Cloudera Runtime 7.2.7

Managing Apache Impala

Date published: 2019-11-01
Date modified: 2021-02-02

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

ACID Operation.....5
Concepts Used in FULL ACID v2 Tables.....5
Key Differences between INSERT-ONLY and FULL ACID Tables.....6
Compaction of Data in FULL ACID Transactional Table.....7

Configuring Impala.....7

Monitoring Impala.....9
Impala Logs.....9
Managing Logs.....9
Impala lineage.....11
Web User Interface for Debugging.....11
Debug Web UI for Impala Daemon.....12
Debug Web UI for StateStore.....15
Debug Web UI for Catalog Server.....16
Configuring Impala Web UI.....18

Securing Impala.....20
Configuring Impala TLS/SSL.....21
Impala Authentication.....22
Configuring Kerberos Authentication.....23
Configuring LDAP Authentication.....24
Impala Authorization.....26
Configuring Authorization.....33

Tuning Impala.....33
Setting Up HDFS Caching.....33
Setting up Data Cache for Remote Reads.....34
Configuring Dedicated Coordinators and Executors.....34

Managing Resources in Impala.....36
Admission Control and Query Queuing.....36
Enabling Admission Control.....40
Creating Static Pools.....40
Configuring Dynamic Resource Pool.....40
Dynamic Resource Pool Settings.....41
Admission Control Sample Scenario.....43
Cancelling a Query.....45

Using HLL Datasketch Algorithms in Impala.....45

Using KLL Datasketch Algorithms in Impala.....49

Managing Metadata in Impala.....52
On-demand Metadata.....52
Automatic Invalidation of Metadata Cache.....53
Automatic Invalidation/Refresh of Metadata.....54
Configuring Event Based Automatic Metadata Sync.....56

READ Support for FULL ACID ORC Tables

FULL ACID v2 transactional tables are readable in Impala without modifying any configurations. You must have Cloudera Runtime 7.2.2 or higher and have connection to Hive Metastore server in order to READ from FULL ACID tables.

There are two types of transactional tables available with Hive ACID.

- INSERT-ONLY
- FULL ACID

Until this release, Impala in CDP supported INSERT-ONLY transactional tables allowing both READ and WRITE operations. The latest version of Impala in CDP now also supports READ of FULL ACID ORC tables.

By default tables created in Impala are INSERT-ONLY managed tables whereas the default tables in Hive are managed tables that are FULL-ACID.

Limitations

- Impala cannot CREATE or WRITE to FULL ACID transactional tables yet. You can CREATE and WRITE FULL ACID transactional tables with transaction scope at the row level via HIVE and use Impala to READ these tables.
- Impala does not support ACID v1.

Concepts Used in FULL ACID v2 Tables

Before beginning to use FULL ACID v2 tables you must be aware of these new concepts like transactions, WriteIds, rowIDs, delta delete directories, locks, etc. that are added to FULL ACID tables to achieve ACID semantics.

Write IDs

For every transaction, both read and write, Hive will assign a globally unique ID. For transactional writes like INSERT and DELETE, it will also assign a table-wise unique ID, a write ID. The write ID range will be encoded in the delta and delete directory names. Results of a DML transactional query are allocated to a location under partition/table. This location is derived by Write ID allocated to the transaction. This provides Isolation of DML queries and such queries can run in parallel without interfering with each other.

New Sub-directories

New data files resulting from a DML query are written to a unique location derived from WriteId of the transaction. You can find the results of an INSERT query in delta directories under partition/table location. Depending on the operation type there can be two types of delta directories:

- Delta Directory: This type is created for the results of INSERT statements and is named delta_<writeId>_<writeId> under partition/table location.
- Delete Delta Directory: This delta directory is created for results of DELETE statements and is named delete_delta_<writeId>_<writeId> under partition/table location.

UPDATE operations create both delete and delta directories.

Row IDs

rowId is the auto-generated unique ID within the transaction and bucket. This is added to each row to identify each row in a table. RowID is used during a DELETE operation. When a record is deleted from a table, the rowId of the deleted row will be written to the delete_delta directory. So for all subsequent READ operations all rows will be read except these rows.

Schematic differences between INSERT-ONLY and FULL ACID tables

INSERT-ONLY tables do not have a special schema. They store the data just like plain original files from the non-ACID world. However, their files are organized differently. For every INSERT statement the created files are put into a transactional directory which has transactional information in its name.

Full ACID tables do have a special schema. They have row identifiers to support row-level DELETEs. So a row in Full ACID format looks like this:

```
{
   "operation": 0,
   "originalTransaction": 1,
   "bucket": 536870912,
   "rowId": 0,
   "currentTransaction": 1,
   "row": {"i": 1}
   }
```

- The green columns are the hidden/system ACID columns.
- Field “row” holds the user data.
- operation 0 means INSERT, 1 UPDATE, and 2 DELETE. UPDATE will not appear because of the split-update technique (INSERT + DELETE).
- originalTransaction is the write ID of the INSERT operation that created this row.
- bucket is a 32-bit integer defined by BucketCodec class.
- rowId is the auto-generated unique ID within the transaction and bucket.
- currentTransaction is the current write ID. For INSERT, it is the same as currentTransaction. For DELETE, it is the write ID when this record is first created.
- row contains the actual data. For DELETE, row will be null.

Key Differences between INSERT-ONLY and FULL ACID Tables

Before beginning to use FULL ACID v2 tables you must be aware of the key differences between the INSERT-ONLY and FULL-ACID tables.

This table highlights some of the differences between the INSERT-ONLY and FULL ACID tables.

	INSERT-ONLY	FULL ACID
Schema	There is no special data schema. They store the data just like plain original files from the non-ACID world.	Data is in special format, i.e. there are synthetic columns with transactional information in addition to actual data.
Transactional information	Transactional information is encoded in directory names.	Full ACID tables also use the same directory structure as INSERT-only tables. Transactional information is encoded in the directory names. Directory name and filename are the source of transactional information.
Table properties	'transactional'='true', 'transactional_properties'='insert_only'	'transactional'='true'
Supported operations	INSERT-ONLY tables only support insertion of data. UPDATES and DELETES are not supported. These tables also provide CREATE TABLE, DROP TABLE, TRUNCATE, INSERT, SELECT operations.	FULL ACID ORC tables can be READ using IMPALA. These tables also provide UPDATE and DELETE operations at the row level using HIVE. This is achieved using transactions like Insert-Only Tables along with changes in ORC Reader to support deletes.
WRITE operation	WRITE operations are atomic and the results of the insert operation are not visible to other query operations until the operation is committed.	WRITE operations are atomic - The operation either succeeds completely or fails; it does not result in partial data.

	INSERT-ONLY	FULL ACID
INSERT operation	For every INSERT statement the created files are added to a transactional directory which has transactional information in its name.	INSERT operation is done through HIVE and this statement is executed in a single transaction. This operation creates a delta directory containing information about this transaction and its data.
DELETE operation	N/A	DELETE operation is done through HIVE and this event creates a special “delete delta” directory.
UPDATE operation	N/A	UPDATE operation is done through HIVE. This operation is split into an INSERT and DELETE operation. This operation creates a delta dir followed by a delete dir.
READ operation	READ operations always read a consistent snapshot of the data.	READ operations always read a consistent snapshot of the data.
Supported file format	Supports any file formats.	Supports only ORC.
Compactions	Minor and major compactions are supported.	Minor compactions can be created, which means several delta and delete directories can be compacted into one delta and delete directory. Major compactions are also supported.

File structure of FULL ACID transational table

Hive 3 achieves atomicity and isolation of operations on transactional tables by using techniques in write, read, insert, create, delete, and update operations that involve delta files, which can provide query status information and help you troubleshoot query problems. For more information, see Hive 3 ACID transactions

Compaction of Data in FULL ACID Transactional Table

As administrator, you need to manage compaction of delta files that accumulate during data ingestion. Compaction is a process that performs critical housekeeping of files.

Hive creates a set of delta files for each transaction that alters a table or partition and stores them in a separate delta directory. When the number of delta and delete directories in the table grow, the read performance will be impacted, since reading is a process of merging the results of valid transactions. To avoid any compromise on the read performance, occasionally Hive performs compaction, namely minor and major. This process merges these directories while preserving the transaction information.

Configuring Impala

You must review and configure the mandatory and recommended settings if you installed Impala without Cloudera Manager or if you want to customize your environment after installing Impala. If you installed Impala using Cloudera Manager, some of these configurations are completed automatically.

In some cases, depending on the level of Impala, CDP, and Cloudera Manager, you might need to add particular component configuration details in one of the free-form fields on the Impala configuration pages in Cloudera Manager.

- You must enable short-circuit reads, whether or not Impala was installed through Cloudera Manager. This setting goes in the Impala configuration settings, not the Hadoop-wide settings.
- If you installed Impala in an environment that is not managed by Cloudera Manager, you must enable block location tracking, and you can optionally enable native checksumming for optimal performance.

Short-Circuit Reads

Enabling short-circuit reads allows Impala to read local data directly from the file system. This removes the need to communicate through the DataNodes, improving performance. This setting also minimizes the number of additional

copies of data. Short-circuit reads requires libhadoop.so (the Hadoop Native Library) to be accessible to both the server and the client. You must install it from an .rpm, .deb, or parcel to use short-circuit local reads.

Note: If you use Cloudera Manager, you can enable short-circuit reads through a checkbox in the user interface and that setting takes effect for Impala as well.

To Enable Short-Circuit Reads

You can enable short-circuit reads through a checkbox available under Configuration tab for both Impala and HDFS.

1. In Cloudera Manager Clusters select the Impala service, for example, IMPALA.
2. On the Configuration tab, search for dfs.client.read.shortcircuit.
3. Accept the default (enabled), or check to enable the dfs.client.read.shortcircuit property to read the HDFS file blocks directly.
4. Do the above for HDFS service too by clicking Cloudera Manager Clusters and by selecting the HDFS service.
5. Save the changes and Restart the service.

To configure DataNodes for short-circuit reads:

1. On all Impala nodes, configure the following attributes from the HDFS service as shown:
   a. In Cloudera Manager Clusters select the HDFS service, for example, HDFS.
   b. On the Configuration tab, search for dfs.domain.socket.path and set the attribute.
   c. Search for and set the attribute, if necessary, dfs.client.file-block-storage-locations.timeout.millis.
   d. Search for and set the attribute, if necessary, dfs.datanode.hdfs-blocks-metadata.enabled
2. After applying these changes, restart all DataNodes.

Block Location Tracking

Enabling block location metadata allows Impala to know which disk data blocks are located on, allowing better utilization of the underlying disks. Impala will not start unless this setting is enabled.

To enable block location tracking:

1. For each DataNode, enable the following attribute dfs.datanode.hdfs-blocks-metadata.enabled file:
   a. In Cloudera Manager Clusters select the HDFS service, for example, HDFS.
   b. On the Configuration tab, search for dfs.datanode.hdfs-blocks-metadata.enabled and enable the attribute if not already enabled.
2. After applying these changes, restart all DataNodes.

Native Checksumming

Enabling native checksumming causes Impala to use an optimized native library for computing checksums, if that library is available.

To enable native checksumming:

If you installed CDP from packages, the native checksumming library is installed and setup correctly, and no additional steps are required.

If you installed by other means, native checksumming may not be available due to missing shared objects. Finding the message "Unable to load native-hadoop library for your platform... using builtin-java classes where applicable" in the Impala logs indicates native checksumming may be unavailable.

To enable native checksumming, you must build and install libhadoop.so (the Hadoop Native Library).

Monitoring Impala

As an administrator, you monitor Impala service and take action when necessary to keep it running smoothly and avoid conflicts with other components running on the same cluster.

Impala Logs

You can review Impala log files on each host, when you have traced an issue back to a specific system. The Impala logs contain information about any errors Impala encountered, jobs Impala has completed, and settings Impala is configured with.

The logs store information about Impala startup options. This information appears once for each time Impala is started and may include:

- Machine name.
- Impala version number.
- Flags used to start Impala.
- CPU information.
- The number of available disks.

A new set of log files is produced each time the associated daemon is restarted. These log files have long names including a timestamp. The .INFO, .WARNING, and .ERROR files are physically represented as symbolic links to the latest applicable log files.

Review Impala log files on each host, when you have traced an issue back to a specific system:

- By using the web interface at http://host-name:25000/logs where host-name is your Cloudera cluster host name.

  The web interface limits the amount of logging information displayed. To view every log entry, access the log files directly through the file system. Impala log files can often be several megabytes in size.
- By examining the contents of the log file

  By default, the Impala logs are stored at /var/log/impalad/, /var/log/catalogd/, and /var/log/statestore/. The most comprehensive log, showing informational, warning, and error messages, is in the file name impalad.INFO.

  For each of impalad, statestored, catalogd:

  - Examine the .INFO files to see configuration settings for the processes.
  - Examine the .WARNING files to see all kinds of problem information, including such things as suboptimal settings and also serious runtime errors.
  - Examine the .ERROR and/or .FATAL files to see only the most serious errors, if the processes crash, or queries fail to complete. These messages are also in the .WARNING file.

Cloudera Manager collects front-end and back-end logs together into a single view and let you do a search across log data for all the managed nodes in DiagnosticsLogs.

Managing Logs

You must configure the Impala log settings to change the default log locations, to rotate logs, or to log verbose levels.

Procedure

1. To change log file locations:
   a) In Cloudera Manager, navigate to Impala serviceConfiguration.
   b) In the search field, type log_dir.
   c) Specify the new log directories for Impala Daemon, Catalog Server, or StateStore in the respective fields:
      - Impala Daemon Log Directory
      - Catalog Server Log Directory
      - StateStore Log Directory
   d) Click Save Changes and restart Impala.
2. To set up log rotation:
   a) In Cloudera Manager, navigate to Impala serviceConfiguration.
   b) In the search field, type max_log_files.
   c) Specify the values for Impala Daemon, Catalog Server, or StateStore in the respective fields:
      - Impala Daemon Maximum Log Files
      - Catalog Server Maximum Log Files
      - StateStore Maximum Log Files
   d) Click Save Changes and restart Impala.

   The above configuration option specifies how many log files to keep at each severity level (INFO, WARNING, ERROR, and FATAL).

   - A value of 0 preserves all log files, in which case you would set up set up manual log rotation using your Linux tool or technique of choice.
   - A value of 1 preserves only the very latest log file.
   - The default value is 10.

   For some log levels, Impala logs are first temporarily buffered in memory and only written to disk periodically. The --logbufsecs setting controls the maximum time that log messages are buffered for. For example, with the default value of 5 seconds, there may be up to a 5 second delay before a logged message shows up in the log file.

   It is not recommended that you set --logbufsecs to 0 as the setting makes the Impala daemon to spin in the thread that tries to delete old log files.
3. To specify how often the log information is written to disk:
   a) In Cloudera Manager, navigate to Impala serviceConfiguration.
   b) In the search field, type logbuflevel.
   c) Specify the values for Impala Daemon, Catalog Server, or StateStore in the respective fields:
      - Impala Daemon Log Buffer Level
      - Catalog Server Log Buffer Level
      - StateStore Log Buffer Level

      The default is 0, meaning that the log is immediately flushed to disk when Impala outputs an important messages such as a warning or an error, but less important messages such as informational ones are buffered in memory rather than being flushed to disk immediately.
   d) Click Save Changes and restart Impala.

4. To set the logging verbose levels:
   a) In Cloudera Manager, navigate to Impala serviceConfiguration.
   b) In the search field, type GLOG_v.
   c) Specify the values for Impala Daemon, Catalog Server, or StateStore in the respective fields:
      - Impala Daemon Verbose Log Level
      - Catalog Server Verbose Log Level
      - StateStore Verbose Log Level
   d) Click Save Changes and restart Impala.

   As logging levels increase, the categories of information logged are cumulative. For example, GLOG_v=2 records everything GLOG_v=1 records, as well as additional information.

   Increasing logging levels imposes performance overhead and increases log size. Cloudera recommends using GLOG_v=1 for most cases: this level has minimal performance impact but still captures useful troubleshooting information.

   Additional information logged at each level of GLOG_v is as follows:

   - 1: The default level. Logs information about each connection and query that is initiated to an impalad instance, including runtime profiles.
   - 2: Everything from the previous level plus information for each RPC initiated. This level also records query execution progress information, including details on each file that is read.
   - 3: Everything from the previous level plus logging of every row that is read. This level is only applicable for the most serious troubleshooting and tuning scenarios, because it can produce exceptionally large and detailed log files, potentially leading to its own set of performance and capacity problems.

   Note: For performance reasons, Cloudera highly recommends not setting the most verbose logging level to 3.

Impala lineage

You can use the Atlas lineage graph to understand the source and impact of data and changes to data over time and across all your data.

Atlas collects metadata from Impala to represent the lineage among data assets. The Atlas lineage graph shows the input and output processes that the current entity participated in. Entities are included if they were inputs to processes that lead to the current entity or they are output from processes for which the current entity was an input. Impala processes follow this pattern.

Note that lineage is not updated between a table and views that the table is a part of when an Impala ALTER TABLE operation runs on the table.

Related Information

Impala metadata collection

Web User Interface for Debugging

You can use the Impala daemons (`impalad`, `statestored`, and `catalogd`) Web UI to display the diagnostic and status information. Each of these Impala daemons includes a built-in web server.

Impala Daemon Web UI

The Impala Daemon (`impalad`) Web UI includes information about configuration settings, running and completed queries, and associated performance and resource usage for queries. In particular, the Details link for each query displays alternative views of the query including a graphical representation of the plan, and the output of the EXPLAIN, SUMMARY, and PROFILE statements from `impala-shell`. Each host that runs the `impalad` daemon has its own instance of the Web UI, with details about those queries for which that host served as the coordinator. The

`impalad` Web UI is primarily used for diagnosing query problems that can be traced to a particular node.

StateStore Web UI

The StateStore (`statestored`) Web UI includes information about memory usage, configuration settings, and ongoing health checks performed by statestored. Because there is only a single instance of the statestored within any Impala cluster, you access the Web UI only on the particular host that serves as the Impala StateStore.

Catalog Server Web UI

The Catalog Server (`catalogd`0 Web UI includes information about the databases, tables, and other objects managed by Impala, in addition to the resource usage and configuration settings of the catalogd. Because there is only a single instance of the catalogd within any Impala cluster, you access the Web UI only on the particular host that serves as the Impala Catalog Server.

Debug Web UI for Impala Daemon

You can use the Impala Daemon (`impalad`) Web UI to view information about configuration settings, running and completed queries, and associated performance and resource usage for queries.

To debug and troubleshoot an `impalad` using a web-based interface, open the URL http ://*impala#server#hostname*:25000/ in a browser. (For secure clusters, use the prefix https:// instead of http://.)

Because each Impala node produces its own set of debug information, you should choose a specific node that you want to investigate an issue on.

Main Page

The main impalad Web UI page at / lists the following information about the impalad:

- The version of the `impalad` daemon

  The Version section also contains other information, such as when Impala was built and what build flags were used.
- Process start time
- Hardware information
- OS information
- Process information
- CGroup information

Admission Controller Page

The Admission Controller impalad debug Web UI is at /admission page under the main impalad Web UI.

Use the /admission page to troubleshoot queued queries and the admission control.

The admission page provides the following information about each resource pool to which queries have been submitted at least once:

- Time since the statestored received the last update
- A warning if this impalad is considered disconnected from the statestored and thus the information on this page could be stale.
- Pool configuration
- Queued queries submitted to this coordinator, in the order of submission
- Running queries submitted to this coordinator
- Pool stats
  - Average of time in queue: An exponential moving average which represents the average time in queue over the last 10 to 12 queries. If a query is admitted immediately, the wait time of 0 is used in calculating this average wait time.

- Histogram of the distribution of peak memory used by queries admitted to the pool

  Use the histogram to figure out settings for the minimum and maximum query MEM_LIMIT ranges for this pool.

  The histogram displays data for all queries admitted to the pool, including the queries that finished, got canceled, or hit an error.

Click on the pool name to only display information relevant to that pool. You can then refresh the debug page to see only the information for that specific pool.

Click Reset informational stats for all pools to reset the stats that keep track of historical data, such as Totals stats, Time in queue (exponential moving average), and the histogram.

The above information is also available as a JSON object from the following HTTP endpoint:

```
http://impala-server-hostname:port/admission?json
```

Known Backends Page

The Known backends page of the impalad debug Web UI is at /backends under the main impalad Web UI.

This page lists the following info for each of the `impalad` nodes in the cluster. Because each `impalad` daemon knows about every other `impalad` daemon through the StateStore, this information should be the same regardless of which node you select.

- Address of the node: Host name and port
- KRPC address: The KRPC address of the impalad. Use this address when you issue the SHUT DOWN command for this impalad.
- Whether acting as a coordinator
- Whether acting as an executor
- Quiescing status: Specify whether the graceful shutdown process has been initiated on this impa lad.
- Memory limit for admission: The amount of memory that can be admitted to this backend by the admission controller.
- Memory reserved: The amount of memory reserved by queries that are active, either currently executing or finished but not yet closed, on this backend.

  The memory reserved for a query that is currently executing is its memory limit, if set. Otherwise, if the query has no limit or if the query finished executing, the current consumption is used.
- Memory of the queries admitted to this coordinator: The memory submitted to this particular host by the queries admitted by this coordinator.

Catalog Page

The Catalog page of the impalad debug Web UI is at /catalog under the main impalad Web UI.

This page displays a list of databases and associated tables recognized by this instance of `impalad`. You can use this page to locate which database a table is in, check the exact spelling of a database or table name, look for identical table names in multiple databases. The primary debugging use case would be to check if an impalad instance has knowledge of a particular table that someone expects to be in a particular database.

Hadoop Configuration

The Hadoop Configuration page of the impalad debug Web UI is at /hadoop-varz under the main impalad Web UI.

This page displays the Hadoop common configurations that Impala is running with.

JMX

The JMX page of the impalad debug Web UI is at /jmx under the main impalad Web UI.

This page displays monitoring information about various JVM subsystems, such as memory pools, thread management, runtime. etc.

Java Log Level

The Change log level page of the impalad debug Web UI is at /log_level under the main impalad Web UI.

This page displays the current Java and backend log levels, and it allows you to change the log levels dynamically without having to restart the impalad.

Logs Page

The INFO logs page of the impalad debug Web UI is at /logs under the main impalad Web UI.

This page shows the last portion of the impalad.INFO log file, including the info, warning, and error logs for the `impalad`. You can see the details of the most recent operations, whether the operations succeeded or encountered errors. This page provides one central place for the log files and saves you from looking around the filesystem for the log files, which could be in different locations on clusters that use cluster management software.

Memz Page

The Memory Usage page of the impalad debug Web UI is at /memz under the main impalad Web UI.

This page displays the summary and detailed information about memory usage by the impalad.

Metrics Page

The Metrics page of the impalad debug Web UI is at /metrics under the main impalad Web UI.

This page displays the current set of metrics, counters and flags representing various aspects of impalad internal operations.

Queries Page

The Queries page of the impalad debug Web UI is at /queries under the main impalad Web UI.

This page lists:

- Currently running queries
- Queries that have completed their execution, but have not been closed yet
- Completed queries whose details still reside in memory

The queries are listed in reverse chronological order, with the most recent at the top. You can control the amount of memory devoted to completed queries by specifying the -#-#query_log_size startup option for impalad.

This page provides:

- How many SQL statements are failing (State value of EXCEPTION)
- How large the result sets are (# rows fetched)
- How long each statement took (Start Time and End Time)

Click the Details link for a query to display the detailed performance characteristics of that query, such as the profile output.

On the query detail page, in the Profile tab, you have options to export the query profile output to the Thrift, text, or Json format.

The Queries page also includes the Query Locations section that lists the number of running queries with fragments on this host.

RPC Services Page

The RPC durations page of the impalad debug Web UI is at /rpcz under the main impalad Web UI.

This page displays information, such as the duration, about the RPC communications of this impa lad with other Impala daemons.

Sessions Page

The Sessions page of the impalad debug Web UI is at /session under the main impalad Web UI.

This page displays information about the sessions currently connected to this `impalad` instance. For example, sessions could include connections from the impala-shell command, JDBC or ODBC applications, or the Impala Query UI in the Hue web interface.

Threadz Page

The Threads page of the impalad debug Web UI is at /threadz under the main impalad Web UI.

This page displays information about the threads used by this instance of impalad, and it shows which categories they are grouped into. Making use of this information requires substantial knowledge about Impala internals.

Varz Page

The Varz page of the impalad debug Web UI is at /varz under the main impalad Web UI.

This page shows the configuration settings in effect when this instance of impalad communicates with other Hadoop components such as HDFS and YARN. These settings are collected from a set of configuration files.

The bottom of this page also lists all the command-line settings in effect for this instance of impa lad.

Prometheus Metrics Page

At /metrics_prometheus, under the main impalad Web UI, the metrics are generated in Prometheus exposition format that Prometheus can consume for event monitoring and alerting.

The /metrics_prometheus is not shown in the Web UI list of pages.

Debug Web UI for StateStore

You can use the StateStore (`statestored`) Web UI to view information about memory usage, configuration settings, and ongoing health checks performed by statestored.

To debug and troubleshoot the statestored daemon using a web-based interface, open the URL http ://*impala#server#hostname*:25010/ in a browser. (For secure clusters, use the prefix https:// instead of http://.)

Main Page

The main statestored Web UI page at / lists the following information about the statestored:

- The version of the statestored daemon
- Process start time
- Hardware information
- OS information
- Process information
- CGroup information

Logs Page

The INFO logs page of the debug Web UI is at /logs under the main statestored Web UI.

This page shows the last portion of the statestored.INFO log file, including the info, warning, and error logs for the statestored. You can refer here to see the details of the most recent operations, whether the operations succeeded or encountered errors. This page provides one central place for the log files and saves you from looking around the filesystem for the log files, which could be in different locations on clusters that use cluster management software.

Memz Page

The Memory Usage page of the debug Web UI is at /memz under the main statestored Web UI.

This page displays summary and detailed information about memory usage by the statestored. You can see the memory limit in effect for the node, and how much of that memory Impala is currently using.

Metrics Page

The Metrics page of the debug Web UI is at /metrics under the main statestored Web UI.

This page displays the current set of metrics: counters and flags representing various aspects of stat estored internal operation.

RPC Services Page

The RPC durations page of the statestored debug Web UI is at /rpcz under the main statestored Web UI.

This page displays information, such as the durations, about the RPC communications of this stat estored with other Impala daemons.

Subscribers Page

The Subscribers page of the debug Web UI is at /subscribers under the main statestored Web UI.

This page displays information about the other Impala daemons that have registered with the stat estored to receive and send updates.

Threadz Page

The Threads page of the debug Web UI is at /threadz under the main statestored Web UI.

This page displays information about the threads used by this instance of statestored, and shows which categories they are grouped into. Making use of this information requires substantial knowledge about Impala internals.

Topics Page

The Topics page of the debug Web UI is at /topics under the main statestored Web UI.

This page displays information about the topics to which the other Impala daemons have registered to receive updates.

Varz Page

The Varz page of the debug Web UI is at /varz under the main statestored Web UI.

This page shows the configuration settings in effect when this instance of statestored communicates with other Hadoop components such as HDFS and YARN. These settings are collected from a set of configuration files.

The bottom of this page also lists all the command-line settings in effect for this instance of statesto red.

Prometheus Metrics Page

At /metrics_prometheus, under the main statestored Web UI, the metrics are generated in Prometheus exposition format that Prometheus can consume for event monitoring and alerting.

The /metrics_prometheus is not shown in the Web UI list of pages.

Debug Web UI for Catalog Server

You can use the Catalog Server (`catalogd`) Web UI to view information about the databases, tables, and other objects managed by Impala, in addition to the resource usage and configuration settings of the catalogd.

The main page of the debug Web UI is at http://*impala#server#hostname*:25020/ (non-secure cluster) or http s://*impala#server#hostname*:25020/ (secure cluster).

Main Page

The main catalogd Web UI page at / lists the following information about the catalogd:

- The version of the catalogd daemon
- Process start time
- Hardware information
- OS information
- Process information
- CGroup information

Catalog Page

The Catalog page of the debug Web UI is at /catalog under the main catalogd Web UI.

This page displays a list of databases and associated tables recognized by this instance of catalogd. You can use this page to locate which database a table is in, check the exact spelling of a database or table name, look for identical table names in multiple databases. The catalog information is represented as the underlying Thrift data structures.

JMX

The JMX page of the catalogd debug Web UI is at /jmx under the main catalogd Web UI.

This page displays monitoring information about various JVM subsystems, such as memory pools, thread management, runtime. etc.

Java Log Level

The Change log level page of the catalogd debug Web UI is at /log_level under the main catalogd Web UI.

The page displays the current Java and backend log levels and allows you to change the log levels dynamically without having to restart the catalogd

Logs Page

The INFO logs page of the debug Web UI is at /logs under the main catalogd Web UI.

This page shows the last portion of the catalogd.INFO log file, including the info, warning, and error logs for the catalogd daemon. You can refer here to see the details of the most recent operations, whether the operations succeeded or encountered errors. This page provides one central place for the log files and saves you from looking around the filesystem for the log files, which could be in different locations on clusters that use cluster management software.

Memz Page

The Memory Usage page of the debug Web UI is at /memz under the main catalogd Web UI.

This page displays summary and detailed information about memory usage by the catalogd. You can see the memory limit in effect for the node, and how much of that memory Impala is currently using.

Metrics Page

The Metrics page of the debug Web UI is at /metrics under the main catalogd Web UI.

This page displays the current set of metrics: counters and flags representing various aspects of cata logd internal operation.

RPC Services Page

The RPC durations page of the catalogd debug Web UI is at /rpcz under the main catalogd Web UI.

This page displays information, such as the durations, about the RPC communications of this cata logd with other Impala daemons.

Threadz Page

The Threads page of the debug Web UI is at /threadz under the main catalogd Web UI.

This page displays information about the threads used by this instance of catalogd, and shows which categories they are grouped into. Making use of this information requires substantial knowledge about Impala internals.

Varz Page

The Varz page of the debug Web UI is at /varz under the main catalogd Web UI.

This page shows the configuration settings in effect when this instance of catalogd communicates with other Hadoop components such as HDFS and YARN. These settings are collected from a set of configuration files.

The bottom of this page also lists all the command-line settings in effect for this instance of cata logd.

Prometheus Metrics Page

At /metrics_prometheus, under the main catalogd Web UI, the metrics are generated in Prometheus exposition format that Prometheus can consume for event monitoring and alerting.

The /metrics_prometheus is not shown in the Web UI list of pages.

Configuring Impala Web UI

As an administrator, you can diagnose issues with each daemon on a particular host, or perform other administrative actions such as cancelling a running query from the built-in web server's UI. The built-in web server is inlcuded within each of the Impala-related daemons. By default, these web servers are enabled. You can turn them off in a high-security configuration where it is not appropriate for users to have access to this kind of monitoring information through a web interface.

Enabling and Disabling Access to Impala Web Servers

By default, these web servers are enabled. You might turn them off in a high-security configuration where it is not appropriate for users to have access to this kind of monitoring information through a web interface.

To enable or disable Impala Web Servers for Web UI in Cloudera Manager:

- Impala Daemon
  1. Navigate to ClustersImpala ServiceConfiguration.
  2. Select ScopeImpala Daemon .
  3. Select CategoryPorts and Addresses.
  4. Select or clear Enable Impala Daemon Web Server.
  5. Click Save Changes, and restart the Impala service.
- Impala StateStore
  1. Navigate to ClustersImpala ServiceConfiguration.
  2. Select ScopeImpala StateStore.
  3. Select CategoryMain.
  4. Select or clear Enable StateStore Web Server.
  5. Click Save Changes, and restart the Impala service.

- Impala Catalog Server
  1. Navigate to ClustersImpala ServiceConfiguration.
  2. Select ScopeImpala Catalog Server.
  3. Select CategoryMain.
  4. Check or clear Enable Catalog Server Web Server.
  5. Click Save Changes, and restart the Impala service.

Configuring Secure Access for Impala Web Servers

Cloudera Manager supports two methods of authentication for secure access to the Impala Catalog Server, Impala Daemon, and StateStore web servers: password-based authentication and SPNEGO authentication.

Authentication for the three types of daemons can be configured independently.

Configuring Password Authentication

1. Navigate to ClustersImpala ServiceConfiguration.
2. Search for "password" using the Search box in the Configuration tab. This should display the password-related properties (Username and Password properties) for the Impala Daemon, StateStore, and Catalog Server. If there are multiple role groups configured for Impala Daemon instances, the search should display all of them.
3. Enter a username and password into these fields.
4. Click Save Changes, and restart the Impala service.

Now when you access the Web UI for the Impala Daemon, StateStore, or Catalog Server, you are asked to log in before access is granted.

Enabling Kerberos HTTP SPNEGO Authentication for Web UI

To provide security through Kerberos, Impala Web UIs support SPNEGO. SPNEGO is a protocol for securing HTTP requests with Kerberos by passing negotiation tokens through HTTP headers.

To enable authorization using SPNEGO in Cloudera Manager:

1. Navigate to ClustersImpala ServiceConfiguration.
2. Select ScopeImpala 1 (Service-Wid).
3. Select the Enable Kerberos Authentication for HTTP Web-Consoles field.

   This setting is effective only Kerberos is enabled for the HDFS service.
4. Click Save Changes, and restart the Impala service.

Configuring TLS/SSL for Web UI

1. Create or obtain an TLS/SSL certificate.
2. Place the certificate, in .pem format, on the hosts where the Impala Catalog Server and StateStore are running, and on each host where an Impala Daemon is running. It can be placed in any location (path) you choose. If all the Impala Daemons are members of the same role group, then the .pem file must have the same path on every host.
3. Navigate to ClustersImpala ServiceConfiguration.
4. Search for "certificate" using the Search box in the Configuration tab. This should display the certificate file location properties for the Impala Catalog Server, Impala Daemon, and StateStore. If there are multiple role groups configured for Impala Daemon instances, the search should display all of them.
5. In the property fields, enter the full path name to the certificate file, the private key file path, and the password for the private key file..
6. Click Save Changes, and restart the Impala service.

Important: If Cloudera Manager cannot find the .pem file on the host for a specific role instance, that role will fail to start.

When you access the Web UI for the Impala Catalog Server, Impala Daemon, and StateStore, https will be used.

Opening Impala Web UIs

Procedure

1. Navigate to ClustersImpala ServiceConfiguration.
2. Open the appropriate Web UI:Select Web UIImpala Catalog Web UI.
   - To open StateStore Web UI, select Web UIImpala StateStore Web UI.
   - To open Catalog Server Web UI, select Web UIImpala Catalog Web UI.
   - To open Impala Daemon Web UI:
     a. Click the Instances tab.
     b. Click an Impala Daemon instance.
     c. Click Impala Daemon Web UI.

Securing Impala

You must be aware of this set of security features Impala provides to protect your critical and sensitive data.

The Impala security features have several objectives.

- At the basic level, security prevents accidents or mistakes that could disrupt application processing, delete or corrupt data, or reveal data to unauthorized users.
- Advanced security features and practices can harden the system against malicious users trying to gain unauthorized access or perform other disallowed operations.
- The auditing feature provides a way to confirm that no unauthorized access occurred, and detect whether any such attempts were made.

Based on the above objectives, Impala security features are divided into the following categories.

Authentication

How does Impala verify the identity of the user to confirm that they really are allowed to exercise the privileges assigned to that user? Impala relies on the Kerberos subsystem for authentication.

Authorization

Which users are allowed to access which resources, and what operations are they allowed to perform? Impala relies on the open source Ranger project for authorization. By default (when authorization is not enabled), Impala does all read and write operations with the privileges of the impala user, which is suitable for a development/test environment but not for a secure production environment. When authorization is enabled, Impala uses the OS user ID of the user who runs `impala-shell` or other client program, and associates various privileges with each user.

TLS/SSL Encryption

This feature encrypts TLS/SSL network encryption, between Impala and client programs, and between the Impala-related daemons running on different nodes in the cluster.

Auditing

What operations were attempted, and did they succeed or not? This feature provides a way to look back and diagnose whether attempts were made to perform unauthorized operations. You use this information to track down suspicious activity, and to see where changes are needed in authorization policies. The audit data produced by this feature is collected by the Cloudera Manager product and then presented in a user-friendly form by the Cloudera Manager product.

The following are the important guidelines to protect a cluster running Impala against accidents and mistakes, or malicious attackers trying to access sensitive data. See the subsequent topics that describe the security features in detail.

- Secure the root account. The root user can tamper with the `impalad` daemon, read and write the data files in HDFS, log into other user accounts, and access other system services that are beyond the control of Impala.
- Restrict membership in the sudoers list (in the /etc/sudoers file). The users who can run the sudo command can do many of the same things as the root user.
- Ensure the Hadoop ownership and permissions for Impala data files are restricted.

  The Impala authorization feature makes use of the HDFS file ownership and permissions mechanism.

  1. Set up users and assign them to groups at the OS level, corresponding to the different categories of users with different access levels for various databases, tables, and HDFS locations (URIs).
  2. Create the associated Linux users using the `useradd` command if necessary.
  3. Add Linux users to the appropriate groups with the `usermod` command.
- Ensure the Hadoop ownership and permissions for Impala log files are restricted.

  If you issue queries containing sensitive values in the WHERE clause, such as financial account numbers, those values are stored in Impala log files in the Linux file system, and you must secure those files.
- Ensure that the Impala web UI (available by default on port 25000 on each Impala node) is password-protected.
- Create a policy file that specifies which Impala privileges are available to users in particular Hadoop groups (which by default map to Linux OS groups). Create the associated Linux groups using the `groupadd` command if necessary.
- Design your databases, tables, and views with database and table structure to allow policy rules to specify simple, consistent rules.

  For example, if all tables related to an application are inside a single database, you can assign privileges for that database and use the * wildcard for the table name. If you are creating views with different privileges than the underlying base tables, you might put the views in a separate database so that you can use the * wildcard for the database containing the base tables, while specifying the specific names of the individual views.
- Enable authorization for all impalad daemons.

  The authorization feature does not apply to the `statestored` daemon, which has no access to schema objects or data files.
- Set up authentication using Kerberos, to make sure users really are who they say they are.

Related Information

Configuring Impala Web UI

Configuring Impala TLS/SSL

To protect sensitive information being transmitted, Impala supports TLS/SSL network encryption, between Impala and client programs, and between the Impala-related daemons running on different nodes in the cluster.

To configure Impala to listen for Beeswax and HiveServer2 requests on TLS/SSL-secured ports:

1. In Cloudera Manager, select the Impala service from the Clusters drop down.
2. In the Configuration tab, select ScopeImpala (Service-Wide).
3. Select CategorySecurity.
4. Edit the following property fields:

Property	Description
Enable TLS/SSL for Impala	Encrypt communication between clients (like ODBC, JDBC, and the Impala shell) and the Impala daemon using Transport Layer Security (TLS) (formerly known as Secure Socket Layer (SSL)).
Impala TLS/SSL Server Certificate File (PEM Format)	Local path to the X509 certificate that identifies the Impala daemon to clients during TLS/SSL connections. This file must be in PEM format.
Impala TLS/SSL Server Private Key File (PEM Format)	Local path to the private key that matches the certificate specified in the Certificate for Clients. This file must be in PEM format.
Impala TLS/SSL Private Key Password	The password for the private key in the Impala TLS/SSL Server Certificate and Private Key file. If left blank, the private key is not protected by a password.

Property	Description
Impala TLS/SSL CA Certificate	The location on disk of the certificate, in PEM format, used to confirm the authenticity of SSL/TLS servers that the Impala daemons might connect to. Because the Impala daemons connect to each other, this should also include the CA certificate used to sign all the SSL/TLS Certificates. Without this parameter, SSL/TLS between Impala daemons will not be enabled.

Using TLS/SSL with Business Intelligence Tools

You can use Kerberos authentication, TLS/SSL encryption, or both to secure connections from JDBC and ODBC applications to Impala.

Configuring TLS/SSL Communication for the Impala Shell

Typically, a client program has corresponding configuration properties in Cloudera Manager to verify that it is connecting to the right server. For example, with SSL enabled for Impala, you use the following options when starting the `impala-shell` interpreter:

- --ssl: enables TLS/SSL for `impala-shell`.
- --ca_cert: the local pathname pointing to the third-party CA certificate, or to a copy of the server certificate for self-signed server certificates.

If --ca_cert is not set, `impala-shell` enables TLS/SSL, but does not validate the server certificate. This is useful for connecting to a known-good Impala that is only running over TLS/SSL, when a copy of the certificate is not available (such as when debugging customer installations).

For `impala-shell` to successfully connect to an Impala cluster that has the minimum allowed TLS/SSL version set to 1.2 (--ssl_minimum_version=tlsv1.2), the cluster that `impala-shell` runs on must have Python version 2.7.9 or higher (or a vendor-provided Python version with the required support. Some vendors patched Python 2.7.5 versions on Red Hat Enterprise Linux 7 and derivatives).

Specifying TLS/SSL Minimum Allowed Version and Ciphers

Depending on your cluster configuration and the security practices in your organization, you might need to restrict the allowed versions of TLS/SSL used by Impala. Older TLS/SSL versions might have vulnerabilities or lack certain features. You can use startup options for the `impalad`, `catalogd`, and `statestored` daemons to specify a minimum allowed version of TLS/SSL.

Specify one of the following values for the --ssl_minimum_version configuration setting:

- tlsv1: Allow any TLS version of 1.0 or higher. This setting is the default when TLS/SSL is enabled.
- tlsv1.1: Allow any TLS version of 1.1 or higher.
- tlsv1.2: Allow any TLS version of 1.2 or higher.

Along with specifying the version, you can also specify the allowed set of TLS ciphers by using the --ssl_cipher_list configuration setting. The argument to this option is a list of keywords, separated by colons, commas, or spaces, and optionally including other notation. For example:

```
--ssl_cipher_list="RC4-SHA,RC4-MD5"
```

By default, the cipher list is empty, and Impala uses the default cipher list for the underlying platform. See the output of `man ciphers` for the full set of keywords and notation allowed in the argument string.

Impala Authentication

Authentication is the mechanism to ensure that only specified hosts and users can connect to Impala.

Authentication feature verifies that when clients connect to Impala, they are connected to a legitimate server. The feature prevents spoofing such as *impersonation* (setting up a phony client system with the same account and group

names as a legitimate user) and *man-in-the-middle attacks* (intercepting application requests before they reach Impala and eavesdropping on sensitive information in the requests or the results).

Impala automatically handles both Kerberos and LDAP authentication. Each `impalad` daemon can accept both Kerberos and LDAP requests through the same port. No special actions need to be taken if some users authenticate through Kerberos and some through LDAP.

You can also make proxy connections to Impala through Apache Knox.

Regardless of the authentication mechanism used, Impala always creates HDFS directories and data files owned by the same user (typically impala). Once you are finished setting up authentication, set up authorization to control user-level access to databases, tables, columns, partitions when they connect through Impala.

Configuring Kerberos Authentication

Impala supports an enterprise-grade authentication system called Kerberos. Kerberos provides strong security benefits including capabilities that render intercepted authentication packets unusable by an attacker. It virtually eliminates the threat of impersonation by never sending a user's credentials in cleartext over the network.

Requirements for Using Impala with Kerberos

Important:

- If you plan to use Impala in your cluster, you must configure your KDC to allow tickets to be renewed, and you must configure krb5.conf to request renewable tickets. Typically, you can do this by adding the max_renewable_life setting to your realm in kdc.conf, and by adding the renew_lifetime parameter to the libdefaults section of krb5.conf.

Impala supports the Cloudera ODBC driver and the Kerberos interface provided. To use Kerberos through the ODBC driver, the host type must be set depending on the level of the ODBD driver:

- SecImpala for the ODBC 1.0 driver.
- SecBeeswax for the ODBC 1.2 driver.
- Blank for the ODBC 2.0 driver or higher, when connecting to a secure cluster.
- HS2NoSasl for the ODBC 2.0 driver or higher, when connecting to a non-secure cluster.

Enabling Kerberos in Impala-shell

To enable Kerberos in the Impala shell, start the `impala-shell` command using the -k flag.

Configuring Impala to Support Kerberos Security

Enabling Kerberos authentication for Impala involves steps that can be summarized as follows:

- Creating service principals for Impala and the HTTP service. Principal names take the form: *serviceName/fully.qualified.domain.name@KERBEROS.REALM*
- Creating, merging, and distributing key tab files for these principals.
- Editing /etc/default/impala (in cluster not managed by Cloudera Manager), or editing the Security settings in the Cloudera Manager interface, to accommodate Kerberos authentication.

Enabling Access to Internal Impala APIs for Kerberos Users

For applications that need direct access to Impala APIs, without going through the HiveServer2 or Beeswax interfaces, you can specify a list of Kerberos users who are allowed to call those APIs. By default, the impala and hdfs users are the only ones authorized for this kind of access. Any users not explicitly authorized through the internal _principals_whitelist configuration setting are blocked from accessing the APIs. This setting applies to all the Impala-related daemons, although currently it is primarily used for HDFS to control the behavior of the catalog server.

Mapping Kerberos Principals to Short Names for Impala

Impala can support the additional mapping rules that will be inserted before rules generated from the list of trusted realms and before the default rule. The support is disabled by default in Impala.

To enable mapping Kerberos principals to short names:

1. In Cloudera Manager, select the Impala service.
2. In the Configuration tab, select Impala (Service-Wide) in Scope and Advanced in Category.
3. Select the Use HDFS Rules to Map Kerberos Principals to Short Names field.
4. Click Save Changes, and restart the Impala service.

Configuring LDAP Authentication

Impala uses LDAP for authentication, verifying the credentials of each user who connects through `impala-shell`, Hue, a Business Intelligence tool, JDBC or ODBC applications.

About this task

Authentication is the process of allowing only specified named users to access the server (in this case, the Impala server). This feature is crucial for any production deployment, to prevent misuse, tampering, or excessive load on the server.

Only the connections between clients and Impala can be authenticated by LDAP.

Procedure

To enable and configure LDAP:

1. In Cloudera Manager, select the Impala service.
2. In the Configuration tab, type ldap in the search box. The fields for LDAP configuration will be listed.
3. Set the following fields to enable LDAP.

   Enable LDAP Authentication (enable_ldap_auth)

   Enables LDAP-based authentication between the client and Impala.

   LDAP URL (ldap_uri)

   Sets the URI of the LDAP server to use. Typically, the URI is prefixed with ldap://. You can specify secure SSL-based LDAP transport by using the prefix ldaps://. The URI can optionally specify the port, for example: ldap://ldap_server.example.com:389 or ldaps://ldap_server.example.com:636. (389 and 636 are the default ports for non-SSL and SSL LDAP connections, respectively.)

4. Set the following fields to support custom bind strings.

   When Impala connects to LDAP, it issues a bind call to the LDAP server to authenticate as the connected user. Impala clients, including the Impala-shell, provide the short name of the user to Impala. This is necessary so that Impala can use role-based access, which uses short names.

   However, LDAP servers often require more complex, structured usernames for authentication. Impala supports three ways of transforming the short name (for example, 'henry') to a more complicated string. If necessary, specify one of the following configuration options when starting the `impalad` daemon.

   Active Directory Domain (--ldap_domain)

   Replaces the username with a string *username@ldap_domain*.

   LDAP BaseDN (--ldap_baseDN)

   Replaces the username with a “distinguished name” (DN) of the form: uid=*userid*,ldap_baseDN. (This is equivalent to a Hive option).

   LDAP Pattern (--ldap_bind_pattern)

   This is the most general option, and replaces the username with the string *ldap_bind_pattern* where all instances of the string #UID are replaced with *userid*. For example, an ldap_bind_pattern of "use r=#UID,OU=foo,CN=bar" with a username of henry will construct a bind name of "user=henry, OU=foo,CN=bar".

   The above options are mutually exclusive, and Impala does not start if more than one of these options are specified.

5. Set the following fields for secure LDAP connections.

   To avoid sending credentials over the wire in cleartext, you must configure a secure connection between both the client and Impala, and between Impala and the LDAP server. The secure connection could use SSL or TLS.

   Secure LDAP connections through SSL:

   For SSL-enabled LDAP connections, specify a prefix of ldaps:// instead of ldap://. Also, the default port for SSL-enabled LDAP connections is 636 instead of 389.

   Secure LDAP connections through TLS:

   TLS, the successor to the SSL protocol, is supported by most modern LDAP servers. Unlike SSL connections, TLS connections can be made on the same server port as non-TLS connections. To secure all connections using TLS, specify the following flags as startup options to the `impalad` daemon:

   Enable LDAP TLS (--ldap_tls_

   Tells Impala to start a TLS connection to the LDAP server, and to fail authentication if it cannot be done.

   LDAP Server CA Certificate (--ldap_ca_certificate)

   Specifies the location of the certificate in standard .PEM format. Store this certificate on the local filesystem, in a location that only the impala user and other trusted users can read.

6. Click Save Changes and restart the Impala service.

Enabling LDAP in Hue

To connect to Impala using LDAP authentication through Hue, enable the LDAP setting in Hue.

Procedure

1. Go to the Hue service.
2. Click the Configuration tab.
3. Select ScopeHue Server.
4. Select CategoryAdvanced.

5. Add the following properties to the Hue Server Advanced Configuration Snippet (Safety Valve) for hue_safety_valve_server.ini property

```
[impala]
auth_username=<LDAP username of Hue user to be authenticated>
auth_password=<LDAP password of Hue user to be authenticated>
```

6. Click Save Changes.

Enabling LDAP Authentication for impala-shell

To connect to Impala using LDAP authentication, you specify command-line options to the `impala-shell` command interpreter and enter the password when prompted.

-l

Enables LDAP authentication.

-u

Sets the user. Per Active Directory, the user is the short username, not the full LDAP distinguished name. If your LDAP settings include a search base, use the --ldap_bind_pattern on the `impalad` daemon to translate the short user name from `impala-shell` automatically to the fully qualified name.

`impala-shell` automatically prompts for the password.

Impala Authorization

Authorization determines which users are allowed to access which resources, and what operations they are allowed to perform. You use Apache Ranger to enable and manage authorization in Impala.

Supported Ranger Features in Impala

- Resource-based and tag-based authorization policies
- Resource-based and tag-based column masking

Unsupported Ranger Features in Impala

- Row-level filtering

Using Resource-based Authorization Policies for Impala

In the Ranger Service Manager, you can use the Hadoop SQL preloaded resource-based and tag-based services and policies to authorize access to Impala:

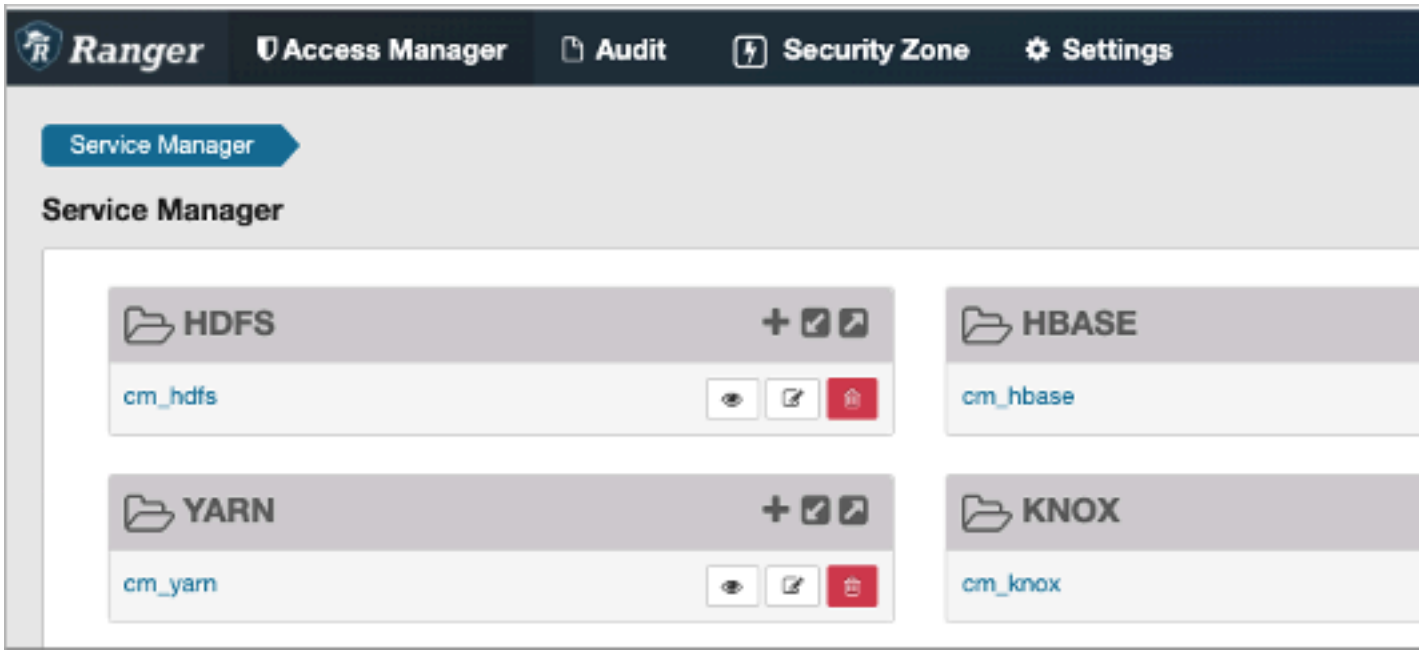

You can configure services for Impala as described in Using Ranger to Provide Authorization in CDP.

For example, you can edit the all-global policy for an Impala user:

For information about using tag-based policies, see Tag-based column masking in Hive with Ranger policies.

Privilege Model

You set up privileges through the GRANT and REVOKE statements in either Impala or Hive. Then both components use those same privileges automatically.

By default, when authorization is not enabled, Impala does all read and write operations with the privileges of the impala user, which is suitable for a development/test environment but not for a secure production environment. When authorization is enabled, Impala uses the OS user ID of the user who runs `impala-shell` or other client programs, and associates various privileges with each user.

Privileges can be granted on different resources in the schema and are associated with a level in the resource hierarchy. A privilege on a particular resource automatically inherits the same privilege of its parent.

The resource hierarchy is:

```
Server
    URI
    Database
        Table
            Column
        Function
```

The table-level privileges apply to views as well. Anywhere you specify a table name, you can specify a view name instead.

You can specify privileges for individual columns.

The table below lists the minimum level of privileges and the scope required to execute SQL statements in Impala. The following notations are used:

- The SERVER resource type in Ranger implies all databases, all tables, all columns, all UDFs, and all URLs.
- ANY denotes the CREATE, ALTER, DROP, SELECT, INSERT, or REFRESH privilege.
- ALL privilege denotes the SELECT, INSERT, CREATE, ALTER, DROP, and REFRESH privileges.
- VIEW_METADATA privilege denotes the SELECT, INSERT, or REFRESH privileges.
- The parent levels of the specified scope are implicitly supported. For example, if a privilege is listed with the TABLE scope, the same privilege granted on DATABASE and SERVER will allow the user to execute that specific SQL statement on TABLE.

Note: For an "impala" user to access any managed tables under the warehouse root directory or external tables with directories outside the warehouse root directory, the "impala" user must have HDFS RW access to the underlying paths that are referenced in a query.

For example, to be able to execute CREATE VIEW, you need the CREATE privilege on the database and the SELE CT privilege on the source table.

SQL Statement	Privileges	Object Type / Resource Type
SELECT	SELECT	TABLE
WITH SELECT	SELECT	TABLE
EXPLAIN SELECT	SELECT	TABLE
INSERT	INSERT	TABLE
EXPLAIN INSERT	INSERT	TABLE
TRUNCATE	INSERT	TABLE
LOAD	INSERT	TABLE
	ALL	URI
CREATE DATABASE	CREATE	SERVER
CREATE DATABASE LOCATION	CREATE	SERVER
	ALL	URI
CREATE TABLE	CREATE	DATABASE
CREATE TABLE LIKE	CREATE	DATABASE
	VIEW_METADATA	TABLE
CREATE TABLE AS SELECT	CREATE	DATABASE
	INSERT	DATABASE
	SELECT	TABLE
EXPLAIN CREATE TABLE AS SELECT	CREATE	DATABASE
	INSERT	DATABASE
	SELECT	TABLE
CREATE TABLE LOCATION	CREATE	TABLE
	ALL	URI
CREATE VIEW	CREATE	DATABASE
	SELECT	TABLE

ALTER DATABASE SET OWNER	ALL WITH GRANT	DATABASE
ALTER TABLE	ALL	TABLE
ALTER TABLE SET LOCATION	ALL	TABLE
	ALL	URI
ALTER TABLE RENAME	CREATE	DATABASE
	ALL	TABLE
ALTER TABLE SET OWNER	ALL WITH GRANT	TABLE
ALTER VIEW	ALL	TABLE
	SELECT	TABLE
ALTER VIEW RENAME	CREATE	DATABASE
	ALL	TABLE
ALTER VIEW SET OWNER	ALL WITH GRANT	VIEW
DROP DATABASE	ALL	DATABASE
DROP TABLE	ALL	TABLE
DROP VIEW	ALL	TABLE
CREATE FUNCTION	CREATE	DATABASE
	ALL	URI
DROP FUNCTION	ALL	DATABASE
COMPUTE STATS	ALL	TABLE
DROP STATS	ALL	TABLE
INVALIDATE METADATA	REFRESH	SERVER
INVALIDATE METADATA <table>	REFRESH	TABLE
REFRESH <table>	REFRESH	TABLE
REFRESH AUTHORIZATION	REFRESH	SERVER
REFRESH FUNCTIONS	REFRESH	DATABASE
COMMENT ON DATABASE	ALL	DATABASE
COMMENT ON TABLE	ALL	TABLE
COMMENT ON VIEW	ALL	TABLE
COMMENT ON COLUMN	ALL	TABLE
DESCRIBE DATABASE	VIEW_METADATA	DATABASE
DESCRIBE <table/view>	VIEW_METADATA	TABLE
If the user has the SELECT privilege at the COLUMN level, only the columns the user has access will show.	SELECT	COLUMN
USE	ANY	TABLE
SHOW DATABASES	ANY	TABLE
SHOW TABLES	ANY	TABLE
SHOW FUNCTIONS	VIEW_METADATA	DATABASE
SHOW PARTITIONS	VIEW_METADATA	TABLE
SHOW TABLE STATS	VIEW_METADATA	TABLE
SHOW COLUMN STATS	VIEW_METADATA	TABLE

SHOW FILES	VIEW_METADATA	TABLE
SHOW CREATE TABLE	VIEW_METADATA	TABLE
SHOW CREATE VIEW	VIEW_METADATA	TABLE
SHOW CREATE FUNCTION	VIEW_METADATA	DATABASE
SHOW RANGE PARTITIONS (Kudu only)	VIEW_METADATA	TABLE
UPDATE (Kudu only)	ALL	TABLE
EXPLAIN UPDATE (Kudu only)	ALL	TABLE
UPSERT (Kudu only)	ALL	TABLE
WITH UPSERT (Kudu only)	ALL	TABLE
EXPLAIN UPSERT (Kudu only)	ALL	TABLE
DELETE (Kudu only)	ALL	TABLE
EXPLAIN DELETE (Kudu only)	ALL	TABLE

The privileges not listed in the table above will be silently ignored by Impala.

Changing Privileges in Impala

Privileges are managed via the GRANT and REVOKE SQL statements that require the Ranger service enabled.

Privileges can be also managed in Ranger UI. Especially, for attribute-based access control, Ranger UI is required to manage authorization.

Impala authorization policies are listed in the Hive service section in Ranger UI.

REFRESH AUTHORIZATION is not required when you make the changes to privileges within Impala. The changes are automatically propagated.

Changing Privileges from Outside of Impala

If you make a change to privileges in Ranger from outside of Impala, e.g. adding a user, removing a user, modifying privileges, there are two options to propagate the change:

- Use the ranger.plugin.hive.policy.pollIntervalMs property to specify how often to do a Ranger refresh. The property is specified in ranger-impala-security.xml in the conf directory under your Impala home directory.
- Run the REFRESH AUTHORIZATION statement to force a refresh.

Warning: As INVALIDATE METADATA is an expensive operation, you should use it judiciously.

Granting Privileges on URI

URIs represent the file paths you specify as part of statements such as CREATE EXTERNAL TABLE and LOAD DATA. Typically, you specify what look like UNIX paths, but these locations can also be prefixed with hdfs:// to make clear that they are really URIs. To set privileges for a URI, specify the name of a directory, and the privilege applies to all the files in that directory and any directories underneath it.

URIs must start with hdfs://, s3a://, adl://, or file://. If a URI starts with an absolute path, the path will be appended to the default filesystem prefix. For example, if you specify:

```
GRANT ALL ON URI '/tmp';
```

The above statement effectively becomes the following where the default filesystem is HDFS.

```
GRANT ALL ON URI 'hdfs://localhost:20500/tmp';
```

When defining URIs for HDFS, you must also specify the NameNode. For example:

```
GRANT ALL ON URI file:///path/to/dir TO <principal>
          GRANT ALL ON URI hdfs://namenode:port/path/to/dir TO <principal>
```

Warning: Because the NameNode host and port must be specified, it is strongly recommended that you use High Availability (HA). This ensures that the URI will remain constant even if the NameNode changes. For example:

```
GRANT ALL ON URI hdfs://ha-nn-uri/path/to/dir TO <principal>
```

Note:

```
<principal> can be a user, or a group
```

Note:

- If a user uses Impala SQL engine to access the resource of the specified URL/URI provided as part of the SQL statement and as an admin if you must deny access permissions to this location for this particular user, you must add the permission "All" in the "HADOOP SQL" policy.
- However, if the user is using Hive as the execution engine and to deny access permission to a location, then you must add both the “Read” and “Write” permissions in the field of ‘Permissions’ for the corresponding deny conditions.
- If the same Ranger policy is shared by both Hive and Impala, then you must add “All”, “Read”, and “Write” to the field of ‘Permissions’ to enforce the policy.

Object Ownership

Object ownership for tables, views and databases is enabled by default in Impala.

To define owner specific privileges, go to Ranger UI and define appropriate policies on the {OWNER} user.

The CREATE statements implicitly make the user running the statement the owner of the object. For example, if *User A* creates a database, *foo*, via the CREATE DATABASE statement, *User A* now owns the *foo* database and is authorized to perform any operation on the *foo* database.

An ownership can be transferred to another user or role via the ALTER DATABASE, ALTER TABLE, or ALTER VIEW with the SET OWNER clause.

Note: Currently, due to a known issue (IMPALA-8937), until the ownership information is fully loaded in the coordinator catalog cache, the owner of a table might not be able to see the table when executing the SHOW TABLES statement. The owner can still query the table.

Ranger Column Masking for Impala

Ranger column masking hides sensitive columnar data in Impala query output. For example, you can define a policy that reveals only the first or last four characters of column data. Column masking is enabled by default. To disable column masking, modify the following configuration property in all coordinators:

```
--enable_column_masking=false
```

This flag might be removed in a future release. The Impala behavior mimics Hive behavior with respect to column masking.

The following table lists all Impala-supported, built-in mask types for defining column masking in a policy using the Ranger REST API or Ranger UI:

Type	Name	Description	Transformer
MASK	Redact	Replace lowercase with 'x', uppercase with 'X', digits with '0'	mask({col})

Type	Name	Description	Transformer
MASK_SHOW_LAST_4	Partial mask: show last 4	Show last 4 characters; replace rest with 'x'	mask_show_last_n({col}, 4, 'x', 'x', 'x', -1, '1')
MASK_SHOW_FIRST_4	Partial mask: show first 4	Show first 4 characters; replace rest with 'x'	mask_show_first_n({col}, 4, 'x', 'x', 'x', -1, '1')
MASK_HASH	Hash	Hash the value	mask_hash({col})
MASK_NULL	Nullify	Replace with NULL	N/A
MASK_NONE	Unmasked (retain original value)	No masking	N/A
MASK_DATE_SHOW_YEAR	Date: show only year	Date: show only year	mask({col}, 'x', 'x', 'x', -1, '1', 1, 0, -1)
CUSTOM	Custom	Custom	N/A

The table includes the mask name as it appears in the Ranger UI.

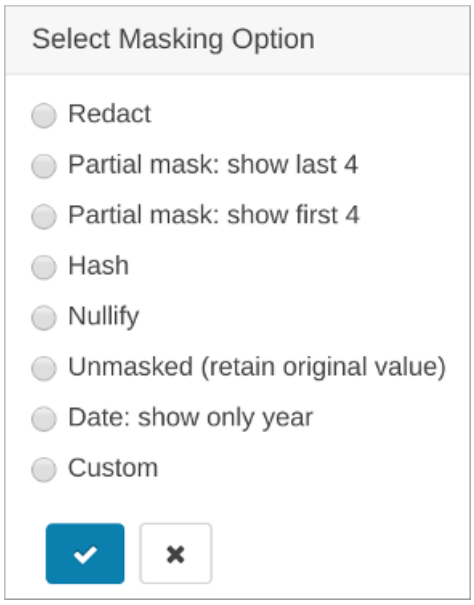

Ranger Column Masking Limitations in Impala

- Column masking introduces unused columns during the query analysis and, consequently, additional SELECT privileges checks on all columns of the masked table.
- Impala might produce more than one audit log entry for a column involved in a column masking policy under all of these conditions: the column appears in multiple places in a query, the query is rewritten, or the query is re-analyzed.
- Column masking policies, shared between Hive and Impala, might be affected by SQL/UDF differences between Hive and Impala, as shown in the following example.

  For instance, UTF-8 strings containing non-ASCII characters are not guaranteed to work properly. Suppose a column masking policy masks the last two characters of a string: s => mask_last_n(s, 2, 'x', 'x', 'x', 'x'). Applying this policy, Hive properly masks SQL## to SQLxx, but the Impala masking produces SQL##xx because each Chinese character is encoded in 3 bytes. Impala and Hive do not handle different lengths in the same way.

Limitations on Mask Functions

The mask functions in Hive are implemented through GenericUDFs. Even though Impala users can call Hive UDFs, Impala does not yet support Hive GenericUDFs, so you cannot use Hive's mask functions in Impala. However, Impala has builtin mask functions that are implemented through overloads. In Impala, when using mask functions, not all parameter combinations are supported. These mask functions are introduced in Impala 3.4

The following list includes all the implemented overloads.

- Overloads used by Ranger default masking policies,
- Overloads with simple arguments,
- Overload with all arguments in int type for full functionality. Char argument needs to be converted to their ASCII value.

To list the available overloads, use the following query:

```
show functions in _impala_builtins like "mask*";
```

Note: An error message that states "No matching function with signature: mask..." implies that Impala does not contain the corresponding overload.

Related Information

GRANT statement
REVOKE statement
Resource-based column masking in Hive with Ranger policies
Tag-based column masking in Hive with Ranger policies
Apache Ranger documentation

Configuring Authorization

Enable Ranger authorization in Impala.

Procedure

1. In Cloudera Manager, navigate to the Impala service.
2. In the Configuration tab, type ranger in the search field.
3. Specify the values in the following fields.
   - Ranger Service: Select the Ranger service you are using. This is the server name to use when granting server level privileges.
   - Ranger Policy Cache Directory: Specify the directory where Ranger security policies are cached locally.
   - Ranger DFS Audit Path: Specify the DFS path on which Ranger audits are written. The special placeholder '${ranger_base_audit_url}' should be used as the prefix, in order to use the centralized location defined in the Ranger service.
   - Ranger Audit DFS Spool Dir: Specify the spool directory for Ranger audits being written to DFS.
   - Ranger Audit Solr Spool Dir: Specify the spool directory for Ranger audits being written to Solr.
4. Restart Impala and Hive.

Tuning Impala

The following sections explains procedures for tuning Impala queries and other SQL operations.

Setting Up HDFS Caching

Set up HDFS caching with Impala for improved performance.

Before you begin

Decide how much memory to devote to the HDFS cache on each host. The total memory available for cached data is the sum of the cache sizes on all the hosts. By default, any data block is only cached on one host although you can cache a block across multiple hosts by increasing the replication factor.

Procedure

1. Enable or disable HDFS caching through Cloudera Manager, using the configuration setting Maximum Memory Used for Caching for the HDFS service.

   This control sets the HDFS configuration parameter dfs_datanode_max_locked_memory, which specifies the upper limit of HDFS cache size on each node. Set up the HDFS caching for your Hadoop cluster.

   - All the other manipulation of the HDFS caching settings, such as what files are cached, is done through the command line, either Impala DDL statements or the Linux `hdfs cacheadmin` command.
2. Using the `hdfs cacheadmin` command, set up one or more pools owned by the same user as the `impalad` daemon (typically impala).

   For example:

   ```
   hdfs cacheadmin -addPool four_gig_pool -owner impala -limit 4000000000
   ```
3. Once HDFS caching is enabled and one or more pools are available, on the Impala side, you specify the cache pool name defined by the hdfs cacheadmin command in the Impala DDL statements that enable HDFS caching for a table or partition, such as CREATE TABLE ... CACHED IN *pool* or ALTER TABLE ... SET CACHED IN *pool*.
4. You can use `hdfs cacheadmin -listDirectives` to get a list of existing cache pools.
5. You can use `hdfs cacheadmin -listDirectives -stats` to get detailed information about the pools.

Setting up Data Cache for Remote Reads

When Impala compute nodes and its storage are not co-located, the network bandwidth requirement goes up as the network traffic includes the data fetch as well as the shuffling exchange traffic of intermediate results. To mitigate the pressure on the network, you can enable the compute nodes to cache the working set read from remote filesystems, such as, remote HDFS data node, S3, ABFS, ADLS.

About this task

To enable remote data cache as follows.

Procedure

1. In Cloudera Manager, navigate to ClustersImpala Service.
2. In the Configuration tab, select Enable Local Data Cache to enable the local Impala Daemon data cache for caching of remote reads.
3. In Impala Daemon Data Cache Directories, add the directories Impala Daemon will use for caching of remote read data.
4. In Impala Daemon Data Cache Per Directory Capacity, specify the maximum amount of local disk space Impala will use per daemon in each of the configured directrories for caching of remote read data.
5. Click Save Changes and restart the Impala service.

Configuring Dedicated Coordinators and Executors

Configure a dedicated coordinator and a dedicated executor roles to improve scalability of Impala.

Guidelines for Dedicated Coordinators and Executors

- Dedicated coordinator:
  - Should be on an edge node with no other services running on it.
  - Does not need large local disks but still needs some that can be used for Spilling.
  - Require at least the same or even larger memory allocation than executors.
- (Dedicated)Executors:
  - Should be co-located with DataNodes.
  - The number of hosts with dedicated executors typically increases as the cluster grows larger and handles more table partitions, data files, and concurrent queries.

Procedure

1. In Cloudera Manager, navigate to ClustersImpalaConfigurationRole Groups.
2. Click Create to create two role groups with the following values.
   a) Group for Coordinators
      1. Group Name: Coordinators
      2. Role Type: Impala Daemon
      3. Copy from:
         - Select Impala Daemon Default Group if you want the existing configuration gets carried over to the Coordinators.
         - Select None if you want to start with a blank configuration.

   b) Group for Executors
      1. Group Name: Executors
      2. Role Type: Impala Daemon
      3. Copy from:
         - Select Impala Daemon Default Group if you want the existing configuration gets carried over to the Executors.
         - Select None if you want to start with a blank configuration.
3. In the Role Groups page, click Impala Daemon Default Group.
   a) Select the set of nodes intended to be coordinators.
      1. Click Action for Selected and select Move To Different Role Group….
      2. Select the Coordinators.

   b) Select the set of nodes intended to be Executors.
      1. Click Action for Selected and select Move To Different Role Group….
      2. Select the Executors.
4. Click Configuration. In the search field, type Impala Daemon Specialization.
5. Click Edit Individual Values.
6. For Coordinators role group, select COORDINATOR_ONLY.
7. For Executors role group, select EXECUTOR_ONLY.
8. Click Save Changes and then restart the Impala service.

Related Information

Dedicated Coordinator

Managing Resources in Impala

Impala includes the features that balance and maximize resources to improve query performance and scalability of your Cloudera cluster.

A typical deployment uses the following resource management features:

- Static service pools

  Use the static service pools to allocate dedicated resources for Impala to manage and prioritize workloads on clusters.
- Admission control

  Within the constraints of the static service pool, you can further subdivide Impala's resources using dynamic resource pools and admission control.

Related Information

Creating Static Pools

Dynamic Resource Pool Settings

Admission Control and Query Queuing

Admission control is an Impala feature that imposes limits on concurrent SQL queries to avoid resource usage spikes and out-of-memory conditions on busy Cloudera clusters.

The admission control feature lets you set an upper limit on the number of concurrent Impala queries and on the memory used by those queries. Any additional queries are queued until the earlier ones finish, rather than being cancelled or running slowly and causing contention. As other queries finish, the queued queries are allowed to proceed.

You can specify these limits and thresholds for each pool or globally so that you can balance the resource usage and throughput among steady well-defined workloads, rare resource-intensive queries, and ad-hoc exploratory queries.

In addition to the threshold values for currently executing queries, you can place limits on the maximum number of queries that are queued (waiting) and a limit on the amount of time they might wait before returning with an error. These queue settings let you ensure that queries do not wait indefinitely so that you can detect and correct “starvation” scenarios.

Queries, DML statements, and some DDL statements, including CREATE TABLE AS SELECT and COMPUTE STATS are affected by admission control.

On a busy Cloudera cluster, you might find there is an optimal number of Impala queries that run concurrently. For example, when the I/O capacity is fully utilized by I/O-intensive queries, you might not find any throughput benefit in running more concurrent queries. By allowing some queries to run at full speed while others wait, rather than having all queries contend for resources and run slowly, admission control can result in higher overall throughput.

For another example, consider a memory-bound workload such as many large joins or aggregation queries. Each such query could briefly use many gigabytes of memory to process intermediate results. Because Impala by default cancels queries that exceed the specified memory limit, running multiple large-scale queries at once might require re-running some queries that are cancelled. In this case, admission control improves the reliability and stability of the overall workload by only allowing as many concurrent queries as the overall memory of the cluster can accommodate.

Concurrent Queries and Admission Control

One way to limit resource usage through admission control is to set an upper limit on the number of concurrent queries. This is the initial technique you might use when you do not have extensive information about memory usage for your workload. The settings can be specified separately for each dynamic resource pool.

Max Running Queries

Maximum number of concurrently running queries in this pool. The default value is unlimited. (optional)

The maximum number of queries that can run concurrently in this pool. The default value is unlimited. Any queries for this pool that exceed Max Running Queries are added to the admission control queue until other queries finish. You can use Max Running Queries in the early stages of resource management, when you do not have extensive data about query memory usage, to determine if the cluster performs better overall if throttling is applied to Impala queries.

For a workload with many small queries, you typically specify a high value for this setting, or leave the default setting of “unlimited”. For a workload with expensive queries, where some number of concurrent queries saturate the memory, I/O, CPU, or network capacity of the cluster, set the value low enough that the cluster resources are not overcommitted for Impala.

Once you have enabled memory-based admission control using other pool settings, you can still use Max Running Queries as a safeguard. If queries exceed either the total estimated memory or the maximum number of concurrent queries, they are added to the queue.

If Max Running Queries Multiple is set, the Max Running Queries setting is ignored.

Max Running Queries Multiple

This floating point number is multiplied by the current total number of executors at runtime to give the maximum number of concurrently running queries allowed in the pool. The effect of this setting scales with the number of executors in the resource pool.

This calculation is rounded up to the nearest integer, so the result will always be at least one.

If set to zero or a negative number, the setting is ignored.

Max Queued Queries

Maximum number of queries that can be queued in this pool. The default value is 200. (optional).

If Max Queued Queries Multiple is set, the Max Queued Queries setting is ignored.

Max Queued Queries Multiple

This floating point number is multiplied by the current total number of executors at runtime to give the maximum number of queries that can be queued in the pool. The effect of this setting scales with the number of executors in the resource pool.

This calculation is rounded up to the nearest integer, so the result will always be at least one.

If set to zero or a negative number, the setting is ignored.

Queue Timeout

The amount of time, in milliseconds, that a query waits in the admission control queue for this pool before being canceled. The default value is 60,000 milliseconds.

In the following cases, Queue Timeout is not significant, and you can specify a high value to avoid canceling queries unexpectedly:

- In a low-concurrency workload where few or no queries are queued
- In an environment without a strict SLA, where it does not matter if queries occasionally take longer than usual because they are held in admission control

You might also need to increase the value to use Impala with some business intelligence tools that have their own timeout intervals for queries.

In a high-concurrency workload, especially for queries with a tight SLA, long wait times in admission control can cause a serious problem. For example, if a query needs to run in 10 seconds, and you have tuned it so that it runs in 8 seconds, it violates its SLA if it waits in the admission control queue longer than 2 seconds. In a case like this, set a low timeout value and monitor how many queries are cancelled because of timeouts. This technique helps you to discover capacity,

tuning, and scaling problems early, and helps avoid wasting resources by running expensive queries that have already missed their SLA.

If you identify some queries that can have a high timeout value, and others that benefit from a low timeout value, you can create separate pools with different values for this setting.

You can combine these settings with the memory-based approach described below. If either the maximum number of or the expected memory usage of the concurrent queries is exceeded, subsequent queries are queued until the concurrent workload falls below the threshold again.

Memory Limits and Admission Control

Each dynamic resource pool can have an upper limit on the cluster-wide memory used by queries executing in that pool.

Use the following settings to manage memory-based admission control.

Max Memory

The maximum amount of aggregate memory available across the cluster to all queries executing in this pool. This should be a portion of the aggregate configured memory for Impala daemons, which will be shown in the settings dialog next to this option for convenience. Setting this to a non-zero value enables memory based admission control.

Impala determines the expected maximum memory used by all queries in the pool and holds back any further queries that would result in Max Memory being exceeded.

If you specify Max Memory, you should specify the amount of memory to allocate to each query in this pool. You can do this in two ways:

- By setting Maximum Query Memory Limit and Minimum Query Memory Limit. This is preferred as it gives Impala flexibility to set aside more memory to queries that are expected to be memory-hungry.
- By setting Default Query Memory Limit to the exact amount of memory that Impala should set aside for queries in that pool.

Note that in the following cases, Impala will rely entirely on memory estimates to determine how much memory to set aside for each query. This is not recommended because it can result in queries not running or being starved for memory if the estimates are inaccurate. And it can affect other queries running on the same node.

- Max Memory, Maximum Query Memory Limit, and Minimum Query Memory Limit are not set, and the MEM_LIMIT query option is not set for the query.
- Default Query Memory Limit is set to 0, and the MEM_LIMIT query option is not set for the query.

If Max Memory Multiple is set, the Max Memory setting is ignored.

Max Memory Multiple

This number of bytes is multiplied by the current total number of executors at runtime to give the maximum memory available across the cluster for the pool. The effect of this setting scales with the number of executors in the resource pool.

If set to zero or a negative number, the setting is ignored.

Minimum Query Memory Limit and Maximum Query Memory Limit

These two options determine the minimum and maximum per-host memory limit that will be chosen by Impala Admission control for queries in this resource pool. If set, Impala admission control will choose a memory limit between the minimum and maximum values based on the per-host memory estimate for the query. The memory limit chosen determines the amount of memory that Impala admission control will set aside for this query on each host that the query is running on. The

aggregate memory across all of the hosts that the query is running on is counted against the pool's Max Memory.

Minimum Query Memory Limit must be less than or equal to Maximum Query Memory Limit and Max Memory.

A user can override Impala's choice of memory limit by setting the MEM_LIMIT query option. If the Clamp MEM_LIMIT Query Option setting is set to TRUE and the user sets MEM_LIMIT to a value that is outside of the range specified by these two options, then the effective memory limit will be either the minimum or maximum, depending on whether MEM_LIMIT is lower than or higher the range.

For example, assume a resource pool with the following parameters set:

- Minimum Query Memory Limit = 2GB
- Maximum Query Memory Limit = 10GB

If a user tries to submit a query with the MEM_LIMIT query option set to 14 GB, the following would happen:

- If Clamp MEM_LIMIT Query Option = true, admission controller would override MEM_LIMIT with 10 GB and attempt admission using that value.
- If Clamp MEM_LIMIT Query Option = false, the admission controller will retain the MEM_LIMIT of 14 GB set by the user and will attempt admission using the value.

Clamp MEM_LIMIT Query Option

If this field is not selected, the MEM_LIMIT query option will not be bounded by the Maximum Query Memory Limit and the Minimum Query Memory Limit values specified for this resource pool. By default, this field is selected. The field is disabled if both Minimum Query Memory Limit and Maximum Query Memory Limit are not set.

You can combine the memory-based settings with the upper limit on concurrent queries. If either the maximum number of or the expected memory usage of the concurrent queries is exceeded, subsequent queries are queued until the concurrent workload falls below the threshold again.

Monitoring Admission Control

To see how admission control works for particular queries, examine the profile output or the summary output for the query.

Profile

The information is available through the PROFILE statement in `impala-shell` immediately after running a query in the shell, on the queries page of the Impala debug web UI, or in the Impala log file (basic information at log level 1, more detailed information at log level 2).

The profile output contains details about the admission decision, such as whether the query was queued or not and which resource pool it was assigned to. It also includes the estimated and actual memory usage for the query, so you can fine-tune the configuration for the memory limits of the resource pools.

Summary

The summary output includes the queuing status consisting of whether the query was queued and what was the latest queuing reason.

The information is available in `impala-shell` when the LIVE_PROGRESS or LIVE_SUMMARY query option is set to TRUE.

You can also start an impala-shell session with the --live_progress or --live_summary flags to monitor all queries in that impala-shell session.

Enabling Admission Control

Enable admission control on all production clusters to alleviate possible capacity issues. The capacity issues could be because of a high volume of concurrent queries, because of heavy-duty join and aggregation queries that require large amounts of memory, or because Impala is being used alongside other Hadoop data management components and the resource usage of Impala must be constrained to work well.

About this task

Procedure

1. Navigate to ClustersImpala.
2. In the Configuration tab, navigate to CategoryAdmission Control.
3. Select or clear both the Enable Impala Admission Control and the Enable Dynamic Resource Pools.
4. Enter a Reason for change, and then click Save Changes to commit the changes.
5. Restart the Impala service.
6. After completing this task, for further configuration settings, customize the configuration settings for the dynamic resource pools.

Creating Static Pools

To manage and prioritize workloads on clusters, use the static service pools to allocate dedicated resources for Impala for predictable resource availability. When static service pools are used, Cloudera Manager creates a cgroup in which Impala runs. This cgroup limits memory, CPU and Disk I/O according to the static partitioning policy.

About this task

Create static resource pools for the services running in your Cloudera cluster.

Procedure

1. In Cloudera Manager, navigate to Clusters Static service pools .
2. In the Configuration tab, allocate a portion of resource to each component.
   - HDFS always needs to have a minimum of 5-10% of the resources.
   - Generally, YARN and Impala split the rest of the resources.
     - For mostly batch workloads, you might allocate YARN 60%, Impala 30%, and HDFS 10%.
     - For mostly ad-hoc query workloads, you might allocate Impala 60%, YARN 30%, and HDFS 10%.
3. Click Continue.
4. Review the changes and click Continue.
5. Click Restart Now.

Configuring Dynamic Resource Pool

Admission control and dynamic resource pools are enabled by default. However, until you configure the settings for the dynamic resource pools, the admission control feature is effectively not enabled.

About this task

There is always a resource pool designated as root.default. By default, all Impala queries run in this pool when the dynamic resource pool feature is enabled for Impala. You create additional pools when your workload includes identifiable groups of queries (such as from a particular application, or a particular group within your organization)

that have their own requirements for concurrency, memory use, or service level agreement (SLA). Each pool has its own settings related to memory, number of queries, and timeout interval.

Procedure

1. In Cloudera Manager, navigate to ClustersDynamic Resource Pool Configuration. If the cluster has an Impala service, the Resource Pools tab displays under the Impala Admission Control tab.
2. In the Impala Admission Control tab, click Create Resource Pool.
3. Specify a name and resource limits for the pool:
   - In the Resource Pool Name field, type a unique name containing only alphanumeric characters.
   - Optionally, in the Submission Access Control tab, specify which users and groups can submit queries. By default, anyone can submit queries. To restrict this permission, select the Allow these users and groups option and provide a comma-delimited list of users and groups in the Users and Groups fields respectively.
4. Click Create.
5. Click Refresh Dynamic Resource Pools.

Dynamic Resource Pool Settings

Use the following settings to configure your dynamic resource pools for Impala.

Max Memory

Maximum amount of aggregate memory available across the cluster to all queries executing in this pool. This should be a portion of the aggregate configured memory for Impala daemons, which will be shown in the settings dialog next to this option for convenience. Setting this to a non-zero value enables memory based admission control.

Impala determines the expected maximum memory used by all queries in the pool and holds back any further queries that would result in Max Memory being exceeded.

If you specify Max Memory, you should specify the amount of memory to allocate to each query in this pool. You can do this in two ways:

- By setting Maximum Query Memory Limit and Minimum Query Memory Limit. Setting them gives Impala flexibility to set aside more memory to queries that are expected to be memory-hungry.
- By setting Default Query Memory Limit to the exact amount of memory that Impala should set aside for queries in that pool.

Note that if you do not set any of the above options, or set Default Query Memory Limit to 0, Impala will rely entirely on memory estimates to determine how much memory to set aside for each query. This is not recommended because it can result in queries not running or being starved for memory if the estimates are inaccurate.

For example, consider the following scenario:

- The cluster is running impalad daemons on five hosts.
- A dynamic resource pool has Max Memory set to 100 GB.
- The Maximum Query Memory Limit for the pool is 10 GB and Minimum Query Memory Limit is 2 GB. Therefore, any query running in this pool could use up to 50 GB of memory (Maximum Query Memory Limit * number of Impala nodes).
- Impala will run varying numbers of queries concurrently because queries may be given memory limits anywhere between 2 GB and 10 GB, depending on the estimated memory requirements. For example, Impala may run up to 10 small queries with 2 GB memory limits or two large queries with 10 GB memory limits because that is what will fit in the 100 GB cluster-wide limit when executing on five hosts.

- The executing queries may use less memory than the per-host memory limit or the Max Memory cluster-wide limit if they do not need that much memory. In general this is not a problem so long as you are able to run enough queries concurrently to meet your needs.

Minimum Query Memory Limit and Maximum Query Memory Limit

These two options determine the minimum and maximum per-host memory limit that will be chosen by Impala Admission control for queries in this resource pool. If set, Impala admission control will choose a memory limit between the minimum and maximum value based on the per-host memory estimate for the query. The memory limit chosen determines the amount of memory that Impala admission control will set aside for this query on each host that the query is running on. The aggregate memory across all of the hosts that the query is running on is counted against the pool's Max Memory.

Minimum Query Memory Limit must be less than or equal to Maximum Query Memory Limit and Max Memory.

You can override Impala's choice of memory limit by setting the MEM_LIMIT query option. If the Clamp MEM_LIMIT Query Option is selected and the user sets MEM_LIMIT to a value that is outside of the range specified by these two options, then the effective memory limit will be either the minimum or maximum, depending on whether MEM_LIMIT is lower than or higher than the range.

Max Running Queries

Maximum number of concurrently running queries in this pool. The default value is unlimited.

The maximum number of queries that can run concurrently in this pool. The default value is unlimited. Any queries for this pool that exceed Max Running Queries are added to the admission control queue until other queries finish. You can use Max Running Queries in the early stages of resource management, when you do not have extensive data about query memory usage, to determine if the cluster performs better overall if throttling is applied to Impala queries.

For a workload with many small queries, you typically specify a high value for this setting, or leave the default setting of "unlimited". For a workload with expensive queries, where some number of concurrent queries saturate the memory, I/O, CPU, or network capacity of the cluster, set the value low enough that the cluster resources are not overcommitted for Impala.

Once you have enabled memory-based admission control using other pool settings, you can still use Max Running Queries as a safeguard. If queries exceed either the total estimated memory or the maximum number of concurrent queries, they are added to the queue.

Max Queued Queries

Maximum number of queries that can be queued in this pool. The default value is 200. (optional)

Queue Timeout

The amount of time, in milliseconds, that a query waits in the admission control queue for this pool before being canceled. The default value is 60,000 milliseconds.

It the following cases, Queue Timeout is not significant, and you can specify a high value to avoid canceling queries unexpectedly:

- In a low-concurrency workload where few or no queries are queued
- In an environment without a strict SLA, where it does not matter if queries occasionally take longer than usual because they are held in admission control

You might also need to increase the value to use Impala with some business intelligence tools that have their own timeout intervals for queries.

In a high-concurrency workload, especially for queries with a tight SLA, long wait times in admission control can cause a serious problem. For example, if a query needs to run in 10 seconds, and you have tuned it so that it runs in 8 seconds, it violates its SLA if it waits in the admission control queue longer than 2 seconds. In a case like this, set a low timeout value and monitor how

many queries are cancelled because of timeouts. This technique helps you to discover capacity, tuning, and scaling problems early, and helps avoid wasting resources by running expensive queries that have already missed their SLA.

If you identify some queries that can have a high timeout value, and others that benefit from a low timeout value, you can create separate pools with different values for this setting.

Clamp MEM_LIMIT Query Option

If this field is not selected, the MEM_LIMIT query option will not be bounded by the Maximum Query Memory Limit and the Minimum Query Memory Limit values specified for this resource pool. By default, this field is selected. The field is disabled if both Minimum Query Memory Limit and Maximum Query Memory Limit are not set.

Admission Control Sample Scenario

You can learn about the factors you must consider when allocating Impala's resources and the process you need to follow to set up admission control for the selected workload.

Anne Chang is administrator for an enterprise data hub that runs a number of workloads, including Impala.

Anne has a 20-node cluster that uses Cloudera Manager static partitioning. Because of the heavy Impala workload, Anne needs to make sure Impala gets enough resources. While the best configuration values might not be known in advance, she decides to start by allocating 50% of resources to Impala. Each node has 128 GiB dedicated to each impalad. Impala has 2560 GiB in aggregate that can be shared across the resource pools she creates.

Next, Anne studies the workload in more detail. After some research, she might choose to revisit these initial values for static partitioning.

To figure out how to further allocate Impala's resources, Anne needs to consider the workloads and users, and determine their requirements. There are a few main sources of Impala queries:

- Large reporting queries executed by an external process/tool. These are critical business intelligence queries that are important for business decisions. It is important that they get the resources they need to run. There typically are not many of these queries at a given time.
- Frequent, small queries generated by a web UI. These queries scan a limited amount of data and do not require expensive joins or aggregations. These queries are important, but not as critical, perhaps the client tries resending the query or the end user refreshes the page.
- Occasionally, expert users might run ad-hoc queries. The queries can vary significantly in their resource requirements. While Anne wants a good experience for these users, it is hard to control what they do (for example, submitting inefficient or incorrect queries by mistake). Anne restricts these queries by default and tells users to reach out to her if they need more resources.

To set up admission control for this workload, Anne first runs the workloads independently, so that she can observe the workload's resource usage in Cloudera Manager. If they could not easily be run manually, but had been run in the past, Anne uses the history information from Cloudera Manager. It can be helpful to use other search criteria (for example, user) to isolate queries by workload. Anne uses the Cloudera Manager chart for Per-Node Peak Memory usage to identify the maximum memory requirements for the queries.

From this data, Anne observes the following about the queries in the groups above:

- Large reporting queries use up to 32 GiB per node. There are typically 1 or 2 queries running at a time. On one occasion, she observed that 3 of these queries were running concurrently. Queries can take 3 minutes to complete.
- Web UI-generated queries use between 100 MiB per node to usually less than 4 GiB per node of memory, but occasionally as much as 10 GiB per node. Queries take, on average, 5 seconds, and there can be as many as 140 incoming queries per minute.
- Anne has little data on ad hoc queries, but some are trivial (approximately 100 MiB per node), others join several tables (requiring a few GiB per node), and one user submitted a huge cross join of all tables that used all system resources (that was likely a mistake).

Based on these observations, Anne creates the admission control configuration with the following pools:

XL_Reporting

Property	Value
Max Memory	1280 GiB
Maximum Query Memory Limit	32 GiB
Minimum Query Memory Limit	32 GiB
Max Running Queries	2
Queue Timeout	5 minutes

This pool is for large reporting queries. To support running 2 queries at a time, the pool memory resources are set to 1280 GiB (aggregate cluster memory). This is for 2 queries, each with 32 GiB per node, across 20 nodes. Anne sets the pool's Maximum Query Memory Limit to 32 GiB so that no query uses more than 32 GiB on any given node. She sets Max Running Queries to 2 (though it is not necessary she do so). She increases the pool's queue timeout to 5 minutes in case a third query comes in and has to wait. She does not expect more than 3 concurrent queries, and she does not want them to wait that long anyway, so she does not increase the queue timeout. If the workload increases in the future, she might choose to adjust the configuration or buy more hardware.

HighThroughput_UI

Property	Value
Max Memory	960 GiB (inferred)
Maximum Query Memory Limit	4 GiB
Minimum Query Memory Limit	2 GiB
Max Running Queries	12
Queue Timeout	5 minutes

This pool is used for the small, high throughput queries generated by the web tool. Anne sets the Maximum Query Memory Limit to 4 GiB per node, and sets Max Running Queries to 12. This implies a maximum amount of memory per node used by the queries in this pool: 48 GiB per node (12 queries * 4 GiB per node memory limit).

Notice that Anne does not set the pool memory resources, but does set the pool's Maximum Query Memory Limit. This is intentional: admission control processes queries faster when a pool uses the Max Running Queries limit instead of the peak memory resources.

This should be enough memory for most queries, since only a few go over 4 GiB per node. For those that do require more memory, they can probably still complete with less memory (spilling if necessary). If, on occasion, a query cannot run with this much memory and it fails, Anne might reconsider this configuration later, or perhaps she does not need to worry about a few rare failures from this web UI.

With regard to throughput, since these queries take around 5 seconds and she is allowing 12 concurrent queries, the pool should be able to handle approximately 144 queries per minute, which is enough for the peak maximum expected of 140 queries per minute. In case there is a large burst of queries, Anne wants them to queue. The default maximum size of the queue is already 200, which should be more than large enough. Anne does not need to change it.

Default

Property	Value
Max Memory	320 GiB
Maximum Query Memory Limit	4 GiB
Minimum Query Memory Limit	2 GiB
Max Running Queries	Unlimited
Queue Timeout	60 Seconds

The default pool (which already exists) is a catch all for ad-hoc queries. Anne wants to use the remaining memory not used by the first two pools, 16 GiB per node (XL_Reporting uses 64 GiB per node, High_Throughput_UI uses 48 GiB per node). For the other pools to get the resources they expect, she must still set the Max Memory resources and the Maximum Query Memory Limit. She sets the Max Memory resources to 320 GiB (16 * 20). She sets the Maximum Query Memory Limit to 4 GiB per node for now. That is somewhat arbitrary, but satisfies some of the ad hoc queries she observed. If someone writes a bad query by mistake, she does not actually want it using all the system resources. If a user has a large query to submit, an expert user can override the Maximum Query Memory Limit (up to 16 GiB per node, since that is bound by the pool Max Memory resources). If that is still insufficient for this user's workload, the user should work with Anne to adjust the settings and perhaps create a dedicated pool for the workload.

Cancelling a Query

Various client applications let you interactively cancel queries submitted or monitored through those applications.

Setting a Time Limit on Query Execution

An Impala administrator can set a default value of the EXEC_TIME_LIMIT_S query option for a resource pool. If a user accidentally runs a large query that executes for longer than the limit, it will be automatically terminated after the time limit expires to free up resources.

You can override the default value per query or per session if you do not want to apply the default EXEC_TIME_LIMIT_S value to a specific query or a session.

Interactively Cancelling a Query

- In the Impala Web UI for the `impalad` host (on port 25000 by default), cancel a query: In the /queriestab, click Cancel for a query in the queries in flight list.
- In `impala-shell`, press ^C
- In Hue, click Cancel from the Watch page.

You can manually cancel a query in the Impala Web UI for the `impalad` host (on port 25000 by default):

Using HLL Datasketch Algorithms in Impala

You can use Datasketch algorithms (HLL) for queries that take too long to calculate exact results due to very large data sets (e.g. number of distinct values).

You may use data sketches (HLL algorithms) to generate approximate results that are much faster to retrieve. HLL is an algorithm that gives approximate answers for computing the number of distinct values in a column. The value returned by this algorithm is similar to the result of COUNT(DISTINCT col) and the NDV function integrated with Impala. However, HLL algorithm is much faster than COUNT(DISTINCT col) and the NDV function and is less memory-intensive for columns with high cardinality.

These HLL algorithms create "sketches", which are data structures containing an approximate calculation of some facet of the data and which are compatible regardless of which system reads or writes them. For e.g., This particular algorithm works in two phases. First it creates a sketch from the dataset provided for the algorithm and then as the second step in the process you can read the estimate from the sketch. Using this approach it's also possible to create granular sketches for a specific dataset (e.g. one sketch per partition) and later on use them to provide answers for different count(distinct) queries without reading the actual data as they can also be merged together. Using HLL you can create sketches by Hive and read it with Impala for getting the estimate and vice versa.

Usage Notes

Query results produced by the HLL algorithm are approximate but within well defined error bounds. Because the number of distinct values in a column returned by this HLL algorithm is an estimate, it might not reflect the precise number of different values in the column, especially if the cardinality is very high.

Data Types

The following is a list of currently supported data types of the datasets that these HLL functions are compatible with.

Supported:

1. Numeric Types:

   TINYINT, INT, BIGINT, FLOAT, DOUBLE
2. String types:

   STRING, CHAR, VARCHAR
3. Other types:

   BINARY

Unsupported:

1. Complex types
2. Some of the scalar types: DECIMAL, TIMESTAMP, BOOLEAN, SMALLINT, DATE

File Formats

The currently supported file formats are ORC and Parquet. Since Impala does not yet have read support for materialized views, it is required for Hive to write the sketches into regular tables (i.e. not into a materialized view), using the BINARY type. Then Impala can be used to read these sketches and calculate the estimate as Hive does.

Limitations:

You can write HLL sketches into a text table, however it may not work as the serialized sketches can contain characters that are special for text format.

Internal Configuration

The configuration is hard coded and the level of compression is HLL_4 and the number of buckets in the HLL array is k=12 (2 power of 12).

Using HLL Sketches

This datasketch algorithm works in two phases.

- First it creates a sketch from the dataset provided for the algorithm.
- As the second step in the process the function provides an estimate from the sketch.

The following example shows how to create a sketch and to create an estimate from the sketch.

Creating a sketch

This query creates a sketch using the ds_hll_sketch function from a column date_string_col containing the supported data type.

```
Select ds_hll_sketch(date_string_col) from tableName;
```

This query receives a primitive expression and returns a serialized HLL sketch that is not in a readble format.

Estimating the number of distinct values in a column

The following query returns a cardinality estimate (similarly to ndv()) for that particular column (date_string_col) by using the HLL function ds_hll_estimate.

```
Select ds_hll_estimate(ds_hll_sketch(date_string_col)) from tableName;
```

This query receives a primitive expression (column name) and then creates a sketch from the column, and then passes the sketch to the outer function that gives a cardinality estimate.

```
+-----------------------------------------------+
| ds_hll_estimate(ds_hll_sketch(date_string_col))|
+-----------------------------------------------+
| 729                                           |
+-----------------------------------------------+
Fetched 1 row(s) in 0.25s
```

The above result is similar to the results received using the count(distinct col) or the ndv() approach.

```
Select count(distinct date_string_col) from tableName;
+-----------------------------------------------+
| count (distinct date_string_col)              |
+-----------------------------------------------+
| 730                                           |
+-----------------------------------------------+
```

```
Select ndv(distinct date_string_col) from tableName;
+-----------------------------------------------+
| ndv (date_string_col)                         |
+-----------------------------------------------+
| 736                                           |
+-----------------------------------------------+
```

Inserting the derived sketch into a table

You can create a sketch and save it for later use. If you save the sketches to a table or view with the same granularity (e.g. one sketch for each partition in a table) then later you can simply read the sketch and get the estimate almost for free as the real cost here is just the sketch creation.

The following example shows the creation of a sketch for the column (date_string_col) from the table “tableName” and saving the created sketch into another table called sketch_table and later reading the sketch and its estimate from the saved table called sketch_table.

```
   select year, month, ds_hll_estimate(ds_hll_sketch(date_string_col)) from
tableName group by year, month;
   +-------------------------------------------------------------+
   | year | month | ds_hll_estimate(ds_hll_sketch(date_string_col))|
   +------+-------+-----------------------------------------------+
   | 2010 | 5     | 31                                            |
   | 2010 | 11    | 30                                            |
   | 2010 | 6     | 30                                            |
   | 2010 | 2     | 28                                            |
   | 2010 | 10    | 31                                            |
   | 2009 | 11    | 30                                            |
   | 2009 | 7     | 31                                            |
   | 2009 | 10    | 31                                            |
   | 2010 | 11    | 31                                            |
   | 2009 | 7     | 31                                            |
   | 2010 | 10    | 31                                            |
   | 2010 | 8     | 30                                            |
   | 2010 | 5     | 31                                            |
   | 2009 | 7     | 30                                            |
```

```
| 2010 | 4     | 31                                                |
| 2010 | 12    | 30                                                |
| 2009 | 9     | 31                                                |
| 2009 | 8     | 30                                                |
| 2010 | 6     | 28                                                |
| 2010 | 3     | 31                                                |
| 2010 | 4     | 31                                                |
| 2010 | 2     | 30                                                |
| 2009 | 1     | 31                                                |
| 2009 | 12    | 30                                                |
+-----------------------------------------------------------------+
```

Insert the created sketch into another table:

```
insert into sketch_table select year, month, ds_hll_sketch(date_string_
col) from tableName group by year, month;
```

Verify the number of rows modified using count function in the table:

```
select count(1) from sketch_table;
```

You can also verify it by getting the estimates from the sketches stored in the table using the following query:

```
select year, month, ds_hll_estimate(sketch_col) from sketch_table;
```

Combining Sketches

You can combine sketches to allow complex queries to be computed from a few number of simple sketches.

```
select ds_hll_estimate(ds_hll_union(sketch_col)) from sketch_table;
```

Non-deterministic vs Deterministic Results

Based on the characteristics of a query, DataSketches HLL might return results in a non-deterministic manner. During the cache populating phase it returns exact results but in the approximating phase the algorithm can be sensitive to the order of data it receives. As a result you might see the results differ in consecutive runs but all of them will be within the error range of the algorithm.

Due to the characteristics of the algorithm, and to how it was implemented in Impala, the results might be non-deterministic when the query is executed on one node and the input data set is big enough that the algorithm starts approximating instead of giving exact results.

Inter-opreability between Hive and Impala

From the above examples you will notice that the NDV function is much faster than the Datasketches. But NDV is used only by Impala and not by Hive whereas the Datasketches are implemented in Hive as well as in Impala. In order to maintain the inter-operability between Hive and Impala we recommend to use Datasketches to calculate the estimate. A sketch created by Hive using Datasketches can be fed to Impala to produce an estimate.

Using KLL Datasketch Algorithms in Impala

You can use the stochastic streaming algorithm (KLL) that uses the percentile/quantile functions to statistically analyze the approximate distribution of comparable data from a very large stream.

Use the ds_kll_quantiles() or its inverse functions the Probability Mass Function ds_kll_PMF() and the Cumulative Distribution Function ds_kll_CDF() to obtain the analysis. Query results produced by this KLL algorithm are approximate but within well defined error bounds.

Data Types and File Formats

KLL function currently supports only floats. If you sketch an int data the KLL function converts it into float.

You can sketch data from any file formats. However it is recommended to save the sketches into parquet tables. It is not recommended to save the sketches to text format.

Using KLL Sketches

This datasketch algorithm works in two phases.

- First it creates a sketch from the dataset provided for the algorithm.
- As the second step in the process you can read the estimate from the sketch.

Using this approach you can create granular sketches for a specific dataset (e.g. one sketch per partition) and later use them to provide answers for different quantile queries without reading the actual data as they can also be merged together. Using KLL you can create sketches by Hive and read it with Impala for getting the estimate and vice versa.

Creating a sketch

The following example shows how to create a sketch using the ds_kll_sketch function from a column (float_col) containing the supported data type. This created sketch is a data structure containing an approximate calculation of some facet of the data and which are compatible regardless of which system reads or writes them.

```
select ds_kll_sketch(float_col) from tableName;
```

where ds_kll_sketch() is an aggregate function that receives a float parameter (e.g. a float column of a table) and returns a serialized Apache DataSketches KLL sketch of the input data set wrapped into STRING type. Since the serialized data might contain unsupported characters this query may error out. So another approach to this method is to insert the derived sketch into a table or a view and later use for quantile approximations.

Inserting the derived sketch into a table

If you save the sketches to a table or view with the same granularity (e.g. one sketch for each partition in a table) then later you can simply read the sketch for quantile approximation.

```
insert into table_name select ds_kll_sketch(float_col) from tableName;
```

Debugging the created sketch

To troubleshoot the sketch created by the ds_kll_sketch function, use ds_kll_stringify on the sketch. ds_kll_stringify () receives a string that is a serialized Apache DataSketches sketch and returns its stringified format by invoking the related function on the sketch's interface.

```
select ds_kll_stringify(ds_kll_sketch(float_col)) from tableName;
```

```
    +-------------------------------------------+
    | ds_kll_stringify(ds_kll_sketch(float_col))|
    +-------------------------------------------+
```

```
    | ### KLL sketch summary:                    |
    |    K                     : 200             |
    |    min K                 : 200             |
    |    M                     : 8               |
    |    N                     : 100             |
    |    Epsilon               : 1.33%           |
    |    Epsilon PMF           : 1.65%           |
    |    Empty                 : false           |
    |    Estimation mode       : false           |
    |    Levels                : 1               |
    |    Sorted                : false           |
    |    Capacity items        : 200             |
    |    Retained items        : 100             |
    |    Storage bytes         : 432             |
    |    Min value             : 0               |
    |    Max value             : 9               |
    | ### End sketch summary                     |
    +--------------------------------------------+
```

Determining the number of datasets sketched

To determine the number of records sketched into this sketch, use ds_kll_n function.

```
select ds_kll_n(ds_kll_sketch(float_col)) from tableName;
```

```
    +-------------------------------------+
    | ds_kll_n(ds_kll_sketch(float_col)) |
    +-------------------------------------+
    | 100                                 |
    +-------------------------------------+
```

Calculate Quantile

This function ds_kll_quantile() function receives two parameters, a STRING parameter that contains a serialized KLL sketch and a DOUBLE (0.5) that represents the rank of the quantile in the range of [0,1]. E.g. rank=0.5 means the approximate median of all the sketched data.

```
    +-------------------------------------------------+
    | ds_kll_quantile(ds_kll_sketch(float_col), 0.5) |
    +-------------------------------------------------+
    | 4.0                                             |
    +-------------------------------------------------+
```

This query returns a data (4.0) that is bigger than the 50% of the data.

Calculating quantiles

To calculate the quantiles for multiple rank parameters, use the function ds_kll_quantiles_as_string () that is very similar to ds_kll_quantile() but this function receives any number of rank parameters and returns a comma separated string that holds the results for all of the given ranks.

```
    select ds_kll_quantiles_as_string(ds_kll_sketch(float_col), 0.5, 0.6, 1)
 from tableName;
```

```
+-----------------------------------------------------------------+
| ds_kll_quantiles_as_string(ds_kll_sketch(float_col), 0.5, 0.6, 1) |
+-----------------------------------------------------------------+
| 4, 5, 9                                                         |
+-----------------------------------------------------------------+
```

Calculate Rank

This rank function ds_kll_rank() receives two parameters, a STRING that represents a serialized DataSketches KLL sketch and a float to provide a probing value in the sketch.

```
select ds_kll_rank(ds_kll_sketch(float_col), 4) from tableName;
```

```
+-------------------------------------------------+
| ds_kll_rank(ds_kll_sketch(float_col), 4)        |
+-------------------------------------------------+
| 0.48                                            |
+-------------------------------------------------+
```

This query returns a DOUBLE that is the rank of the given probing value in the range of [0,1]. E.g. a return value of 0.48 means that the probing value given as parameter is greater than 48% of all the values in the sketch.

Note: This is only an approximate calculation.

Calculate Probability Mass Function (PMF)

This Probabilistic Mass Function (PMF) ds_kll_pmf_as_string()receives a serialized KLL sketch and one or more float values to represent ranges in the sketched values. In the following example, the float values [1, 3, 4, 8] represent the following ranges: (-inf, 1), [1, 3), [3, 4), [4, 8) [8, +inf)

Note: The input values for the ranges have to be unique and monotonically increasing.

```
select ds_kll_pmf_as_string(ds_kll_sketch(float_col), 1, 3, 4, 8) from table
Name
```

```
+-------------------------------------------------------------+
| ds_kll_pmf_as_string(ds_kll_sketch(float_col), 1, 3, 4, 8) |
+-------------------------------------------------------------+
| 0.12, 0.24, 0.12, 0.36, 0.16                                |
+-------------------------------------------------------------+
```

This query returns a comma separated string where each value in the string is a number in the range of [0,1] and shows what percentage of the data is in the particular ranges.

Calculate Cumulative Distribution Function (CDF)

This Cumulative Distribution Function (CDF) ds_kll_cmf_as_string()receives a serialized KLL sketch and one or more float values to represent ranges in the sketched values. In the following example, the float values [1, 3, 4, 8] represents the following ranges: (-inf, 1), (-inf, 3), (-inf, 4), (-inf, 8), (-inf, +inf)

Note: The input values for the ranges have to be unique and monotonically increasing.

```
select ds_kll_cdf_as_string(ds_kll_sketch(float_col), 1, 3, 4, 8) from table
Name;
```

```
    +-------------------------------------------------------------+
    | ds_kll_cdf_as_string(ds_kll_sketch(float_col), 1, 3, 4, 8) |
    +-------------------------------------------------------------+
    | 0.12, 0.36, 0.48, 0.84, 1                                   |
    +-------------------------------------------------------------+
```

This query returns a comma separated string where each value in the string is a number in the range of [0,1] and shows what percentage of the data is in the particular ranges.

Calculate the Union

To take advantage of the UNION function, create granular sketches for a specific dataset (one sketch per partition), write these sketches to a separate table and based on the partition you are interested in, you can UNION the relevant sketches together to get an estimate.

```
insert into sketch_tbl select ds_kll_sketch(float_col) from tableName;
```

```
select ds_kll_quantile(ds_kll_union(sketch_col), 0.5) from sketch_tbl
    where partition_col=1 OR partition_col=5;
```

This function ds_kll_union() receives a set of serialized Apache DataSketches KLL sketches produced by ds_kll_s ketch() and merges them into a single sketch.

Managing Metadata in Impala

This section describes various knobs you can use to control how Impala manages its metadata in order to improve performance and scalability.

On-demand Metadata

With the on-demand metadata feature, the Impala coordinators pull metadata as needed from catalogd and cache it locally. The cached metadata gets evicted automatically under memory pressure.

The granularity of on-demand metadata fetches is at the partition level between the coordinator and catalogd. Common use cases like add/drop partitions do not trigger unnecessary serialization/deserialization of large metadata.

The feature can be used in either of the following modes.

Metadata on-demand mode

In this mode, all coordinators use the metadata on-demand.

Set the following on catalogd:

```
--catalog_topic_mode=minimal
```

Set the following on all impalad coordinators:

```
--use_local_catalog=true
```

Mixed mode

In this mode, only some coordinators are enabled to use the metadata on-demand.

We recommend that you use the mixed mode only for testing local catalog's impact on heap usage.

Set the following on catalogd:

```
--catalog_topic_mode=mixed
```

Set the following on impalad coordinators with metdadata on-demand:

```
--use_local_catalog=true
```

Limitation:

HDFS caching is not supported in On-demand metadata mode coordinators.

Note:

In Impala 3.4.0 and above, global INVALIDATE METADATA statement is supported when On-demand feature is enabled.

INVALIDATE METADATA Usage Notes:

To return accurate query results, Impala needs to keep the metadata current for the databases and tables queried. Through "automatic invalidation" or "HMS event polling" support, Impala automatically picks up most changes in metadata from the underlying systems. However there are some scenarios where you might need to run INVALIDA TE METADATA or REFRESH.

- if some other entity modifies information used by Impala in the metastore, the information cached by Impala must be updated via INVALIDATE METADATA or REFRESH,
- if you have "local catalog" enabled without "HMS event polling" and need to pick up metadata changes that were done outside of Impala in Hive and other Hive client, such as SparkSQL,
- and so on.
- **Note:** As this is a very expensive operation compared to the incremental metadata update done by the REFRESH statement, when possible, prefer REFRESH rather than INVALIDATE METADATA.

Related Information

Invalidate Metadata Statement

Automatic Invalidation of Metadata Cache

To keep the size of metadata bounded, the Impala Catalog Server periodically scans all the tables and invalidates those not recently used.

There are two types of configurations in Catalog Server that control the automatic invalidation of metadata in the Catalog Server Command Line Argument Advanced Configuration Snippet (Safety Valve) field in Cloudera Manager.

Time-based cache invalidation

Catalogd invalidates tables that are not recently used in the specified time period (in seconds).

The ##invalidate_tables_timeout_s flag needs to be applied to both impalad and catalogd.

Memory-based cache invalidation

When the memory pressure reaches 60% of JVM heap size after a Java garbage collection in cata logd, Impala invalidates 10% of the least recently used tables.

The ##invalidate_tables_on_memory_pressure flag needs to be applied to both impalad and cata logd.

Automatic invalidation of metadata provides more stability with lower chances of running out of memory, but the feature could potentially cause performance issues and may require tuning.

Automatic Invalidation/Refresh of Metadata

In this release, you can invalidate or refresh metadata automatically after changes to databases, tables or partitions render metadata stale. You control the synching of tables or database metadata by basing the process on events. You learn how to access metrics and state information about the invalidate event processor.

When tools such as Hive and Spark are used to process the raw data ingested into Hive tables, new HMS metadata (database, tables, partitions) and filesystem metadata (new files in existing partitions/tables) are generated. In previous versions of Impala, in order to pick up this new information, Impala users needed to manually issue an INVALIDATE or REFRESH commands.

When automatic invalidate/refresh of metadata is enabled,, the Catalog Server polls Hive Metastore (HMS) notification events at a configurable interval and automatically applies the changes to Impala catalog.

Impala Catalog Server polls and processes the following changes.

- Invalidates the tables when it receives the ALTER TABLE event.
- Refreshes the partition when it receives the ALTER, ADD, or DROP partitions.
- Adds the tables or databases when it receives the CREATE TABLE or CREATE DATABASE events.
- Removes the tables from catalogd when it receives the DROP TABLE or DROP DATABASE events.
- Refreshes the table and partitions when it receives the INSERT events.

  If the table is not loaded at the time of processing the INSERT event, the event processor does not need to refresh the table and skips it.
- Changes the database and updates catalogd when it receives the ALTER DATABASE events. The following changes are supported. This event does not invalidate the tables in the database.
  - Change the database properties
  - Change the comment on the database
  - Change the owner of the database
  - Change the default location of the database

    Changing the default location of the database does not move the tables of that database to the new location. Only the new tables which are created subsequently use the default location of the database in case it is not provided in the create table statement.

This feature is controlled by the ##hms_event_polling_interval_s flag. Start the catalogd with the ##hms_event_polling_interval_s flag set to a positive integer to enable the feature and set the polling frequency in seconds. We recommend the value to be less than 5 seconds.

The following use cases are not supported:

- When you bypass HMS and add or remove data into table by adding files directly on the filesystem, HMS does not generate the INSERT event, and the event processor will not invalidate the corresponding table or refresh the corresponding partition.

  It is recommended that you use the LOAD DATA command to do the data load in such cases, so that event processor can act on the events generated by the LOAD command.

- The Spark API that saves data to a specified location does not generate events in HMS, thus is not supported. For example:

```
Seq((1, 2)).toDF("i", "j").write.save("/user/hive/warehouse/spark_etl.db/
customers/date=01012019")
```

This feature is turned off by default with the ##hms_event_polling_interval_s flag set to 0.

Disable Event Based Automatic Metadata Sync

When the ##hms_event_polling_interval_s flag is set to a non-zero value for your catalogd, the event-based automatic invalidation is enabled for all databases and tables. If you wish to have the fine-grained control on which tables or databases need to be synced using events, you can use the impala.disableHmsSync property to disable the event processing at the table or database level.

When you add the DBPROPERTIES or TBLPROPERTIES with the impala.disableHmsSync key, the HMS event based sync is turned on or off. The value of the impala.disableHmsSync property determines if the event processing needs to be disabled for a particular table or database.

- If 'impala.disableHmsSync'='true', the events for that table or database are ignored and not synced with HMS.
- If 'impala.disableHmsSync'='false' or if impala.disableHmsSync is not set, the automatic sync with HMS is enabled if the ##hms_event_polling_interval_s global flag is set to non-zero.
- To disable the event based HMS sync for a new database, set the impala.disableHmsSync database properties in Hive as currently, Impala does not support setting database properties:

```
CREATE DATABASE <name> WITH DBPROPERTIES ('impala.disableHmsSync'='true');
```

- To enable or disable the event based HMS sync for a table:

```
CREATE TABLE <name> ... TBLPROPERTIES ('impala.disableHmsSync'='true' |
'false');
```

- To change the event based HMS sync at the table level:

```
ALTER TABLE <name> SET TBLPROPERTIES ('impala.disableHmsSync'='true' | '
false');
```

When both table and database level properties are set, the table level property takes precedence. If the table level property is not set, then the database level property is used to evaluate if the event needs to be processed or not.

If the property is changed from true (meaning events are skipped) to false (meaning events are not skipped), you need to issue a manual INVALIDATE METADATA command to reset event processor because it doesn't know how many events have been skipped in the past and cannot know if the object in the event is the latest. In such a case, the status of the event processor changes to NEEDS_INVALIDATE.

Metrics for Event Based Automatic Metadata Sync

You can use the web UI of the catalogd to check the state of the automatic invalidate event processor.

By default, the debug web UI of catalogd is at http://*impala-server-hostname*:25020 (non-secure cluster) or https://*impala-server-hostname*:25020 (secure cluster).

Under the web UI, there are two pages that presents the metrics for HMS event processor that is responsible for the event based automatic metadata sync.

- /metrics#events
- /events

  This provides a detailed view of the metrics of the event processor, including min, max, mean, median, of the durations and rate metrics for all the counters listed on the /metrics#events page.

The /metrics#events page provides the following metrics about the HMS event processor.

Name	Description
events-processor.avg-events-fetch-duration	Average duration to fetch a batch of events and process it.
events-processor.avg-events-process-duration	Average time taken to process a batch of events received from the Metastore.
events-processor.events-received	Total number of the Metastore events received.
events-processor.events-received-15min-rate	Exponentially weighted moving average (EWMA) of number of events received in last 15 min. This rate of events can be used to determine if there are spikes in event processor activity during certain hours of the day.
events-processor.events-received-1min-rate	Exponentially weighted moving average (EWMA) of number of events received in last 1 min. This rate of events can be used to determine if there are spikes in event processor activity during certain hours of the day.
events-processor.events-received-5min-rate	Exponentially weighted moving average (EWMA) of number of events received in last 5 min. This rate of events can be used to determine if there are spikes in event processor activity during certain hours of the day.
events-processor.events-skipped	Total number of the Metastore events skipped. Events can be skipped based on certain flags are table and database level. You can use this metric to make decisions, such as: • If most of the events are being skipped, see if you might just turn off the event processing. • If most of the events are not skipped, see if you need to add flags on certain databases.
events-processor.status	Metastore event processor status to see if there are events being received or not. Possible states are: • PAUSED The event processor is paused because catalog is being reset concurrently. • ACTIVE The event processor is scheduled at a given frequency. • ERROR • The event processor is in error state and event processing has stopped. • NEEDS_INVALIDATE The event processor could not resolve certain events and needs a manual INVALIDATE command to reset the state. • STOPPED The event processing has been shutdown. No events will be processed. • DISABLED The event processor is not configured to run.

Configuring Event Based Automatic Metadata Sync

As the first step to use the HMS event based metadata sync, enable and configure HMS notifications in Cloudera Manager.

Procedure

1. In Cloudera Manager, navigate to ClustersHive.
2. Navigate to ConfigurationFiltersSCOPEHive Metastore Server.

3. In Hive Metastore Server Advanced Configuration Snippet (Safety Valve) for hive-site.xml, click + to expand and enter the following:
   - Name: hive.metastore.notifications.add.thrift.objects
   - Value: true
   - Name: hive.metastore.alter.notifications.basic
   - Value: false
   - Name: hive.metastore.dml.events
   - Value: true
4. Click Save Changes.
5. Navigate to FiltersSCOPEHive-1 (Service-Wide).
6. Select Enable Metastore Notifications for DML Operations.
7. Click Save Changes.
8. If you want the INSERT events are generated when the Spark (and other non-Hive) applications insert data into existing tables and partitions:
   a. Navigate to FiltersSCOPEGateway.
   b. In Hive Client Advanced Configuration Snippet (Safety Valve) for hive-site.xml, click + to expand and enter the following:
      - Name: hive.metastore.dml.events
      - Value: true
   c. Click Save Changes.
9. Restart stale services.