Cloud Search Service

Troubleshooting

Issue	03
Date	2020-11-19

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 Clusters

1.1 What Should I Do If the Access to Kibana Fails?

Symptom

After I click **Kibana** in the **Operation** column in the row where cluster **Es-event** resides on the **Clusters** page of the CSS management console, the Kibana page fails to be loaded and access to Kibana fails.

Fault Locating

The browser cache is not cleared.

Procedure

1. Log in to the CSS management console.
2. In the left navigation pane, click **Clusters**.
3. On the displayed **Clusters** page, locate the row where cluster **Es-event** resides and click **Kibana** in the **Operation** column.
4. On the displayed **Kibana** page, press **F12**.
5. Click **Network**, right-click **data:image**, and choose **clear browser cache** from the shortcut menu. In the displayed dialog box, click **OK**. Close the Kibana window.
6. Switch to the **Clusters** page, locate the row where cluster **Es-event** resides and click **Kibana** in the **Operation** column.

1.2 Filebeat Configuration Optimization

Symptom

Filebeat is a high-performance file collection tool. By default, one core is allocated to Filebeat, and it writes 1 MB data to Elasticsearch per second. However, in real practice, when a large number of service logs are generated, Filebeat cannot

promptly collect and write them to Elasticsearch. In this case, you can optimize parameter settings in the **filebeat.yml** file to improve the Filebeat performance.

Fault Locating

For Filebeat, the default configuration of the **filebeat.yml** file cannot deliver optimal performance in handling a large number of logs. In such scenarios, modify parameter settings in the **filebeat.yml** file to meet your demands.

Procedure

1. Optimize the parameters involved in **input** of the **filebeat.yml** configuration file.

   # Increase the value of **harvester_buffer_size** based on actual requirements. This parameter defines the buffer size used by every **harvester**.

   harvester_buffer_size: 40,960,000

   # Increase the value of **filebeat.spool_size** based on actual requirements. This parameter defines the number of log records that can be uploaded by the **spooler** at a time.

   filebeat.spool_size: 250,000

   # Adjust the value of **filebeat.idle_timeout** according to the actual requirements. This parameter defines how often the **spooler** is flushed. After the **idle_timeout** is reached, the **spooler** is flushed regardless of whether the **spool_size** has been reached.

   filebeat.idle_timeout: 1s

2. Optimize the parameters involved in **output.elasticsearch** in the **filebeat.yml** configuration file.

   # Set the value of **worker** to the number of Elasticsearch clusters according to the actual situation. The **worker** parameter indicates the number of Elasticsearch clusters. The default value is **1**.

   worker: 1

   # Increase the value of **bulk_max_size** based on the actual requirements. This parameter defines the maximum number of events to bulk in a single Elasticsearch bulk API index request. The default is **50**.

   bulk_max_size: 15,000

   # Adjust the value of **flush_interval** based on the actual requirements. This parameter defines the number of seconds to wait for new events between two bulk API index requests. If **bulk_max_size** is reached before this interval expires, additional bulk index requests are made.

   flush_interval: 1s

1.3 Why Do I Fail to Access CSS Using TransportClient?

Issue

Accessing CSS fails using the spring data elasticsearch mode and an error message "None of the configured nodes are available" is reported.

Symptom

Accessing CSS fails using the spring data elasticsearch mode and an error message is reported.

Possible Causes

Generally, **cluster.name** needs to be configured when you access clusters using TransportClient. The possible cause of access failure is that the **Settings.EMPTY** option is used or the setting is incorrect.

Procedure

For details on how to access clusters using a client, see **Accessing a Cluster > Accessing a Cluster Using the Java API in Security Mode with Elasticsearch**.

1.4 What Should I Do When a Cluster Is Unavailable?

Issue

A CSS cluster is in the unavailable status.

Symptom

A cluster is displayed as **Unavailable** in the **Cluster Status** column, as shown in the following figure.

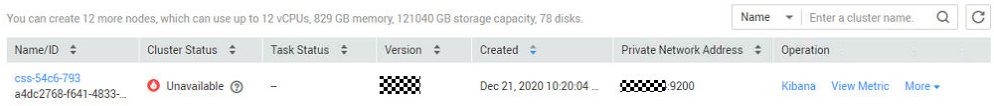

Possible Causes

When a cluster is unavailable, the CSS backend will report it to the console. The possible causes are as follows:

- The cluster is abnormal or faulty.
- The value of parameter **status** is **red**.

Procedure

- **When you can log in to Kibana or Cerebro,**
  a. Log in to Kibana or Cerebro of the target cluster.
  b. Click **Dev Tools** in the navigation tree on the left of Kibana, as shown in the following figure.

Figure 1-1 Clicking Dev Tools

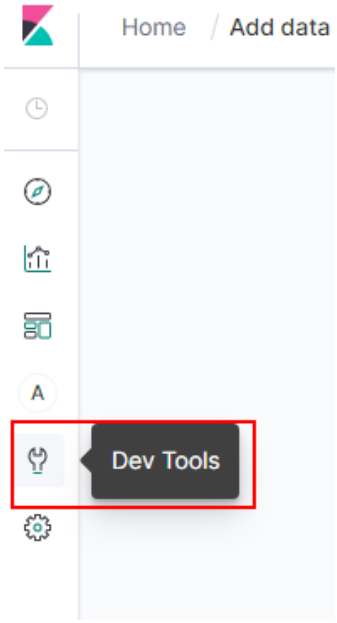

c. Run the **GET _cluster/health?pretty** command on the **Dev Tools** page. Check the value of parameter **status**.

Figure 1-2 Checking the value of parameter status

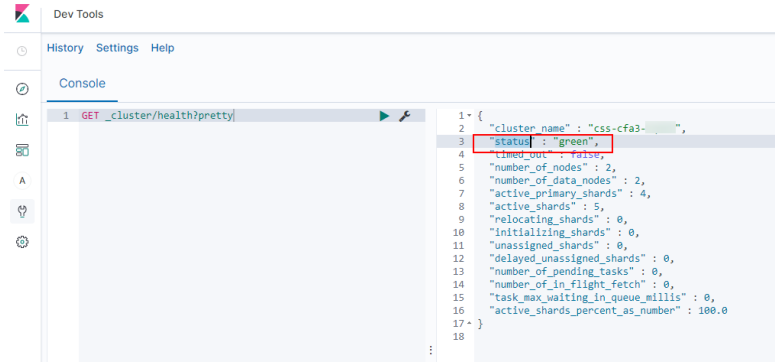

- If the value of **status** is **green**, the cluster is normal. The cluster status is detected once every minute, which may cause a time error. You can wait for several minutes and then check the status. If the cluster is still unavailable, it is an error. In this case, submit a service ticket or contact technical support.
- If the value of **status** is **red**, shards are not appropriately assigned. In this case, note the number of initializing and unassigned shards, as shown in the following figure:


```
{
  "cluster_name" : "css-cfa3-cqian",
  "status" : "green",
  "timed_out" : false,
  "number_of_nodes" : 2,
  "number_of_data_nodes" : 2,
  "active_primary_shards" : 4,
  "active_shards" : 5,
  "relocating_shards" : 0,
  "initializing_shards" : 0,
  "unassigned_shards" : 0,
  "delayed_unassigned_shards" : 0,
  "number_of_pending_tasks" : 0,
  "number_of_in_flight_fetch" : 0,
  "task_max_waiting_in_queue_millis" : 0,
  "active_shards_percent_as_number" : 100.0
}
```

NOTE

A cluster has the following three statuses:

- **green** indicates a cluster is normal.
- **yellow** indicates the replica shard is abnormal.
- **red** indicates the primary shard is abnormal, which often causes cluster unavailability.

d. Run the **GET /_cluster/allocation/explain?pretty** command to check the reason why shards are unassigned.

You can also view the number of unassigned or initializing shards in Cerebro. You can also run commands on the **rest** page of Cerebro. For details on how to log in to Cerebro, see **Cerebro**.

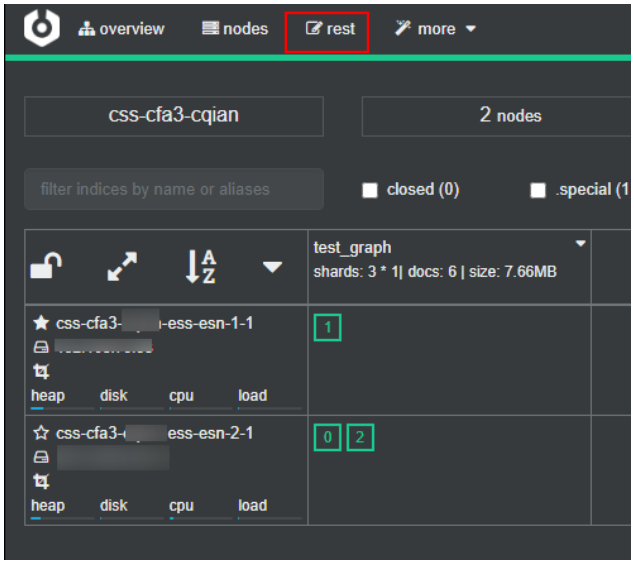

In the following scenarios, a primary shard of a cluster may be unavailable:

- **Shards are not assigned. Possible reasons are as follows:**
  - INDEX_CREATED: An API for creating an index is called. In this case, check whether the disk usage exceeds 85%. If it exceeds 85%, CSS will not assign new shards to the node.
  - CLUSTER_RECOVERED: full data restoration of the cluster is performed.

- INDEX_REOPENED: An index is opened or closed.
- DANGLING_INDEX_IMPORTED: The dangling index results are imported.
- NEW_INDEX_RESTORED: Data is restored to a new index.
- EXISTING_INDEX_RESTORED: Data is restored to disabled indices.
- REPLICA_ADDED: Replica shards are added explicitly.
- ALLOCATION_FAILED: Shard assignment fails.
- NODE_LEFT: The node that carries the shard is located outside of the cluster.
- REINITIALIZED: Misoperations (such as the use of the shadow replica shard) exist in the process from moving the shard to the shard initialization.
- REROUTE_CANCELLED: The assignment is canceled because the routing is canceled explicitly.
- REALLOCATED_REPLICA: It is determined that a better replica location will be used, and the existing replica assignment is canceled.

Shards are initializing. Possible causes are as follows:

This situation rarely occurs. Generally, translog may result in initializing primary shards. When a primary shard starts, it loads the translog file in the folder. If the translog file is too large, this process lasts for a long time.

- **When you cannot log in to Kibana or Cerebro,**

  CSS greatly improves cluster robustness. When a node is faulty, the daemon process of the node attempts to rectify the fault. Only when the rectification fails, an alarm indicating that the node is unavailable is reported. The possible causes for a rectification failure are as follows:

  a. The cluster load is too heavy and nodes frequently go offline. Go to the Cloud Eye console and view the cluster monitoring metrics, such as current and previous CPU usage, memory, and load. Focus on the trend and check whether there are sharp increases or the metric remains high for a long time. The surges may be caused by the sudden increase of access to the cluster. You can view the number of HTTP connections to learn about cluster access. If the load, CPU, or memory usage is high, the node may go offline.

  b. There are too many shards in the cluster. When shards are started, the metadata related to shards needs to be loaded to the memory. If there are too many shards, the memory usage is high. In addition, if a node is offline or a new index is created, the master node will consume more resources when calculating shard assignment because the number of shards is too large, increasing the cluster pressure.

  c. If the fault persists after several attempts, submit a service ticket to get technical support.

2 Data Import and Export

2.1 Why Is An Error Reported When I Import Data by Calling an API for Cluster 7.1.1?

Issue

When calling an API to import data, an error is reported.

Symptom

Failed to import data.

Possible Causes

The command used to create an index is incorrect.

Procedure

1. Run the following command to create an index:

   curl -X PUT -u admin:gh_-Ut=zm7H7A8E https://IP:9200/zxh_store -H 'Content-Type: application/json' -d '{ "settings": { "number_of_shards": 1 }, "mappings": { "properties": {"products": { "properties": { "productName": { "type": "text" }, "size": { "type": "keyword" } } } } } }' -k

2. If the error persists after an index is created, contact technical support.

2.2 What Should I Do When Logs Cannot Be Written to CSS and A High CPU Usage Is Displayed?

Issue

The CSS CPU usage is high, an error message "Elasticsearch Unreachable" is displayed on Logstash, and logs cannot be written to CSS.

Symptom

Logs cannot be written to CSS.

Possible Causes

The customer index has a single shard and the single node is overwhelmed. As a result, after the queuing is full, new jobs are rejected.

Procedure

1. Log in to the CSS management console.
2. Locate the target cluster and choose **More** > **Cerebro** in the **Operation** column of the cluster.

   If the cluster is in security mode, you need to enter the login username and password.
3. In Cerebro, view the number of shards in the cluster and metrics such as the CPU, load, head, and dis of each node.
4. Analyze the possible causes based on metrics and conduct optimization accordingly.
   a. Increase the number of queues and reduce rejected jobs by changing the value of **write.queue_size**.
      i. Log in to the CSS management console, click the target cluster name to go to the **Basic Information** page.
      ii. Click **Parameter Configurations**, search for **write.queue_size**, and change its value.

      If this parameter does not exist, add it in the **Customize** area. For details, see **Configuring Parameters**.
   b. Create indices again to ensure that the number of shards is greater than that of nodes in the cluster.
5. If the number of shards and queues are appropriate but the CPU usage and load are still high, you are advised to scale out the node.

2.3 Why Is An Error Reported When Pushing Data to CSS After Deploying Logstash on an ECS?

Issue

After Logstash is deployed on an ECS, an error is reported when data is pushed to CSS.

Symptom

After Logstash is deployed on an ECS, an error is reported when data is pushed to CSS. The error message is as follows:

LogStash::Outputs::ElasticSearch::HttpClient::Pool::BadResponseCodeError: Got response code '500' contacting Elasticsearch at URL 'https://192.168.xx.xx:9200/_xpack'.

Possible Causes

CSS currently does not integrate the x-pack plugin. When you access CSS after deploying Logstash, the system will check whether x-pack is enabled for CSS.

Procedure

1. Delete the x-pack directory in Logstash.
2. Add the configuration item **ilm_enabled => false** to **elasticsearch** under the **output** tag in the Logstash configuration file.
3. Push data to CSS again.

3 Functions

3.1 Why Cannot I Perform Index Backup?

Index backup is implemented by creating cluster snapshots. If you cannot perform index backup, perform the following steps:

Check the Time When the Target Cluster Is Created

1. Log in to the CSS management console.
2. In the left navigation pane, click **Clusters**.
3. On the **Clusters** page, view the creation time of the target cluster in the **Created** column.
   - If the creation time is earlier than March 10, 2018, no further action is required because the index backup and restoration function was unavailable at that time.
   - If the creation time is later than March 10, 2018, check whether the account or IAM user used for logging in to the management console has the index backup permission. For details, see **Check Whether the Account or IAM User Has the Index Backup Permission**.

Check Whether the Account or IAM User Has the Index Backup Permission

1. Log in to the IAM management console.
2. Check the user group, to which the account or the IAM user belongs.

   For details, see **Viewing and Modifying User Information** in the *Identity and Access Management User Guide*.
3. Check whether the permissions assigned to the user group contain the following: **Tenant Administrator** for project **OBS** in region **Global service** and **Elasticsearch Administrator** in the current region.

   For details, see **Viewing and Modifying User Group Information** in the *Identity and Access Management User Guide*.
   - If neither of the preceding permissions is contained, go to **4**.
   - If both of the preceding permissions are contained, contact the customer service.

4. Add the following permissions to the user group: **Tenant Administrator** for project **OBS** in region **Global service** and **Elasticsearch Administrator** in the current region.

   For details, see **Viewing and Modifying User Group Information** in the *Identity and Access Management User Guide*.

3.2 Why Cannot I Use The Custom Word Dictionary Function?

Perform the following steps to troubleshoot this problem:

Check the Time When the Target Cluster Is Created

Step 1 Log in to the CSS management console.

Step 2 In the left navigation pane, click **Clusters**.

Step 3 On the **Clusters** page, view the creation time of the target cluster in the **Created** column.

- If the creation time is earlier than March 10, 2018, no further action is required because the custom word dictionary function was unavailable at that time.
- If the creation time is later than March 10, 2018, check whether the account or IAM user used for logging in to the management console has the custom word dictionary permission. For details, see **Check Whether the Account or IAM User Has the Custom Word Dictionary Permission**.

----End

Check Whether the Account or IAM User Has the Custom Word Dictionary Permission

1. Log in to the IAM management console.
2. Check the user group, to which the account or the IAM user belongs.

   For details, see **Viewing and Modifying User Information** in the *Identity and Access Management User Guide*.
3. Check whether the permissions assigned to the user group contain the following: **Tenant Administrator** for project **OBS** in region **Global service** and **Elasticsearch Administrator** in the current region.

   For details, see **Viewing and Modifying User Group Information** in the *Identity and Access Management User Guide*.

   – If neither of the preceding permissions is contained, go to **4**.

   – If both of the preceding permissions are contained, contact the customer service personnel.
4. Add the following permissions to the user group: **Tenant Administrator** for project **OBS** in region **Global service** and **Elasticsearch Administrator** in the current region.

For details, see **Viewing and Modifying User Group Information** in the *Identity and Access Management User Guide*.

4 Ports

4.1 Why Do I Fail to Access the 9200 Port?

Symptom

In the scenario where a VPN private line or VPC peering connection is used to access the CSS cluster, no result is returned when the cURL command is used to connect to the CSS cluster.

For example, if you run the following command to connect to the cluster, no result is returned:

```
curl -s 'http://< node private access address >:9200'
```

Possible Cause

In the scenario where the VPN private line or VPC peering connection is used to access the CSS, the client and CSS are not in the same VPC. Therefore, the subnet of the CSS must be in a different network segment from that of the VPC.

For example, for a CSS cluster, the VPC is **vpc-8e28** and its network segment is **192.168.0.0/16**. The subnet **subnet-4a81** of the VPC is selected, whose network segment is **192.168.0.0/16**. This segment is the same as that of the **vpc-8e28**. In this case, if the VPN private line or the VPC peering connection is used to access the CSS cluster, the host created on the subnet does not have the gateway corresponding to the VPC. As a result, the default route of the CSS service is affected and the 9200 port fails to be accessed.

Procedure

When the 9200 port fails to be accessed and the CSS is available, the execution procedures are as follows:

1. Go to the CSS management console. In the cluster list, click the cluster name to view the VPC and subnet used by the cluster.
2. Go to the VPC management console. In the VPC list, click the name of the VPC used by the CSS cluster. The VPC details page is displayed. View the VPC and subnet network segment information.

As shown in **Figure 4-1**, the VPC network segment information is the same as the subnet network segment information. When a VPN private line or a VPC peer connection is used for access, the 9200 port fails to be accessed.

Figure 4-1 Viewing network segment information

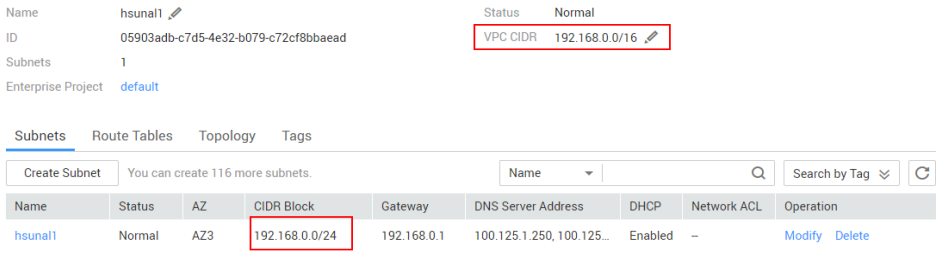

3. If the preceding error occurs, create a cluster again and select a subnet that is different from the VPC subnet. If the subnet does not exist, create another subnet on the VPC management console.

   After a new CSS cluster is created, migrate the data of the old cluster to the new cluster, and then use the VPN private line or the VPC peer to access the cluster.

   NOTE

   If the VPN private line access is required or the VPC peer connection is used to access the CSS cluster, ensure that the VPC and subnet of the newly created CSS have different network segment information.

5 Change History

Release Date	Description
2020-11-19	This is the third official release. • Added: **What Should I Do When a Cluster Is Unavailable?**
2020-08-10	This is the second official release. • Added: **Why Do I Fail to Access CSS Using TransportClient?** **Why Is An Error Reported When Pushing Data to CSS After Deploying Logstash on an ECS?**
2020-04-26	This is the first official release.

