Cloud Search Service

FAQ

Issue 02

Date 2020-07-14

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: https://www.huawei.com

Email: support@huawei.com

Contents

1 General Consulting

1.1 What Is CSS?

Cloud Search Service (CSS) is a fully managed, distributed search service. It supports Elasticsearch and provides users with structured and unstructured data search, statistics, and report capabilities.

CSS can be automatically deployed, where you can quickly create Elasticsearch clusters. It is built with search debugging practices with zero O&M. Additionally, it has a robust monitoring system to present you key metrics, including clusters and query performance so that you can focus on the business logic.

1.2 What Are Regions and AZs?

Regions and AZs

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- Regions are divided from the dimensions of geographical location and network latency. Public services, such as ECS, EVS, OBS, VPC, Elastic IP (EIP), and Image Management Service (IMS), are shared within the same region. Regions are classified as universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides services of the same type only or for specific tenants.
- An AZ contains one or multiple physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. Within an AZ, computing, network, storage, and other resources are logically divided into multiple clusters. AZs within a region are interconnected using high-speed optical fibers to allow users to build cross-AZ high-availability systems.

Figure 1-1 shows the relationship between regions and AZs.

Figure 1-1 Regions and AZs

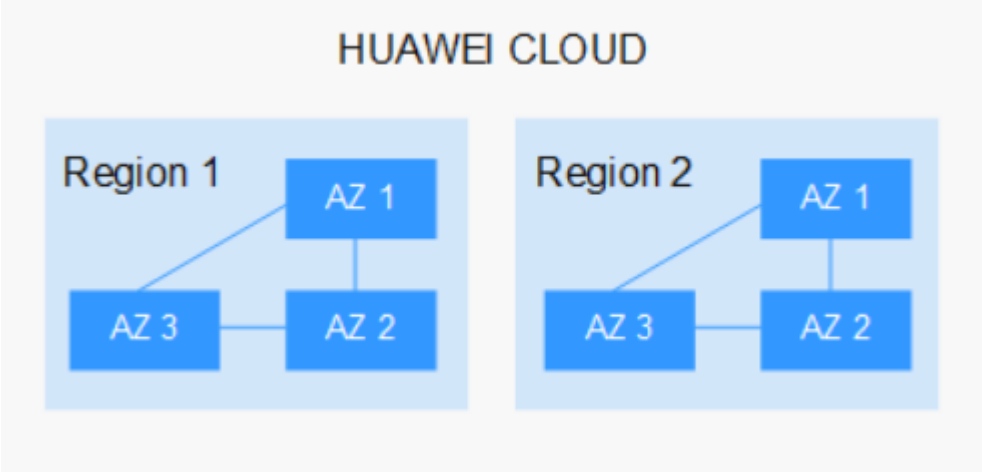

HUAWEI CLOUD provides services in many regions around the world. You can select a region and AZ as needed. For more information, see HUAWEI CLOUD Global Regions.

Region Selection

When selecting a region, consider the following factors:

- Location

  You are advised to select a region close to you or your target users. This reduces network latency and improves access rate. Regions within the Chinese mainland provide the same infrastructure, BGP network quality, as well as resource operations and configurations. Therefore, if your target users are in the Chinese mainland, you do not need to consider the network latency differences when selecting a region.

  - If you or your target users are in Asia Pacific except the Chinese mainland, select the **AP-Hong Kong**, **AP-Bangkok**, or **AP-Singapore** region.
  - If you or your target users are in Africa, select the **AF-Johannesburg** region.
  - If you or your target users are in Europe, select the **EU-Paris** region.

- Resource price

  Resource prices may vary in different regions. For details, see **Product Pricing Details**.

AZ Selection

When determining whether to deploy resources in the same AZ, consider your application's requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs in the same region.
- For low network latency, deploy resources in the same AZ.

Regions and Endpoints

Before using an API to call resources, specify its region and endpoint. For details about HUAWEI CLOUD regions and endpoints, see **Regions and Endpoints**.

1.3 Which Scenarios Can CSS Be Applied To?

CSS applies to scenarios such as log search and analysis, time and space retrieval, time sequence retrieval and report generation, and intelligent search.

1.4 How Does CSS Ensure Secure Running of Data and Services?

CSS ensures secure running of data and services from the following aspects:

- Network isolation

  The entire network is divided into two planes: service plane and management plane. The two planes are deployed and isolated physically to ensure the security of the service and management networks.

  - Service plane: refers to the network plane of the cluster. It provides service channels for users and delivers data definition, index, and search capabilities.
  - Management plane: refers to the management console. It is used to manage CSS.

- Host security

  CSS provides the following security measures:

  - The VPC security group ensures the security of hosts in a VPC.
  - Using the network access control list (ACL), you can permit or deny the network traffic entering and exiting the subnets.
  - Internal security infrastructure (including the network firewall, intrusion detection system, and protection system) can monitor all network traffic that enters or exits the VPC through the IPsec VPN.

- Data security

  Multiple replicas, cross-AZ deployment of clusters, and third-party (OBS) backup of index data ensure user data security.

1.5 Which CSS Metrics Should I Focus On?

The metrics that you need to focus on include the disk usage and cluster health status. You can log in to Cloud Eye and configure alarm prompts according to actual conditions. If alarms are reported, clear them by taking related measures. For details about how to configure alarms, see **Creating Alarm Rules** in the *Cloud Search Service User Guide*.

Configuration examples:

- Alarms are reported if the disk usage is higher than or equal to a specified value (for example, 85%) and has reached this value multiple times (for example, 5 times) within a specified time period (for example, 5 minutes).
- Alarms are reported if the value of the cluster health status metric exceeds 0 for multiple times (for example, 5 times) within a specified time period (for example, 5 minutes).

Measures:

- Upon receiving alarms related to the disk usage, view disk space consumption, check whether data can be deleted from cluster nodes or archived to other systems to release space, or expand the disk capacity.
- If an alarm related to the cluster health status is received, check whether shard allocation is normal, whether shards are lost, and check whether the process is restarted on Cerebro.

1.6 Which Storage Options Does CSS Provide?

CSS uses EVS and local disks to store your indices. During cluster creation, you can specify the EVS disk type and specifications (that is, the EVS disk size).

- Supported EVS disk types include common I/O, high I/O, and ultra-high I/O.
- For details about the sizes of EVS disks for different ECSs, see **Restrictions** in the *Cloud Search Service Overview*.

1.7 What Is the Upper Limit for the Storage Capacity of CSS?

You can create 1 to 200 nodes during cluster creation. A certain number of disks are mounted to each node, which corresponds to an ECS. You can calculate the total storage capacity of CSS based on the sizes of EVS disks attached to different ECSs. For details about the limitation on EVS disk sizes, see **Restrictions** in the *Cloud Search Service Overview*.

1.8 Which Tools Can I Use to Manage CSS?

You can use any of the following three methods to manage CSS or to use search engine APIs. You can initiate requests based on constructed request messages.

- curl

  cURL is a command-line tool used to perform URL operations and transfer information. It serves as the HTTP client that can send HTTP requests to the HTTP server and receive response messages. cURL is applicable to API debugging. For more information about cURL, visit **https://curl.haxx.se/**.
- Code

  You can call APIs through code to assemble, send, and process request messages.
- REST client

  Both Mozilla Firefox and Google Chrome provide a graphical browser plugin, that is, REST client, to send and process requests.

  - For Mozilla Firefox, see **Firefox REST Client**.
  - For Google Chrome, see **Postman**.

1.9 Which Elasticsearch Versions Does CSS Support?

In CSS, Elasticsearch versions 5.5.1, 6.2.3, 6.5.4, 7.1.1, and 7.6.2 and Kibana versions 5.5.1, 6.2.3, 6.5.4, 7.1.1, and 7.6.2 are supported.

1.10 Which Methods Can I Use to Access CSS?

You can access CSS using either of the following methods:

- RESTful API
- Transport Client

If you use Transport Client to access the cloud service, ensure that versions of both the client and server are the same.

1.11 Can I Upgrade a Cluster from an Earlier Version to a Later Version?

Currently, the cluster version cannot be upgraded directly. You can purchase a cluster of a later version.

1.12 What Can Be the Disk Space of the Requested Cluster Used For?

The following logs and files can be stored in the disks:

- Log files: Elasticsearch logs
- Data files: Elasticsearch index files
- Other files: cluster configuration files
- OS: 5% storage space reserved for the OS by default

2 Billing

2.1 How Is a Monthly Package Billed?

Within the validity period of your purchased discount package, the system deducts the quota from the discount package first. You are billed for the amount exceeding the upper limit. Assume you purchase a 500 GB storage package for one month. If you use less than 500 GB of storage capacity within one month, you are not billed. If you use more than 500 GB of storage capacity, you are billed for the part beyond 500 GB on a pay-per-use basis.

Discount packages and clusters are created separately. The node specifications of the package correspond to those when you create a cluster. If the discount package you purchase has the same node specifications with those you select when you create a cluster, the specifications in the package are preferentially deducted.

3 Functions

3.1 Why Do I Fail to Create a Cluster?

The possible causes of a cluster creation failure are as follows:

- Insufficient resource quota. You are advised to apply for sufficient resource quotas. For details, see **How Can I Apply for a Higher Quota?**
- The value of **Port Range/ICMP Type** in **Security Group** does not include port **9200**. Modify the security group information or select another available security group.

3.2 How Do I Use Custom Scoring of Elasticsearch to Query Data?

You can use Elasticsearch to score the matched documents. This section describes how to use the custom scoring query function of Elasticsearch.

Procedure

1. Log in to the CSS management console.
2. In the left navigation pane, click **Clusters** to switch to the **Clusters** page.
3. In the cluster list, locate the row where the target cluster resides and click **Kibana** in the **Operation** column.
4. In the left navigation pane of Kibana, click **Dev Tools**. Click **Get to work** to switch to the **Console** page.
5. Run the following command to create an index and specify a user-defined mapping to define the data type:

   For example, the **tv.json** data file is available and contains data as follows:

   ```
   {
   "tv":[
   { "name": "tv1", "description": "USB, DisplayPort", "vote": 0.98 }
   { "name": "tv2", "description": "USB, HDMI", "vote": 0.99 }
   { "name": "tv3", "description": "USB", "vote": 0.5 }
   { "name": "tv4", "description": "USB, HDMI, DisplayPort", "vote": 0.7 }
   ```



```
]
}
```

Run the following command to create the **mall** index and specify the user-defined mapping to define the data type:

```
PUT /mall?pretty
{
  "mappings": {
    "tv": {
      "properties": {
        "description": {
          "type": "text",
          "fields": {
            "keyword": {
              "type": "keyword"
            }
          }
        },
        "name": {
          "type": "text",
          "fields": {
            "keyword": {
              "type": "keyword"
            }
          }
        },
        "vote": {
          "type": "float"
        }
      }
    }
  }
}
```

6. Run the following command to import data in the **tv.json** file to the **mall** index:

```
POST /mall/tv/_bulk?pretty
{ "index": {"_id": "1"}}
{ "name": "tv1", "description": "USB, DisplayPort", "vote": 0.98 }
{ "index": {"_id": "2"}}
{ "name": "tv2", "description": "USB, HDMI", "vote": 0.99 }
{ "index": {"_id": "3"}}
{ "name": "tv3", "description": "USB", "vote": 0.5 }
{ "index": {"_id": "4"}}
{ "name": "tv4", "description": "USB, HDMI, DisplayPort", "vote": 0.7 }
```

7. Query data by using custom scoring.
   - Compute the total score of each TV based on the absolute praise rate and sort the searched products in a descending order of the total score.

     The following sample code illustrates how to query TVs with the USB, HDMI, or DisplayPort port. In this example, give score 0 for a TV if it has none of the three ports, 1 if it has one, 2 if it has two, and 3 if it has three. Then, multiply the obtained score by the TV's absolute praise rate to obtain its total score. Finally, sort the TVs in a descending order of their total scores.

     ```
     GET /mall/tv/_search?pretty
     {
       "query":{
         "function_score":{
           "query":{
             "bool":{
               "should":[
                 {"constant_score":{
                     "query":{"match":{"description":"USB"}}
                 }},
                 {"constant_score":{
     ```



```
                "query":{"match":{"description":"HDMI"}}
              }},
              {"constant_score":{
                  "query":{"match":{"description":"DisplayPort"}}
              }}
            ]
          }
        },
        "field_value_factor":{
          "field":"vote",
          "factor":1
        },
        "boost_mode":"multiply",
        "max_boost":10
      }
    }
}
```

In the preceding example, the total score is calculated as follows: new_score = query_score x (factor x vote)

The command output is similar to the following:

```
{
  "took": 13,
  "timed_out": false,
  "_shards": {
    "total": 5,
    "successful": 5,
    "failed": 0
  },
  "hits": {
    "total": 4,
    "max_score": 2.1,
    "hits": [
      {
        "_index": "mall",
        "_type": "tv",
        "_id": "4",
        "_score": 2.1,
        "_source": {
          "name": "tv4",
          "description": "USB, HDMI, DisplayPort",
          "vote": 0.7
        }
      },
      {
        "_index": "mall",
        "_type": "tv",
        "_id": "2",
        "_score": 1.98,
        "_source": {
          "name": "tv2",
          "description": "USB, HDMI",
          "vote": 0.99
        }
      },
      {
        "_index": "mall",
        "_type": "tv",
        "_id": "1",
        "_score": 1.96,
        "_source": {
          "name": "tv1",
          "description": "USB, DisplayPort",
          "vote": 0.98
        }
      },
      {
        "_index": "mall",
        "_type": "tv",
```



```
      "_id": "3",
      "_score": 0.5,
      "_source": {
        "name": "tv3",
        "description": "USB",
        "vote": 0.5
      }
    }
  ]
 }
}
```

The preceding command output shows that Elasticsearch computes the total score of each TV based on the absolute praise rate and then sorts the TVs in a descending order of the total score.

- Compute the total score of each TV based on the relative praise rate and sort the searched products in a descending order of the total score.

  The following sample code illustrates how to query TVs with the USB, HDMI, or DisplayPort port. In this example, give score 0 for a TV if it has none of the three ports, 1 if it has one, 2 if it has two, and 3 if it has three. Then, multiply the obtained score by the TV's relative praise rate to obtain its total score. The relative praise rate is calculated as follows: Specify a praise rate threshold like 0.8. If a TV's absolute praise rate is higher than 0.8, take the praise rate as 1 in the calculation formula for the TV. Otherwise, take the praise rate as 0.5. Finally, sort the TVs in a descending order of their total scores.

  ```
  GET /mall/tv/_search?pretty
  {
   "query":{
    "function_score":{
     "query":{
      "bool":{
       "should":[
        {"constant_score":{
            "query":{"match":{"description":"USB"}}
        }},
        {"constant_score":{
            "query":{"match":{"description":"HDMI"}}
        }},
        {"constant_score":{
            "query":{"match":{"description":"DisplayPort"}}
        }}
       ]
      }
     },
     "script_score":{
      "script":{
       "params":{"threshold":0.8},
       "inline":"if (doc[\"vote\"].value > params.threshold) {return 1;} return 0.5;"
      }
     },
     "boost_mode":"multiply",
     "max_boost":10
    }
   }
  }
  ```

  In the preceding example, the total score is calculated as follows: new_score = query_score x vote (If the value of **vote** is greater than 0.8, take its value as 1 in the formula. If the value of **vote** is not greater than 0.8, take its value as 0.5 in the formula.)

  The command output is similar to the following:


```
{
  "took": 634,
  "timed_out": false,
  "_shards": {
    "total": 5,
    "successful": 5,
    "failed": 0
  },
  "hits": {
    "total": 4,
    "max_score": 2,
    "hits": [
      {
        "_index": "mall",
        "_type": "tv",
        "_id": "2",
        "_score": 2,
        "_source": {
          "name": "tv2",
          "description": "USB, HDMI",
          "vote": 0.99
        }
      },
      {
        "_index": "mall",
        "_type": "tv",
        "_id": "1",
        "_score": 2,
        "_source": {
          "name": "tv1",
          "description": "USB, DisplayPort",
          "vote": 0.98
        }
      },
      {
        "_index": "mall",
        "_type": "tv",
        "_id": "4",
        "_score": 1.5,
        "_source": {
          "name": "tv4",
          "description": "USB, HDMI, DisplayPort",
          "vote": 0.7
        }
      },
      {
        "_index": "mall",
        "_type": "tv",
        "_id": "3",
        "_score": 0.5,
        "_source": {
          "name": "tv3",
          "description": "USB",
          "vote": 0.5
        }
      }
    ]
  }
}
```


The preceding command output shows that Elasticsearch computes the total score of each TV based on the relative praise rate and then sorts the TVs in a descending order of the total score.

3.3 How Do I Migrate the Entire Elasticsearch Database to CSS?

For details about how to migrate the entire Elasticsearch database to CSS, see ***Cloud Data Migration User Guide***.

3.4 Can I Migrate CSS Data to Other Tenants?

Data cannot be directly migrated. However, you can back up and restore data using OBS buckets. For details, see **Index Backup and Restoration**.

3.5 Does CSS Have a Data Collection Client Available?

CSS integrates Elasticsearch and Kibana. You can install Logstash to collect data.

3.6 Does CSS Support Data Migration Between VPCs?

Data cannot be directly migrated on CSS. However, you can back up and restore data by backing up and restoring data. For details, see **Index Backup and Restoration**. Alternatively, you can create a VPC peering connection and then use Logstash to migrate data.

3.7 How Can I Find All Snapshots in an OBS Bucket When I Create a Snapshot for a CSS Cluster?

After a snapshot is created for a CSS cluster, it is incrementally created based on the original snapshot. **index-***x* records the information about each snapshot. **snap-***xxxx***.dat** and **meta-***xxxx***.dat** record the metadata of the snapshot. **indices** stores the index data.

3.8 Will an EIP Be Automatically Used After It Is Associated with a CSS Cluster?

Yes.

3.9 Can I Set the Maximum Data Volume to Be Queried on CSS?

Add **http.max_connect_length :100MB** to custom parameters and change its value. For details, see **Configuring Parameters**.

3.10 Can I Use the CSS Cluster Data as a Data Source and Obtain It?

Yes, CSS is fully compatible with all open-source APIs of Elasticsearch.

3.11 Can a Tool on HUAWEI CLOUD Synchronize RDS Data to CSS in Real Time?

Currently, only Cloud Data Migration (CDM) supports data migration. However, it does not support data import to CSS. For details about data import, see **Using CDM to Import Data from OBS to Elasticsearch**.

4 Clusters in Security Mode

4.1 What Is the Relationship Between the Filebeat Version and Cluster Version?

- Clusters in non-security mode: not limited.
- Cluster in security mode: The Filebeat OSS version must match the cluster version. For details on how to download the Filebeat OSS version, see **Past Releases of Elastic Stack Software**.

4.2 Can I Enable Security Mode Later to Add Authentication and Use SSL Communication If I Did Not Enable It When I Purchase a CSS Cluster for the First Time?

Security mode cannot be enabled after a cluster is created. To enable security mode, you need to create a cluster in security mode. For details, see **Clusters in Security Mode**.

5 Resource Usage and Change

5.1 How Do I Clear Expired Data to Release Storage Space?

- Run the following command to delete a single index data record.

  curl -XDELETE http://IP:9200/*Index name*

- Run the following command to delete all Logstash data of a day. For example, delete all data on June 19, 2017:

  curl -XDELETE 'http://IP:9200/logstash-2017.06.19*'

5.2 How Do I Configure a Two-Replica CSS Cluster?

1. Run the **GET _cat/indices?v** command in Kibana to check the number of cluster replicas. If the value of **rep** is **1**, the cluster has two replicas.

   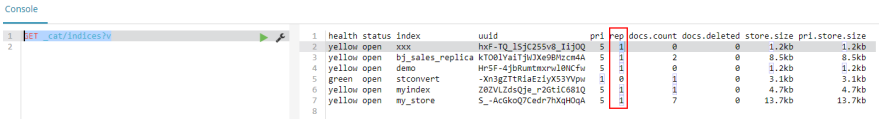

2. If the value of **rep** is not **1**, run the following command to set the number of replicas:

   PUT /index/_settings

   {

   "number_of_replicas":1 //Number of replicas

   }

   NOTE

   index specifies the index name. Set this parameter based on site requirements.

5.3 How Do I Delete Index Resources?

- Manual deletion: Run the **DELETE /my_index** command in Kibana.

- Automatic periodic deletion: Create scheduled tasks to call the index deletion request and periodically execute the tasks.

5.4 Can I Change the Number of Shards to Four with Two Replicas When There Is One Shard Set in the JSON File?

Once an index is created, the number of primary shards cannot be changed.

You can run the following command in Kibana to change the number of replicas:

```
PUT /indexname/_settings
{
"number_of_replicas":1 //Number of replicas
}
```

NOTE

index specifies the index name. Set this parameter based on site requirements.

6 Components

6.1 Does CSS Support APM?

CSS can integrate Application Performance Management (APM), but you need to install and deploy the APM server.

6.2 Does the CSS Cluster 6.5.4 Support Filebeat 6.1.1 (amd)?

CSS clusters and Filebeat of versions 6.*x.x* are compatible with each other. The actual situation depends on site requirements and features. You are advised to use the same version.

6.3 Can I Install Search Guard on CSS?

CSS does not support installation of Search Guard currently.

CSS provides clusters in security mode, which have the same functions as Search Guard. For details about clusters in security mode, see **Clusters in Security Mode**.

7 Kibana

7.1 What Are the Username and Password for Logging in to Kibana and Cerebro?

If you create a cluster in security mode, the username for logging in to Kibana and Cerebro is **admin**, and the password is the one set when you create the cluster. If you forget the password, you can reset it on the cluster details page and then log in. For details on how to change the password, see **Clusters in Security Mode**.

8 Clusters

8.1 Does CSS Support APIs or Functions of Open-Source Elasticsearch?

Yes. Versions 5.5.1, 6.2.3, 6.5.4, 7.1.1, and 7.6.2 are supported. The Elasticsearch Cloud mode is adopted. Kibana can interconnect with Elasticsearch, and functions of open-source Elasticsearch 5.5.1, 6.2.3, 6.5.4, 7.1.1 and 7.6.2 are supported.

8.2 Can CSS Interconnect with Logstash?

Yes. Logstash 5.5.1, 6.2.3, 6.5.4, 7.1.1, and 7.6.2 are recommended. You need to apply for an ECS to install and configure Logstash.

8.3 What Should I Do If an ECS Cannot Connect to a Cluster?

Perform the following steps:

1. Check whether the ECS instance and cluster are in the same VPC.
   - If yes, go to **2**.
   - If not, create an ECS instance and ensure that the ECS instance is in the same VPC as the cluster.
2. View the security group rule setting of the cluster to check whether port **9200** (TCP protocol) is allowed or port **9200** is included in the port range allowed in both the outbound and inbound directions.
   - If yes, go to **3**.
   - If not, switch to the VPC management console and set the security group rule of the cluster to allow port **9200** in both the outbound and inbound directions.
3. Check whether the ECS instance is added to a security group.
   - If added to a security group, check whether the security group configuration rules meet the requirements. For details, see the description

of **Security Group** in the cluster information table in **Clusters**. Then go to **4**.

Figure 8-1 Viewing security group information

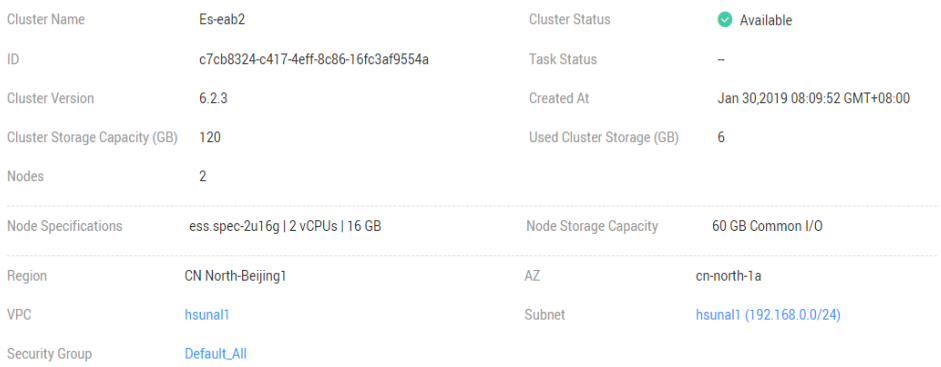

- If not, go to the VPC page from the ECS instance details page, select a security group, and add the security group.

4. Check whether the ECS instance can connect to the cluster.

   ssh *<Private network address and port number of a node>*

   NOTE

   If the cluster contains multiple nodes, check whether the ECS can be connected to each node in the cluster.

   - If the connection is normal, the network is running properly.
   - If the port is unreachable, contact the customer service.

8.4 Which Search Functions Does CSS Support?

CSS supports the following search functions: full-text search, highlighting, facet search, near-real-time search, dynamic clustering, processing of documents of various types like word and PDF, and geographic information search.

For details about the search function supported by Elasticsearch, see section "Search in Depth" in the ***Elasticsearch Reference***.

8.5 Can a New Cluster Use the IP Address of the Old Cluster?

No.

8.6 Can I Associate My EIP If I Want to Access the Cluster from the Internet?

No. To access a cluster from the Internet, see **Public IP Address Access**.

8.7 Can I Use Self-built Kibana and Cerebro to Access CSS Clusters?

Yes.

8.8 Can I Use x-pack-sql-jdbc to Access CSS Clusters and Query Data?

No. Currently, CSS does not integrate the x-pack component.

9 Ports

9.1 Do Ports 9200 and 9300 Both Open?

Yes. Port 9200 is used by external systems to access CSS clusters, and port 9300 is used for communication between nodes.

The methods for accessing port 9300 are as follows:

- If your client is in the same VPC and subnet with the CSS cluster, you can directly access port 9300.
- If your client is in the same VPC with but different subnet from the CSS cluster, apply for a route separately.
- If your client is in the different VPCs and subnets from the CSS cluster, create a VPC peering connection to enable communication between the two VPCs, and then apply for routes to connect the two subnets.

10 Network

10.1 How Do I Use a NAT Gateway to Access CSS from the Internet?

Operation process

1.**Obtaining CSS Information**

2.**Configuring a NAT Gateway**

3.**Modifying Security Group Rules**

4.**Accessing CSS from the Internet**

CAUTION

If your CSS clusters are in non-security mode, do not access CSS through the NAT gateway to prevent data of the cluster from being exposed to the Internet.

Obtaining CSS Information

Step 1 Obtain the private network address for accessing CSS.

Log in to the CSS management console, create a cluster, and then obtain the private network address of the cluster on the **Clusters** page.

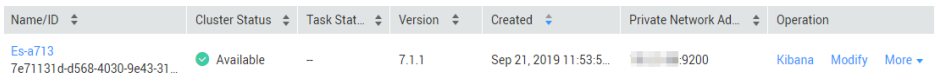

Step 2 Obtain the VPC and subnet information.

Click the desired cluster name/ID to go to the **Basic Information** page and obtain the VPC and subnet information.

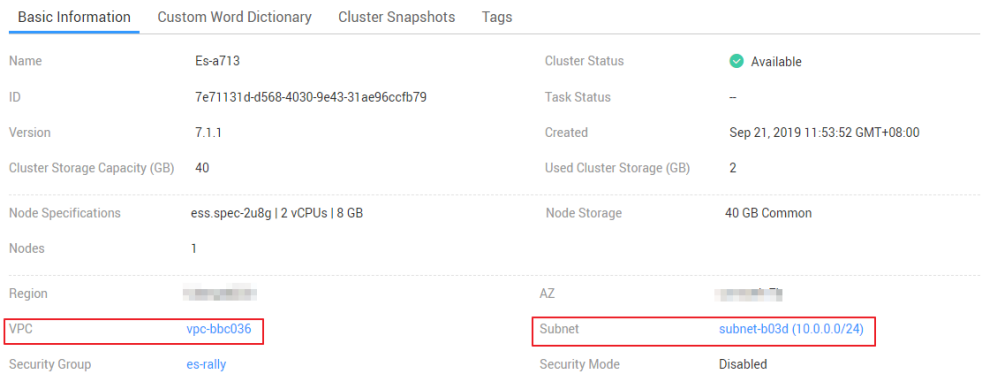

----End

Configuring a NAT Gateway

Step 1 Create a NAT gateway.

1. Log in to the console and choose **Service List** > **Network** >**NAT Gateway**. The **Network Console** page is displayed.
2. Click **Buy NAT Gateway**. On the displayed page, configure related parameters of the NAT gateway. For details, see section "Buying a NAT Gateway" in the *NAT Gateway User Guide*.

   NOTE

   Set **VPC** and **Subnet** to the values obtained in **Step 2**.
3. Click **Next**, confirm the configurations, and click **Submit**.

Step 2 Add DNAT rules.

1. On the **NAT Gateway** console, click the name of the NAT gateway you purchased. The details page is displayed.
2. Choose **DNAT Rules** > **Add DNAT Rule**. For details, see section "Adding a DNAT Rule" in the *NAT Gateway User Guide*.

   NOTE

   - **EIP**: Create an EIP on the **EIPs** page based on your service requirements.
   - **Outside Port**: Custom.
   - **Private IP Address**: private network IP address of CSS, which is **Private Network Address** obtained in **Step 1**.
   - **Inside Port**: 9200.
   - If your cluster contains multiple private IP addresses, add one DNAT rule for each address.
3. Click **OK**.

----End

Modifying Security Group Rules

Step 1 Log in to the CSS management console. Switch to the **Clusters** page. Click the name/ID of the target cluster to switch to the **Basic Information** page.

Step 2 On the displayed page, click the value of the **Security Group** parameter.

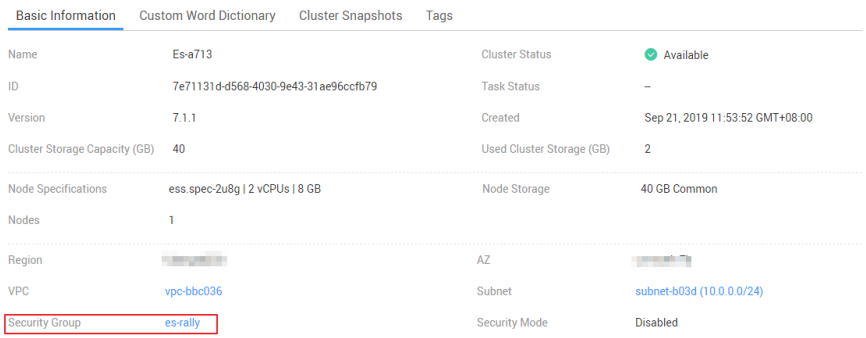

Step 3 On the displayed page, click **Inbound Rules**.

Step 4 Click **Add Rule** to add an inbound rule for port 9200.

Step 5 After the configuration is completed, click **OK**.

----End

Accessing CSS from the Internet

Step 1 Enter **https://*IP*:*port*** or **http://*IP*:*port*** in the address box of the browser.

- ***IP*** refers to EIP and ***port*** refers to the port number, both of which were set when adding DNAT rules.
- If **Security Mode** is enabled for the cluster, enter **https://*IP*:*port*** and then enter the username and password set for security mode on the displayed page.
- If **Security Mode** is not enabled for the cluster, enter **https://*IP*:*port***.

----End

10.2 Can I Change the Security Group of a CSS Cluster?

No.

11 Change History

Release Date	Description
2020-07-14	This is the second official release.
2020-04-26	This is the first official release.

