

Cloudera Runtime 7.3.1

Troubleshooting Apache Impala

Date published: 2020-11-30

Date modified: 2024-12-10

CLOUDERA

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Troubleshooting common issues in Impala.....	4
Using Breakpad Minidumps for Crash Reporting.....	4
Performance Issues Related to Data Encryption.....	5
Troubleshooting Crashes Caused by Memory Resource Limit.....	6

Troubleshooting common issues in 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 <code>impalad</code> nodes and cached.	Adjust timeout and synchronicity settings.
Query rejected with the default pool-defined memory limit settings.	Some complex queries fail because the minimum memory reservation per host is greater than the memory available to the query for buffer reservations.	Increase VW t-shirt size so that there are more hosts in the executor group and less memory is needed per host.
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 <code>COMPUTE STATS</code> statement for each table involved in the join. Consider specifying the <code>[SHUFFLE]</code> 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 <code>REFRESH</code> statement.
Attempts to complete Impala tasks such as executing <code>INSERT SELECT</code> 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 <code>INSERT SELECT</code> on the table. Because the table was created using one user and the <code>INSERT SELECT</code> 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 <code>impalad</code> logs referring to errors connecting to the statestore 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 statestore 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.

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 `statedored` 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.

Performance Issues Related to Data Encryption

To troubleshoot any performance issues related to data encryption, run the `PROFILE` statement directly after running a query to view the encrypted bytes read.

Viewing the encrypted bytes read using query profile

When Impala executes any query on an Ozone encrypted file, the query `PROFILE` captures the runtime details of the execution including the total number of encrypted bytes read from Ozone by the query.



Note: The `PROFILE` statement should be run immediately after running the query. You cannot revisit later and view an earlier session.

You can view query profiles from the command line by logging in to the Impala shell and running the profile; command. You can also collect from Cloudera Manager Web UI, by navigating to CM > Impala > Queries , locate the query you just ran and click on Query Details . Look for the excerpts “BytesReadEncrypted” in the profile output.

Troubleshooting Crashes Caused by Memory Resource Limit

Default settings for `max_map_count` can be insufficient for Impala installations that run with many concurrent queries. We recommend increasing `max_map_count` to avoid potential failures due to exhausting memory mapping limits under heavy load.

Increasing the virtual memory parameter

`max_map_count` is the OS virtual memory parameter and defines the maximum number of memory map areas that a process can use.

To increase the `max_map_count` parameter in your OS and to make the above setting durable, refer to your OS documentation.

- For example, if you are using RHEL 9, add the following line to `/etc/sysctl.conf`:

```
/etc/sysctl.conf:vm.max_map_count=8000000
```

- Reload the config as root: `sysctl -p`.
- Check the new value: `cat /proc/sys/vm/max_map_count`.
- Restart.