Cloudera Runtime 7.1.8

Troubleshooting Apache Impala

Date published: 2020-11-30
Date modified: 2022-08-25

CLOUDERA

Legal Notice

Contents

Troubleshooting Impala

This topic describes the general troubleshooting procedures to diagnose some of the commonly encountered issues in Impala.

Symptom	Explanation	Recommendation
Impala takes a long time to start.	Impala instances with large numbers of tables, partitions, or data files take longer to start because the metadata for these objects is broadcast to all impalad nodes and cached.	Adjust timeout and synchronicity settings.
Joins fail to complete.	There may be insufficient memory. During a join, data from the second, third, and so on sets to be joined is loaded into memory. If Impala chooses an inefficient join order or join mechanism, the query could exceed the total memory available.	Start by gathering statistics with the COMPUTE STATS statement for each table involved in the join. Consider specifying the [SHUFFLE] hint so that data from the joined tables is split up between nodes rather than broadcast to each node. If tuning at the SQL level is not sufficient, add more memory to your system or join smaller data sets.
Queries return incorrect results.	Impala metadata may be outdated after changes are performed in Hive.	After inserting data, adding a partition, or other operation in Hive, refresh the metadata for the table with the REFRESH statement.
Attempts to complete Impala tasks such as executing INSE RT SELECT statements fail. The Impala logs include notes that files could not be opened due to permission denied.	This can be the result of permissions issues. For example, you could use the Hive shell as the hive user to create a table. After creating this table, you could attempt to complete some action, such as an INSERT SELECT on the table. Because the table was created using one user and the INSERT SELECT is attempted by another, this action may fail due to permissions issues.	Ensure the Impala user has sufficient permissions to the table that the Hive user created.
Impala fails to start up, with the impalad logs referring to errors connecting to the statestored service and attempts to re-register.	A large number of databases, tables, partitions, and so on can require metadata synchronization, particularly on startup, that takes longer than the default timeout for the stat estore service.	Configure the statestore timeout value and possibly other settings related to the frequency of statestore updates and metadata loading.

Using Breakpad Minidumps for Crash Reporting

The breakpad project is an open-source framework for crash reporting. Impala can use breakpad to record stack information and register values when any of the Impala-related daemons crash due to an error such as SIGSEGV or unhandled exceptions. The dump files are much smaller than traditional core dump files. The dump mechanism itself uses very little memory, which improves reliability if the crash occurs while the system is low on memory.

Using the Minidump Files for Problem Resolution

You can see in the Impala log files or in the Cloudera Manager charts for Impala when crash events occur that generate minidump files. Because each restart begins a new log file, the "crashed" message is always at or near the bottom of the log file. (There might be another later message if core dumps are also enabled.)

⚠️ **Important:** If an Impala-related daemon experiences a crash due to an out-of-memory condition, it does not generate a minidump for that error.

Using the Minidump Files for Problem Resolution

Typically, you provide minidump files to Cloudera Support as part of problem resolution, in the same way that you might provide a core dump. The Send Diagnostic Data under the Support menu in Cloudera Manager guides you through the process of selecting a time period and volume of diagnostic data, then collects the data from all hosts and transmits the relevant information for you.

Procedure

1. In Cloudera Manager, navigate to Impala serviceConfiguration.
2. In the search field, type minidump.
3. Set the following fields to configure breakpad minidumps.

   - minidump_path: Turn on or off generation of the minidump files.

     By default, a minidump file is generated when an Impala-related daemon crashes.
   - minidump_path: Specify the Location for minidump files.

     By default, all minidump files are written to the following location on the host where a crash occurs: /var/log/impala-minidumps/*daemon_name*

     The minidump files for `impalad`, `catalogd`, and `statestored` are each written to a separate directory.

     If you specify a relative path for this setting, the value is interpreted relative to the default minidump_path directory.
   - max_minidumps: Specify the number of minidump files.

     Like any files used for logging or troubleshooting, consider limiting the number of minidump files, or removing unneeded ones, depending on the amount of free storage space on the hosts in the cluster.

     Because the minidump files are only used for problem resolution, you can remove any such files that are not needed to debug current issues.

     The default for this setting is 9. A zero or negative value is interpreted as "unlimited".
4. Click Save Changes and restart Impala.
5. To provide minidump files to Cloudera Support as part of problem resolution, in Cloudera Manager, navigate to SupportSend Diagnostic Data and follow the steps.

Dynamic memory for Impala control service

Understand and configure memory allocation for Impala's control service queue to avoid errors and ensure consistent behavior across daemons with varying memory limits.

Condition

Impala daemon logs report errors when the control_service_queue_mem_limit is set to a smaller value:

```
impala-service-pool.cc: EndDataStream request on impala.DataStreamService fr
om xx.xx.xx.xx:xxxx dropped due to backpressure.
The service queue contains 0 items out of a maximum of xxxxxxxxxx; memory c
onsumption is 0.
```

Cause

Before the implementation of IMPALA-10343, the default value of control_service_queue_mem_limit was 50MB. After IMPALA-10343 (starting from versions 7.1.7 SP2, 7.1.8, and 7.1.9), the value became dynamic and is now set to 1% of mem_limit.

The control service queue memory limit became dynamic, adjusting based on the daemon's memory. To avoid unintentional reductions in memory allocation, use the --control_service_queue_mem_limit_floor_bytes flag. This ensures consistent behavior for daemons with smaller memory limits, while increasing the limit for those with larger memory.

Remedy

Procedure

1. Log in to **Cloudera Manager** as an administrator.
2. Navigate to Clusters Impala Configuration and search for Impala Daemon Command Line   Argument Adva nced Configuration Snippet (Safety   Valve).
3. Add the below value:

```
--control_service_queue_mem_limit=XXXMB
--control_service_queue_mem_limit_floor_bytes=XXMB (default is 50MB)
```

4. The memory limit can be adjusted using the previously mentioned configurations to accommodate higher workloads. Ensure that you restart the Impala service for the changes to take effect.