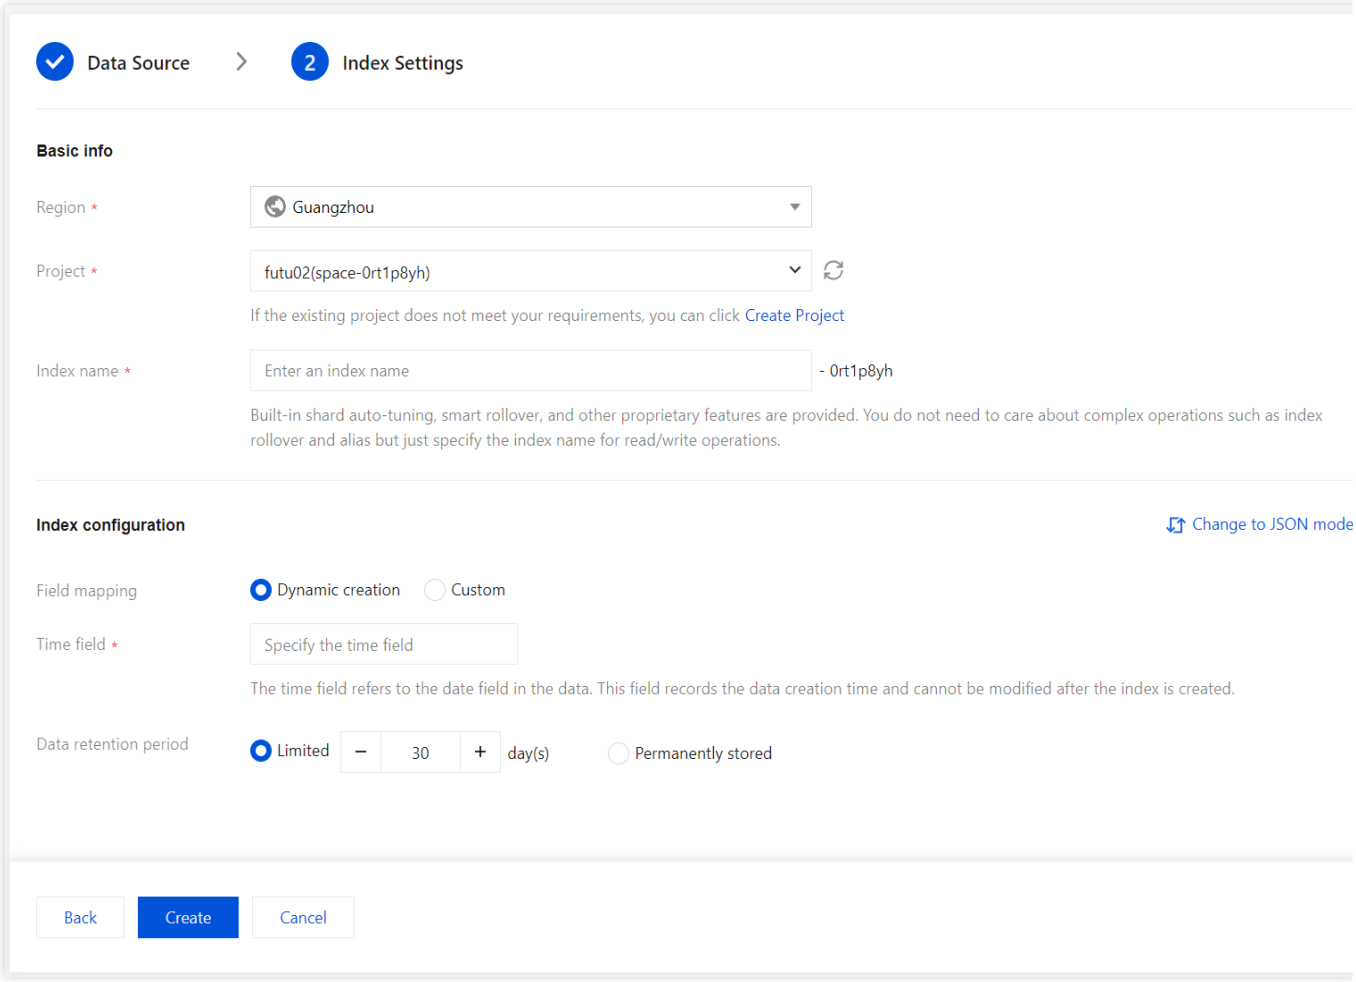


# Data Writing

1. In the project list, click the name of the desired project to enter the Project Management page.

**Note:**

Kibana's relevant modules are embedded directly into the Tencent Cloud Console, allowing you to use search and analysis features directly. **Search and Analysis** corresponds to **Discover**, and **Development Tools** corresponds to **Dev Tools**. This embedded feature requires third-party cookies to be enabled in your browser; if you experience issues, please enable third-party cookies. To access Kibana externally for data writing, see Writing Data.

Enter the **Development Tools** page.

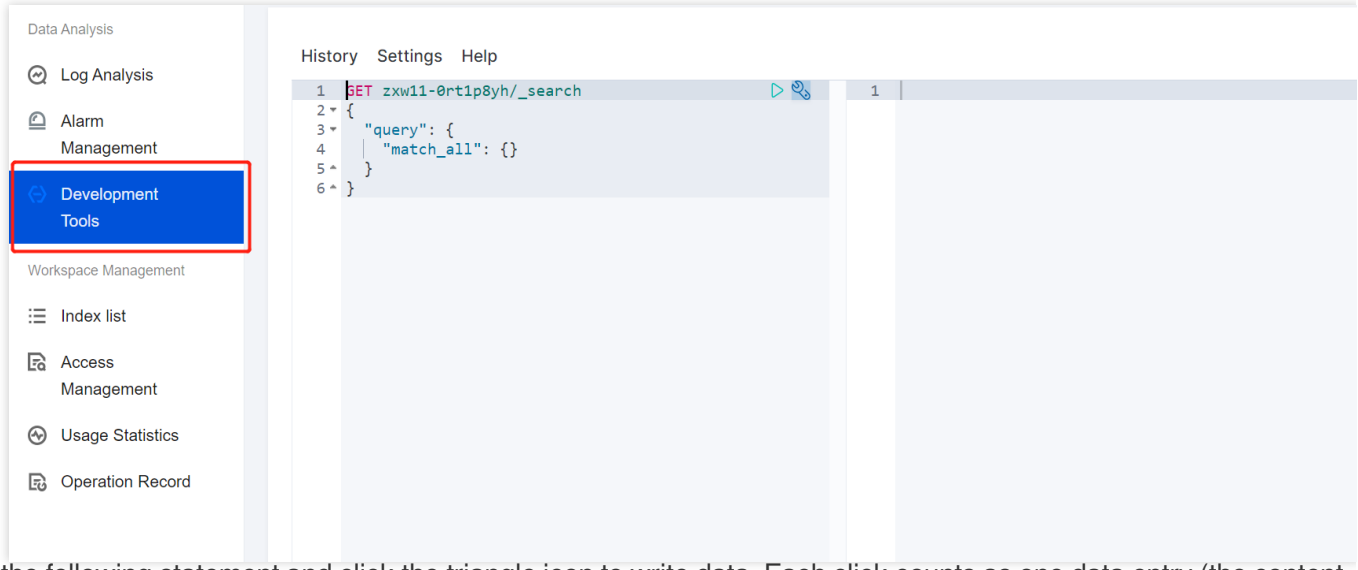


Enter the following statement and click the triangle icon to write data. Each click counts as one data entry (the content within {} represents a complete log entry). You may click several times to generate enough entries for the upcoming data retrieval demonstration.


Sample statement:

```
POST index name/_doc
{
    "id":"090798",
    "routing_no":"4087",
    "region":"10002424",
    "user_name":"user-Ufa9Yee1P",
    "user_type":"01",
    "ip":"119.147.10.191",
    "now_local":"gz",
    "@timestamp":1705648983762
}
```

**Note:**

Replace **Index Name** in the statement with your specific index name.

If your time field is not **@timestamp**, modify **@timestamp** in the figure to match your custom time field.

For batch data entries, see Writing Data.

# Retrieval and Analysis

With the data successfully written into the ES Serverless service, the following steps demonstrate how to query this data.

### Method 1: Using DSL

1. Copy the example statement below and click the triangle icon to execute a query on the written data.

```
GET index name/_search
{
    "query":{
      "match_all":{}
      }
}
```

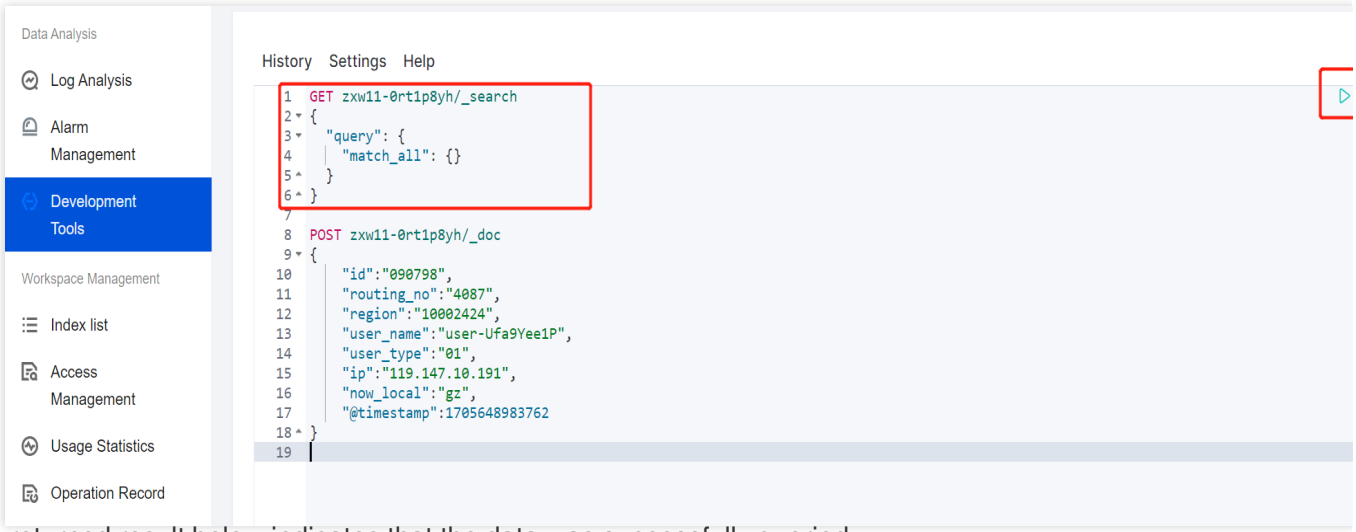


2. The returned result below indicates that the data was successfully queried.

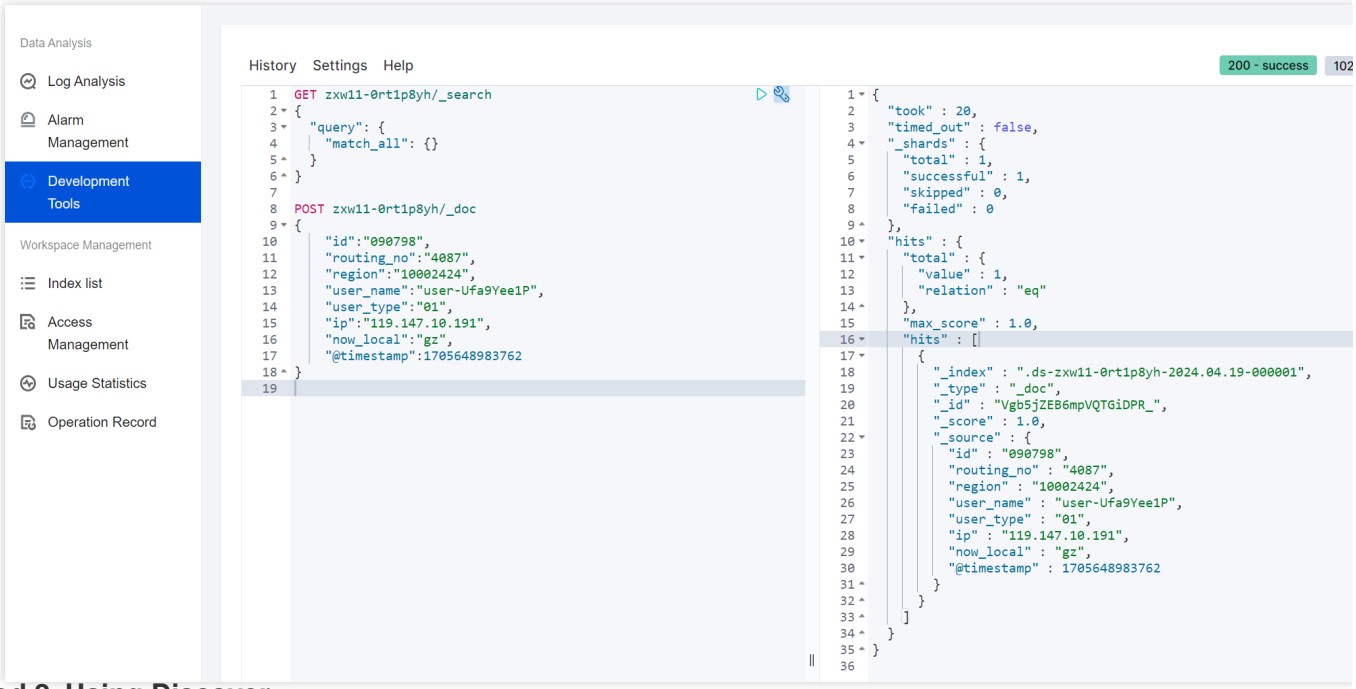


**Method 2: Using Discover**

1. Click **Log Analysis, and select the index just written from the index drop-down list**.

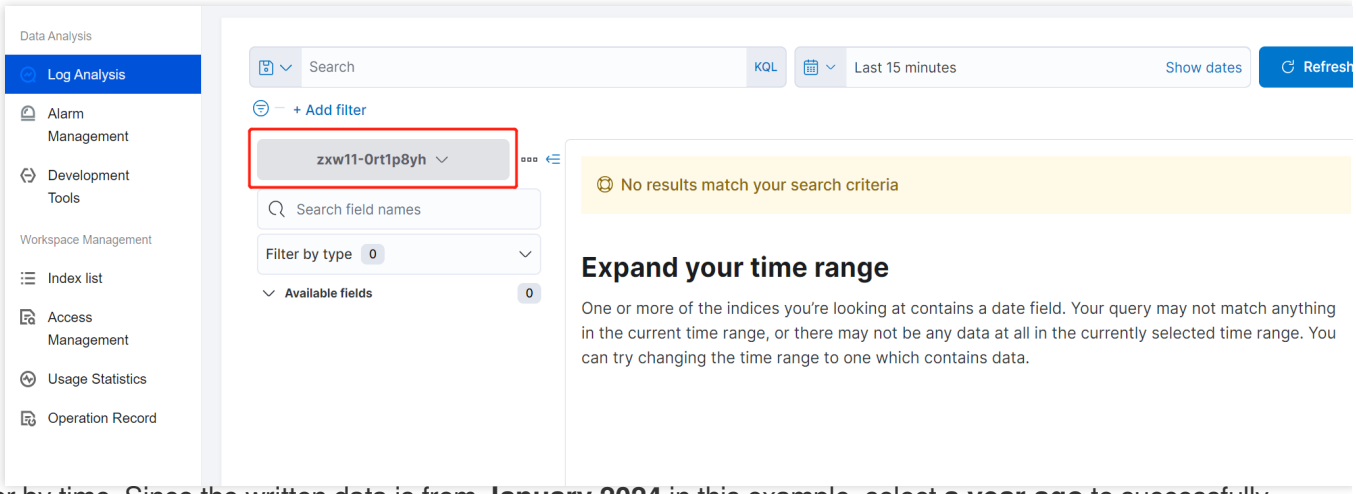


2. Filter by time. Since the written data is from **January 2024** in this example, select **a year ago** to successfully retrieve data from the past year.

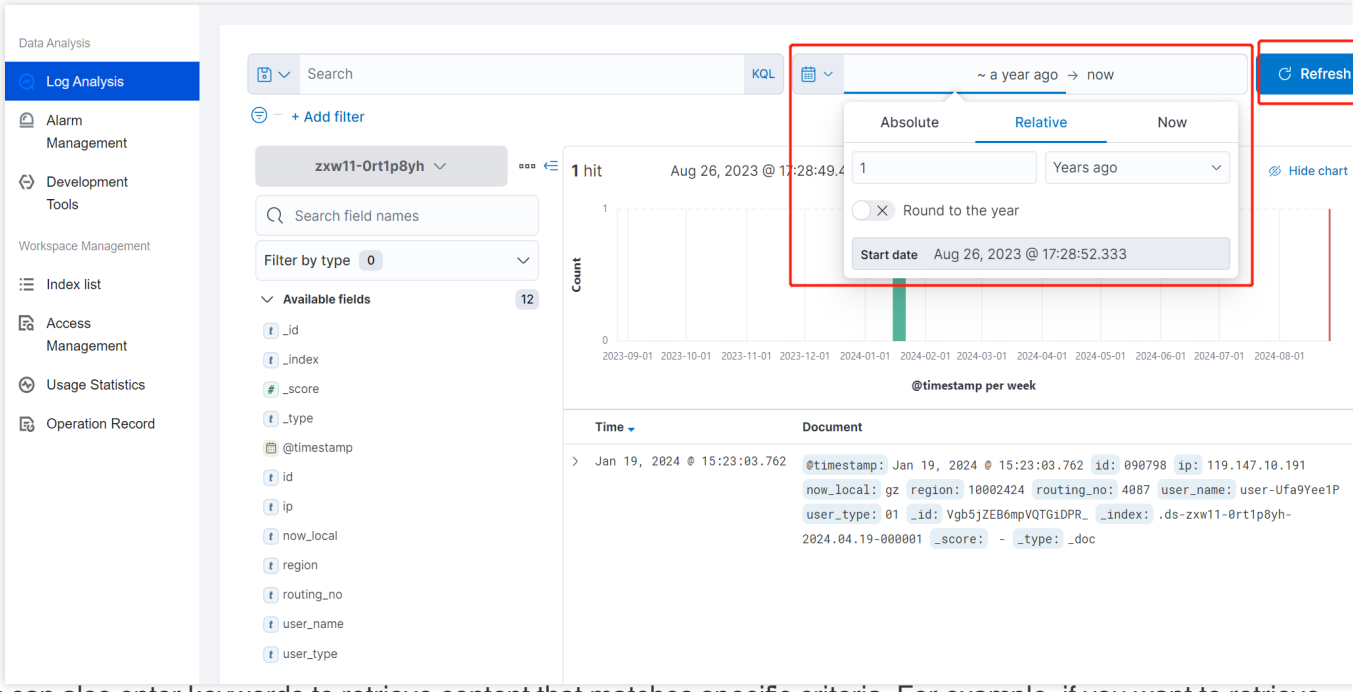


3. You can also enter keywords to retrieve content that matches specific criteria. For example, if you want to retrieve entries where the field **now_local** has the value **gz**.

Click **now_local** as shown below:

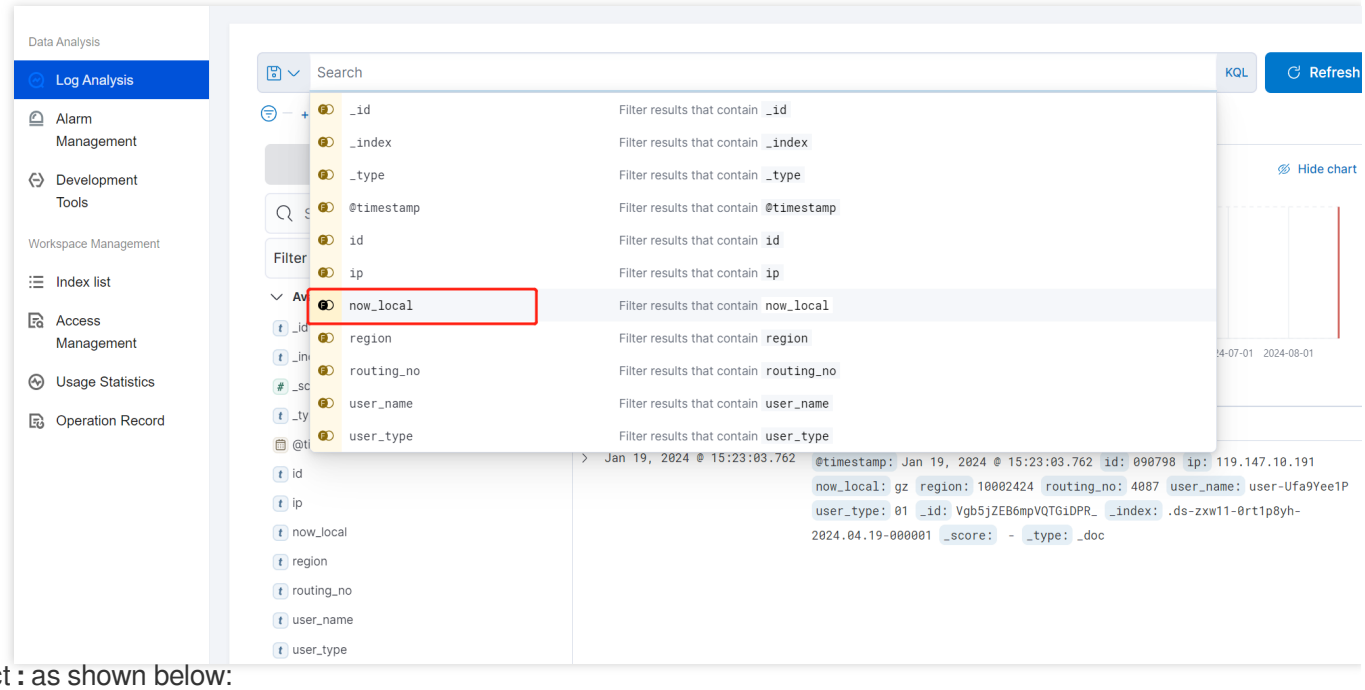


Select **:** as shown below:

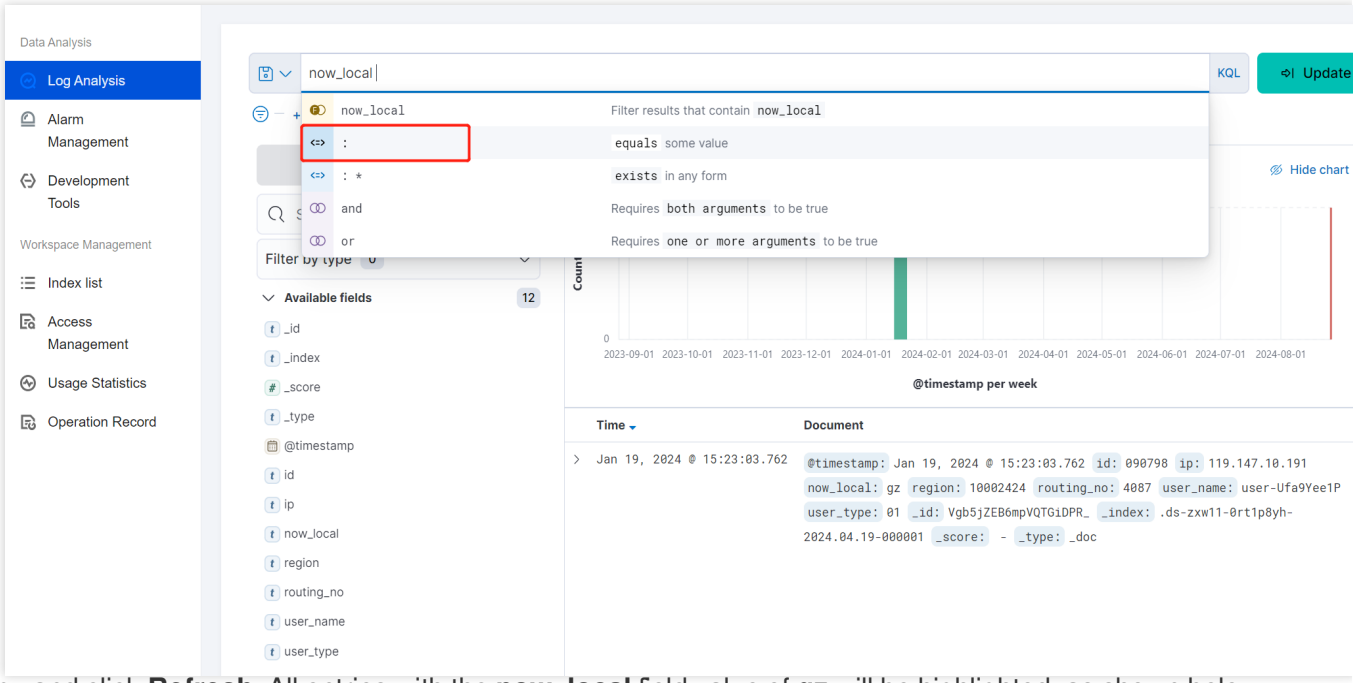


Enter gz and click **Refresh**. All entries with the **now_local** field value of **gz** will be highlighted, as shown below:

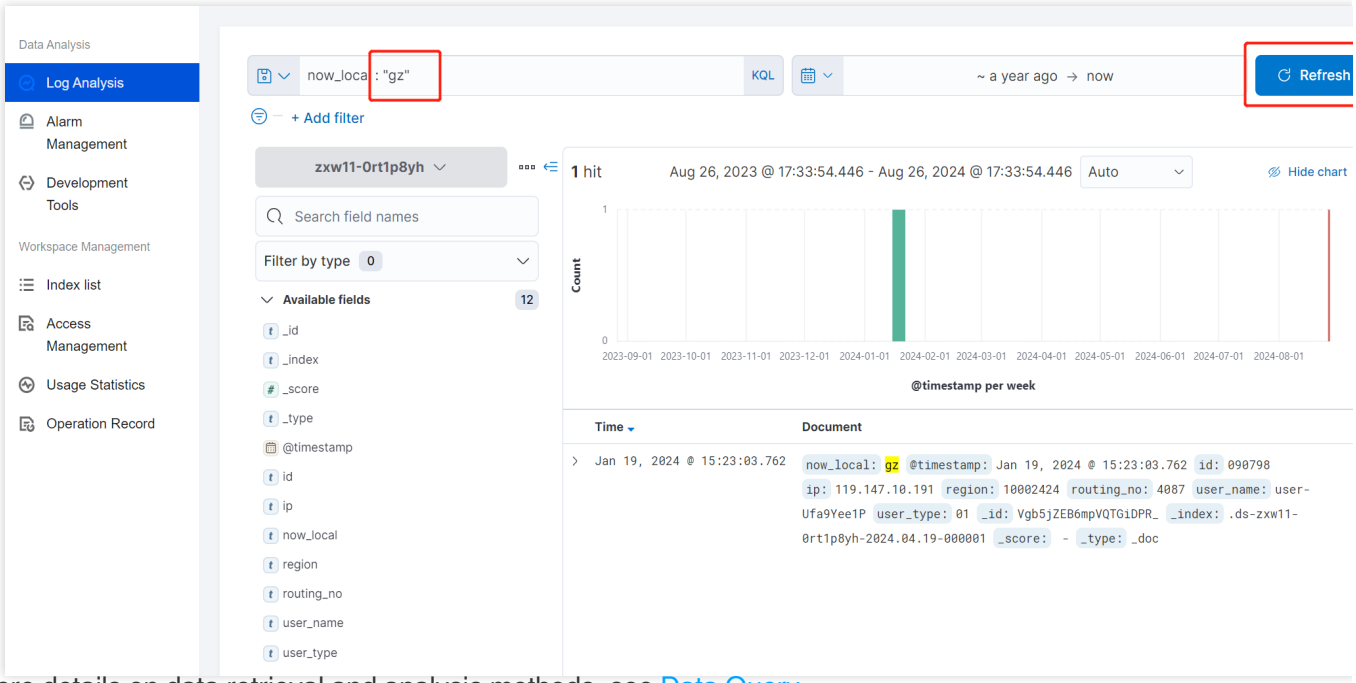


For more details on data retrieval and analysis methods, see Data Query.

# Quick Start
# Creating Indexes

Last updated：2024-12-04 15:59:13

## Prerequisites

A Tencent Cloud account has been created. For account creation, see Signing Up.
If logging in with a sub-account, ensure that the account has read and write permissions for ES.

## Operation Steps

### Logging In to Console

1. Log in to the Elasticsearch console.
2. In the top menu bar, select the region. Currently supported regions include Beijing, Shanghai, Guangzhou, Nanjing, Hong Kong (China), Singapore, Tokyo, and Virginia.
3. In the left sidebar, choose **Log Analysis** under the Serverless mode.

### Creating a Project Space

1. Click **Create Project**.
2. Enter the project space name, which can include 1 - 20 characters, including Chinese characters, letters, digits, underscores, or delimiters (-).
3. Click **Confirm**. Once validated, the project space will be created.

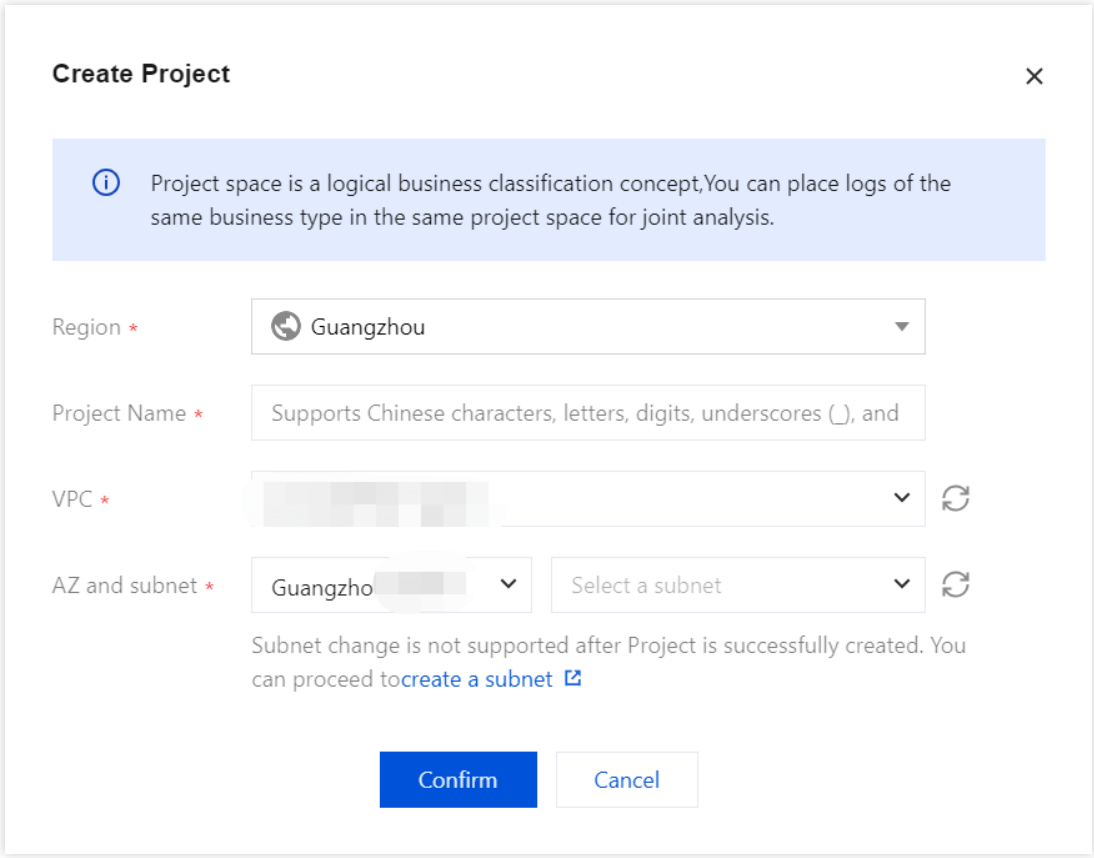


**Note**

In Elasticsearch Serverless Log Analysis, you can create an index and subsequently write data via API or access data sources such as CVM and TKE via the Data Access tab of the corresponding index. You can also set up data access during index creation, enabling one-stop access for CVM Log Access, TKE Log Access, and more. The following section explains the index creation process for API-based data writing.

## Creating an Index

1. On the ES Serverless Log Analysis homepage, click the **Project Name** to enter the Index List page, then select **Create Index And Integrate Data.**

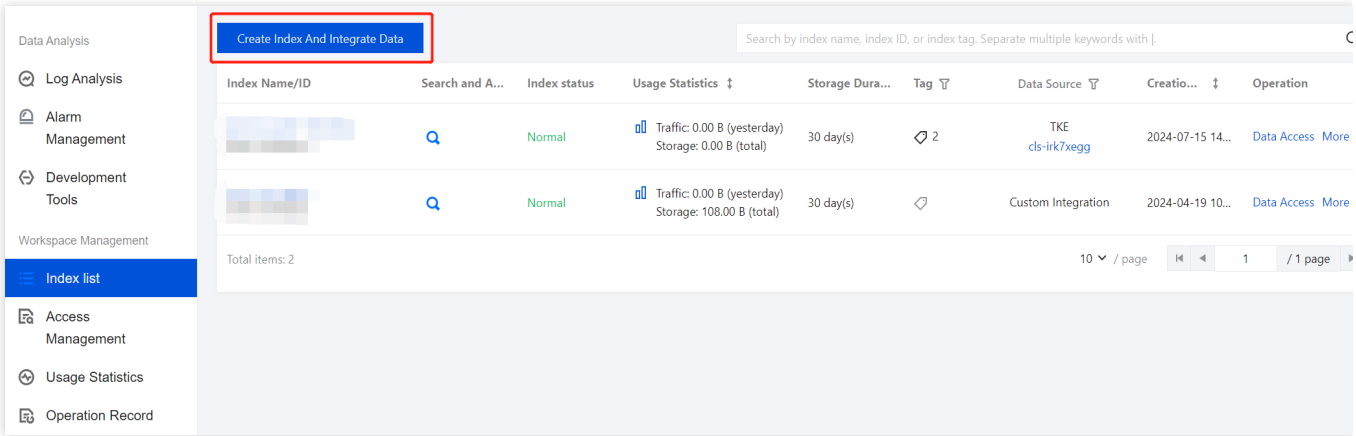


2. On the create index page, select **API writing.**

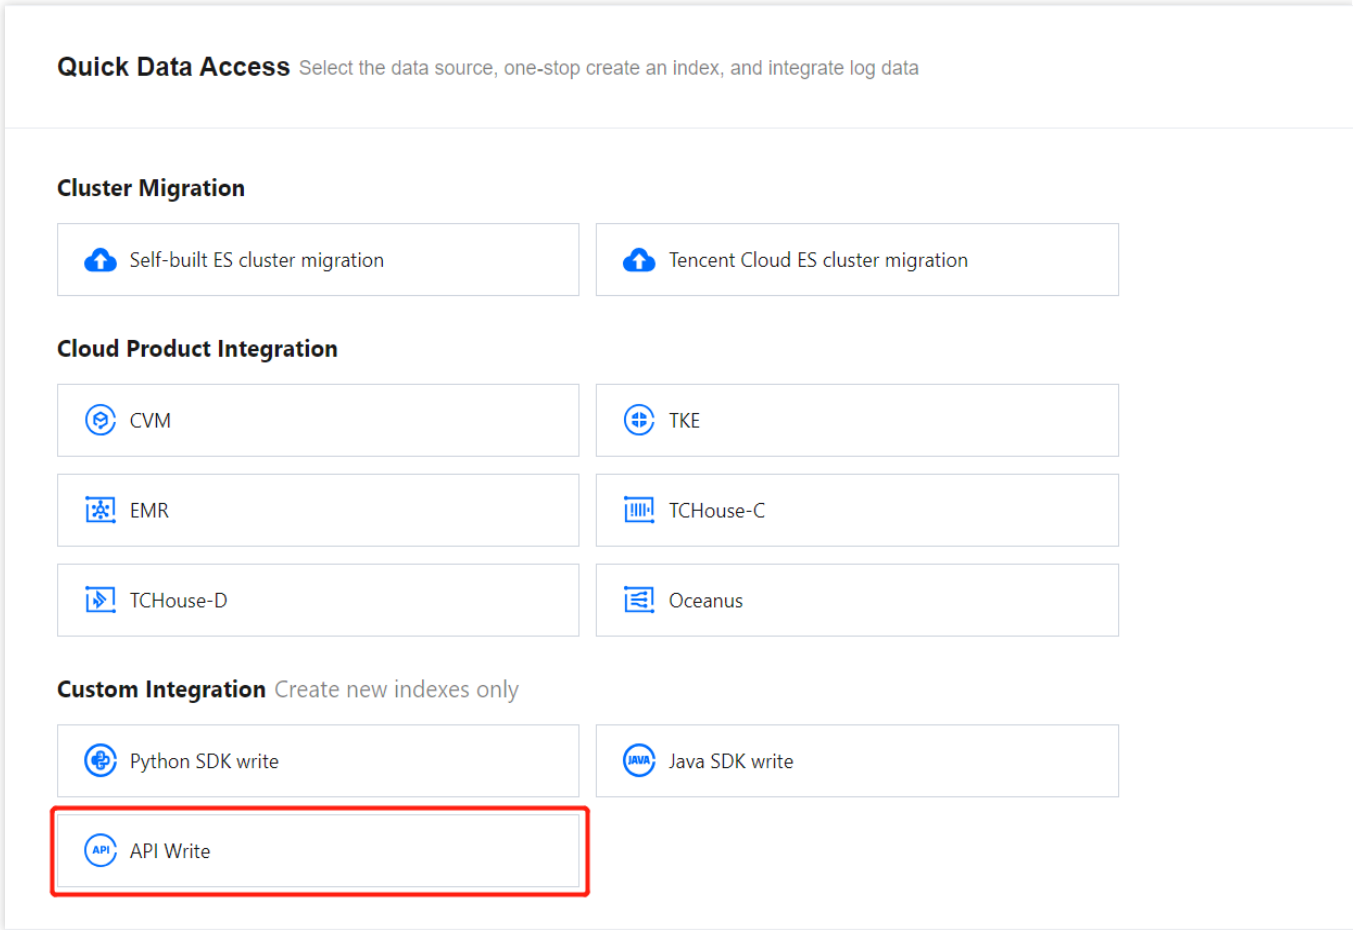


3. If you want to view how to write data to ES Serverless via API, click View Documentation. Then, click **Next**.

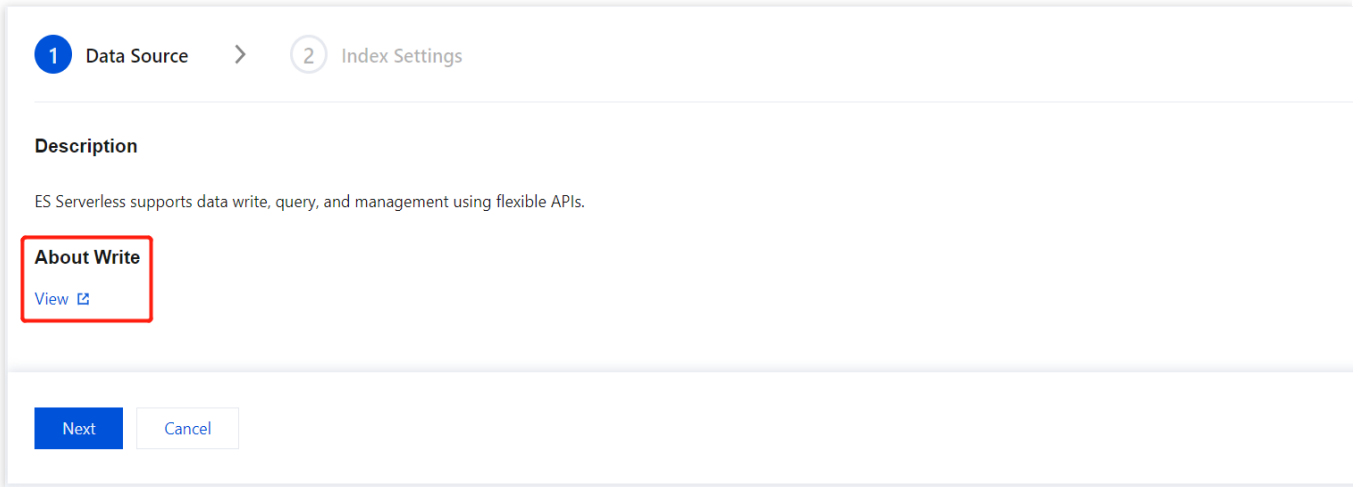


4. Enter the index settings page and fill in the basic information.

Region: Aligns with the region of project space.

Project Space: Defaults to the current project space.

Index Name: Length of 1 - 100 characters. Lowercase letters, digits, and the following symbols: -, _, ;, @, &, =, !, ', %, $, ., +, (, ) are supported.

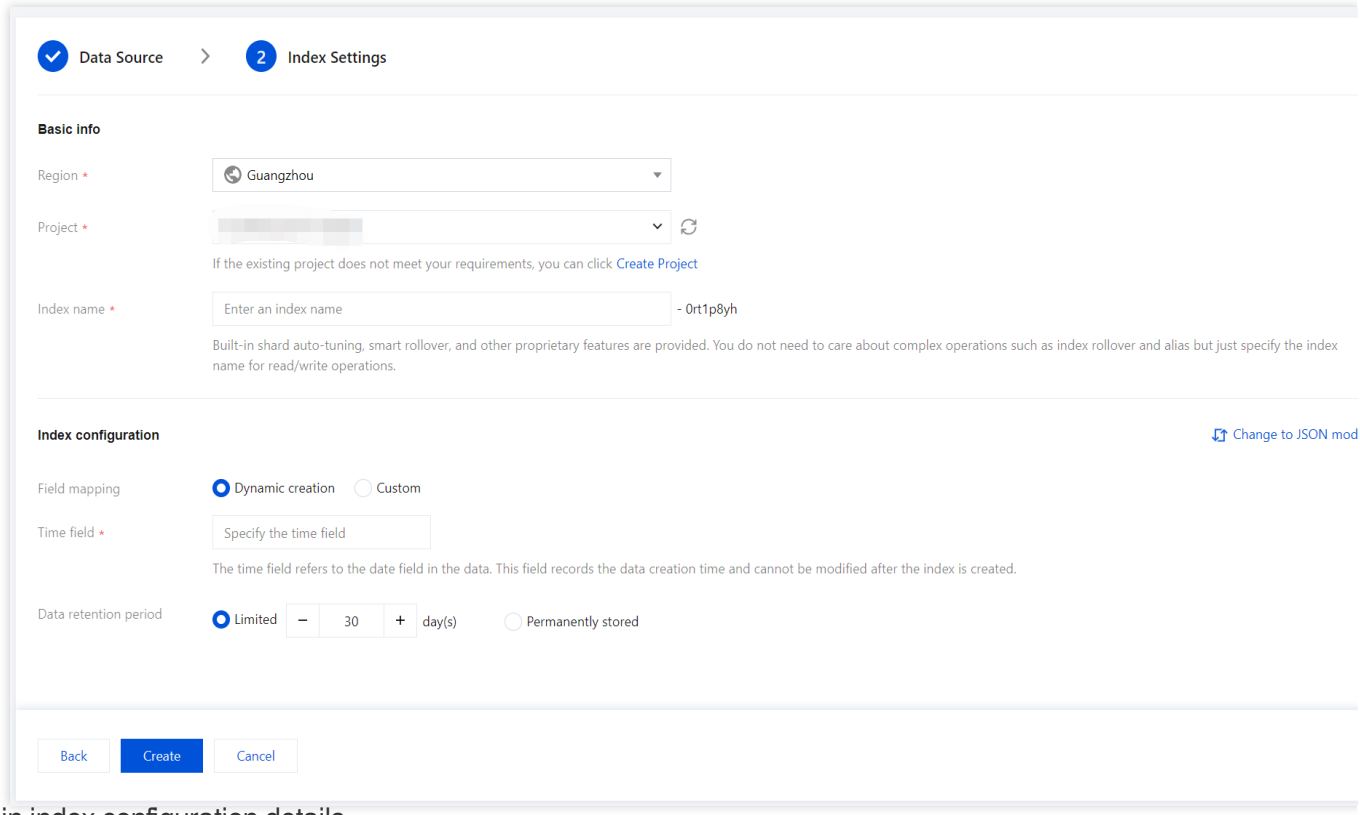


5. Fill in index configuration details.

**Field Mapping**

Dynamic Generation: Enabled by default. When enabled, it automatically parses written data and generates field settings for the index.

Input Sample Auto-Configuration: If **Dynamic Generation** is disabled, you can use **Input Sample Auto-Configuration** to generate field mappings for the index by entering a JSON-formatted data sample. After confirmation, the platform will validate the input; if the validation is successful, the relevant fields will be mapped in the field mapping table.

Field mapping divides the original data into multiple tokens based on fields (key:value) for indexing. Retrieval relies on this mapping, as detailed below:

| Parameter | Description |
| --- | --- |
| Field name | The field name in the written data. |
| Field type | The data type of the field. The interface supports 9 types: text, date, boolean, keyword, long, double, integer, ip, and geo_point. Additional field types are supported in **JSON mode**. For details, see Official Documentation. |
| Includes Chinese | Enable this option if the field includes Chinese text and requires Chinese retrieval. When it is enabled, the ik_max_word tokenizer is applied to the text field by default. |
| Enable indexing | When it is enabled, this field will be indexed for search. |
| | |

| | |
|---|---|
| Enable statistics | When it is enabled, this field's values can be analyzed statistically, which will increase index storage. |

**Time Field**

The time field refers to a field with the type date in the actual data. Once the index is created, this field cannot be modified.

**Note:**

By default, the time field enables **indexing** and **statistics**, and these settings cannot be disabled.

**Data Storage Duration**

You can set the data retention period, with a default of 30 days, or select an option for permanent storage.

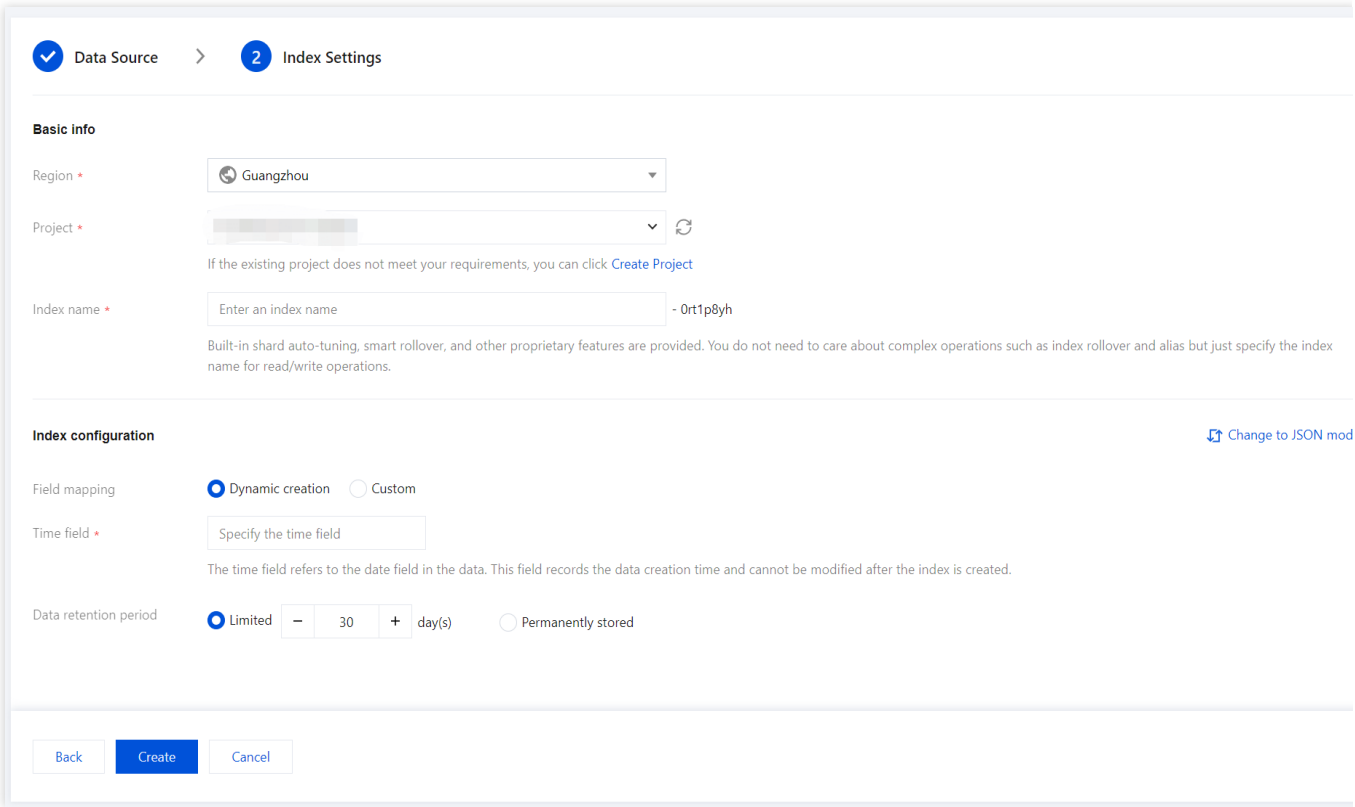


6. Once the information is entered correctly, click **Create** to complete the index creation. For the instructions on data writing, see documentation.

# CVM Log Access

Last updated：2024-12-04 16:01:35

## Prerequisites

A Tencent Cloud account has been created. For account creation, see Signing up for Tencent Cloud.
If logging in with a sub-account, ensure that the account has read and write permissions for ES.

## Operation Steps

### Logging in to the Console

1. Log in to the Elasticsearch Console.
2. In the top menu bar, select **Region**. Currently supported regions include Beijing, Shanghai, Guangzhou, Nanjing, and Hong Kong (China).
3. In the left sidebar, choose **Log Analysis** under the Serverless mode.

### Creating a Project Space

1. Click **Create a project**.
2. Enter a **Project Name** for the project, which can include 1–20 characters, consisting of Chinese characters, letters, digits, underscores, or delimiters (-).
3. Click **Confirm**. If the validation is successful, the project space will be created.

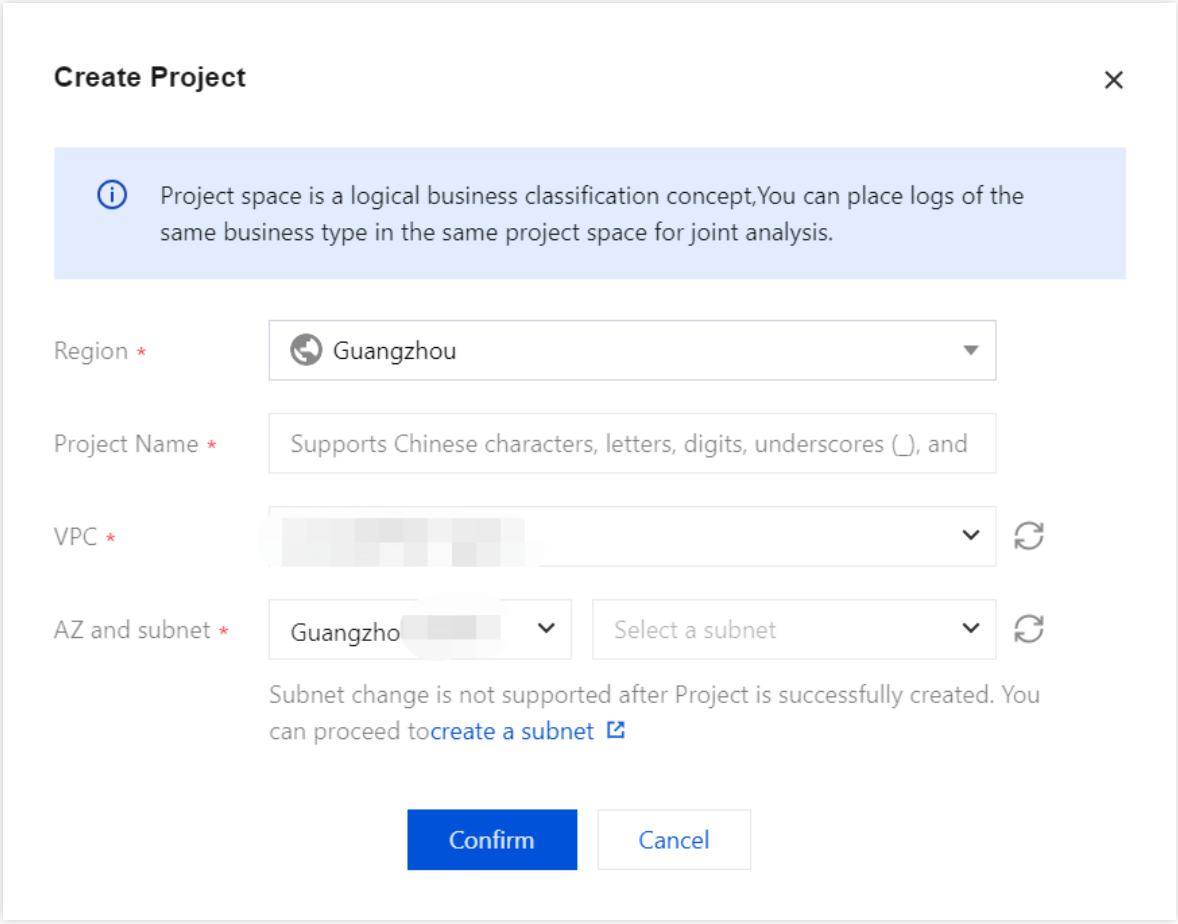


**Note:**

In ES Serverless Log Analysis, you can simply create an index, then use the API for data writing or access data sources such as CVM or TKE via the Data Access tab of the corresponding index. You can also set up data access during index creation for one-stop CVM and TKE log access. The following introduces the one-stop CVM log access process.

## CVM Log Access

On the ES Serverless Log Analysis homepage, select **CVM** to enter the CVM Log Access page.

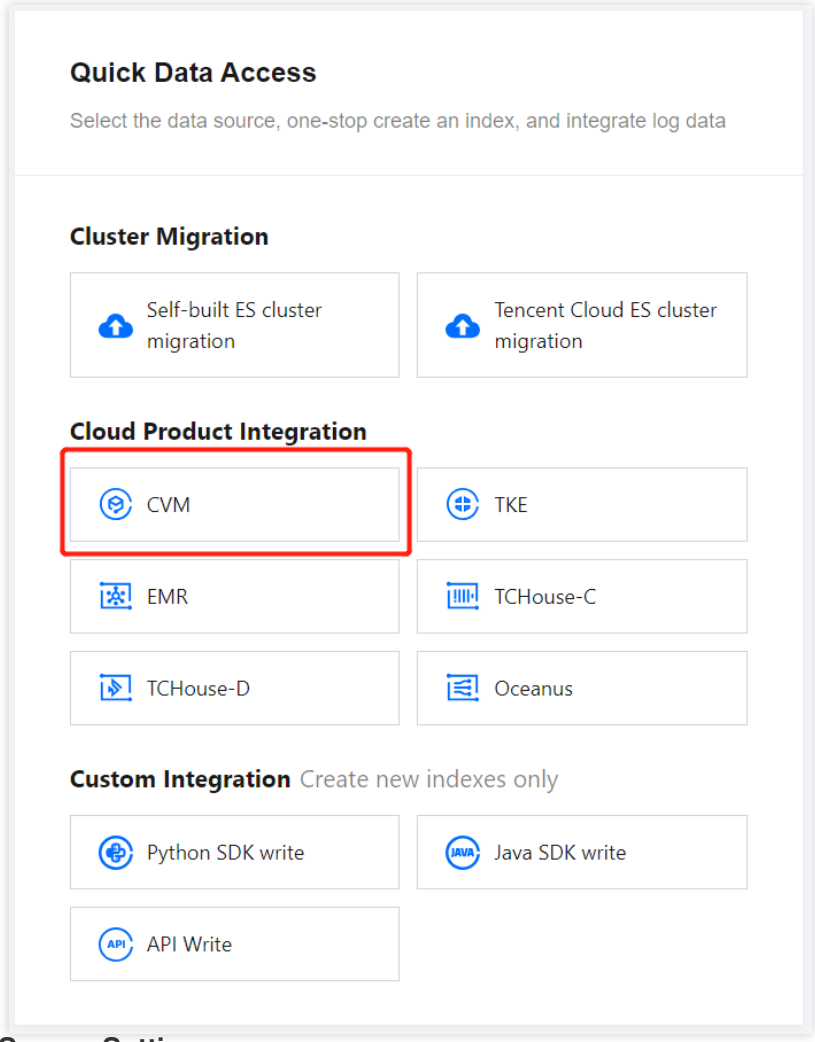


### Data Source Settings

**Region**: Required. It represents the region where the CVM is located.

**VPC**: Required. It represents the private network where the CVM is located. After confirmation, the servers under this VPC will be pulled in.

**Select CVM**: Select the CVM instance for log collection. Currently, only Linux-based CVMs are supported, and data collection requires Installing TAT Agent.

**Collection Path**: Set the log directory and file names based on the location of logs on the server. Supports one or more paths. Directory and file names can be specified using exact names or wildcard patterns.

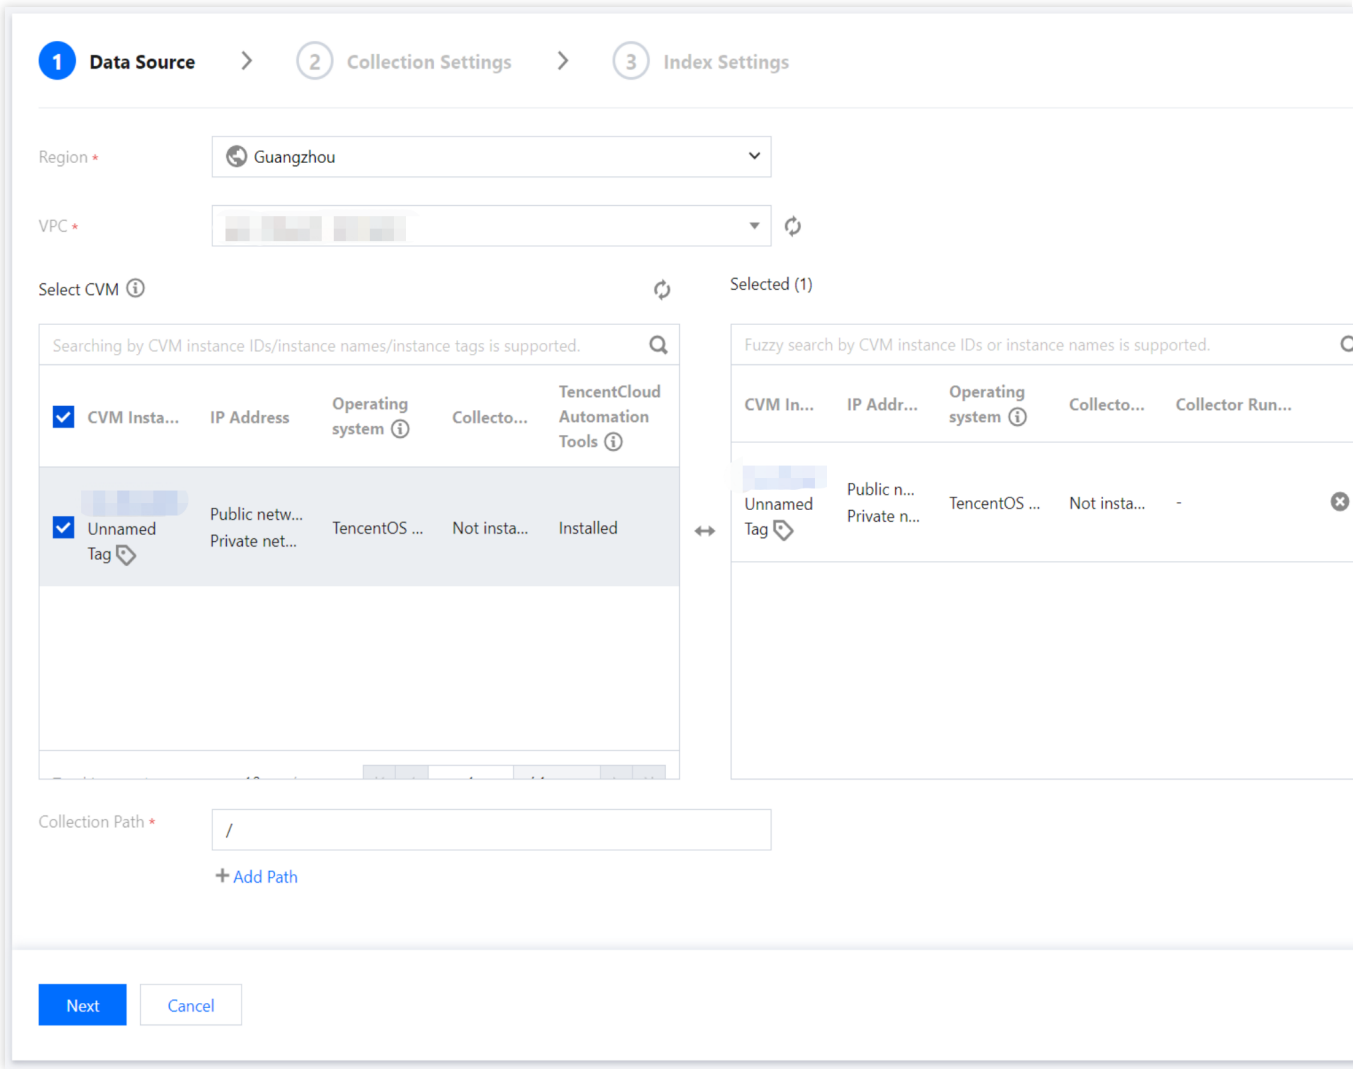


## Collection Settings

### Basic Settings

Collection policy: Supports both full and incremental collection. Once created, the collection policy cannot be modified. Full collection gathers historical log files as well as any logs generated after the Filebeat configuration takes effect; incremental collection only gathers logs generated after the Filebeat configuration takes effect.

### Collection and Parsing

**Collection Template**: If you need a quick setup or trial, select a collection template based on your log output format. After confirming, you can return to the interface and replace the log sample with actual log data to quickly complete the collection parsing setup.

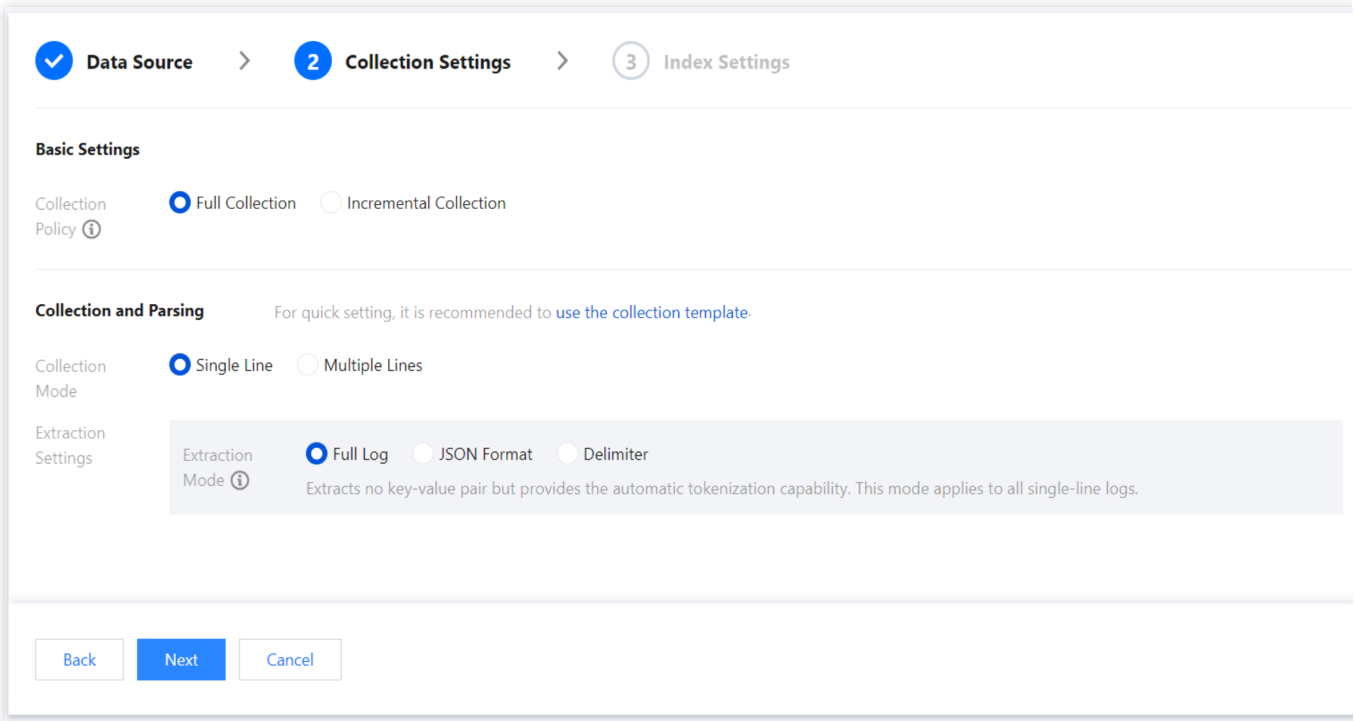


**Collection Mode**: Supports single-line and multi-line modes. Once created, the collection mode cannot be modified.
Single-line text log: Each line in the log file represents one log entry, with each log separated by a line break.
Multi-line text log: Each log entry consists of multiple lines, such as Java stack trace logs. In this mode, you need to configure a log sample and a line-start regular expression. Filebeat uses the line-start regular expression to identify the beginning of each log entry, treating unmatched parts as part of the current log until the next line start appears. After you enter a log sample, the system automatically generates a line-start regular expression by default. You can also customize the expression, with highlighted content in the input box indicating the matched line-start information.

**Note:**

Be sure to use logs from the actual scenarios to facilitate automatic extraction of the line-start regular expression.

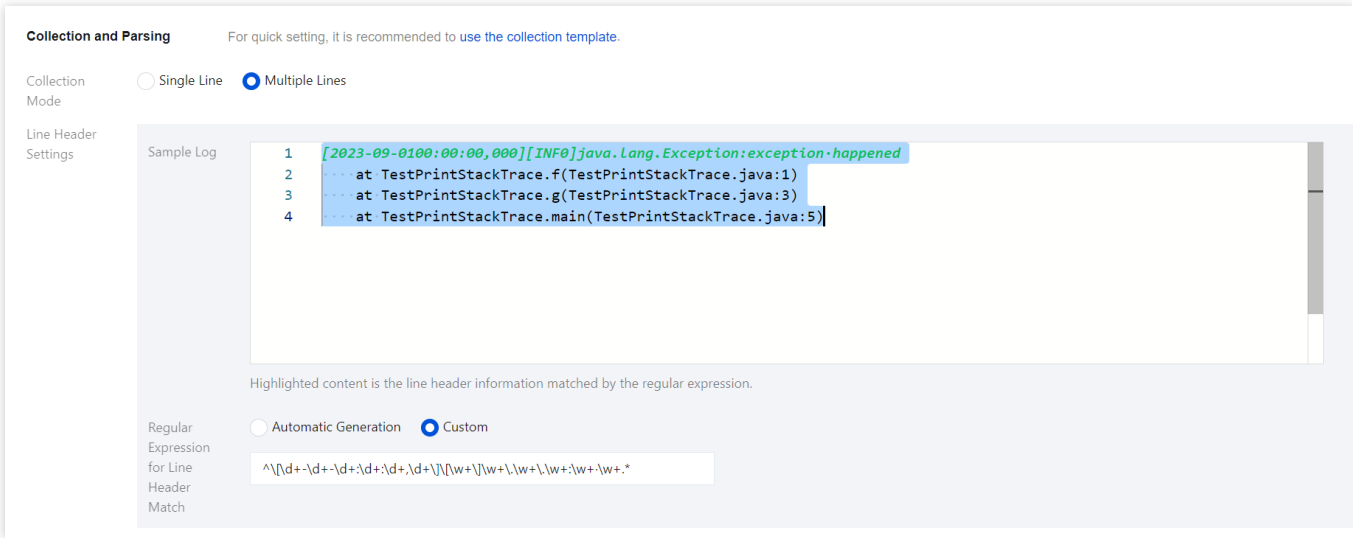


**Extraction Settings**: You can set the extraction mode to full text log, JSON format, or delimiter. Once created, the extraction mode cannot be modified. Details are as follows:

Full Text Log

JSON Format

Delimiter

No key-value extraction is performed on log data, and log content is stored in a field named message. You can perform retrieval and analysis using automatic word segmentation.

For example, a single-line log entry in its original format might be:

```
Tue Jan 01 00:00:00 CST 2023 Running: Content of processing something
```

The data collected in the index would be:

```
massage:Tue Tue Jan 01 00:00:00 CST 2023 Running: Content of processing
something
```

For logs in standard JSON format, fields can be extracted based on the Key: Value pairs within the log.

Suppose your original JSON log entry is:

```
{"pid":321,"name":"App01","status":"WebServer is up and running"}
```

After structuring, this log entry will be transformed as follows:

```
{
    "pid":321,
    "name":"App01",
    "status":"WebServer is up and running"
}
```

For logs with content separated by a fixed delimiter, you can extract key-value pairs based on the specified delimiter. The delimiter can be a single character or a string and can be selected or entered in the console.

Suppose your original log entry is:

```
321 - App01 - WebServer is up and running
```

By specifying the delimiter as -, this log will be split into three fields. You can define a unique key to these fields in the extraction results, as shown below:

```
pid: pid
name: App01
status: WebServer is up and running
```

**Extraction Results**: If the extraction mode is set to JSON format or delimiter, you can enter a log sample, and the system will automatically extract information from it:

For JSON format, the system will automatically populate the extracted Key-Value pairs. If you deselect a field, it will not be written to the index.
For delimiter mode, the system will automatically populate the extracted Values. You can define a unique Key for each Value. If you deselect a field, it will not be written to the index.
Built-in fields: When you configure CVM log collection in the console, Filebeat writes information such as log source and timestamp into the logs as Key-Value pairs. These fields are considered built-in. If a Key name in your business log matches a built-in field name, the content from the business log field will take priority, and the corresponding built-in field will not be written to the index. The meanings of the built-in fields are as follows:

| Built-in Field Name | Meaning |
| --- | --- |
| log.file.path | Path where the log is stored |
| host.name | Name of the server hosting the log |
| host.ip | IP address of the server hosting the log |
| @timestamp | Time when the log entry was collected |

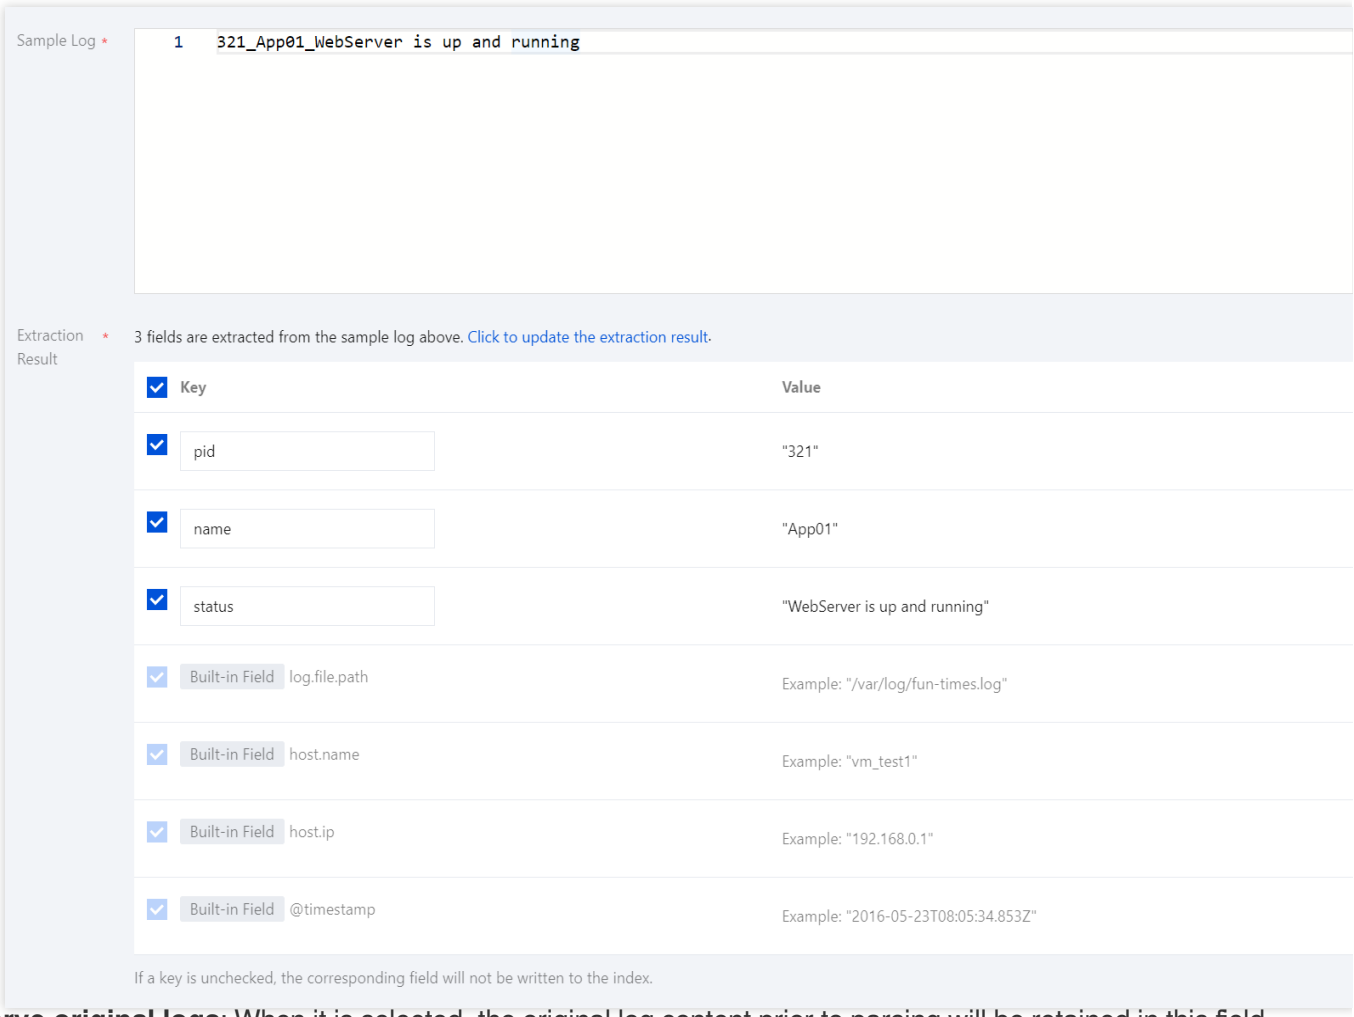


**Preserve original logs**: When it is selected, the original log content prior to parsing will be retained in this field.
**Record parsing errors**: If the extraction mode is set to Delimiter, you can choose whether to log parsing errors. When it is selected, error messages will be uploaded to this field as values in case of parsing failures.

### Index Settings

**Project Space**: You can assign indexes for the same business to a specific project space for easier management.
**Index Name**: Length of 1 - 100 characters. Lowercase letters, digits, and the following symbols: -, _, ;, @, &, =, !, ', %, $, ., +, (, ) are supported.

**Field Map**

Dynamic Generation: Enabled by default. When enabled, it automatically parses and generates field settings for the index based on written data.

Input Sample Auto-Configuration: When **Dynamic Generation** is disabled, you can use **Input Sample Auto-Configuration** to generate field mappings for the index by entering a JSON-formatted data sample. After confirmation, the platform will validate the input; if the validation is successful, the relevant fields will be mapped in the field mapping table.

Field mapping divides the original data into individual tokens based on fields (key:value) for indexing. Retrieval relies on this mapping, as detailed below:

| Parameter | Description |
| --- | --- |
| Field name | The field name in the written data. |
| Field type | The data type of the field; the interface supports 9 types: text, date, boolean, keyword, long, double, integer, ip, and geo_point. Additional field types are supported in **JSON Editing Pattern**. For more details, see Official Documentation. |
| Include Chinese | Enable this option if the field contains Chinese text and requires Chinese retrieval. When it is enabled, the ik_max_word tokenizer is applied to the text field by default. |
| Enable indexing | When it is enabled, an index will be built for this field, allowing it to be searchable. |
| Enable statistics | When it is enabled, statistical analysis can be performed on the field values, which will increase index storage. |

**Time Field**

The time field refers to a field with the date type in the actual data. Once the index is created, this field cannot be modified.

**Note:**

By default, the time field enables **indexing** and **statistics**, and these settings cannot be disabled.

**Data Storage Duration:**

1.1 You can set the data storage duration, with a default of 30 days, or select an option for permanent storage.

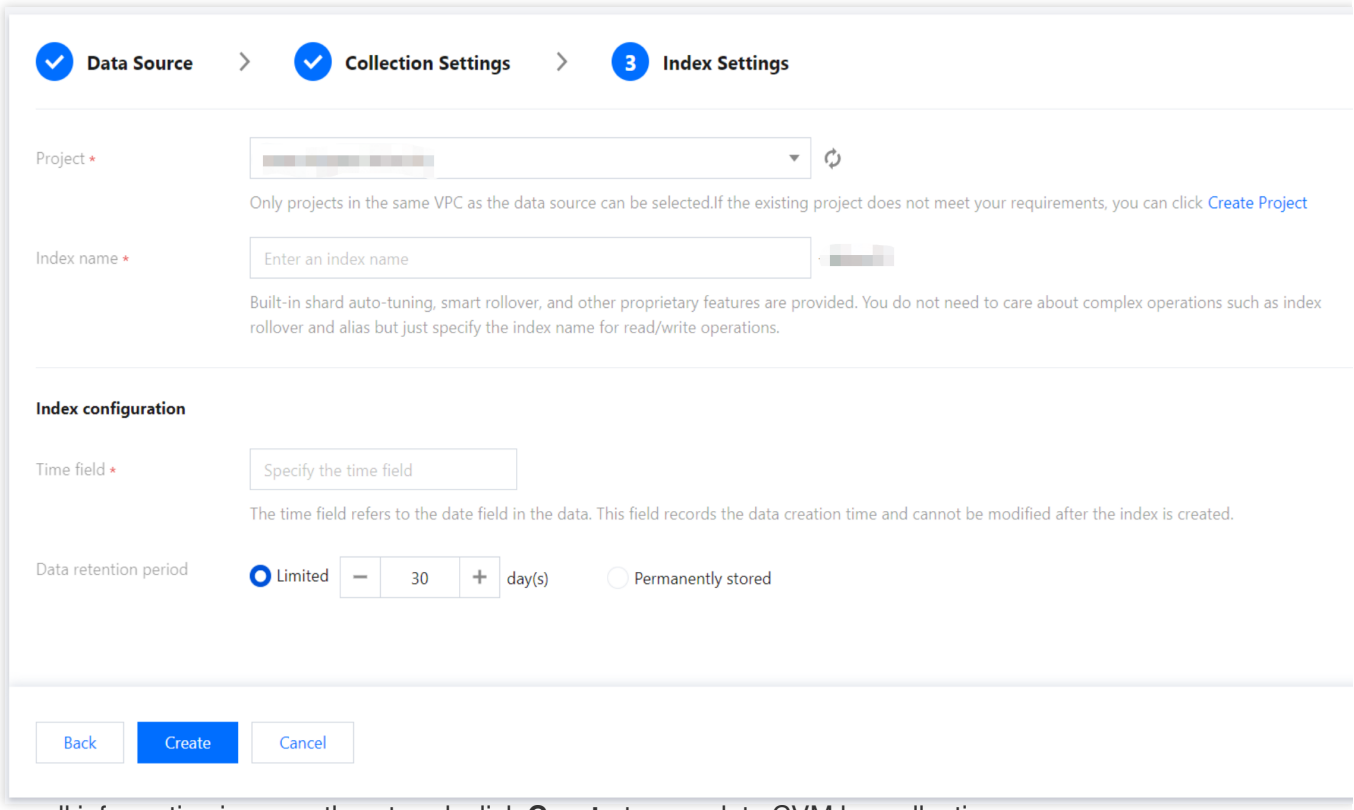


1.2 Once all information is correctly entered, click **Create** to complete CVM log collection.

# TKE Log access

Last updated：2024-12-04 16:06:59

## Prerequisites

A Tencent Cloud account has been created. For account creation, see signing up.
If you log in with a sub-user account, ensure it has read and write permissions on ES.

## Operation Steps

### Logging in to the Console

1. Log in to the ES console.
2. In the top menu, select **Region**. Currently supported regions include Beijing, Shanghai, Guangzhou, Nanjing, and Hong Kong (China).
3. In the left sidebar, choose **Log Analysis** under the Serverless mode.

### Creating a Project Space

1. Click **Create a project**.
2. Enter the **Project Name**, which can include 1 - 20 characters, including Chinese characters, letters, digits.
3. Click **Confirm**. Once validated, the project space will be created.

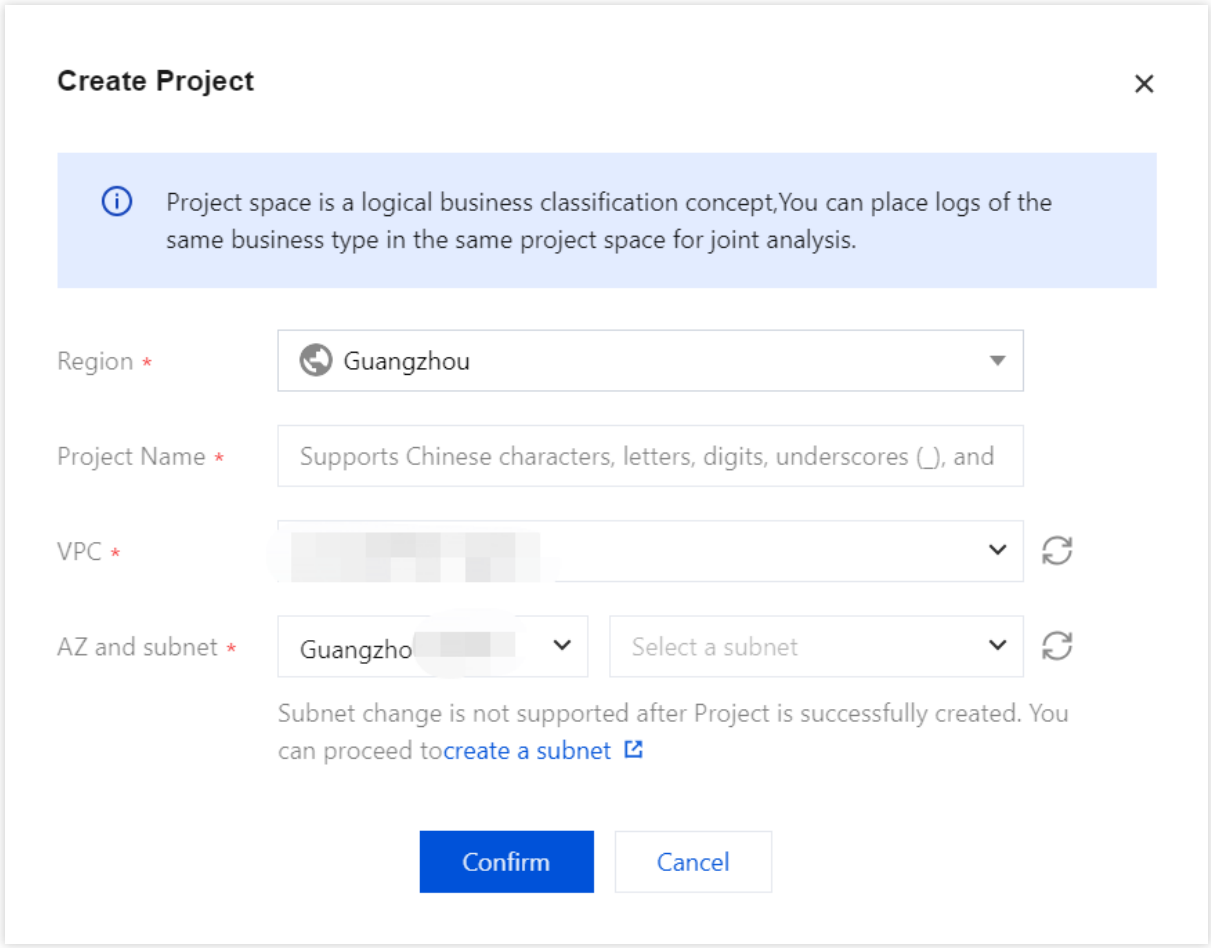


**Note**

In Elasticsearch Serverless Log Analysis, you can create an index and subsequently write data via API or access data sources such as CVM and TKE via the Data Access tab of the corresponding index. You can also set up data access during index creation, enabling one-stop access for CVM and TKE log access. Below is the instruction for one-step TKE log access setup.

## TKE Log Access

On the ES Serverless Log Analysis homepage, select **TKE** to enter the TKE log access page.

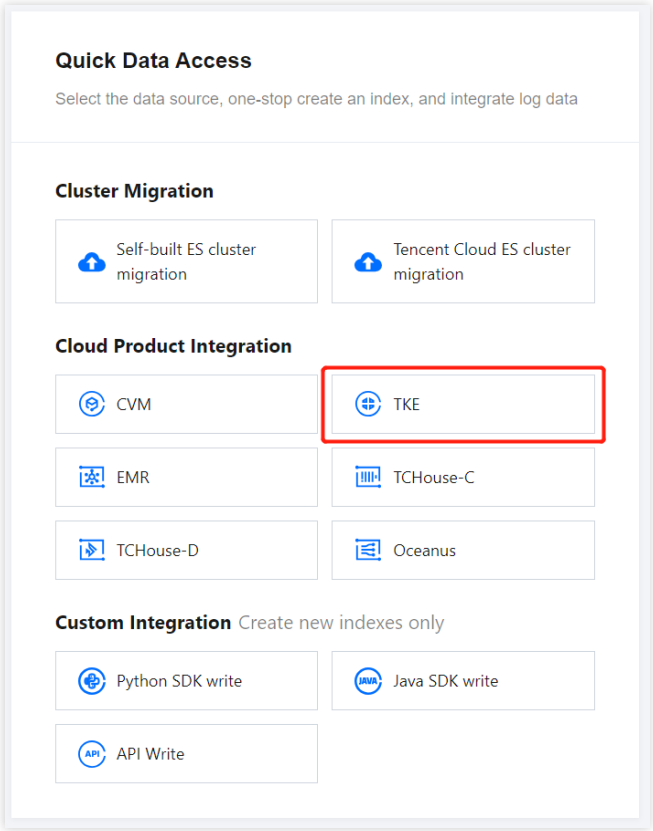


### Data Source Settings

**Region**: The region where the TKE cluster is located.

**VPC**: Required. The VPC where the TKE cluster is located.

**TKE Cluster ID to Be Collected**: Required. The ID of the TKE cluster to collect logs from, which should be in a running status and a standard cluster. If you need log collection for Serverless clusters (EKS), contact us via submiting a ticket.

**Based on Namespace/Host Path**: Required. For **Namespace**, select **Include**/**Exclude** from the first dropdown, and select one or more namespaces from the second dropdown (multi-select supported, but excluding all namespaces is not allowed). For host path-based collection, enter the **Absolute Path** on the host, for example, **/var/log/*.log**.

**Pod Tag**: Optional. You can create multiple Pod labels, which are logically connected using AND.

**Container Name**: Optional. The specified container name should be within the target cluster and namespace. If it is left empty, Filebeat will collect all containers within the namespace that match the specified Pod tags.

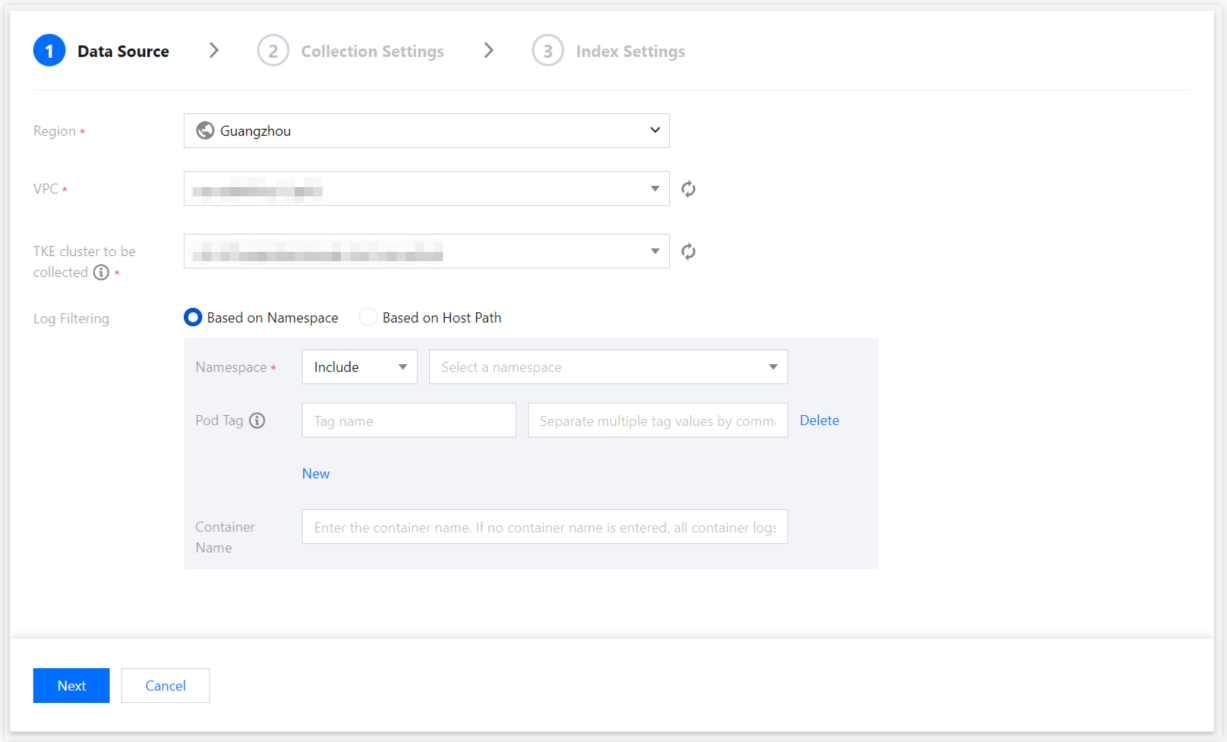


## Collection Settings

### Basic Settings

**Collection Policy:** Supports full collection and incremental collection. Once created, the collection policy cannot be modified. Full collection will collect historical logs as well as logs generated after the Filebeat configuration takes effect. Incremental collection will only collect logs generated after the Filebeat configuration becomes active.

### Collection Parsing

**Collection Template**: If you need a quick setup or are testing, you can select a collection template based on your log output format. After confirmation, return to the interface to modify the log sample with actual log data, enabling a fast completion of the collection parsing settings.

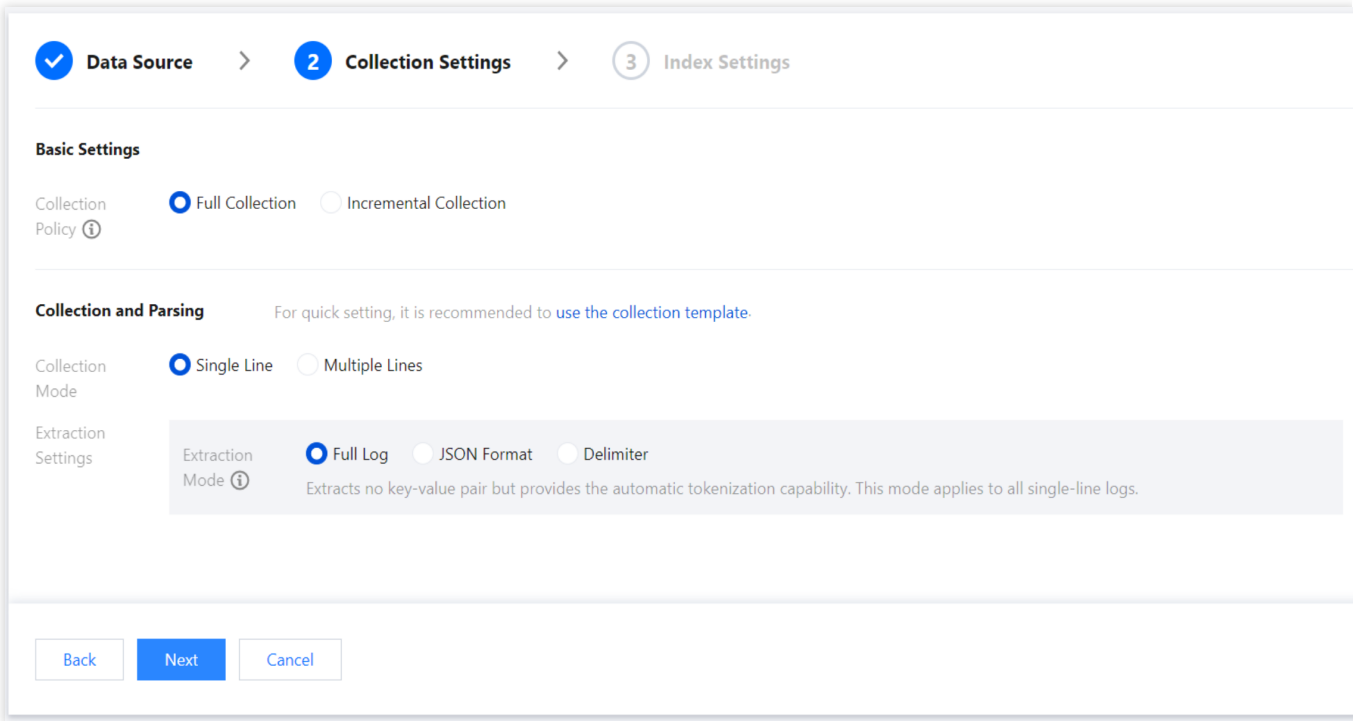


**Collection Mode**: Supports both single-line and multi-line modes; once set, the mode cannot be modified.

Single-line text log: Each line in the log file represents a single log entry, separated by a newline character.

Multi-line text log: Each log entry spans multiple lines, such as Java stack traces. In this mode, you need to configure a log sample and a regex pattern for line beginnings. Filebeat uses the regex to identify the start of each log entry, treating unmatched lines as part of the current log entry until the next matched line beginning appears. Once you enter a log sample, the system automatically generates a default regex pattern for line beginnings. You can also customize this pattern, with highlighted text in the input box indicating the matched line beginnings.

**Note:**

Ensure that actual scenario logs are used to facilitate automatic extraction of the leading line regular expression.

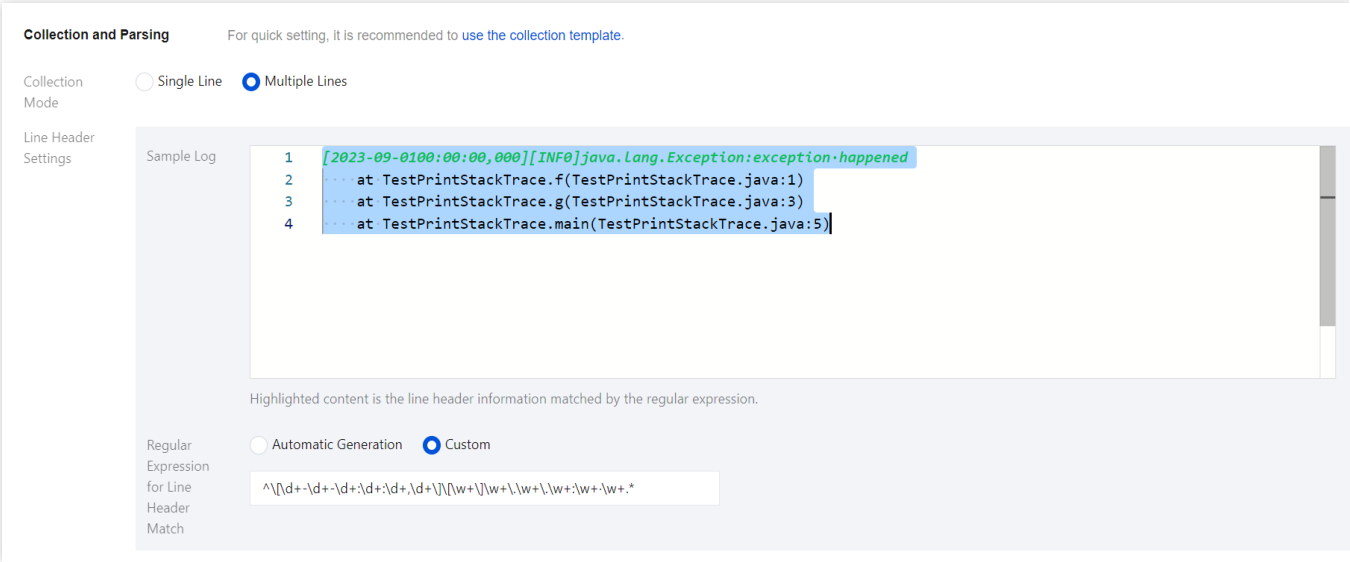


**Extraction Settings**: Extraction mode can be set to full-text log, JSON format, or delimiter-based. Once set, the extraction mode cannot be modified. Details are as follows:

Full-text log

JSON format

Delimiter

No key-value extraction is performed. The log content is stored in a field named message, which can be retrieved and analyzed using automatic tokenization.

For example, a single-line raw log might look like:

```
Tue Jan 01 00:00:00 CST 2023 Running: Content of processing something
```

When collected in the index, this data would appear as:

```
massage:Tue Tue Jan 01 00:00:00 CST 2023 Running: Content of processing
something
```

For logs in standard JSON format, we can extract fields based on the Key: Value pairs within the log.

For example, suppose a JSON log entry is as follows:

```
{"pid":321,"name":"App01","status":"WebServer is up and running"}
```

After structuring, the log entry will appear as follows:

```
{
    "pid":321,
    "name":"App01",
    "status":"WebServer is up and running"
}
```

For logs with content separated by a fixed delimiter, we can extract key-value pairs based on the specified delimiter. The delimiter can be a single character or a string, which can be selected or entered in the console.

For example, if a log entry is as follows:

```
321 - App01 - WebServer is up and running
```

With the delimiter set to -, this log entry will be split into three fields. Unique keys can then be assigned to these fields in the extraction results, as shown below:

```
pid: pid
name: App01
status: WebServer is up and running
```

**Extraction Results**: When the extraction mode is set to JSON format or Delimiter, a sample log can be provided for automatic extraction:
If the extraction mode is JSON format, the system will automatically populate the extracted Keys and Values. If it is deselected, the respective fields will not be written to the index.
If the extraction mode is Delimiter, the system will automatically populate the extracted Values, allowing you to assign unique Keys to each Value. If it is deselected, the corresponding fields will not be written to the index.
Built-in Fields: When you configure TKE log collection in the console, Filebeat will add information such as the log source and timestamp as Key-Value pairs in the logs. These fields are considered built-in fields. If a Key in your business log matches a built-in field name, the business log field content takes precedence, and the corresponding built-in field will not be added to the index. The meanings of the built-in fields are as follows:

| Built-in Field Name | Meaning |
| --- | --- |
| log.file.path | Path where the log is stored |
| kubernetes.pod.ip | IP address of the Pod containing the log |
| kubernetes.pod.name | Name of the Pod containing the log |
| kubernetes.node.hostname | Name of the host containing the log |
| @timestamp | Timestamp of when the log was collected |

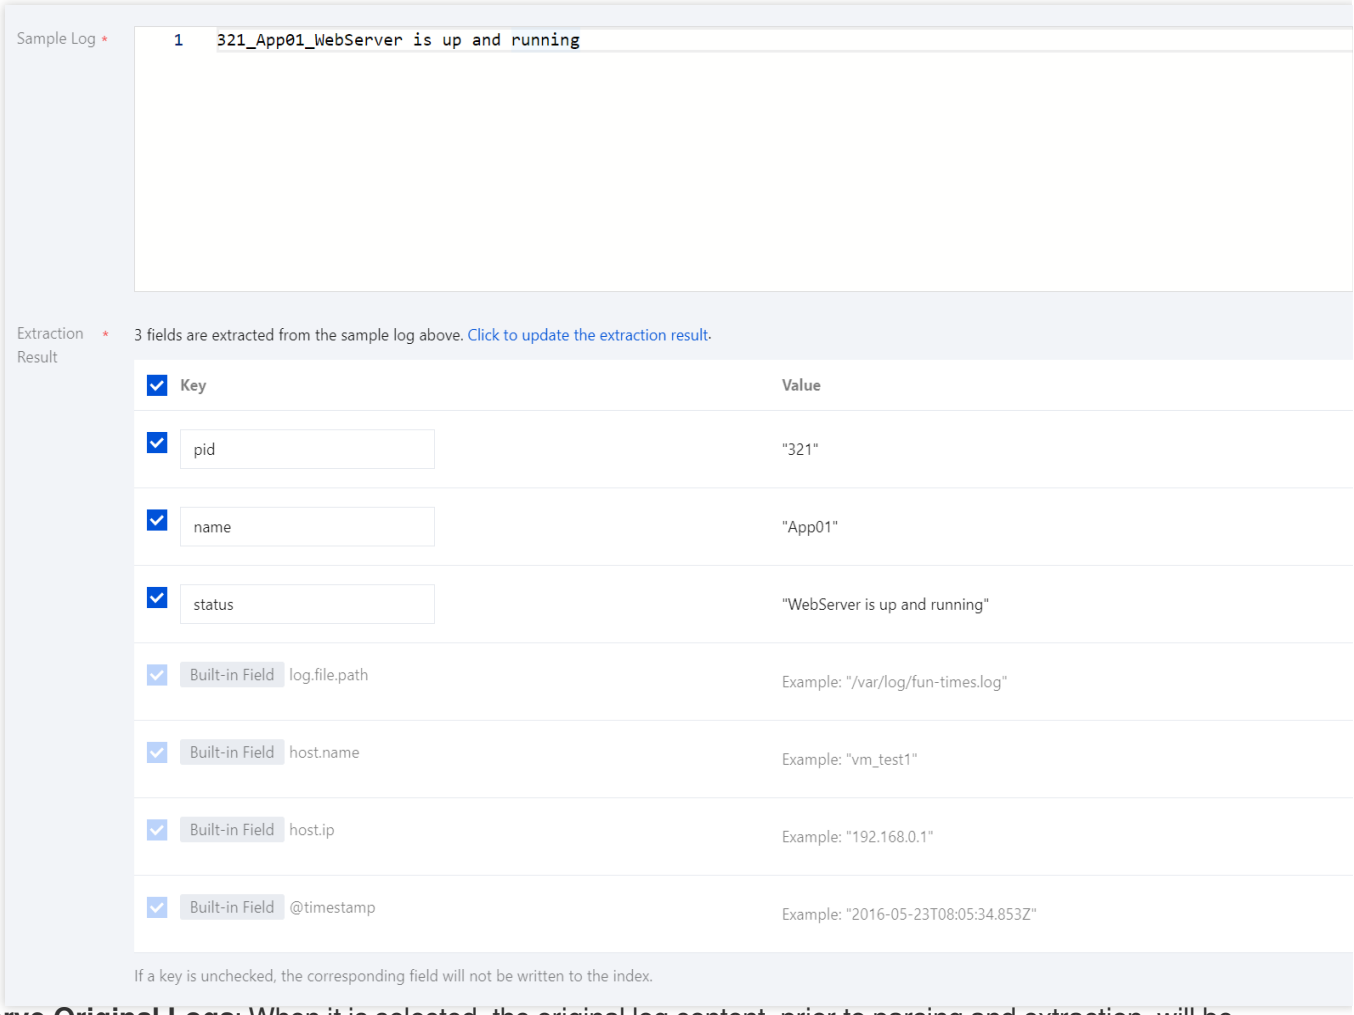


**Preserve Original Logs**: When it is selected, the original log content, prior to parsing and extraction, will be preserved in this field.

**Record Parsing Errors**: If the extraction mode is set to Delimiter, you can choose whether to log parsing errors. When it is selected, any errors encountered during parsing will be uploaded to this field as values.

### Index Settings

**Project Space**: You can assign the index to a specific project space for easier management of related business indexes.

**Index Name**: Supports a length of 1 to 100 characters, including lowercase letters, numbers, -, _, ;, @, &, =, !, ', %, $, ., +, (, ).

### Field Mapping

Dynamic generation: Enabled by default. When it is enabled, the system automatically parses the incoming data to generate the field mappings for the index.

Input sample auto-configuration: If **Dynamic Generation** is disabled, you can use **Input Sample Auto-Configuration** to generate the field mappings. Input a sample in JSON format, and the system will automatically validate it. Once validated, the relevant fields will be mapped in the field mapping table.

 

The field mapping divides the original data into distinct terms by fields (key:value) to construct the index, enabling retrieval based on this mapping. Specific details are as follows:

| Parameter | Feature Description |
| --- | --- |
| Field name | The name of the field within the data being written. |
| Field type | The data type of the field. The interface supports 9 types: text, date, boolean, keyword, long, double, integer, ip, and geo_point. Additional field types can be supported in **JSON Editing Pattern**. For more details, see Official Documentation. |
| Includes Chinese | Enable this option if the field includes Chinese characters that need to be retrieved. When it is enabled, the text field will use the ik_max_word tokenizer by default. |
| Enable index | When it is enabled, an index is built for this field to facilitate retrieval. |
| Enable statistics | When it is enabled, this field's values can be analyzed statistically, which will increase index storage. |

**Time Field**

The time field refers to a field of type date in the actual data. Once the index is created, this field cannot be modified.

**Note:**

The time field has **indexing** and **statistics** enabled by default, and these settings cannot be disabled.

**Data Storage Duration:**

1.1 You can set the storage duration of the data. By default, it is set to retain data for 30 days, though you also have the option to set it to permanent storage.

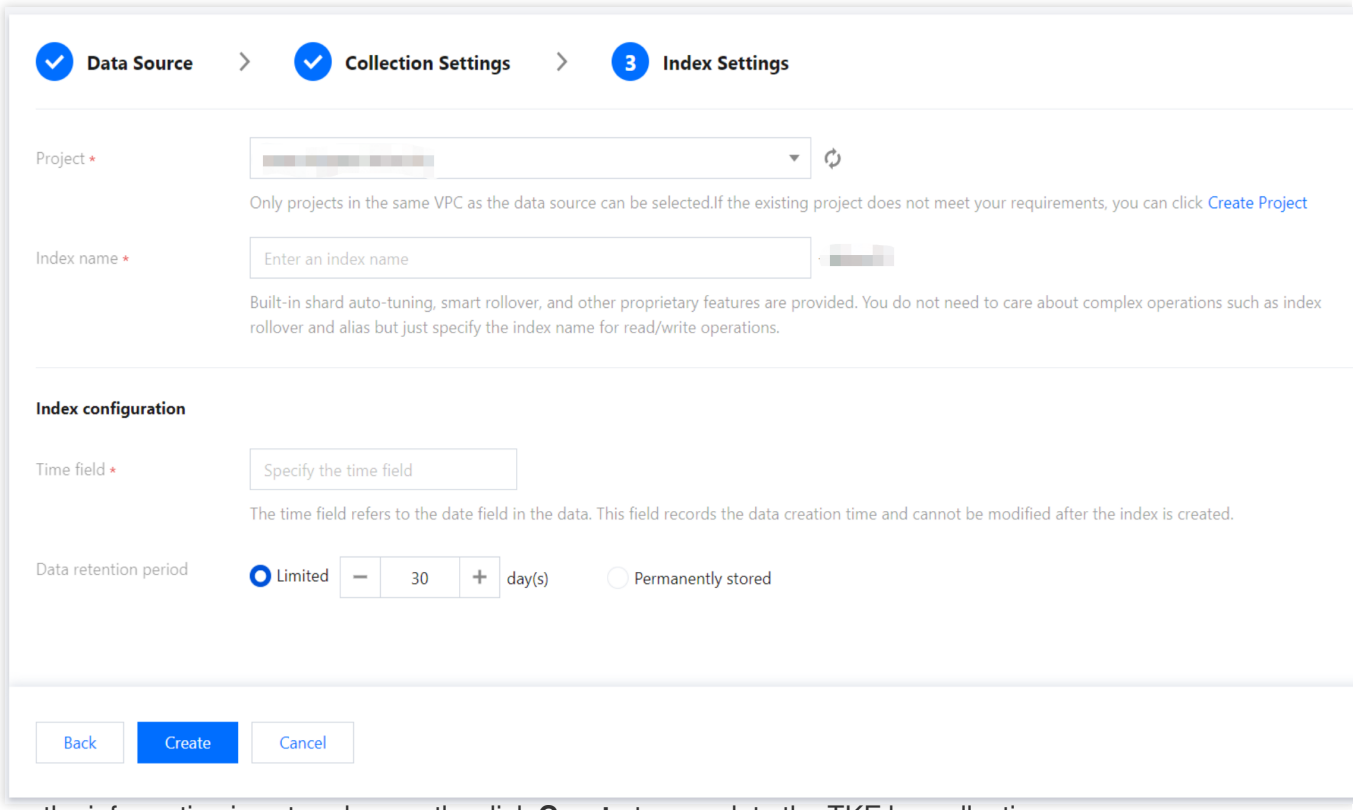


1.2 Once the information is entered correctly, click **Create** to complete the TKE log collection.

# Elastic MapReduce log access

Last updated：2024-12-04 16:09:26

## Prerequisites

A Tencent Cloud account has been created. For account creation, see Signing Up.
If logging in with a sub-account, ensure that the account has read and write permissions for ES.

## Operation Steps

### Logging in to the Console

1. Log in to the ES console.
2. In the left sidebar, choose **Log Analysis** under the Serverless mode.

### Creating a Project Space

1. click **Create a project**.
2. Enter a **Project Name** for the project, which can include 1 - 20 characters, consisting of Chinese characters, letters, digits, underscores, or delimiters (-).
3. Click **Confirm**. Once validated, the project space will be successfully created.

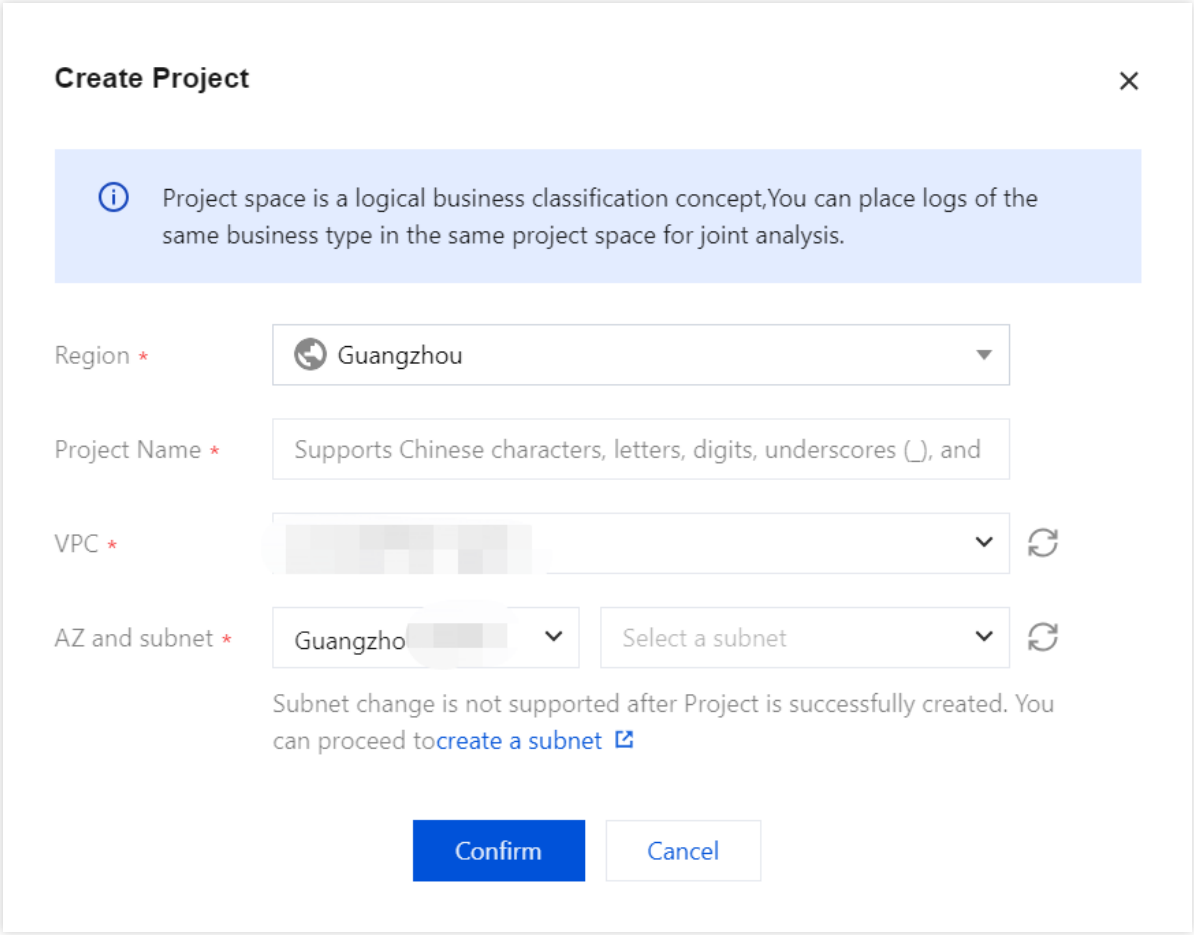


**Note**

In ES Serverless Log Analysis, you can simply Create an index, then use the API for data writing or access data sources such as CVM or TKE via the Data Access tab of the corresponding index. Alternatively, you can set up data access during index creation, enabling one-stop access for CVM logs, TKE logs, and Elastic MapReduce (EMR) logs. The following introduces the one-stop EMR log access process.

## Elastic MapReduce (EMR) Log Access

1. On the ES Serverless Log Analysis homepage, select **EMR** to enter the EMR Log Access page.

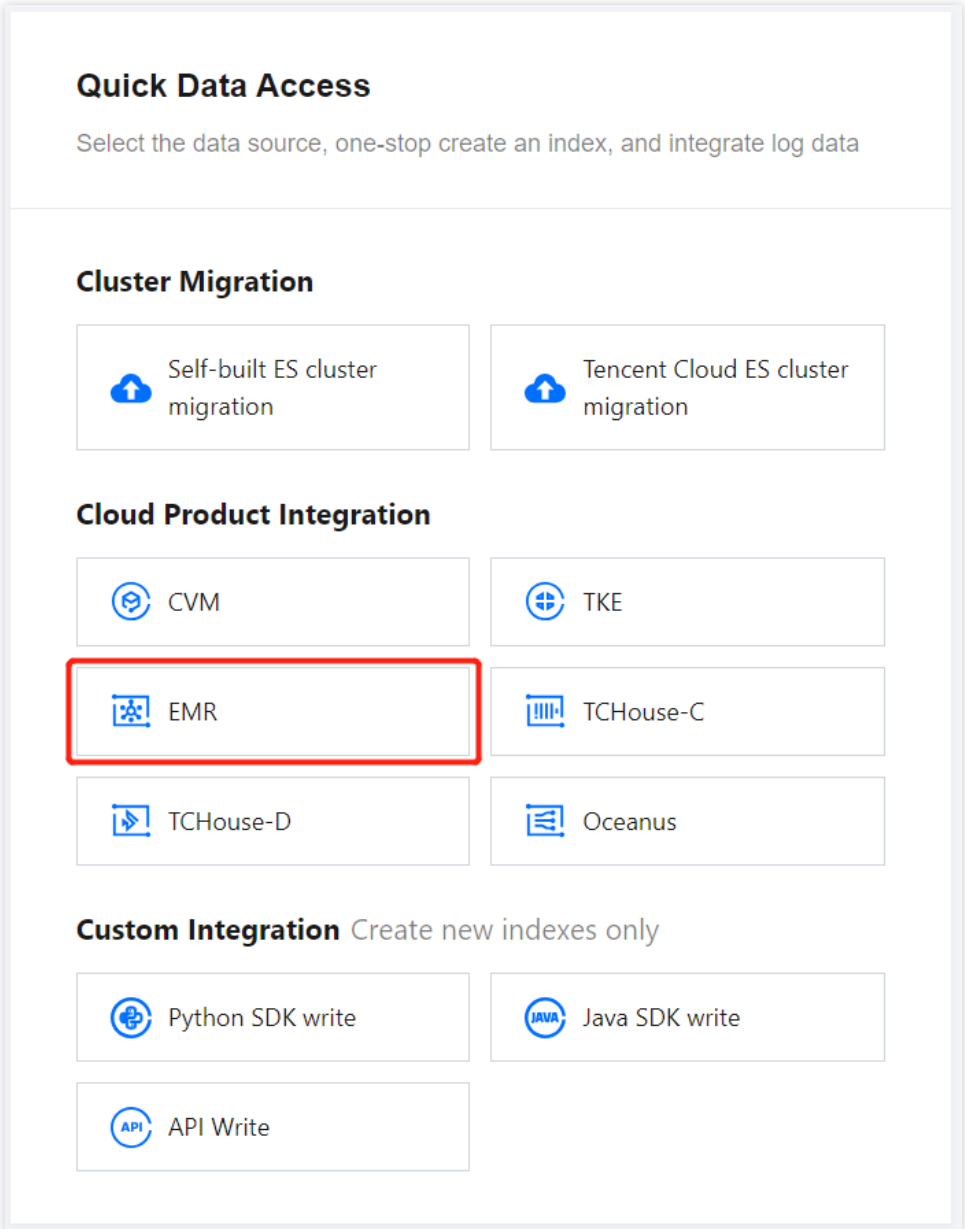


2. Enter the Data Source settings page, configure the data source, and click **Next** once setup is complete.

**Region**: Select the region where the EMR cluster is located. If you enter this page from a project space details page, the region aligns with the region of project space by default.

**VPC**: The Virtual Private Cloud where the EMR cluster is located.

**EMR Cluster**: The EMR cluster from which logs need to be collected.

**Log type**: Specify the component runtime logs to collect; for example, task logs can be collected if the YARN component exists in the cluster.

**Collection Policy**: Supports both full collection and incremental collection. Selecting incremental collection will only collect logs generated after data access setup.

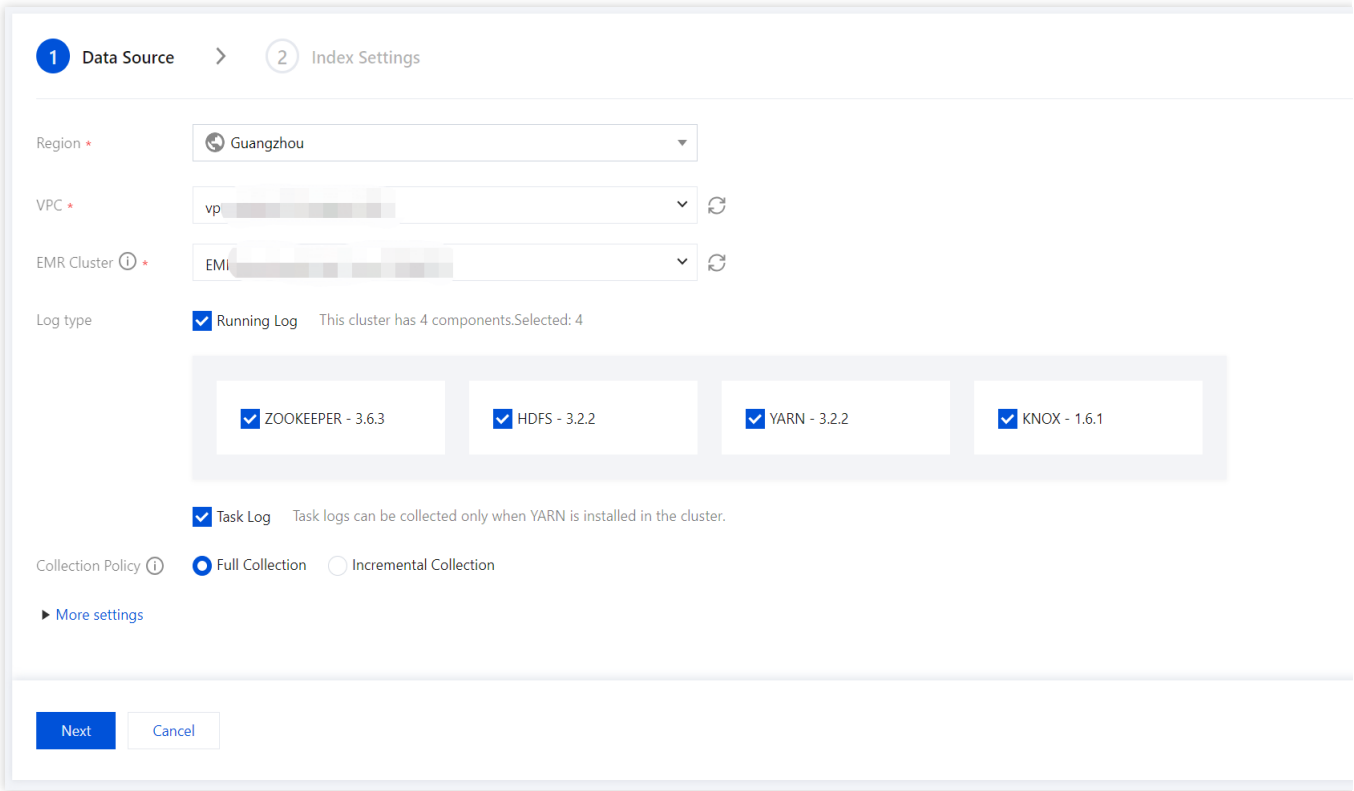


3. Enter the Index Settings page, and configure the index settings.

**Region**: The project space's region, which aligns with the EMR cluster's region by default.

**Project Space**: You can assign indexes for the same business to a specific project space for easier management.

**Index Name**: Length of 1 - 100 characters. Lowercase letters, digits, and the following symbols: -, _, ;, @, &, =, !, ', %, $, ., +, (, ) are supported.

**Field Mapping**

Dynamic Generation: Enabled by default. When enabled, it automatically parses and generates field settings for the index.

Input Sample Auto-Configuration: When **Dynamic Generation** is disabled, you can use **Input Sample Auto-Configuration** to generate field mappings for the index by entering a JSON-formatted data sample. After confirmation, the platform will validate the input; if the validation is successful, the relevant fields will be mapped in the field mapping table.

Field mapping divides the original data into multiple tokens based on fields (key:value) for indexing. Retrieval relies on this mapping, as detailed below:

| Parameter | Description |
| --- | --- |
| Field name | The field name in the written data. |
| Field type | The data type of the field; the interface supports 9 types: text, date, boolean, keyword, long, double, integer, ip, and geo_point. Additional field types are supported in **JSON mode**. For details, see Official Documentation. |
| | |

| | |
|---|---|
| Include Chinese | Enable this option if the field contains Chinese text and requires Chinese retrieval. When it is enabled, the ik_max_word tokenizer is applied to the text field by default. |
| Enable indexing | When it is enabled, this field will be indexed for search. |
| Enable statistics | When it is enabled, statistical analysis can be performed on the field values, which will increase index storage. |

**Time Field**

The time field refers to a field with the type date in the actual data. Once the index is created, this field cannot be modified.

**Note:**

By default, the time field enables **indexing** and **statistics**, and these settings cannot be disabled.

**Data Storage Duration**

You can set the data storage duration, with a default of 30 days, or select permanent storage.

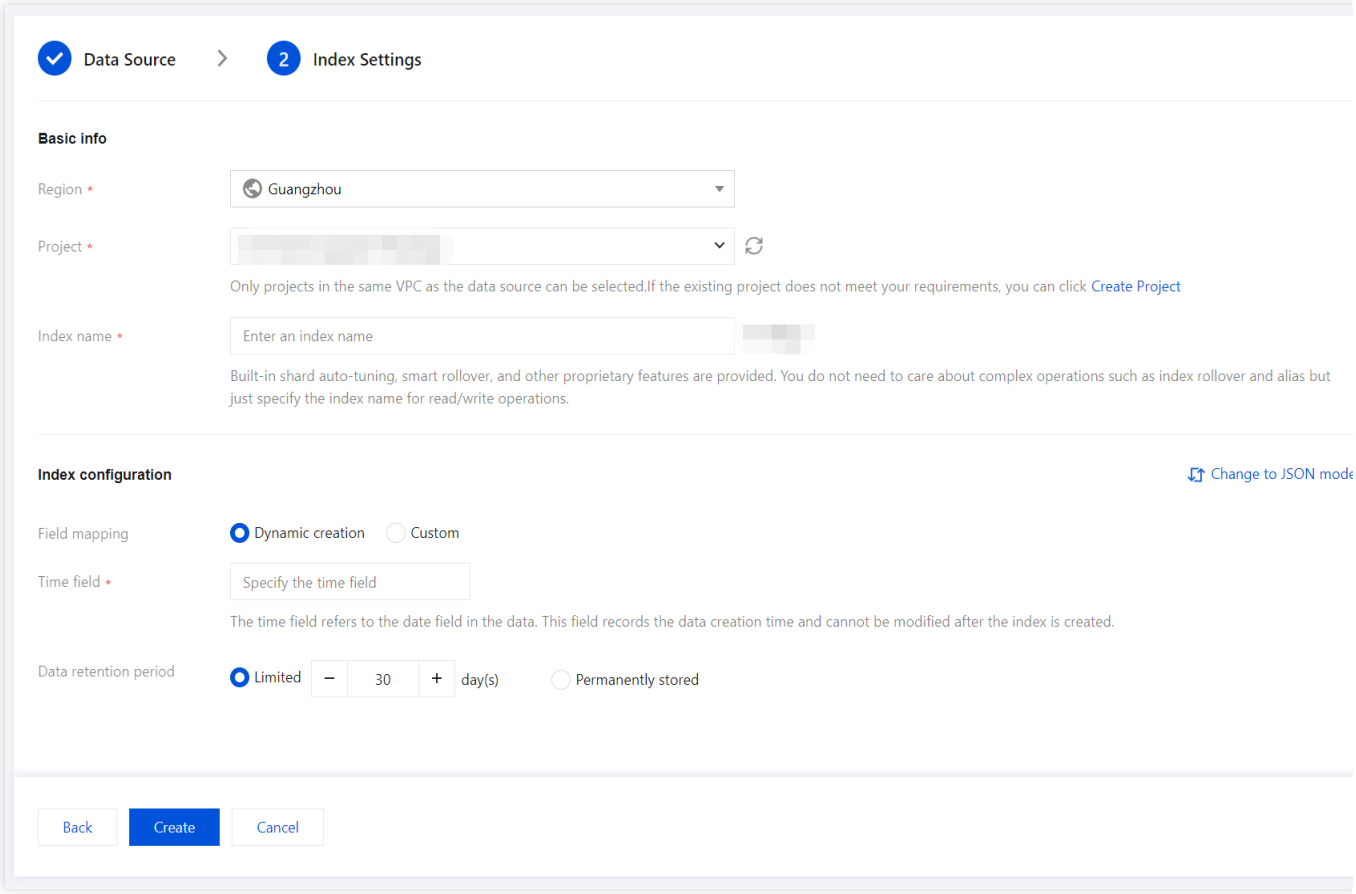

# TCHouse-D Cluster Log Access

Last updated：2024-12-04 16:13:12

The ES Serverless service supports collecting logs from the TCHouse-D nodes to facilitate troubleshooting and issue analysis. For more details, see Log Search.

# Customizing Filebeat Data Access

Last updated：2024-12-04 16:15:17

## Self-Built Filebeat Data Collection

```
Version Support
Only Filebeat versions 7.10.2 or 7.14.2 are supported.
```

| Category | Parameter | Description | Filling Instructions |
| --- | --- | --- | --- |
| Elasticsearch template setting | setup.template.enabled | Index template | Boolean type. It can be set to false; currently, this setting is not supported. |
| | setup.ilm.enabled | Index lifecycle management | Boolean type. It can be set to false; currently, this setting is not supported. |
| | allow_older_versions | Compatibility with ES versions | Boolean type. It can be set to true or false. |
| output | protocol | Data transmission protocol | String type. The default value is http, and it can also be set to https. |
| | hosts | Private network access address for index | Array type.<br>If the protocol is set to http, the port number should be 80. For example, it can be set as ["http://index-xxx.qcloudes.com:80"];<br>If the protocol is set to https, the port number should be 443. For example, it can be set as ["https://index-xxx.qcloudes.com:443"]. |

```
Configuration Description

# ============================== Filebeat inputs
==============================
```

```
filebeat.inputs:
- type: log
  # Change to true to enable this input configuration.
  enabled: true
  # Paths that should be crawled and fetched. Glob based paths.
  paths:
    - /var/log/*.log
# ============================== Filebeat modules
==============================

filebeat.config.modules:
  # Glob pattern for configuration loading
  path: \\${path.config}/modules.d/*.yml

  # Set to true to enable config reloading
  reload.enabled: false

  # Period on which files under path should be checked for changes
  #reload.period: 10s

# ======================== Elasticsearch template setting
========================
setup.template.enabled: false
setup.ilm.enabled: false
  #template setting's value is set to false by default. If you set it to true,
an error will be reported when the configuration is submitted


# ================================== General
==================================

# The name of the shipper that publishes the network data. It can be used to
group
# all the transactions sent by a single shipper in the web interface.
#name:

# The tags of the shipper are included in their own field with each
# transaction published.
#tags: ["service-X", "web-tier"]

# Optional fields that you can specify to add additional information to the
# output.
#fields:
#  env: staging
```

```
# ================================== Processors
==================================
processors:
  - add_host_metadata:
      when.not.contains.tags: forwarded

# ================================== Logging
===================================

# Sets log level. The default log level is info.
# Available log levels are: error, warning, info, debug
#logging.level: debug

# At debug level, you can selectively enable logging only for some components.
# To enable all selectors use ["*"]. Examples of other selectors are "beat",
# "publisher", "service".
#logging.selectors: ["*"]
############################# output ############################################
output.elasticsearch:
  # Array of hosts to connect to.
  allow_older_versions: true
  protocol: "http"
  hosts: ["Private network access address for index"]

  # Authentication credentials - either API key or username/password.
  username: "your index username"
  password: "your index password"
  indices:
  - index: The_index_name
    when.equals:
      fields.type: log
```

# Access Control

Last updated：2024-12-04 16:21:56

1. In the Project list, click the project name/ID to enter the Basic Info page.

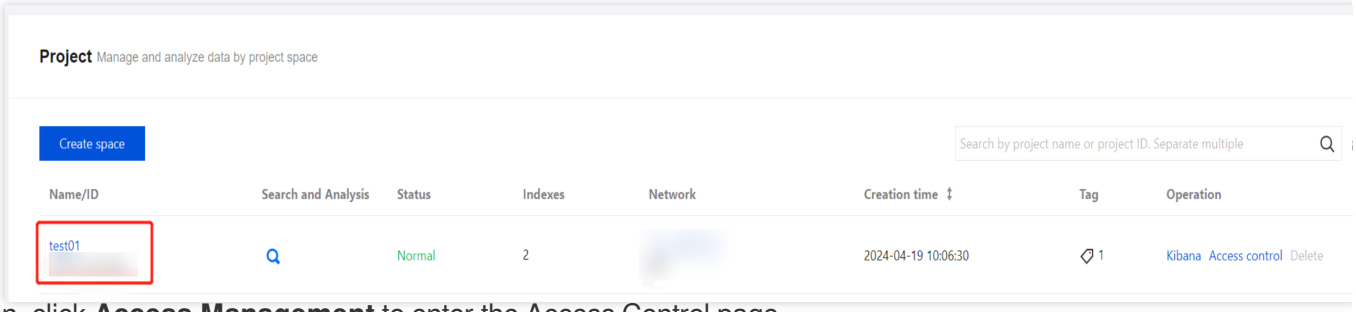


2. Then, click **Access Management** to enter the Access Control page.

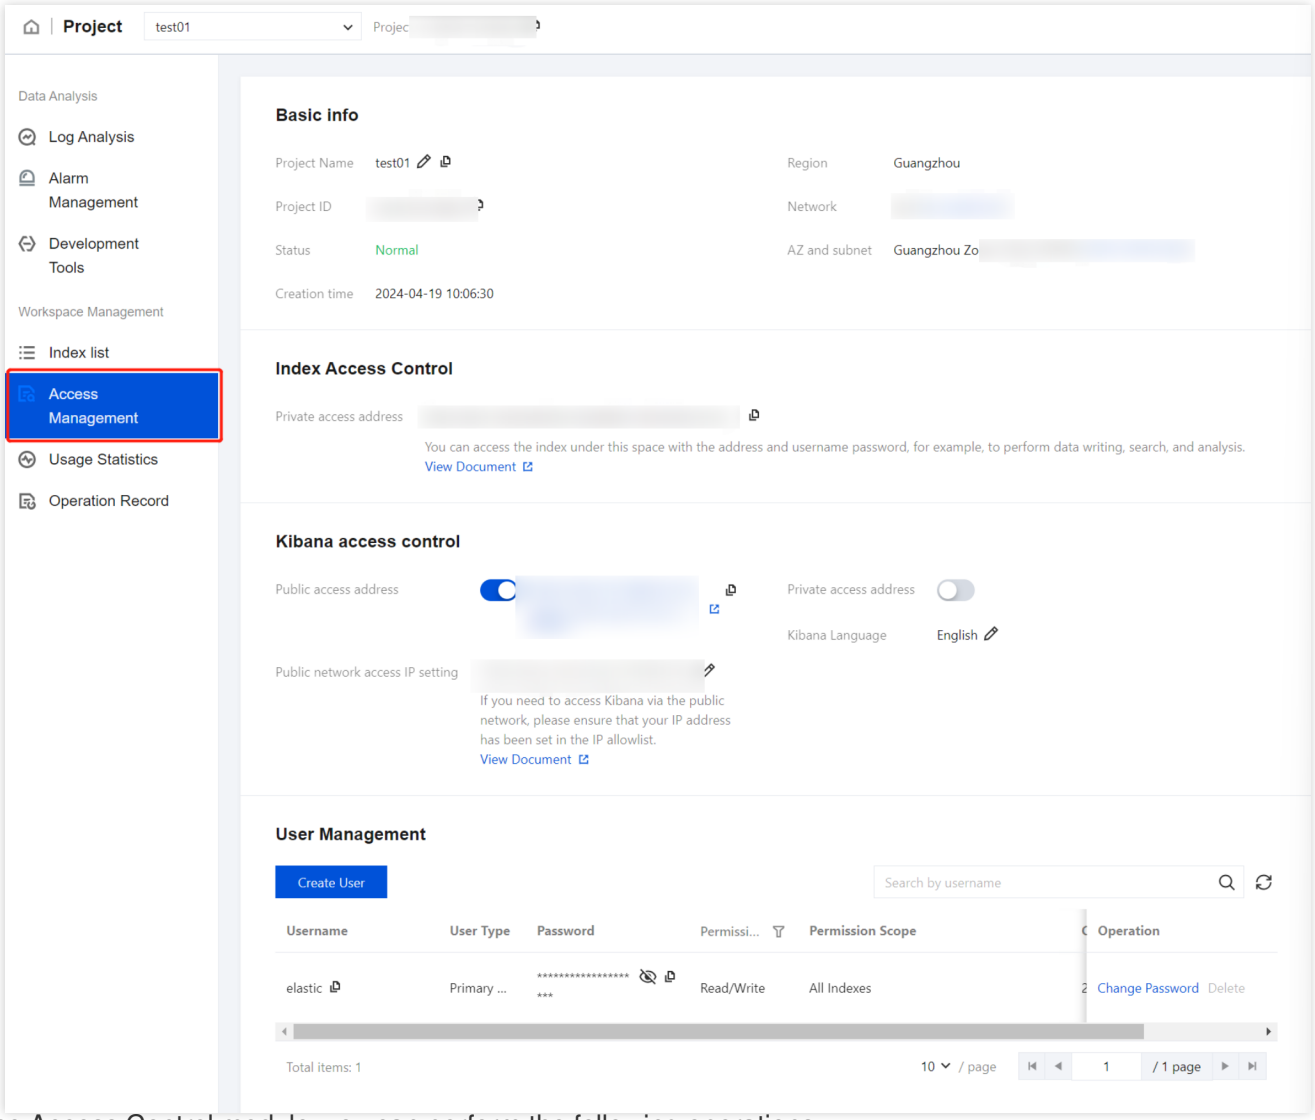


3. In the Access Control module, you can perform the following operations:

View the project space's private network access address, **which can be used for data writing or querying**.

Enable or disable Kibana private network access or public network access.
Modify the allowlist of Kibana public network access addresses. Multiple IP addresses are supported, and are separated by commas, semicolons, or line breaks, in format such as 192.168.0.1,192.168.0.0/24, with a maximum of 50 entries. If you are not aware of the current IP address, click **Click to automatically access the current IP address** to obtain and enter it automatically.

**Note**

Setting 127.0.0.1 means blocking access for all IPv4 addresses. For security, setting the IP allowlist to 0.0.0.0 is not permitted. If you have special requirements, submit a ticket for assistance.

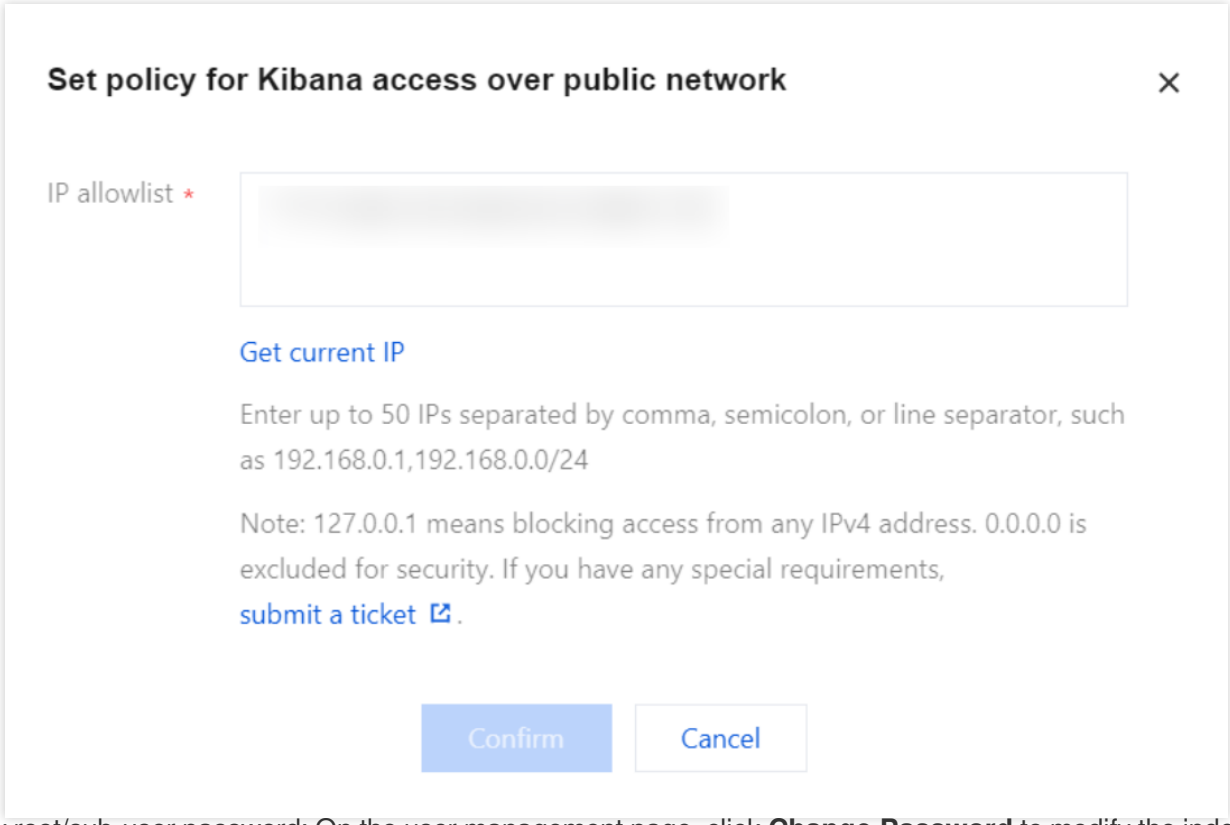


Modify root/sub-user password: On the user management page, click **Change Password** to modify the index access password.

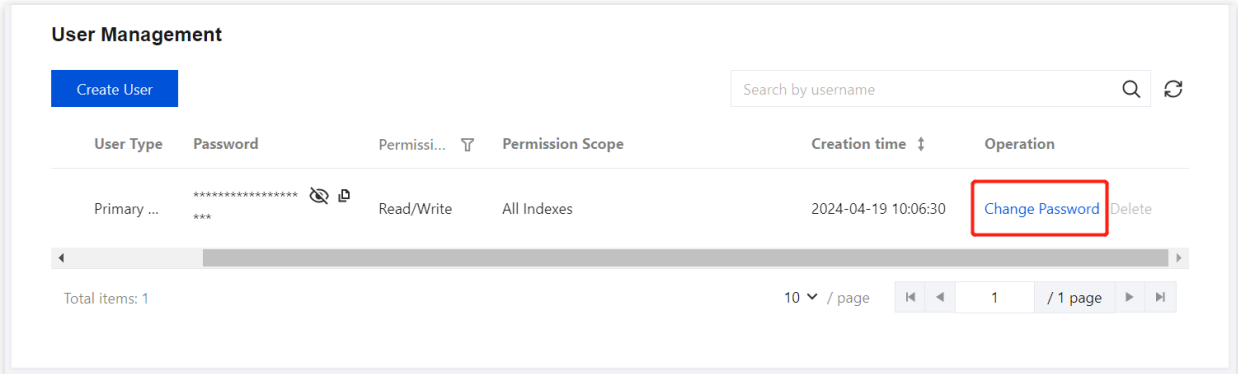


Modify sub-user permissions: On the user management page, click **Modify Permissions**, then select the permission type and scope. Supported permission types include read-only and read-write, and you can select from all indexes

within this space via a dropdown menu.

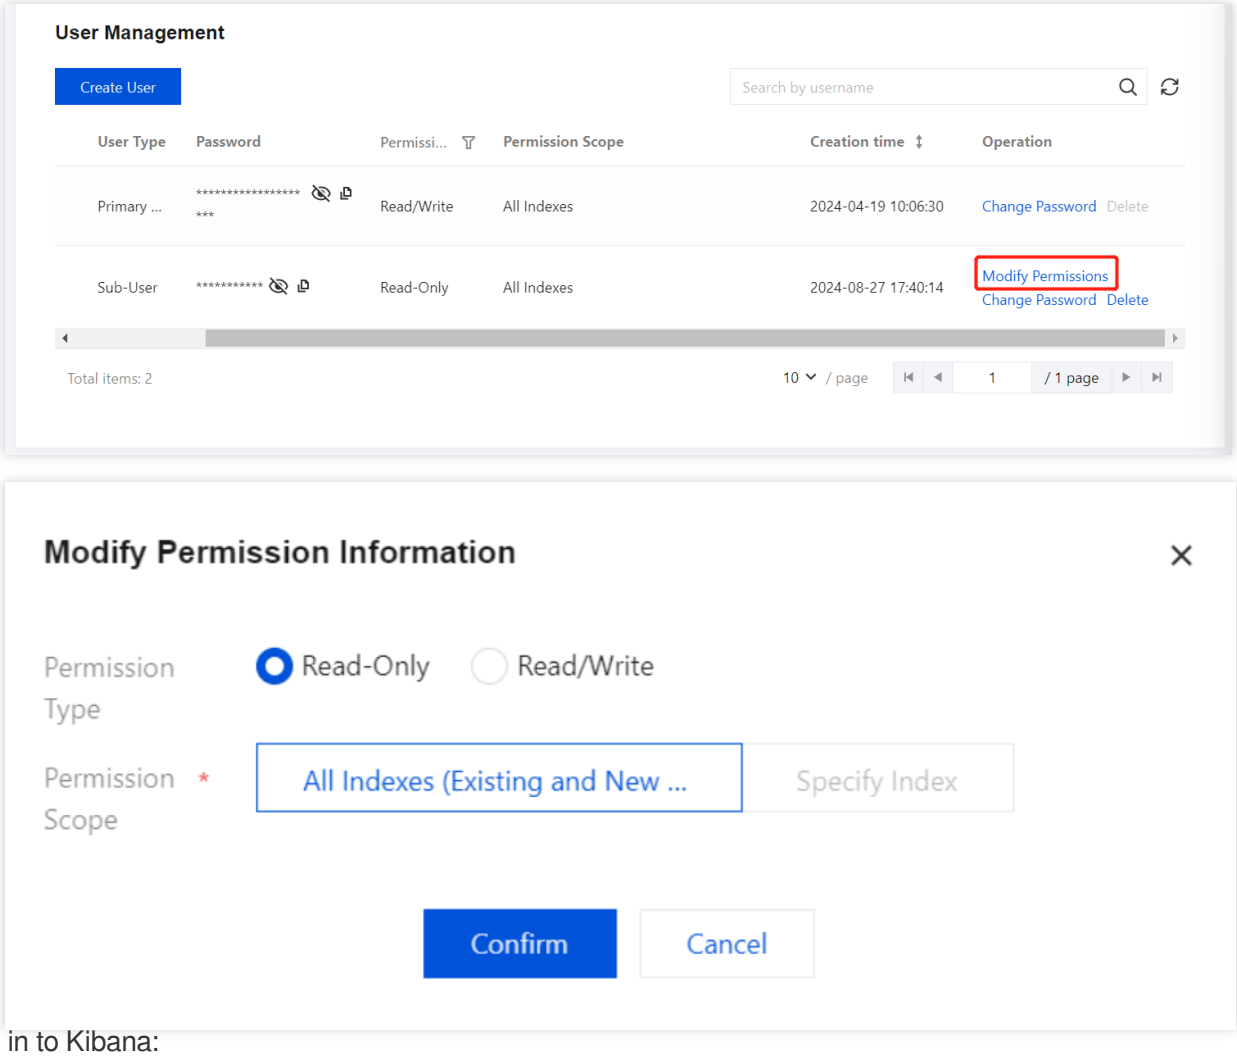


4. Log in to Kibana:

After enabling Kibana public network access and configuring the IP allowlist, click the Kibana public access address to open the Kibana login page. Enter the sub-user's username and password for this space, then click **Log in** to access Kibana.

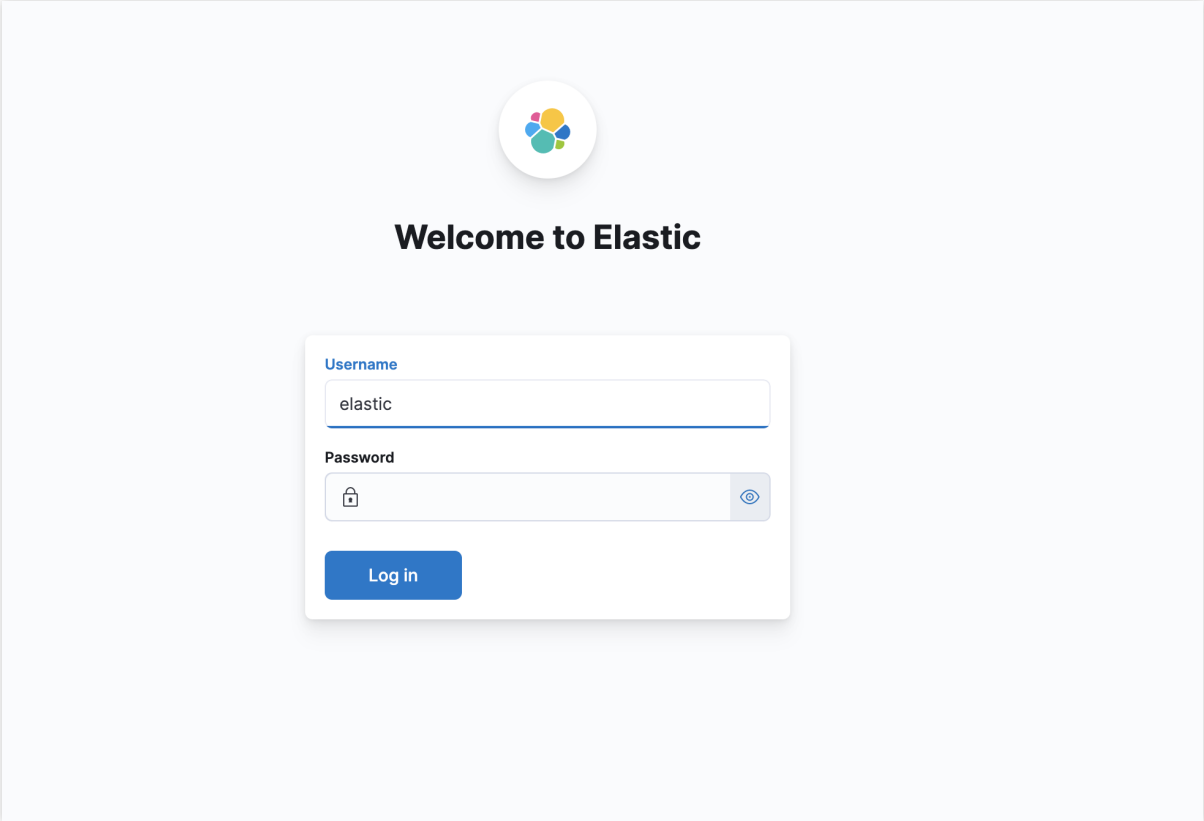

# Writing Data

Last updated：2024-12-04 16:25:16

## Overview

The ES Serverless service supports writing data into indexes through methods such as **ES native APIs, Logstash, Flink, and Kafka**. If you require log collection for services such as CVM, TKE, or TCHouse-C, a one-stop visualized configuration option is also available. By simply setting up data sources and index information, you can collected logs into the indexes for efficient retrieval and analysis. This document provides instructions for writing a single document and writing document in batches using **Kibana** and **Curl commands**.

## Access Control

1. In the Project list, click the corresponding project name to enter the Basic Info page.

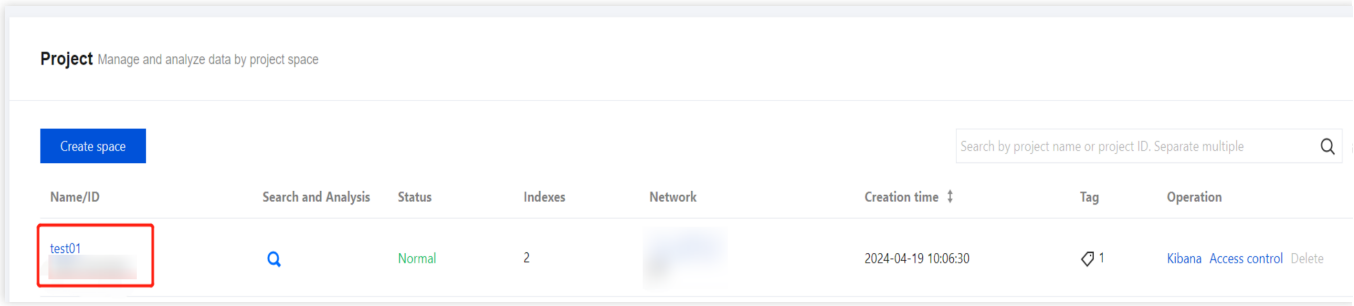


2. In the Access Control module, you can view the username and password for the index, private network access address, Kibana private network access address, and Kibana public network access address. You can also configure the Kibana public network access policy.

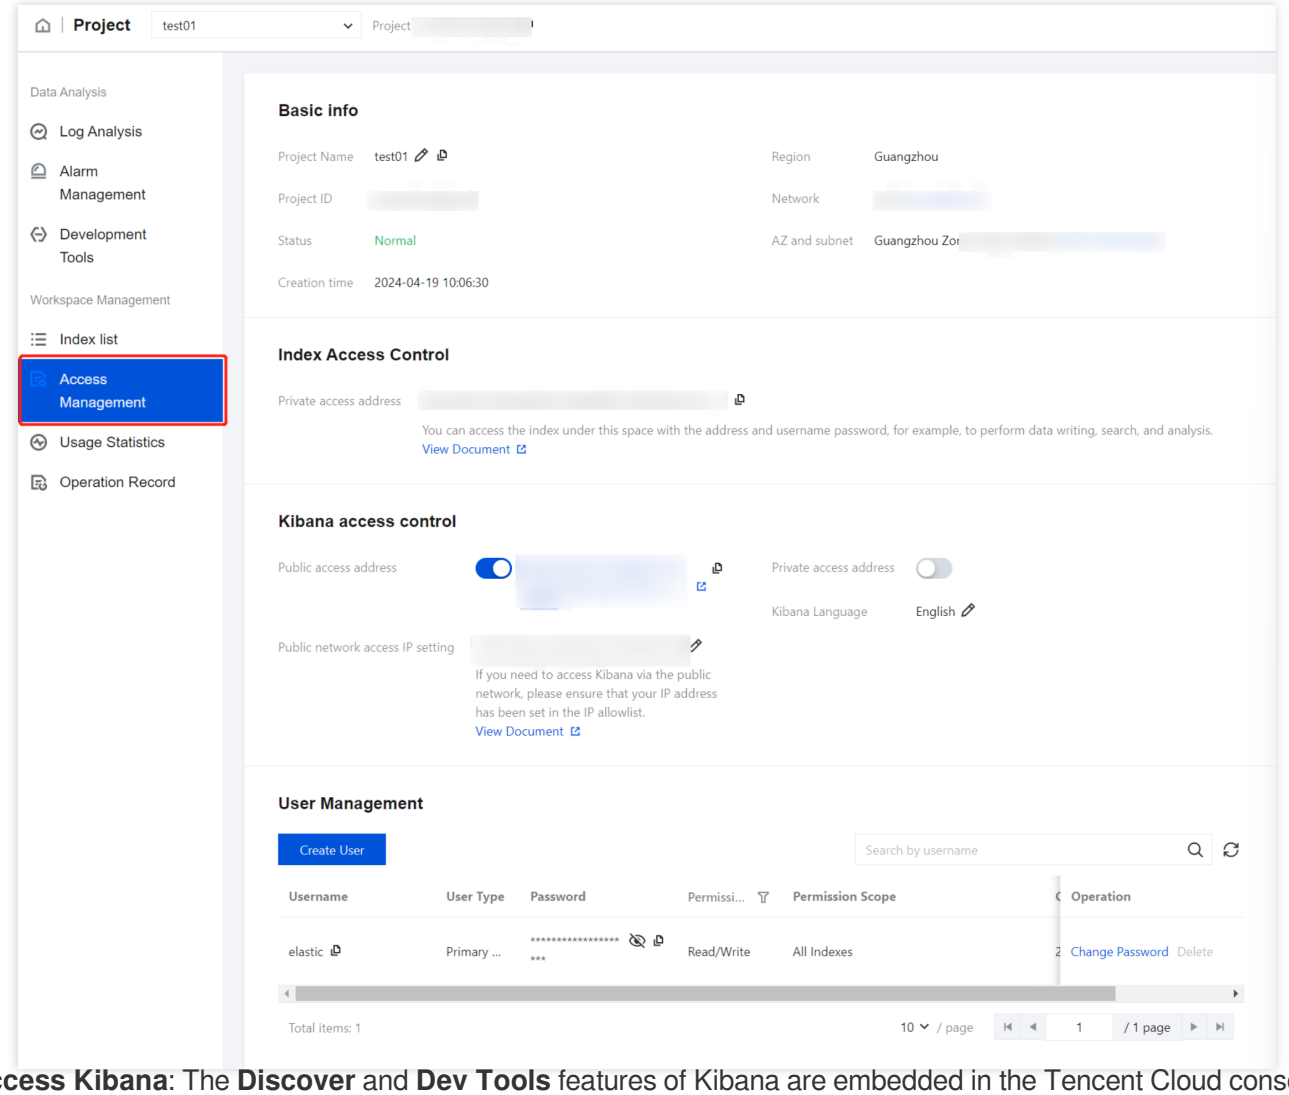


3. **Access Kibana**: The **Discover** and **Dev Tools** features of Kibana are embedded in the Tencent Cloud console, allowing you to use retrieval and analysis capabilities directly within the console or access Kibana via an external link. **Via Console**: In the sidebar of the space details page, click Search and Analysis to enter the relevant page. You can switch between index views by clicking the index pattern dropdown on the left side. **Log Search** corresponds to **Discover**, and **Development Tools** corresponds to **Dev Tools**.

**Note:**

Embedded features require third-party cookies to be enabled in your browser. If you encounter issues, please enable third-party cookies in your browser settings.

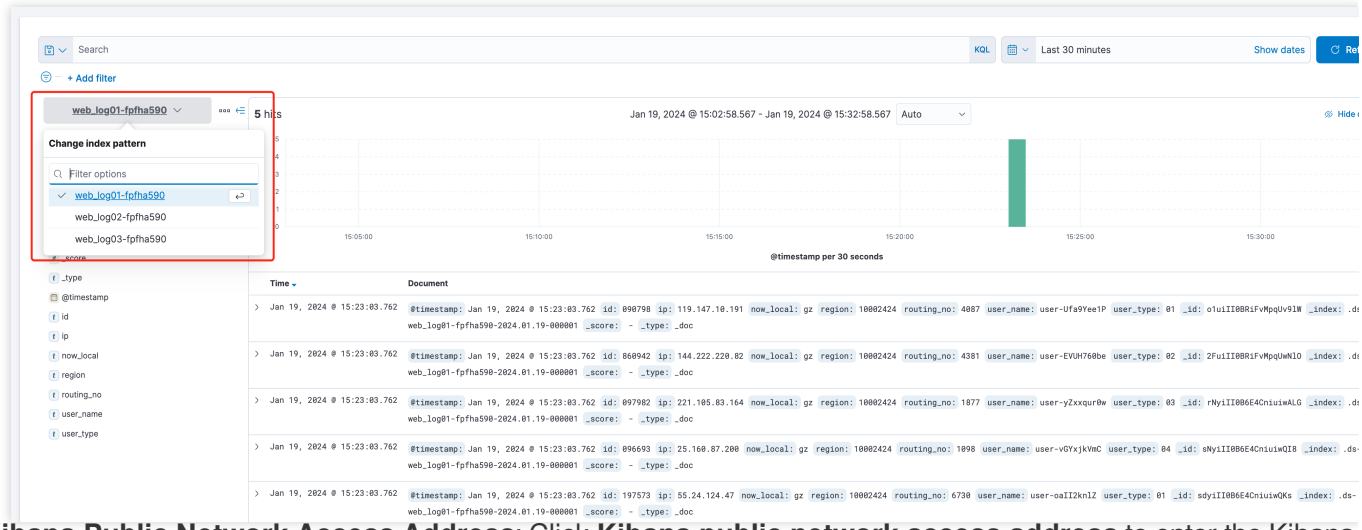


**Via Kibana Public Network Access Address**: Click **Kibana public network access address** to enter the Kibana page.

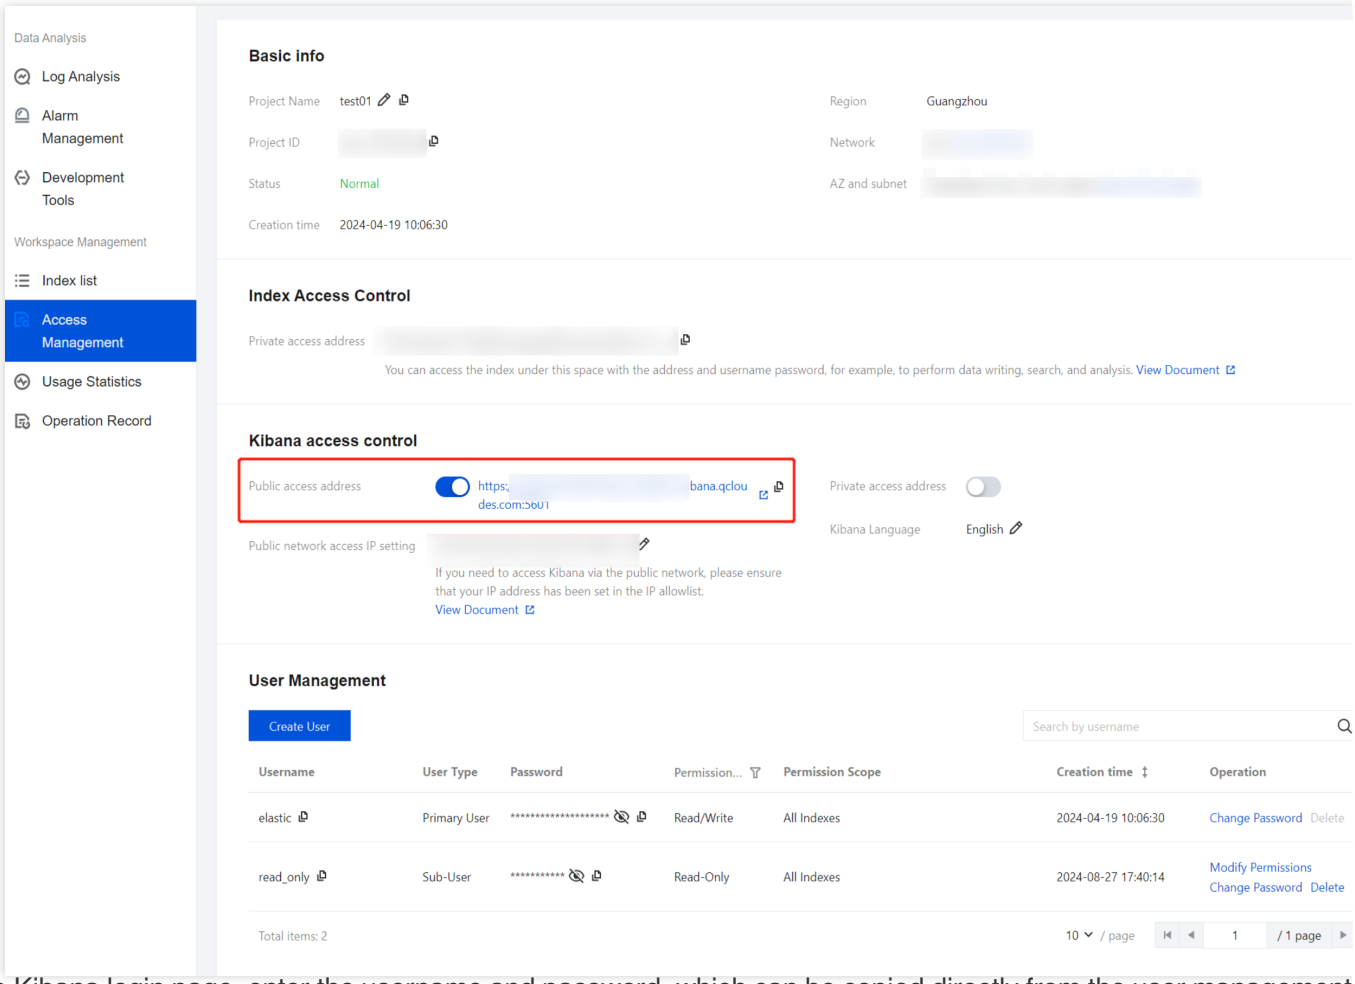


On the Kibana login page, enter the username and password, which can be copied directly from the user management page.

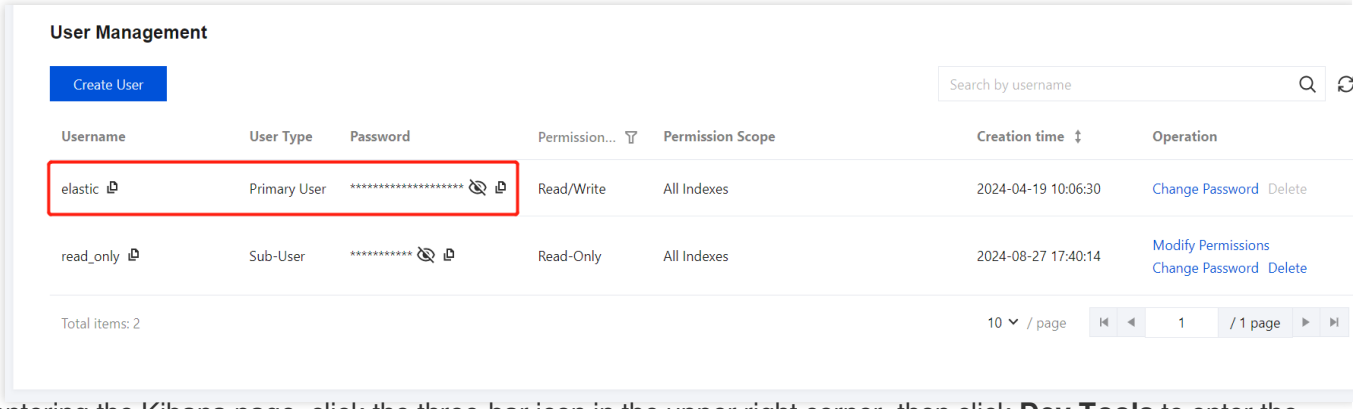


After entering the Kibana page, click the three-bar icon in the upper right corner, then click **Dev Tools** to enter the development tools page.

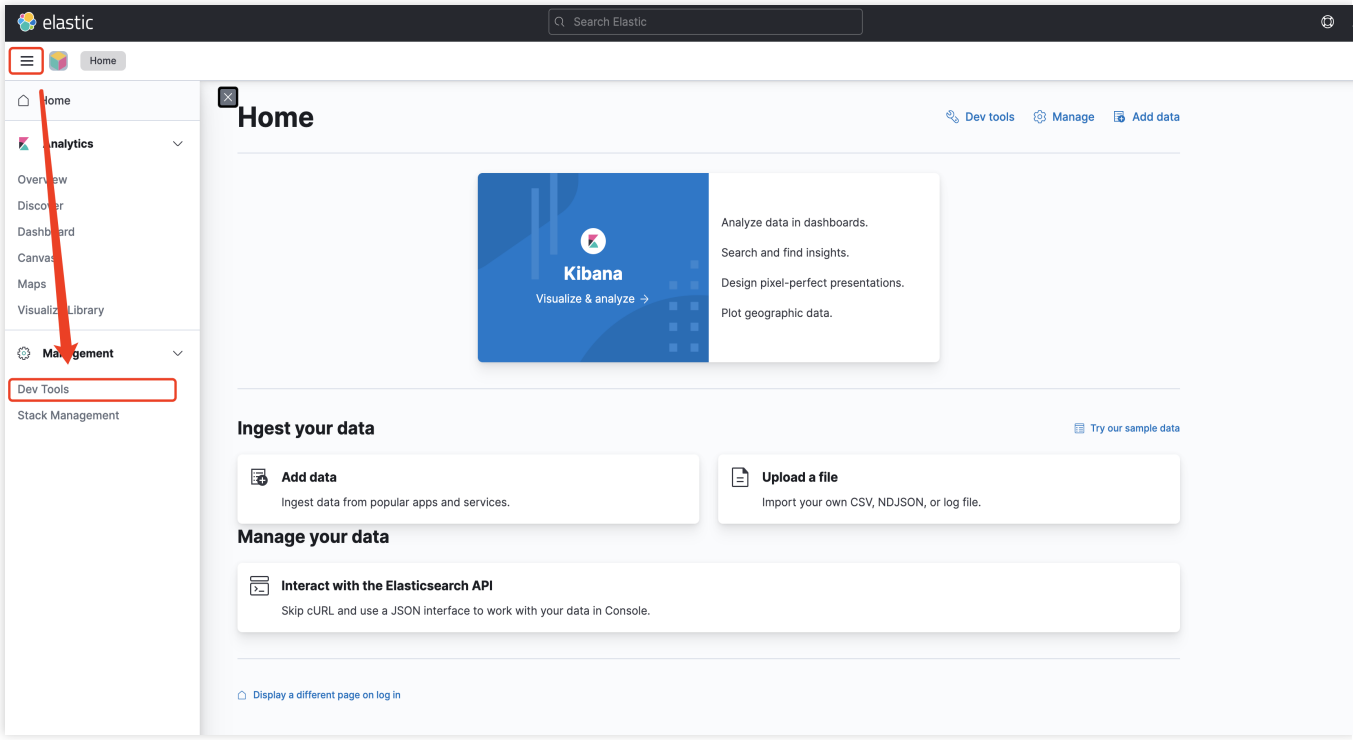


**Note:**

Kibana public network access includes an allowlist mechanism, meaning that IP addresses not included in the access policy cannot access Kibana, enhancing access security. If the page displays Sorry, you do not have permissions to access, click **Kibana Public Network Access Policy** as shown above. In the pop-up window, click **Get current IP** to enter your current IP address to the allowlist.

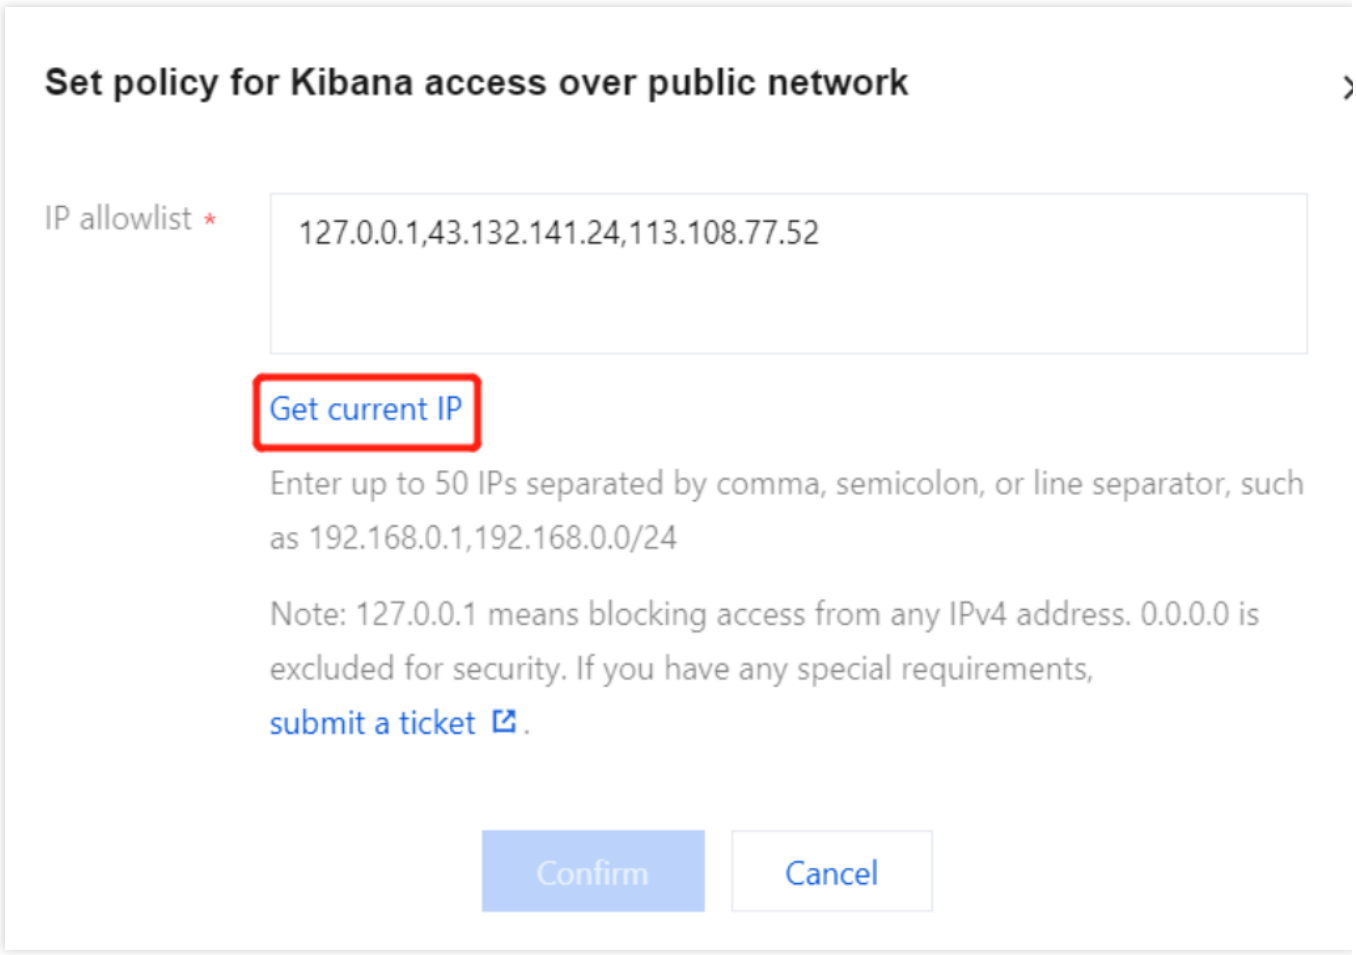


# Writing a Single Document

### Via Kibana Dev Tools

```
POST /index name/_doc
{
"@timestamp": "2023-09-28T11:06:07.000Z",
"user":{
    "id" : "8a4f500"
  },
"message": "Login successful"
}
```

### Via Command Line

```
curl -X POST "project space access address/index name/_doc/?pretty" -H
'Content-Type: application/json' -d'
{
  "@timestamp": "2023-09-28T11:06:07.000Z",
  "user": {
```

```
    "id": "8a4f500d"
  },
  "message": "Login successful"
}
```

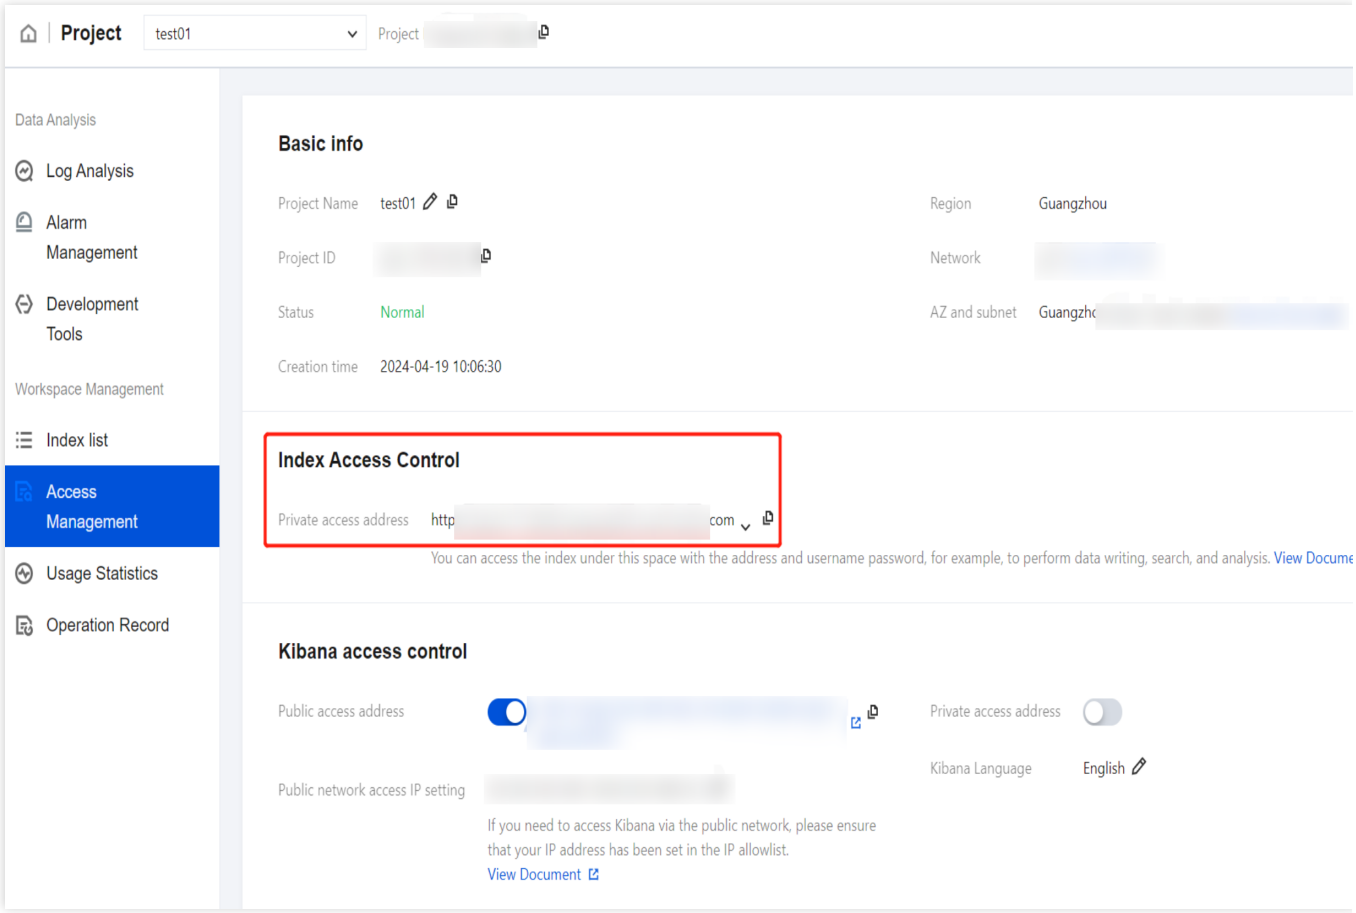


### Caution

The `PUT /index name/_doc/document ID` format cannot be used for writing requests. To specify a document ID, use `PUT /index name/_create/document ID` .

Ensure that the written data includes the **Time Field** set during index creation.

# Writing Document in Batches

### Via Kibana Dev Tools

```
PUT /index name/_bulk?refresh
{"create":{ }}
```

```
{ "@timestamp": "2023-03-28T11:04:05.000Z", "user": { "id": "vlb44hny" },
"message": "Login attempt failed" }
{"create":{ }}
{ "@timestamp": "2023-03-29T11:06:07.000Z", "user": { "id": "8a4f500d" },
"message": "Login successful" }
{"create":{ }}
{ "@timestamp": "2023-03-30T11:07:08.000Z", "user": { "id": "l7gk7f82" },
"message": "Logout successful" }
```

### Via Command Line

```
curl -X PUT "project space access address/index name/_bulk?refresh&pretty" -H
'Content-Type: application/json' -d'
{"create":{ }}
{ "@timestamp": "2023-03-28T11:04:05.000Z", "user": { "id": "vlb44hny" },
"message": "Login attempt failed" }
{"create":{ }}
{ "@timestamp": "2023-03-29T11:06:07.000Z", "user": { "id": "8a4f500d" },
"message": "Login successful" }
{"create":{ }}
{ "@timestamp": "2023-03-30T11:07:08.000Z", "user": { "id": "l7gk7f82" },
"message": "Logout successful" }
'
```

**Caution**

The bulk operation only supports `create` .

Ensure that the written data includes the **Time Field** set during index creation.

# Data Query

Last updated：2024-12-04 16:26:57

## Overview

This document introduces data query operations using the **Kibana** and **Curl command line** methods.

## Access Control

1. In the space list, click the corresponding **Project Name/ID** to enter the Basic Info page.

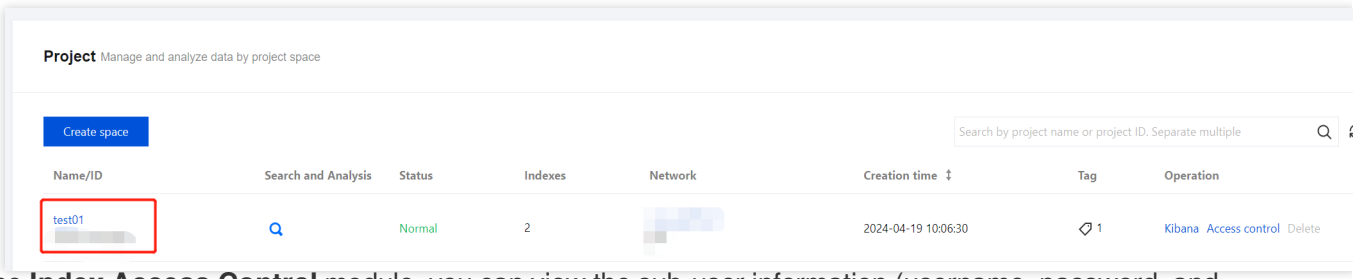


2. In the **Index Access Control** module, you can view the sub-user information (username, password, and permissions), private network access address, Kibana private network access address, and Kibana public network access address. You can also configure the Kibana public network access policy.

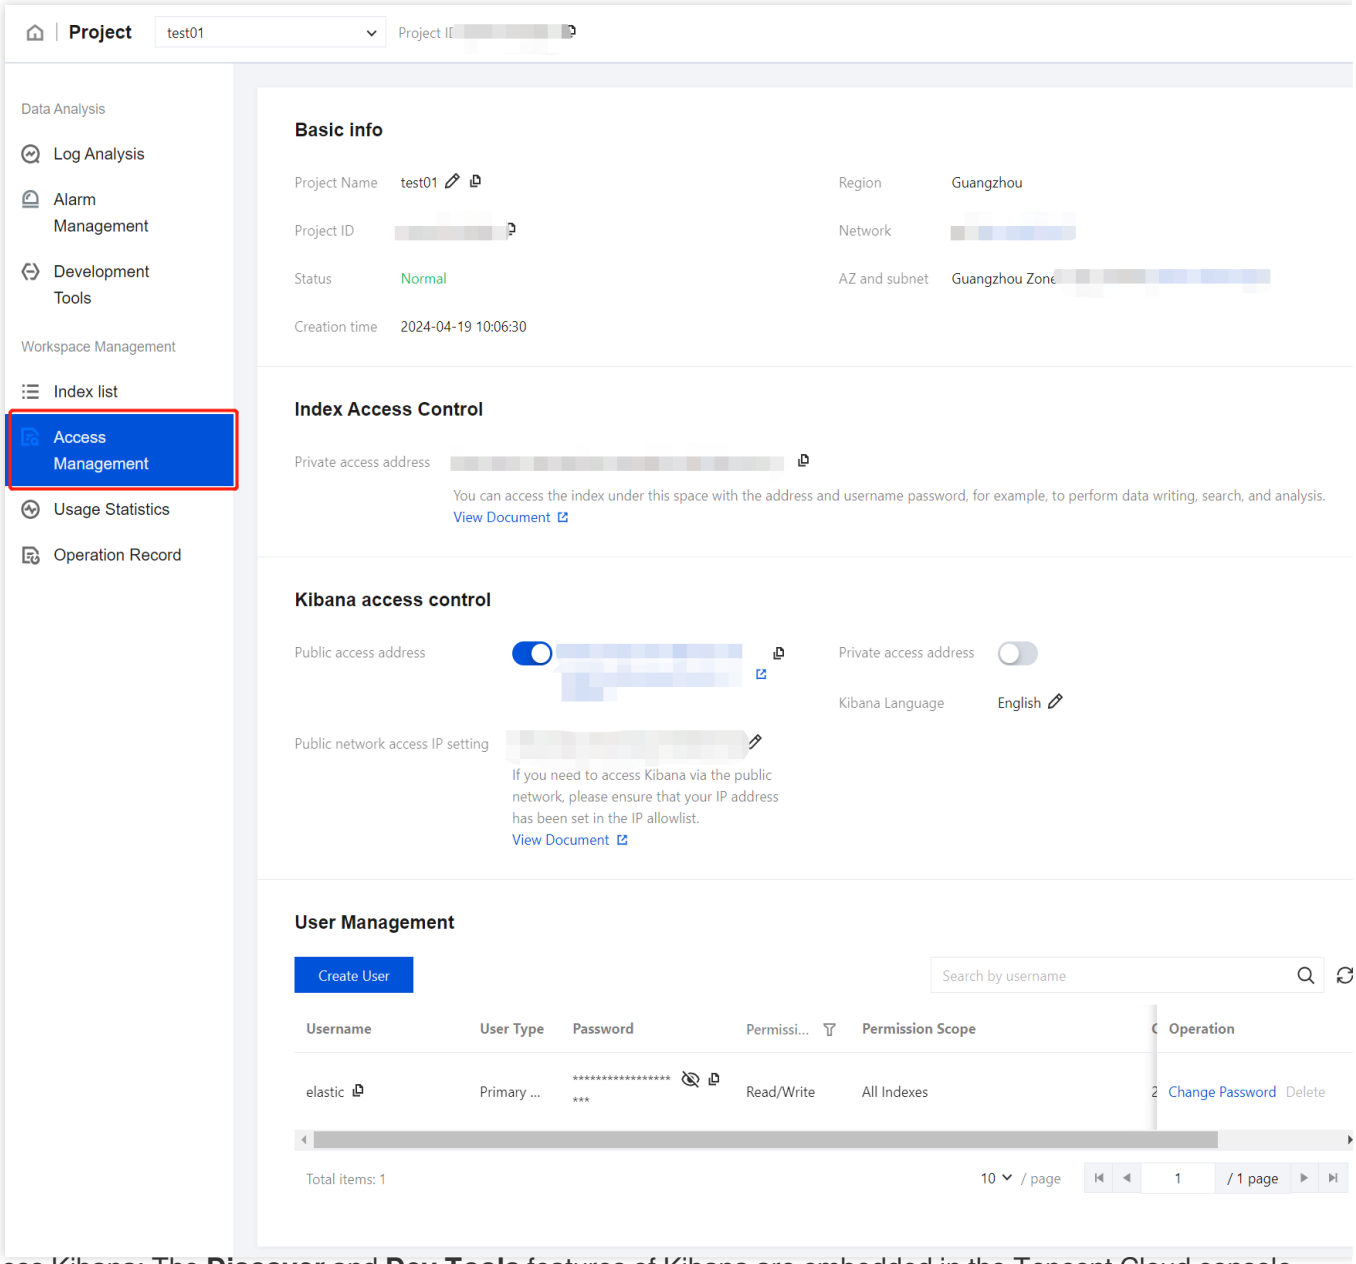


3. Access Kibana: The **Discover** and **Dev Tools** features of Kibana are embedded in the Tencent Cloud console, allowing you to use retrieval and analysis capabilities directly within the console or access Kibana via an external link. **Via Console**: Click **Search and Analysis** in the sidebar to enter the relevant page. You can switch between index views by clicking the index pattern dropdown on the left side. **Log Search** corresponds to **Discover**, and **Development Tools** corresponds to **Dev Tools**.

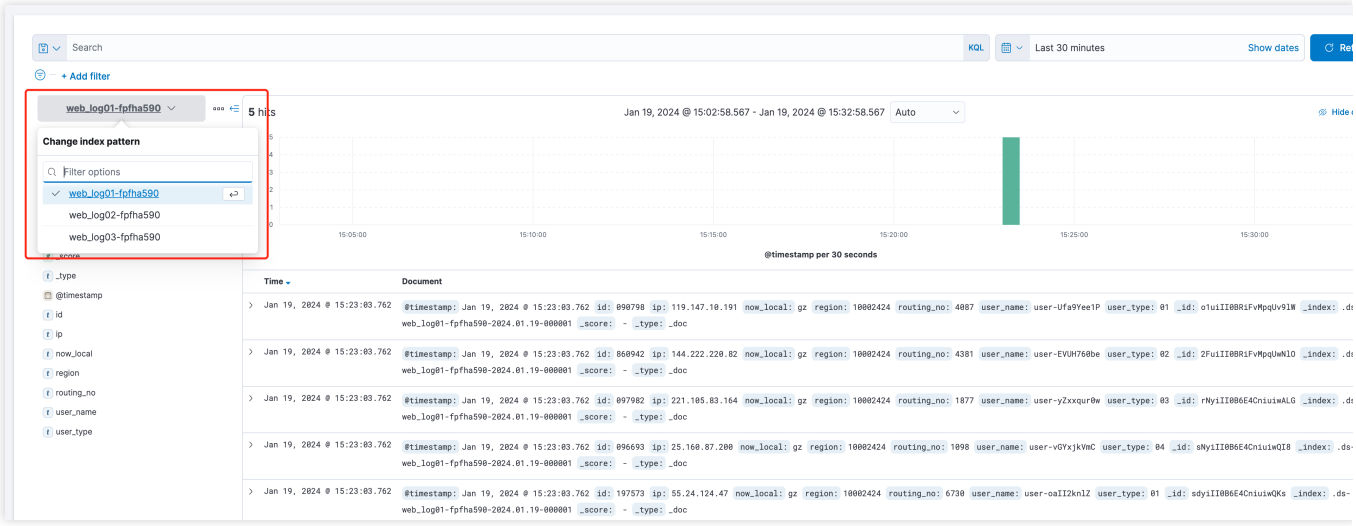


**Note:**

Embedded features require third-party cookies to be enabled in your browser. If you encounter any issue, please enable third-party cookie settings in your browser settings.

**Via Kibana Public Access Address**: Click **Kibana public access address** to enter the Kibana page.

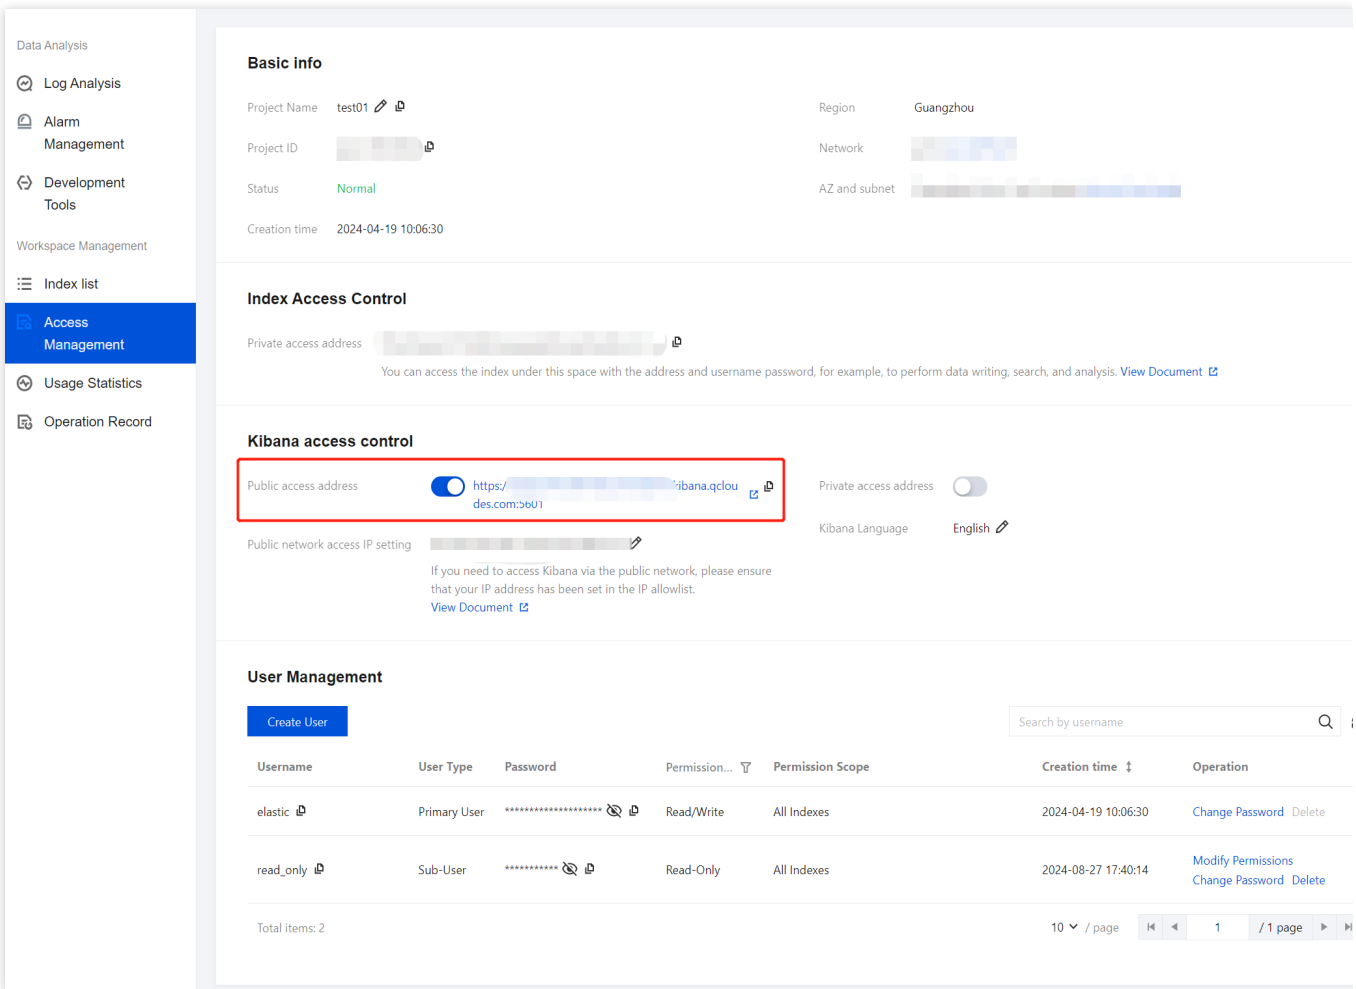


On the Kibana login page, enter the username and password, which can be copied directly from the user management page.

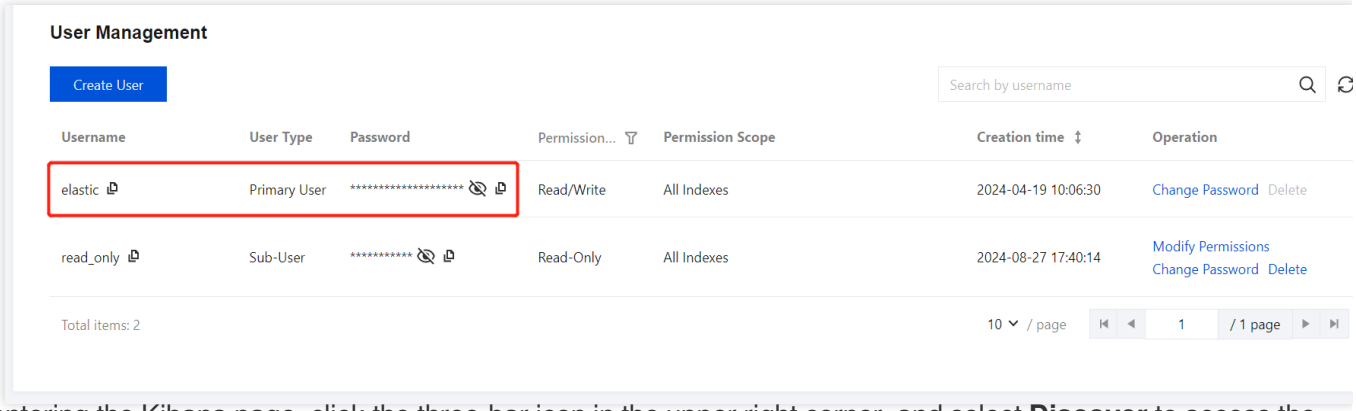


After entering the Kibana page, click the three-bar icon in the upper right corner, and select **Discover** to access the search and analysis page.

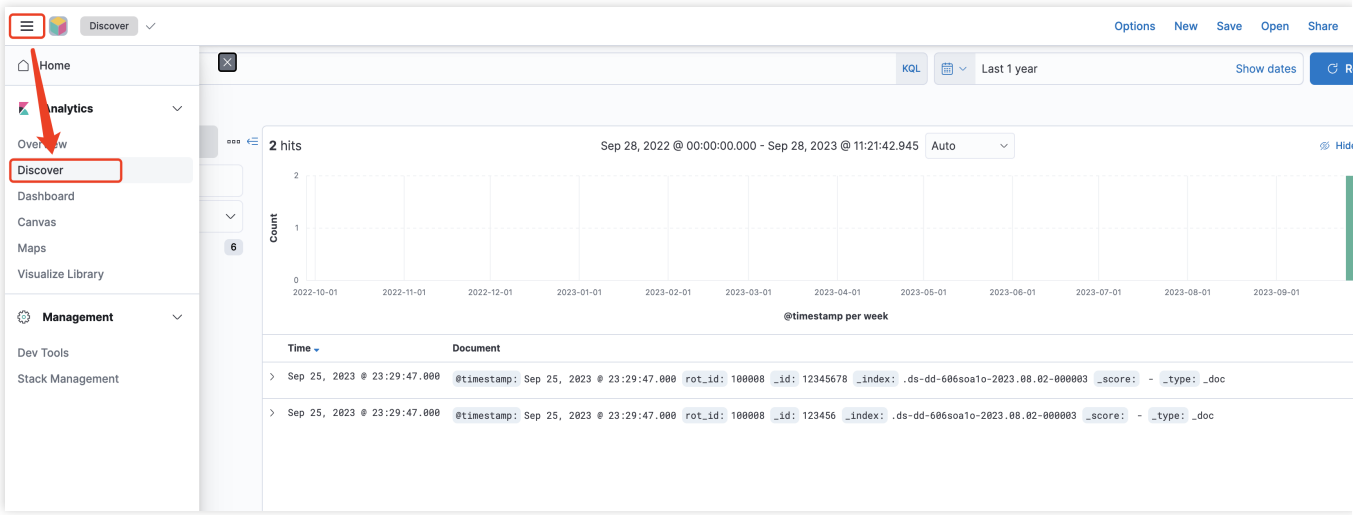


**Note:**

Kibana public network access includes an allowlist mechanism, meaning that IP addresses not included in the access policy cannot access Kibana, enhancing access security. If the page displays Sorry, you do not have permissions to access, you can click **Kibana public network access policy** as shown above. In the pop-up window, click **Get current IP** to enter your current IP address to the allowlist.

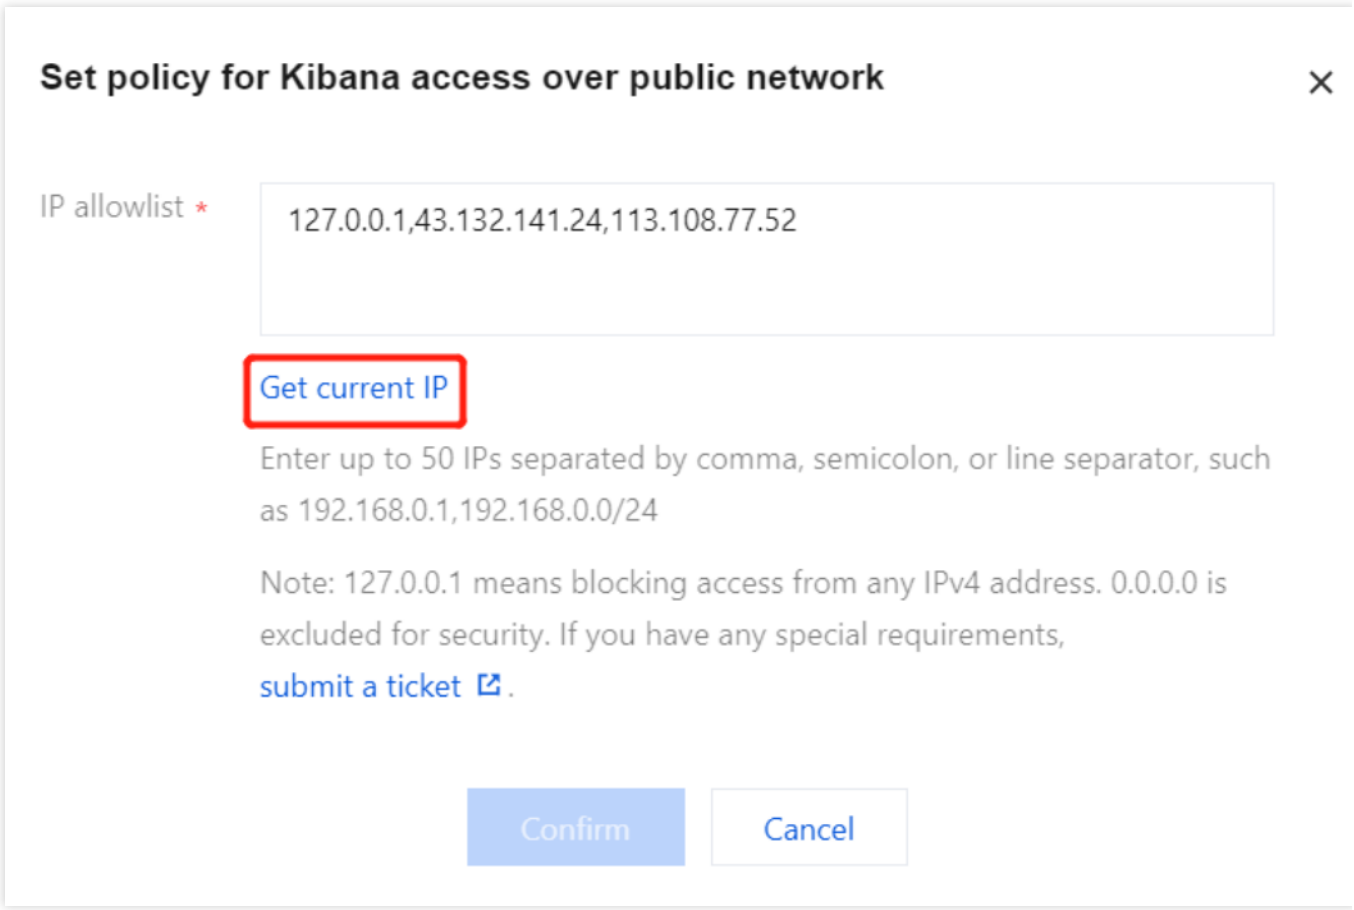


# Retrieval and Analysis

### Via Command Line

```
curl -X GET "index access address/index name/_search?pretty" -H 'Content-Type:
application/json' -d'
{
  "query": {
    "term": {
      "user.id": "kimchy"
    }
  }
}
'
```

### Via Discover

On the Discover page, you can perform time filtering, keyword searches, and other operations:

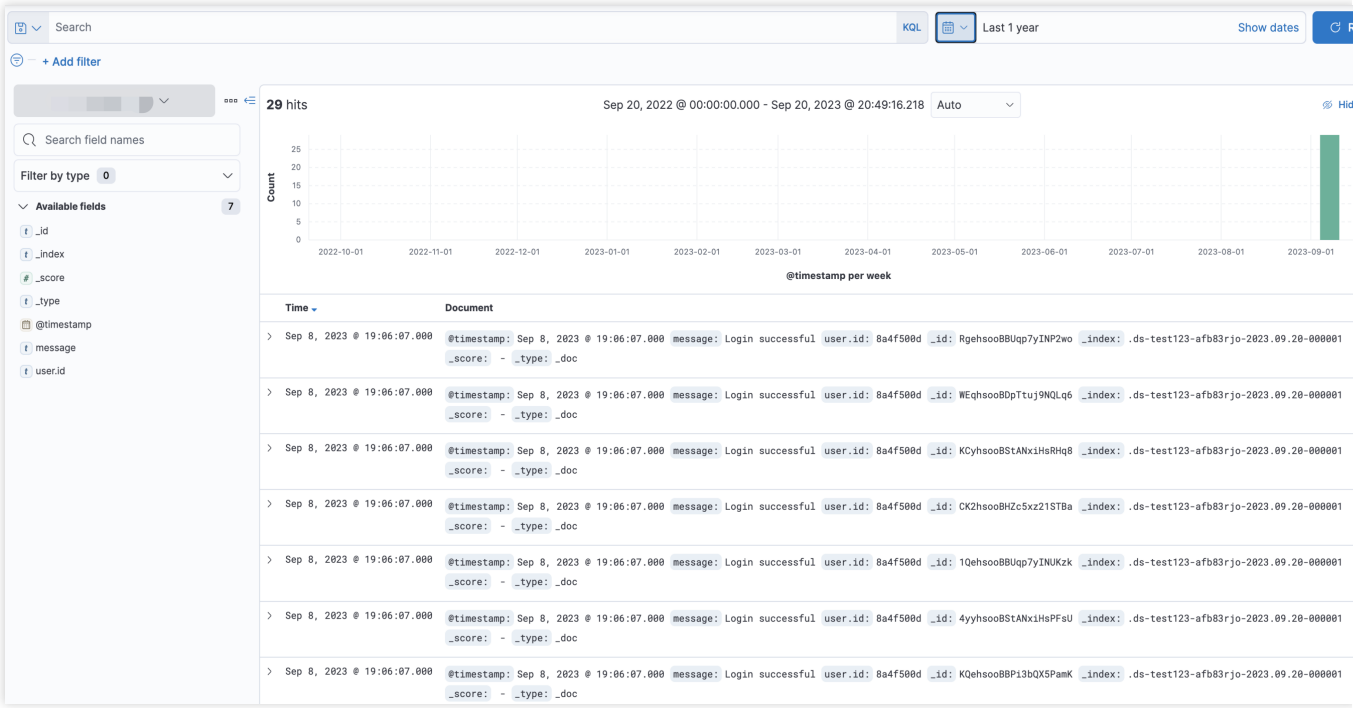


## Via Dev Tools

Perform data queries using DSL. An example is as follows:

```
GET /index_name/_search
{
  "query": {
    "term": {
      "user.id": "kimchy"
    }
  }
}
```

## Via Kibana Dashboard

After entering Kibana, select **Dashboard** in the left sidebar to start data visualization. You can quickly create charts by dragging and dropping elements.

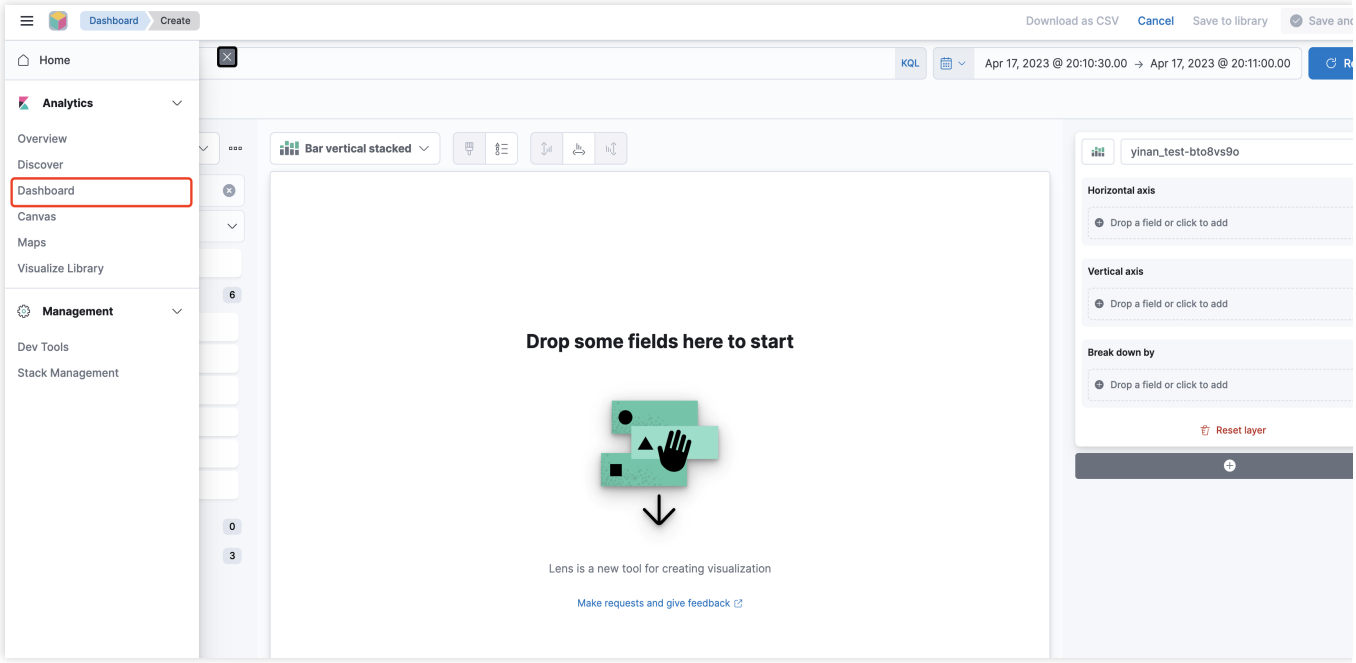
Dashboard
Create
Download as CSV
Cancel
Save to library
Home
Analytics
Overview
Discover
Dashboard
Canvas
Maps
Visualize Library
Management
Dev Tools
Stack Management
KQL
Apr 17, 2023 @ 20:10:30.00 → Apr 17, 2023 @ 20:11:00.00
Bar vertical stacked
yinan_test-bto8vs9o
Horizontal axis
Drop a field or click to add
Vertical axis
Drop a field or click to add
Break down by
Drop a field or click to add
Reset layer
Drop some fields here to start
Lens is a new tool for creating visualization
Make requests and give feedback

# Index Management
# Configuration Management

Last updated：2024-12-04 16:28:21

The Elasticsearch Serverless service provides configuration management features for indexes, allowing you to quickly view an index's configuration on the configuration management page. You can also modify index configurations to quickly adapt to business growth.

## Viewing the Index Configuration

Upon entering this page, the default view mode displays information such as field mappings and data storage duration.

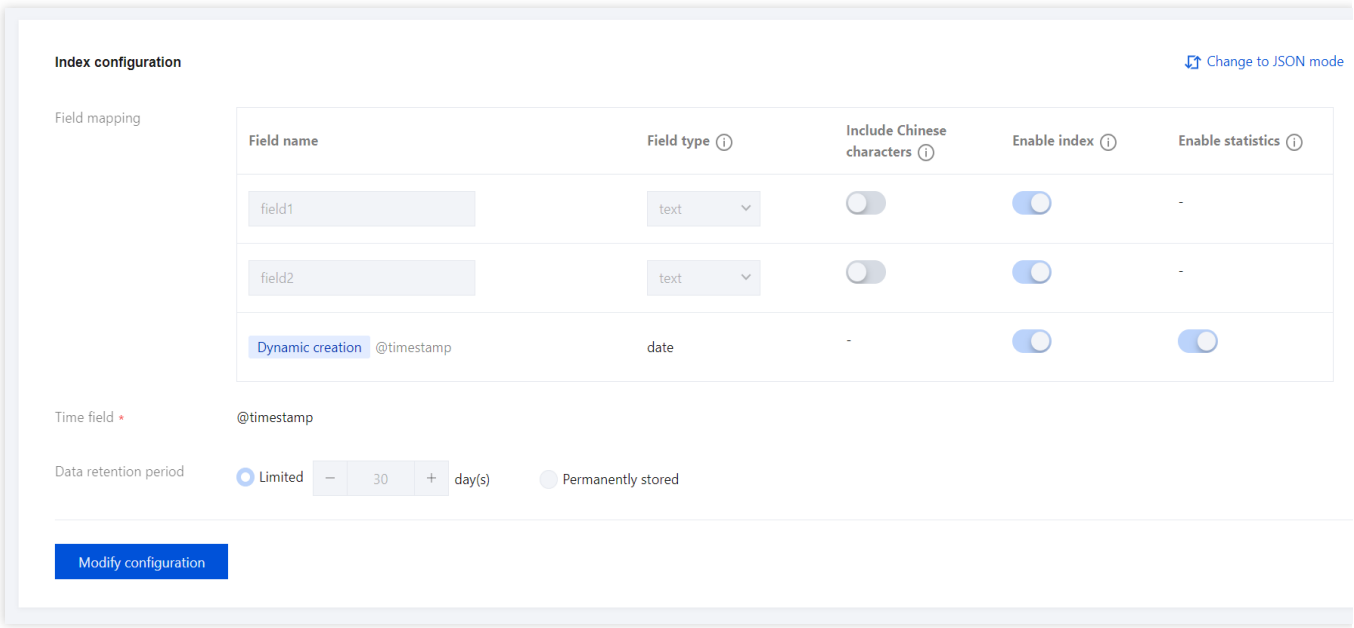


Additionally, you can click **Change to JSON mode** in the upper right corner to view the current index configuration in JSON format.

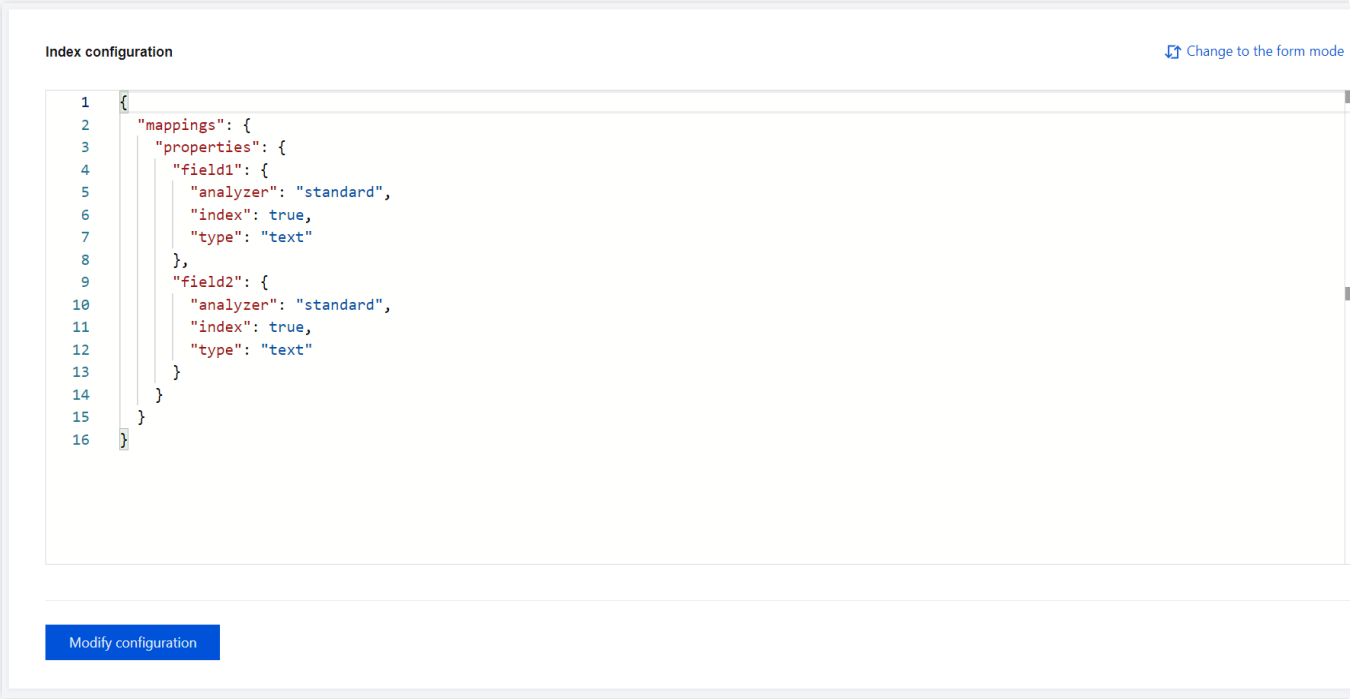


# Modifying the Index Configuration

Click **Modify configuration** in the lower left corner to enter the edit mode, allowing you to adjust index configuration settings, such as field mappings or data storage duration.

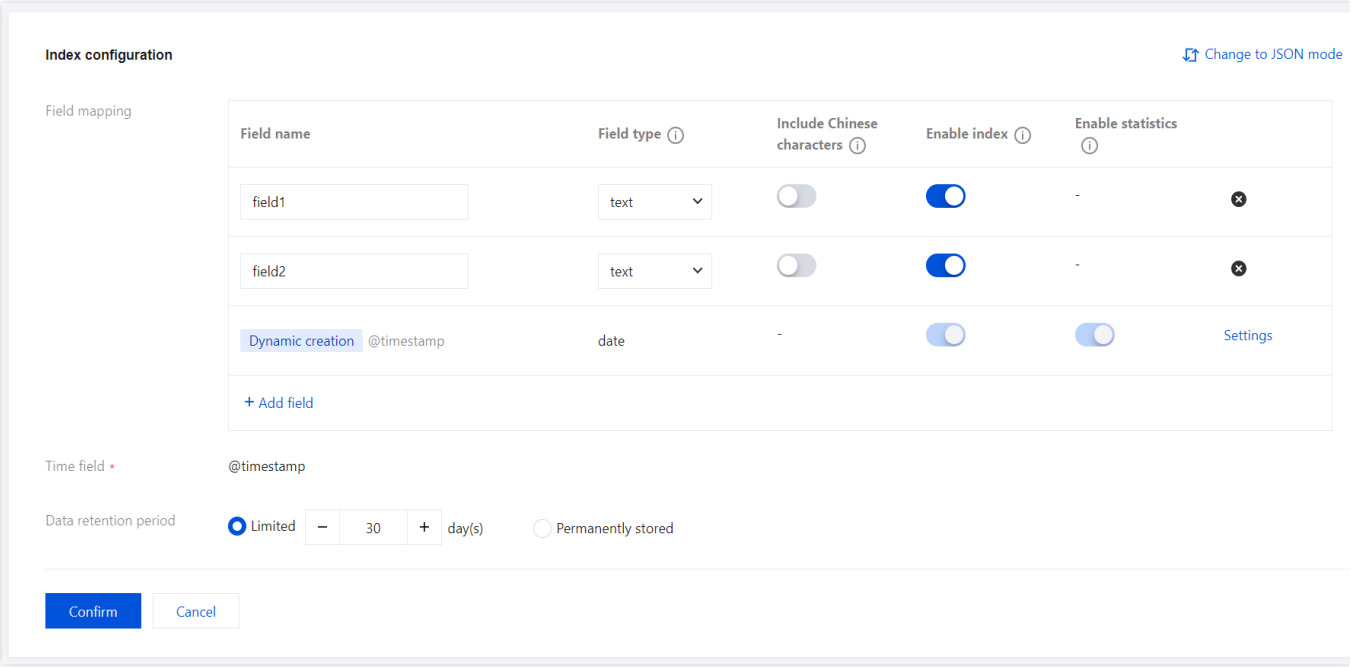


When you change to JSON mode, the left-side panel displays the current configuration for the live index, making it easy to review active settings. The right-side panel provides an input box for modifying configuration settings. Enter

the **corresponding configuration information to be modified** in the input box. Once the modification is successful, the corresponding index configuration items will be updated.

# Alarm Management

Last updated：2024-12-04 16:38:41

The ES Serverless service supports alarm management, allowing you to configure alarm policies for specific objects in the console. These policies periodically perform retrieval and analysis on indexes within monitored objects. When query results meet trigger conditions, an alarm notification is sent (currently supported via email and WeCom), enabling timely detection of issues. This feature supports keyword alarms, such as the number of the term error within logs over a specified time range, and metric monitoring, such as determining whether the maximum value of a numeric field exceeds a set threshold within a specified time range. This capability enhances observability in log analysis scenarios, enabling quick issue detection and resolution.

# Operation Steps

## Prerequisites

1. Log in to the ES Serverless console.
2. In the space list, click the corresponding space name.

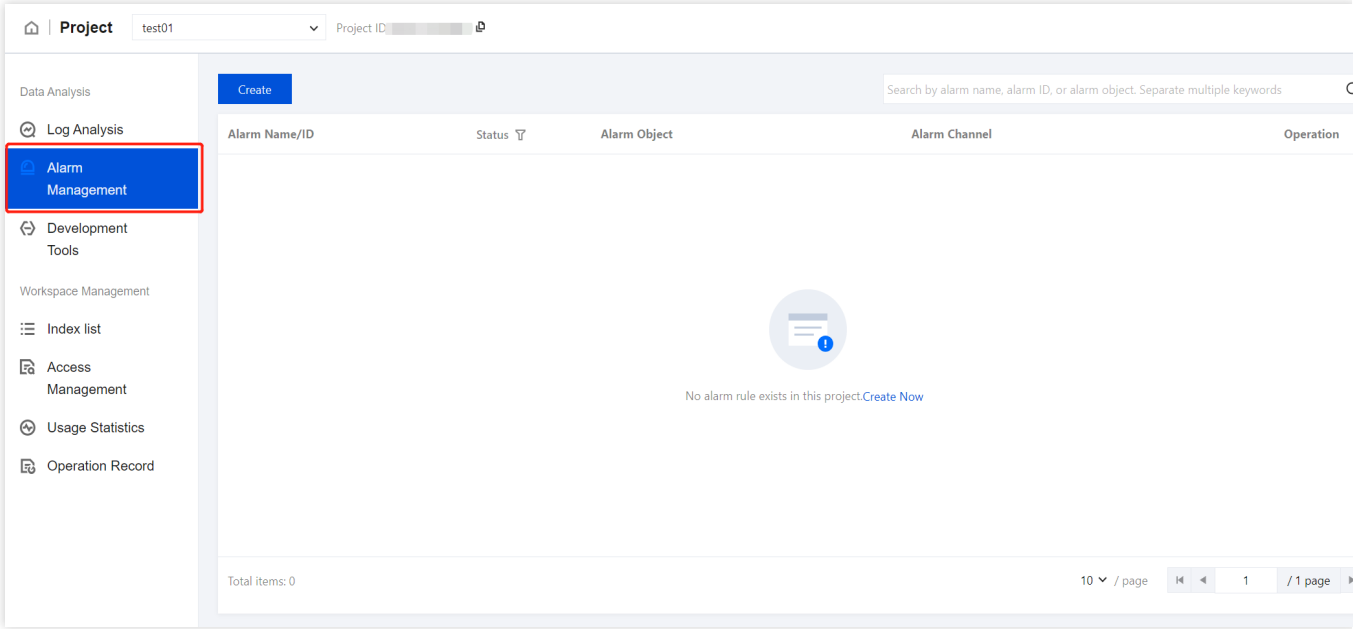


## Creating an Alarm

## Basic Information

1. In the left sidebar, click **Alarm Management**, then click **Create**.
2. Enter an alarm name, with a length of 1–50 characters. Digits, letters, Chinese characters, underscores, and delimiters - are supported.
3. Select an alarm object, with support for indexes within the current space.

**Note:**

Indexes that are still being created cannot be selected.

## Alarm Rules

1. Query statement:

Supported operators include **count, average, sum, max, and min**, with count as the default.

When the operator is **count**, all fields can be selected. The expression supports **equal to, not equal to, belong to, not belong to, existing, and no existing**.

If the expression is **equal to or not equal to**, you need to enter a corresponding value, with support for a single string only.

If the expression is **belong to or not belong to**, you need to enter an array of values, with at least one entry, separated by commas.

When the operator is **average, sum, max, or min**, only numeric fields, such as long, integer, short, double, and float, can be selected.

2. Query range: Defaults to data written within the last 5 minutes. Supports units in minutes and hours.
3. Query frequency: Defaults to querying every 1 minute. Supports units in minutes and hours.
4. Trigger condition: The expression supports **Greater than, Greater than or equal to, equal to, Equal to less than, Less than, and Between**. The default is set to greater than, with a default value of 100.

## Alarm Notification

1. Email:

To ensure the accuracy of the alarm address, enter the email address and complete a Captcha verification.

If the email address is changed, a new Captcha will need to be requested.

2. WeCom: Enter the WeCom bot webhook address.

**Note:**

The WeCom bot webhook address should start with the prefix `https://qyapi.weixin.qq.com` .

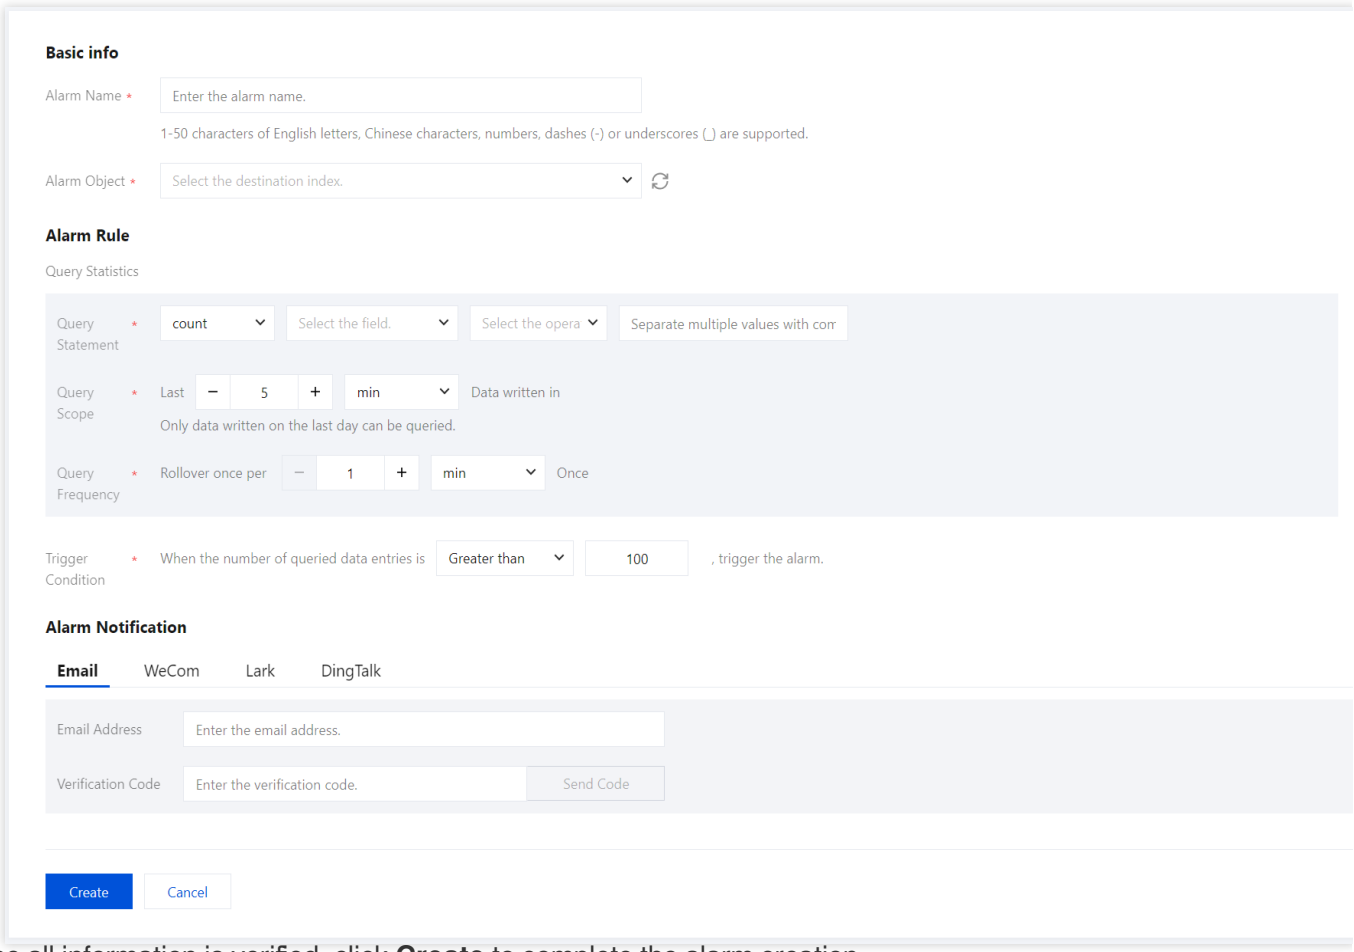


3. Once all information is verified, click **Create** to complete the alarm creation.

### Alarm Content

When an alarm is triggered, you will receive the following information:

Title: Tencent Cloud Elasticsearch Serverless Service Alarm Triggered.

Content:

[Alarm] Dear Tencent Cloud user, your Tencent Cloud account (Account ID: xxx) using the Elasticsearch Serverless service triggered an alarm at {Time} (UTC+8).

Alarm Name: {Corresponding Alarm Name}

Alarm Object: {Corresponding Index Name}

## Alarm Management

1. On the **Alarm Management** page, you can view details and the status of your configured alarm policies.
2. To disable or delete an alarm, click **More** in the operation column.
3. To edit an alarm policy, click **Edit**.

# ES API References

Last updated：2024-12-04 16:40:46

## Using APIs via Command Line or Clients Such as Filebeat

| API URI | Supported Method | Description |
| --- | --- | --- |
| /_bulk | PUT and POST | For more details, see Bulk API. |
| /{index}/_bulk | PUT and POST | For more details, see Bulk API. |
| /{index}/_doc/{id} | PUT and POST | For more details, see Index API. |
| /{index}/_doc | POST | For more details, see Index API. |
| /{index}/_create/{id} | PUT and POST | For more details, see Index API. |
| /{index}/_mapping | GET | For more details, see Get mapping API. |
| /{index}/_msearch | POST and GET | For more details, see Multi search API. |
| /_msearch | POST and GET | For more details, see Multi search API. |
| /{index}/_count | POST and GET | For more details, see Count API. |
| /{index}/_search | POST and GET | For more details, see Search API. |

## Using APIs via Kibana

| API URI | Supported Method | Description |
| --- | --- | --- |
| /{index}/_bulk | PUT and POST | For more details, see Bulk API. |
| /{index}/_doc/{id} | PUT and POST | For more details, see Index API. |
| /{index}/_doc | POST | For more details, see Index API. |
| /{index}/_create/{id} | PUT and POST | For more details, see Index API. |
| /_security/user/_has_privileges | POST and GET | For more details, see Has privileges API. |
| /{index}/_field_caps | POST and GET | For more details, see Field capabilities API. |

| | | |
|---|---|---|
| /{index}/_flush | POST and GET | For more details, see Flush API. |
| /{index}/_mapping | GET | For more details, see Get mapping API. |
| /{index}/_mappings | GET | For more details, see Get mapping API. |
| /{index}/_refresh | POST and GET | For more details, see Refresh API. |
| /_resolve/index/{name} | GET | For more details, see Resolve index API. |
| /{index}/_count | POST and GET | For more details, see Count API. |
| /{index}/_msearch | POST and GET | For more details, see Multi search API. |
| /{index}/_search | POST and GET | For more details, see Search API. |
| /_async_search/{id} | GET | - |
| /{index}/_async_search | POST | - |
| /_security/_authenticate | GET | For more details, see Authenticate API. |

# Related Issues

# Kibana Usage Issues

Last updated：2024-12-04 17:36:10

### How to Set a Field to `geo_point` Type and Draw a Map?

Before writing data, set the type of the specified field to `geo_point` in the mapping. After the data is written, go to **Maps** in the Kibana sidebar to enter the map drawing interface.

**Note:**

Manually set the field to `geo_point` type; otherwise, it may be automatically mapped to an incorrect type, preventing the map from being drawn.

### Where can I Find Kibana's Coordinate Map Feature?

You can use the **Clusters and grids** option in **Maps** to aggregate specific fields.

### How to Distinguish and Display Different Value Ranges for Fields Aggregated by Metrics?

You can adjust the **Fill color** setting: In **Layer settings**, scroll down to find the **Layer Style** module. For **Fill color**, select by value, and select the key to differentiate. Then, in as number, select an appropriate gradient color.

### After Metrics are Displayed, How can I Identify the Value Ranges Represented by Different Colors on the Map?

Click the arrow in the middle of the corresponding layer under LAYERS to display the value ranges (this arrow is hidden by default and only appears when you hover over it).

### After Metrics are Displayed, the Points are Large and Overlap the Base Map Labels, Hiding Place Names. How can I Adjust this?

Click **Road map** in **Layer**, then select **Edit layer settings**. In **Layer settings**, set the base map display priority to top.

# Third-Party Cookie Settings

Last updated：2024-12-04 16:47:23

To use the console retrieval and analysis capabilities, your browser should support third-party cookies. Common browser settings are as follows:

## Chrome

1. Open the Chrome browser.
2. Click More Options in the upper right corner



.
3. Click **Settings** > **Privacy and security** > **Third-party cookies** > **Allow third-party cookies**.

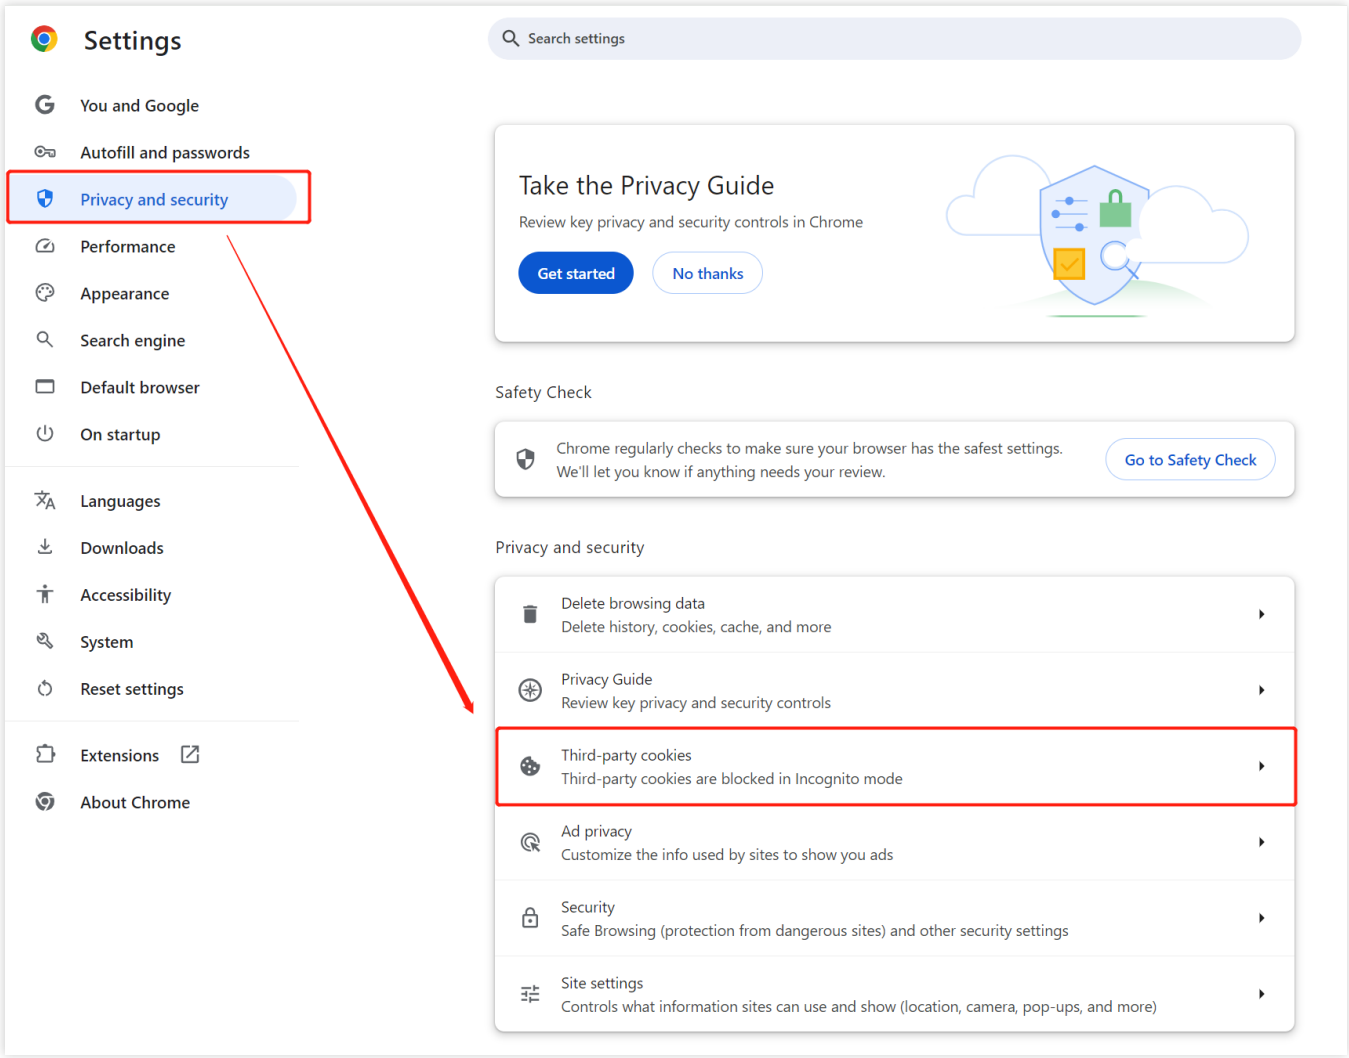
Settings
Search settings
You and Google
Autofill and passwords
Privacy and security
Performance
Appearance
Search engine
Default browser
On startup
Languages
Downloads
Accessibility
System
Reset settings
Extensions
About Chrome
Take the Privacy Guide
Review key privacy and security controls in Chrome
Get started
No thanks
Safety Check
Chrome regularly checks to make sure your browser has the safest settings.
We'll let you know if anything needs your review.
Go to Safety Check
Privacy and security
Delete browsing data
Delete history, cookies, cache, and more
Privacy Guide
Review key privacy and security controls
Third-party cookies
Third-party cookies are blocked in Incognito mode
Ad privacy
Customize the info used by sites to show you ads
Security
Safe Browsing (protection from dangerous sites) and other security settings
Site settings
Controls what information sites can use and show (location, camera, pop-ups, and more)

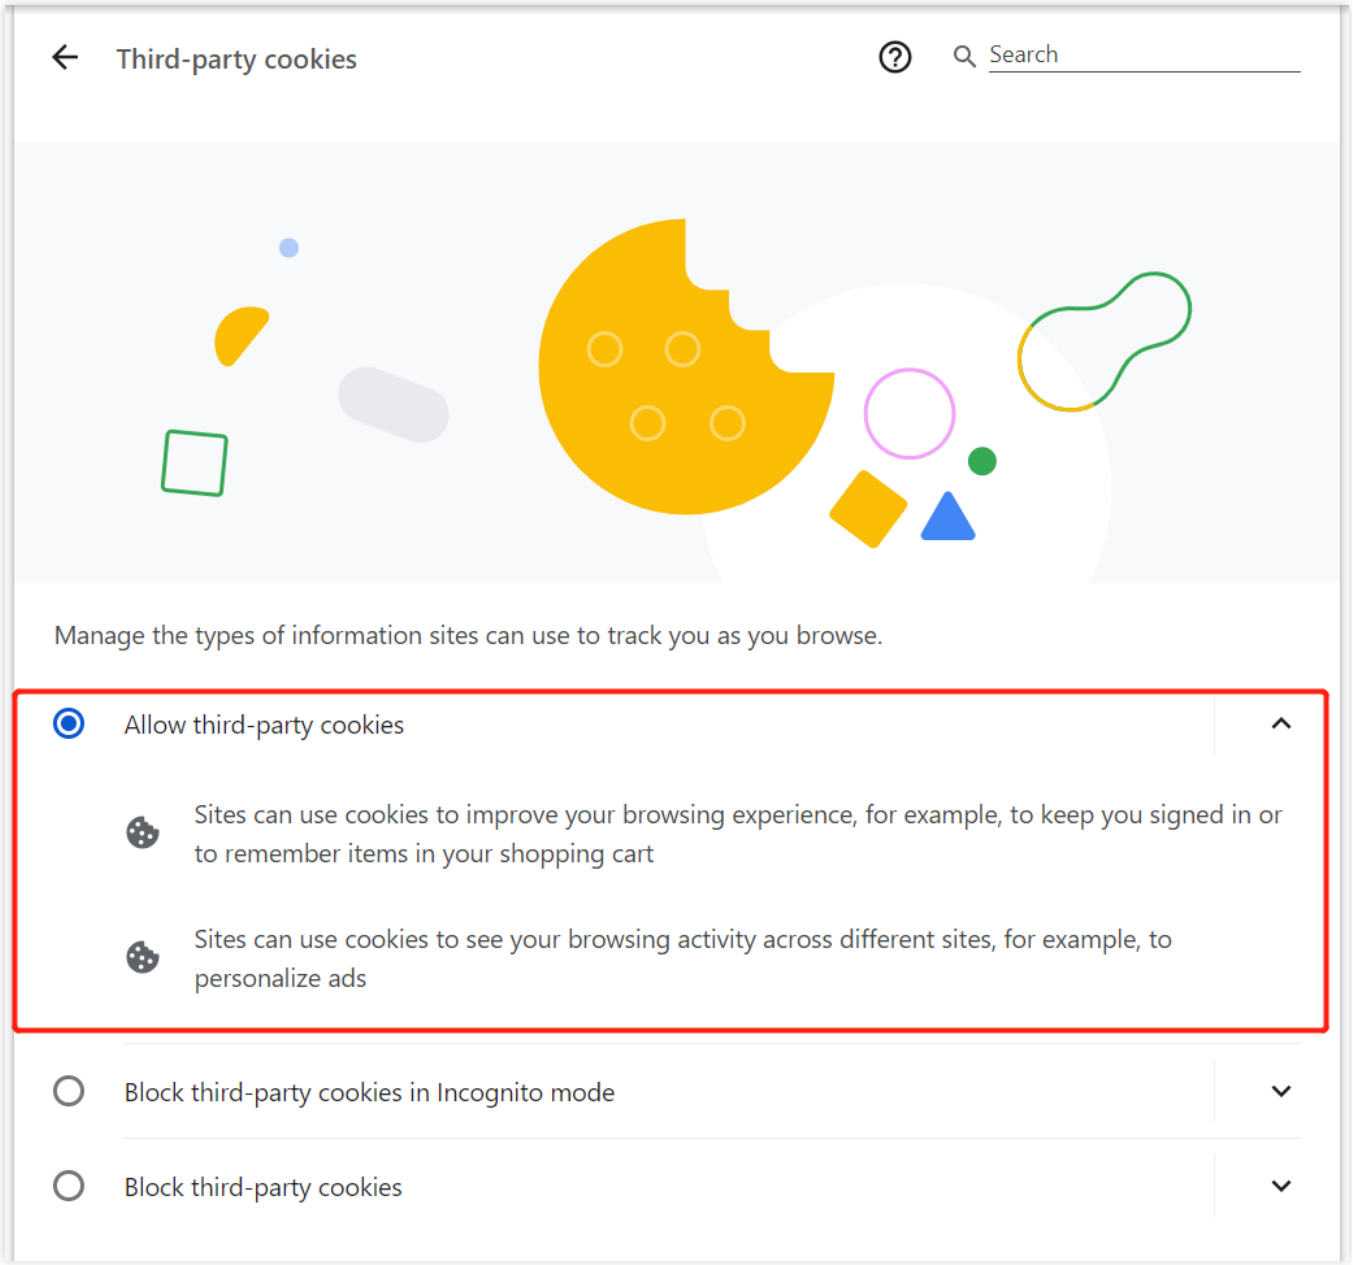


## Safari

1. Open Safari on your Mac computer.
2. Go to **Safari Browser** > **Settings** to open the settings interface.

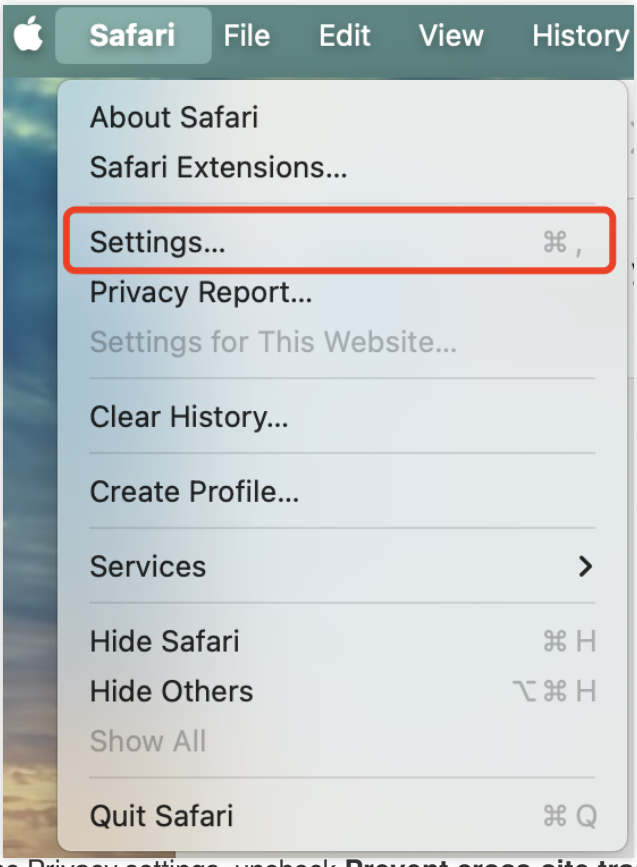


3. In the Privacy settings, uncheck **Prevent cross-site tracking** and **Manage Website Data**.

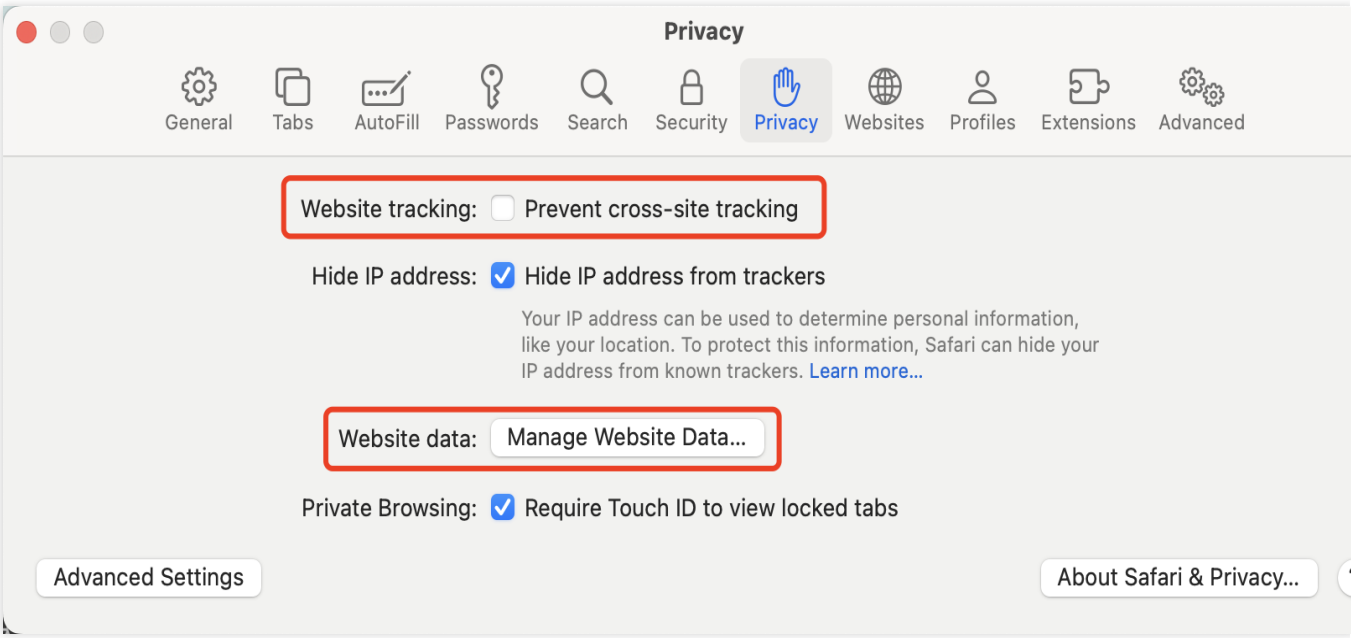

# Field Type Conversion Through Reindex

Last updated：2024-12-04 16:50:23

## Overview

When you create an index in the ES Serverless service, a time field should be specified, and its type should be set to `date`. When you synchronize data from an existing ES cluster to an index in the ES Serverless service, if the field in the data has the same name as the time field but a different type, the write operation will fail. In this case, you can use the Reindex API to convert the field type.

## Process Description

1. Create the target index for reindexing, and set the type of the field to `date` if the field has the same name as the time field in the ES Serverless service index.
2. Use the reindex API to synchronize the existing data to the target index.

## Example

1. Suppose we need to synchronize data from the `source_index` to an index in the ES Serverless service (where the time field is `@timestamp`). Upon checking the field configuration of source_index, we find that in `source_index`, the field `@timestamp` is of type `keyword`. In this case, attempting to synchronize the data will result in a write error.


2. View the number of documents in source_index.

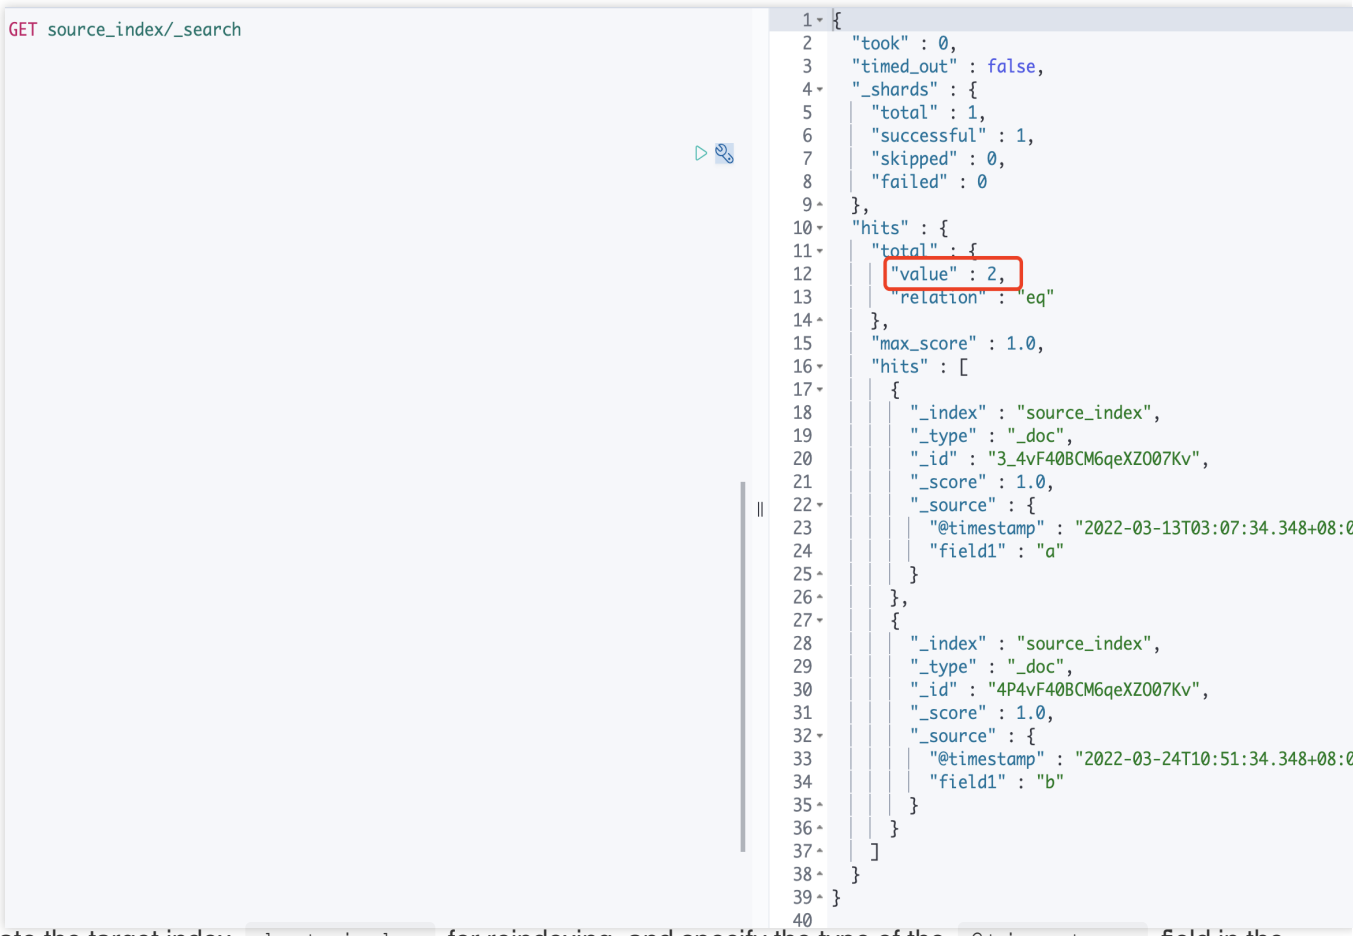


3. Create the target index `dest_index` for reindexing, and specify the type of the `@timestamp` field in the mapping as `date` .


4. Use the reindex API to synchronize data from `source_index` to `dest_index`, and the number of documents in `dest_index` should match exactly the number in the original `source_index`.

```
POST _reindex
{
  "source": {
    "index": "source_index"
  },
  "dest": {
    "index": "dest_index"
  }
}
```

```
POST _reindex
{
  "source": {
    "index": "source_index"
  },
  "dest": {
    "index": "dest_index"
  }
}
```

```
#! [index.search.slowlog.level] setting was deprecated in Elasticsearch and will be removed in a future release
  the breaking changes documentation for the next major version.
#! [index.indexing.slowlog.level] setting was deprecated in Elasticsearch and will be removed in a future releas
  the breaking changes documentation for the next major version.
{
  "took" : 52,
  "timed_out" : false,
  "total" : 2,
  "updated" : 0,
  "created" : 2,
  "deleted" : 0,
  "batches" : 1,
  "version_conflicts" : 0,
  "noops" : 0,
  "retries" : {
    "bulk" : 0,
    "search" : 0
  },
  "throttled_millis" : 0,
  "requests_per_second" : -1.0,
  "throttled_until_millis" : 0,
  "failures" : [ ]
}
```

5. In this case, searching `dest_index` will retrieve the data that was synchronized from `source_index` .

```
GET dest_index/_search
```

```
{
  "took" : 0,
  "timed_out" : false,
  "_shards" : {
    "total" : 1,
    "successful" : 1,
    "skipped" : 0,
    "failed" : 0
  },
  "hits" : {
    "total" : {
      "value" : 2,
      "relation" : "eq"
    },
    "max_score" : 1.0,
    "hits" : [
      {
        "_index" : "dest_index",
        "_type" : "_doc",
        "_id" : "3_4vF40BCM6qeXZO07Kv",
        "_score" : 1.0,
        "_source" : {
          "@timestamp" : "2022-03-13T03:07:34.348+08:
          "field1" : "a"
        }
      },
      {
        "_index" : "dest_index",
        "_type" : "_doc",
        "_id" : "4P4vF40BCM6qeXZO07Kv",
        "_score" : 1.0,
        "_source" : {
          "@timestamp" : "2022-03-24T10:51:34.348+08:
          "field1" : "b"
        }
      }
    ]
  }
}
```