

Operational Database powered by Apache Accumulo 1.0.0

Troubleshooting the OpDB powered by Apache Accumulo installation

Date published: 2021-03-03

Date modified:

The Cloudera logo, featuring the word "CLOUDERA" in a bold, orange, sans-serif font. The letter "E" is stylized with a horizontal bar through its middle.

<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

Under-replicated block exceptions or cluster failure occurs on small clusters.....	4
HDFS storage demands due to retained HDFS trash.....	4
Change the HDFS trash settings in Cloudera Manager.....	4
Disable Operational Database powered by Apache Accumulo's use of HDFS trash.....	5

Under-replicated block exceptions or cluster failure occurs on small clusters

You can troubleshoot the under-replication warning by editing the `dfs.replication.max` and the `table.file.replication` properties.

About this task

By default, Operational Database powered by Apache Accumulo (OpDB) attempts to use a replication factor of five for the HDFS files that make up the metadata table. OpDB ignores the `table.file.replication` setting. This causes under-replication warnings until you correct the number of nodes or until you manually adjust the replication setting on that table.

Normally, Cloudera Manager does not set a maximum replication factor. If it is set, OpDB uses the `table.file.replication` setting only on normal user tables to that maximum.

However, if the `dfs.replication.max` setting for HDFS is adjusted to match the number of cluster nodes, attempts by OpDB to create new files for its internal tables will fail on clusters with fewer than five datanodes.

Procedure

1. Edit the `dfs.replication.max` setting for HDFS to be more or equal to 5.
2. Adjust replication on the metadata and root tables to be less than or equal to the number of datanodes.
3. Lower the `dfs.replication.max` setting value.

For example, to adjust the replication in the Accumulo shell to 3:

```
root@accumulo> config -t accumulo.metadata -s table.file.replication=3
root@accumulo> config -t accumulo.root -s table.file.replication=3
```

HDFS storage demands due to retained HDFS trash

Using Cloudera Manager you can either change the HDFS trash setting or disable the use of the HDFS trash entirely, for your Operational Database powered by Apache Accumulo (OpDB) deployment.

By default, if HDFS trash is enabled, Accumulo uses it for all files it deletes. This includes write-ahead logs and long-term storage files that were blocked due to compaction. By default, the retention period for the HDFS trash is 24 hours. On Accumulo installations with a heavy write workload, this can result in a large amount of data accumulating in the trash folder for the service user.

There are three workarounds:

- Periodically run the `hdfs dfs --expunge` command as the Accumulo service user. The command must be run twice each time you want to purge a backlog of data. On the first run, a trash checkpoint is created. On the second run, the checkpoint created with the first run is removed with immediate effect.
- Change the HDFS trash settings in Cloudera Manager.
- Disable OpDB's use of the HDFS trash entirely.

Change the HDFS trash settings in Cloudera Manager

You can tune the amount of time HDFS retains trash to control how much data Operational Database powered by Apache Accumulo (OpDB) saves.

About this task

This change is HDFS#wide and impacts the ability to recover from accidental deletions unrelated to OpDB.

Procedure

1. In Cloudera Manager, select the HDFS service.
2. Click Configuration.
3. Search for trash.
4. Change the Filesystem Trash Interval to a smaller value. For example, 4 hours.
5. Save your changes.
6. Restart Stale Services with Re-deploy client configuration.

Disable Operational Database powered by Apache Accumulo's use of HDFS trash

Using Cloudera Manager, you can configure Operational Database powered by Apache Accumulo (OpDB) to skip the saving of the HDFS trash.

About this task

In some deployments, changing the system-wide retention period for HDFS trash may not be an option. If that is the case, you can disable Accumulo's use of the HDFS trash entirely. If you do so, any deletions through the Accumulo APIs are unrecoverable.

Procedure

1. In Cloudera Manager, select the Accumulo on CDP service.
2. Click Configuration.
3. Search for accumulo.properties.
4. Find Tablet Server Advanced Configuration Snippet (Safety Valve) for accumulo.properties property.
5. Set the following:

```
gc.trash.ignore=true
```

6. Save your changes.
7. Click Restart Stale Services.