

Cloudera Runtime 7.0.3

Configuring Apache HBase

Date published: 2019-11-22

Date modified:

CLOUdera

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

Legal Notice

© Cloudera Inc. 2025. 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

Use DNS with HBase.....	6
Use the Network Time Protocol (NTP) with HBase.....	6
Configure the graceful shutdown timeout property.....	6
Configure the HBase Thrift Server Role.....	7
Setting user limits for HBase.....	7
Configure ulimit for HBase using Cloudera Manager.....	7
Configure ulimit for HBase using the command line.....	8
Configure ulimit using Pluggable Authentication Modules using the Command Line.....	8
Use dfs.datanode.max.transfer.threads with Hbase.....	9
Configure encryption in HBase.....	9
Using hedged reads.....	10
Enable hedged reads for HBase.....	10
Monitor the performance of hedged reads.....	10
Understanding HBase garbage collection.....	11
Configure HBase garbage collection.....	11
Disable the BoundedByteBufferPool.....	11
Configure the HBase canary.....	12
Using HBase blocksize.....	13
Configure the blocksize for a column family.....	13
Configuring HBase BlockCache.....	14
Contents of the BlockCache.....	14
Size the BlockCache.....	15
Decide to use the BucketCache.....	15
About the Off-heap BucketCache.....	15
Off-heap BucketCache.....	16
BucketCache IO engine.....	18

Configure BucketCache IO engine.....	19
Configure the off-heap BucketCache using Cloudera Manager.....	19
Configure the off-heap BucketCache using the command line.....	20
Cache eviction priorities.....	21
Bypass the BlockCache.....	21
Monitor the BlockCache.....	21
Using quota management.....	22
Configuring quotas.....	22
General Quota Syntax.....	23
Space quotas.....	23
Quota enforcement.....	24
Quota violation policies.....	24
Using HBase scanner heartbeat.....	25
Configure the scanner heartbeat using Cloudera Manager.....	25
Limiting the speed of compactions.....	26
Configure the compaction speed using Cloudera Manager.....	26
Enable HBase indexing.....	27
Using HBase coprocessors.....	27
Add a custom coprocessor.....	27
Disable loading of coprocessors.....	28
Configuring HBase MultiWAL.....	28
Configuring MultiWAL support using Cloudera Manager.....	29
Configuring the storage policy for the Write-Ahead Log (WAL).....	29
Configure the storage policy for WALs using Cloudera Manager.....	29
Configure the storage policy for WALs using the Command Line.....	30
Using RegionServer grouping.....	30
Enable RegionServer grouping using Cloudera Manager.....	30
Configure RegionServer grouping.....	31
Monitor RegionServer grouping.....	31
Remove a RegionServer from RegionServer grouping.....	31
Enable ACL for RegionServer grouping.....	32
Best practices when using RegionServer grouping.....	32
Disable RegionServer grouping.....	32
Optimizing HBase I/O.....	33
HBase I/O components.....	33

Advanced configuration for write-heavy workloads.....	36
--	-----------

Use DNS with HBase

You must configure DNS to resolve RegionServer and Master hostnames in your cluster. HBase uses the local hostname to report its IP address. Both forward and reverse DNS resolving works.

About this task

You must configure a few parameters in the `hbase-site.xml` to use DNS with HBase.

Procedure

If your server has multiple interfaces, HBase uses the interface that the primary hostname resolves to, and you must ensure that the cluster configuration is consistent and every host has the same network interface configuration. If this is your use case:

- Set `hbase.regionserver.dns.interface` in the `hbase-site.xml` file to indicate the primary interface.

If you want to use a different DNS name server than the system-wide default.

- Set `hbase.regionserver.dns.nameserver` in the `hbase-site.xml` file to a different DNS name server.

Use the Network Time Protocol (NTP) with HBase

Ensure that the clocks on all the cluster members are synchronized for your cluster to function correctly. You must configure NTP to synchronize the clock.

About this task

The clocks on cluster members must be synchronized for your cluster to function correctly. Some skew is tolerable, but excessive skew could generate odd behaviors.

Using NTP and DNS ensures that you won't run into odd behaviors when one node A thinks that the time is tomorrow and node B thinks it's yesterday. You will also prevent situations where the master node tells node C to serve a region but node C doesn't know its own name and doesn't answer.

Procedure

1. Run NTP or another clock synchronization mechanism on your cluster.
2. Verify that the system time is synchronized across your cluster nodes.

What to do next

For more information about NTP, see the [NTP website](#)

Configure the graceful shutdown timeout property

You must configure this property to allow enough time for a graceful shutdown of a RegionServer. A graceful shutdown of an HBase RegionServer allows the regions hosted by that RegionServer to be moved to other RegionServers before stopping the RegionServer.

About this task

This timeout only affects a graceful shutdown of the entire HBase service, not individual RegionServers. Therefore, if you have a large cluster with many RegionServers, you should strongly consider increasing the timeout from its default of 180 seconds.

Procedure

1. Go to the HBase service.
2. Click the Configuration tab.
3. Select Scope > HBASE-1 (Service Wide).
4. Use the Search box to search for the Graceful Shutdown Timeout property and edit the value.
5. Click Save Changes to save this setting.

Related Information

[Graceful HBase Shutdown](#)

Configure the HBase Thrift Server Role

You must configure the Thrift Server Role to access certain features such as the Hue HBase browser.

About this task

The Thrift Server role is not added by default when you install HBase, but it is required before you can use certain other features such as the Hue HBase browser. To add the Thrift Server role:

Procedure

1. Go to the HBase service.
2. Click the Instances tab.
3. Click the Add Role Instances button.
4. Select the host(s) where you want to add the Thrift Server role (you only need one for Hue) and click Continue. The Thrift Server role should appear in the instances list for the HBase server.
5. Select the Thrift Server role instance.
6. Select Actions for Selected > Start.

Setting user limits for HBase

You must set user limits to avoid opening many files at the same time. Overloading many files at the same time leads to failure and causes error messages.

Because HBase is a database, it opens many files at the same time. The default setting of 1024 for the maximum number of open files on most Unix-like systems is insufficient. Any significant amount of loading will result in failures and cause error message such as `java.io.IOException...(Too many open files)` to be logged in the HBase or HDFS log files. For more information about this issue, see the [Apache HBase Book](#). You may also notice errors such as:

```
2010-04-06 03:04:37,542 INFO org.apache.hadoop.hdfs.DFSClient: Exception in
createBlockOutputStream java.io.EOFException
2010-04-06 03:04:37,542 INFO org.apache.hadoop.hdfs.DFSClient: Abandoning bl
ock blk_-6935524980745310745_1391901
```

Another setting you should configure is the number of processes a user is permitted to start. The default number of processes is typically 1024. Consider raising this value if you experience `OutOfMemoryException` errors.

Configure ulimit for HBase using Cloudera Manager

You can use Cloudera Manager to configure ulimit.

Procedure

1. Go to the HBase service.
2. Click the Configuration tab.
3. Select Scope Master or Scope RegionServer .
4. Locate the Maximum Process File Descriptors property or search for it by typing its name in the Search box.
5. Edit the property value.
6. Restart the role.
7. Restart the service.

Configure ulimit for HBase using the command line

Follow these command-line instructions to configure ulimit using the command-line interface.

About this task

Cloudera recommends increasing the maximum number of file handles to more than 10,000. Increasing the file handles for the user running the HBase process is an operating system configuration, not an HBase configuration. A common mistake is to increase the number of file handles for a particular user when HBase is running as a different user. HBase prints the ulimit it is using on the first line in the logs. Make sure that it is correct.

Procedure

- To change the maximum number of open files for a user, use the `ulimit -n` command while logged in as that user.
- To set the maximum number of processes a user can start, use the `ulimit -u` command. You can also use the `ulimit` command to set many other limits. For more information, see the online documentation for your operating system, or the output of the `man ulimit` command.
- To make the changes persistent, add the command to the user's Bash initialization file (typically `~/.bash_profile` or `~/.bashrc`). Alternatively, you can configure the settings in the Pluggable Authentication Module (PAM) configuration files if your operating system uses PAM and includes the `pam_limits.so` shared library.

Configure ulimit using Pluggable Authentication Modules using the Command Line

Follow these command-line instructions to configure ulimit using Pluggable Authentication Modules on systems that do not use Cloudera Manager.

About this task

If you are using ulimit, you must make the following configuration changes:

Procedure

1. In the `/etc/security/limits.conf` file, add the following lines, adjusting the values as appropriate. This assumes that your HDFS user is called `hdfs` and your HBase user is called `hbase`.

```
hdfs -      nofile 32768
hdfs -      nproc 2048
hbase -     nofile 32768
```

```
hbase -      nproc    2048
```

**Note:**

- Only the root user can edit this file.
 - If this change does not take effect, check other configuration files in the `/etc/security/limits.d/` directory for lines containing the `hdfs` or `hbase` user and the `nofile` value. Such entries may be overriding the entries in `/etc/security/limits.conf`.
2. To apply the changes in `/etc/security/limits.conf` on Ubuntu and Debian systems, add the following line in the `/etc/pam.d/common-session` file:

```
session required pam_limits.so
```

For more information on the `ulimit` command or per-user operating system limits, refer to the documentation for your operating system.

Use dfs.datanode.max.transfer.threads with Hbase

You must configure the `dfs.datanode.max.transfer.threads` with HBase to specify the maximum number of files that a DataNode can serve at any one time.

Procedure

- A Hadoop HDFS DataNode has an upper bound on the number of files that it can serve at any one time. The upper bound is controlled by the `dfs.datanode.max.transfer.threads` property (the property is spelled in the code exactly as shown here). Before loading, make sure you have configured the value for `dfs.datanode.max.transfer.threads` in the `conf/hdfs-site.xml` file (by default found in `/etc/hadoop/conf/hdfs-site.xml`) to at least 4096 as shown here:

```
<property>
  <name>dfs.datanode.max.transfer.threads</name>
  <value>4096</value>
</property>
```

- Restart HDFS after changing the value for `dfs.datanode.max.transfer.threads`. If the value is not set to an appropriate value, strange failures can occur and an error message about exceeding the number of transfer threads will be added to the DataNode logs. Other error messages about missing blocks are also logged, such as:

```
06/12/14 20:10:31 INFO hdfs.DFSCliet: Could not obtain block blk_XXXXXX
XXXXXXXXXXXXXXXXXX_YYYYYYY from any node:
java.io.IOException: No live nodes contain current block. Will get new blo
ck locations from namenode and retry...
```



Note: The property `dfs.datanode.max.transfer.threads` is a HDFS 2 property which replaces the deprecated property `dfs.datanode.max.xcievers`.

Configure encryption in HBase

You must encrypt the HBase root directory to ensure that you have an additional layer of protection in case the HDFS filesystem is compromised. You can encrypt the HBase root directory within HDFS, using HDFS Transparent Encryption.

About this task

HBase stores all of its data under its root directory in HDFS configured in the `hbase.rootdir`. If you use this feature in combination with bulk-loading of HFiles, you must configure `hbase.bulkload.staging.dir` to point to a location within the same encryption zone as the HBase root directory. Otherwise, you may encounter errors such as:

```
org.apache.hadoop.ipc.RemoteException(java.io.IOException): /tmp/output/f/5  
can't be moved into an encryption zone.
```

Procedure

- Enable HDFS encryption using the HDFS encryption wizard.
- Follow the instructions for setting up HDFS Transparent Encryption.
- Validate and verify that HDFS encryption is enabled and working.

For more information see, HDFS Transparent Encryption.

Using hedged reads

You can enable hedged reads if you want to increase the performance of a read operation from an HDFS block that occasionally takes a long time. This feature helps in situations where a read occasionally takes a long time rather than when there is a systemic problem.

If a read from an HDFS block is slow, the HDFS client starts up another parallel, 'hedged' read against a different block replica. The result of whichever read returns first is used, and the outstanding read is cancelled. Hedged reads can be enabled for HBase when the HFiles are stored in HDFS. This feature is disabled by default.

Enable hedged reads for HBase

You need to enable hedged read if a read operation from a HDFS block is slow.

About this task

The Thrift Server role is not added by default when you install HBase, but it is required before you can use certain other features such as the Hue HBase browser. To add the Thrift Server role:

Procedure

1. Go to the HBase service.
2. Click the Configuration tab.
3. Select Scope > HBASE-1 (Service-Wide).
4. Select Category > Performance.
5. Configure the HDFS Hedged Read Threadpool Size and HDFS Hedged Read Delay Threshold properties. The descriptions for each of these properties on the configuration pages provide more information.
6. Enter a Reason for change, and then click Save Changes to commit the changes.

Monitor the performance of hedged reads

You can monitor the performance of hedged reads using the following metrics provided by Hadoop when hedged reads are enabled.

You can monitor the following properties:

- `hedgedReadOps` - the number of hedged reads that have occurred
- `hedgedReadOpsWin` - the number of times the hedged read returned faster than the original read

Understanding HBase garbage collection

You can configure garbage collection to free up the memory that is no longer referenced by Java objects.



Warning: Configuring the JVM garbage collection for HBase is an advanced operation. Incorrect configuration can have major performance implications for your cluster. Test any configuration changes carefully.

Garbage collection (memory cleanup) by the JVM can cause HBase clients to experience excessive latency.

To tune the garbage collection settings, you pass the relevant parameters to the JVM.

Example configuration values are not recommendations and should not be considered as such. This is not the complete list of configuration options related to garbage collection. See the documentation for your JVM for details on these settings.

- `-XX:+UseG1GC`: Use the 'G1' garbage collection algorithm. You can tune G1 garbage collection to provide a consistent pause time, which benefits long-term running Java processes such as HBase, NameNode, Solr, and ZooKeeper. For more information about tuning G1, see the Oracle documentation on tuning garbage collection.
- `-XX:MaxGCPauseMillis=value`: The garbage collection pause time. Set this to the maximum amount of latency your cluster can tolerate while allowing as much garbage collection as possible. `XX:+ParallelRefProcEnabled` Enable or disable parallel reference processing by using a + or - symbol before the parameter name.
- `-XX:-ResizePLAB`: Enable or disable resizing of Promotion Local Allocation Buffers (PLABs) by using a + or - symbol before the parameter name.
- `-XX:ParallelGCThreads=value`: The number of parallel garbage collection threads to run concurrently.
- `-XX:G1NewSizePercent=value`: The percent of the heap to be used for garbage collection. If the value is too low, garbage collection is ineffective. If the value is too high, not enough heap is available for other uses by HBase.

Related Information

[Tuning Java Garbage Collection for HBase](#)

Configure HBase garbage collection

You must configure garbage collection using Cloudera Manager.

About this task

Procedure

1. Go to the HBase service.
2. Click the Configuration tab.
3. Select Scope > RegionServer.
4. Select Category > Advanced.
5. Locate the Java Configuration Options for HBase RegionServer property or search for it by typing its name in the Search box.
6. Add or modify JVM configuration options.
7. Enter a Reason for change, and then click Save Changes to commit the changes.
8. Restart the role.

Disable the BoundedByteBufferPool

HBase uses a BoundedByteBufferPool to avoid fragmenting the heap. You can disable BoundedByteBufferPool using Cloudera Manager.

About this task

The G1 garbage collector reduces the need to avoid fragmenting the heap in some cases. If you use the G1 garbage collector, you can disable the BoundedByteBufferPool in HBase. This can reduce the number of "old generation" items that need to be collected. This configuration is experimental.

Procedure

1. Go to the HBase service.
2. Click the Configuration tab.
3. Select Scope > RegionServer.
4. Select Category > Advanced.
5. Locate the HBase Service Advanced Configuration Snippet (Safety Valve) for hbase-site.xml property, or search for it by typing its name in the Search box.
6. Add the following XML:

```
<property>
  <name>hbase.ipc.server.reservoir.enabled</name>
  <value>false</value>
</property>
```

7. Enter a Reason for change, and then click Save Changes to commit the changes.
8. Restart the service.

Configure the HBase canary

The HBase canary is an optional service that you can configure to check periodically if a RegionServer is alive. The HBase canary is disabled by default.

About this task

This canary is different from the Cloudera Service Monitoring canary and is provided by the HBase service. After enabling the canary, you can configure several different thresholds and intervals relating to it, as well as exclude certain tables from the canary checks. The canary works on Kerberos-enabled clusters if you have the HBase client configured to use Kerberos.

Procedure

1. Go to the HBase service.
2. Click the Configuration tab.
3. Select Scope > HBase or HBase Service-Wide.
4. Select Category > Monitoring.
5. Locate the HBase Canary property or search for it by typing its name in the Search box. Several properties have Canary in the property name.
6. Select the checkbox.
7. Review other HBase Canary properties to configure the specific behavior of the canary. To apply this configuration property to other role groups as needed, edit the value for the appropriate role group. See [Modifying Configuration Properties Using Cloudera Manager](#).
8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Restart the role.
10. Restart the service.

Using HBase blocksize

You must configure the HBase blocksize to set the smallest unit of data HBase can read from the column family's HFiles.

HBase data is stored in one (after a major compaction) or more (possibly before a major compaction) HFiles per column family per region. The blocksize determines:

- The blocksize for a given column family determines the smallest unit of data HBase can read from the column family's HFiles.
- The basic unit of measure cached by a RegionServer in the BlockCache.

The default blocksize is 64 KB. The appropriate blocksize is dependent upon your data and usage patterns. Use the following guidelines to tune the blocksize size, in combination with testing and benchmarking as appropriate.



Warning: The default blocksize is appropriate for a wide range of data usage patterns, and tuning the blocksize is an advanced operation. The wrong configuration can negatively impact performance.

- Consider the average key/value size for the column family when tuning the blocksize. You can find the average key/value size using the HFile utility:

```
$ hbase org.apache.hadoop.hbase.io.hfile.HFile -f /path/to/HFILE -m -v
...
Block index size as per heapsize: 296
reader=hdfs://srv1.example.com:9000/path/to/HFILE, \
compression=none, inMemory=false, \
firstKey=US6683275_20040127/mimetype:/1251853756871/Put, \
lastKey=US6684814_20040203/mimetype:/1251864683374/Put, \
avgKeyLen=37, avgValueLen=8, \
entries=1554, length=84447
...
```

- Consider the pattern of reads to the table or column family. For instance, if it is common to scan for 500 rows on various parts of the table, performance might be increased if the blocksize is large enough to encompass 500-1000 rows, so that often, only one read operation on the HFile is required. If your typical scan size is only 3 rows, returning 500-1000 rows would be overkill.

It is difficult to predict the size of a row before it is written, because the data will be compressed when it is written to the HFile. Perform testing to determine the correct blocksize for your data.

Configure the blocksize for a column family

You can configure the blocksize of a column family at table creation or by disabling and altering an existing table.

About this task

These instructions are valid whether or not you use Cloudera Manager to manage your cluster.

To configure the blocksize for a column family:

Procedure

1. In the HBase shell, type:

```
hbase> create 'test_table#',{NAME => 'test_cf#', BLOCKSIZE => '262144'}
hbase> disable 'test_table'
hbase> alter 'test_table', {NAME => 'test_cf', BLOCKSIZE => '524288'}
```

```
hbase> enable 'test_table'
```

After changing the blocksize, the HFiles will be rewritten during the next major compaction.

2. To trigger a major compaction, issue the following command in HBase Shell:

```
hbase> major_compact 'test_table'
```

Depending on the size of the table, the major compaction can take some time and have a performance impact while it is running.

3. To view the blocksize metrics, see the `block_cache*` entries in the RegionServer metrics section in the HBase web user interface.

Configuring HBase BlockCache

You can configure BlockCache in two different ways in HBase: the default on-heap LruBlockCache and the BucketCache, which is usually off-heap.

If you have less than 20 GB of RAM available for use by HBase, consider tailoring the default on-heap BlockCache implementation (LruBlockCache) for your cluster.

If you have more than 20 GB of RAM available, consider adding off-heap BlockCache (BucketCache).

In the default configuration, HBase uses a single on-heap cache. If you configure the off-heap BucketCache, the on-heap cache is used for Bloom filters and indexes, and the off-heap BucketCache is used to cache data blocks. This is called the Combined Blockcache configuration. The Combined BlockCache allows you to use a larger in-memory cache while reducing the negative impact of garbage collection in the heap, because HBase manages the BucketCache instead of relying on the garbage collector.

Contents of the BlockCache

In HBase, a block is a single unit of I/O. The block cache keeps data blocks resident in the memory after they are read.

To size the BlockCache correctly, you need to understand what HBase places into it.

- Your data: Each time a Get or Scan operation occurs, the result is added to the BlockCache if it was not already cached there. If you use the BucketCache, data blocks are always cached in the BucketCache.
- Row keys: When a value is loaded into the cache, its row key is also cached. This is one reason to make your row keys as small as possible. A larger row key takes up more space in the cache.
- hbase:meta: The hbase:meta catalog table keeps track of which RegionServer is serving which regions. It can consume several megabytes of cache if you have a large number of regions, and has in-memory access priority, which means HBase attempts to keep it in the cache as long as possible.
- Indexes of HFiles: HBase stores its data in HDFS in a format called HFile. These HFiles contain indexes which allow HBase to seek for data within them without needing to open the entire HFile. The size of an index is a factor of the block size, the size of your row keys, and the amount of data you are storing. For big data sets, the size can exceed 1 GB per RegionServer, although the entire index is unlikely to be in the cache at the same time. If you use the BucketCache, indexes are always cached on-heap.
- Bloom filters: If you use Bloom filters, they are stored in the BlockCache. If you use the BucketCache, Bloom filters are always cached on-heap.

The sum of the sizes of these objects is highly dependent on your usage patterns and the characteristics of your data. For this reason, the HBase Web UI and Cloudera Manager each expose several metrics to help you size and tune the BlockCache.

Size the BlockCache

When you use the LruBlockCache, the blocks needed to satisfy each read are cached, old blocks are evicted to make room for new blocks using a Least-Recently-Used algorithm. Set the size of the BlockCache to satisfy your read requirements.

The size cached objects for a given read may be significantly larger than the actual result of the read. For instance, if HBase needs to scan through 20 HFile blocks to return a 100 byte result, and the HFile blocksize is 100 KB, the read will add 20 * 100 KB to the LruBlockCache.

Because the LruBlockCache resides entirely within the Java heap, the amount of which is available to HBase and what percentage of the heap is available to the LruBlockCache strongly impact performance. By default, the amount of HBase heap reserved for LruBlockCache (hfile.block.cache.size) is .40, or 40%. To determine the amount of heap available for the LruBlockCache, use the following formula. The 0.99 factor allows 1% of heap to be available as a "working area" for evicting items from the cache. If you use the BucketCache, the on-heap LruBlockCache only stores indexes and Bloom filters, and data blocks are cached in the off-heap BucketCache.

```
number of RegionServers * heap size * hfile.block.cache.size * 0.99
```

To tune the size of the LruBlockCache, you can add RegionServers or increase the total Java heap on a given RegionServer to increase it, or you can tune hfile.block.cache.size to reduce it. Reducing it will cause cache evictions to happen more often, but will reduce the time it takes to perform a cycle of garbage collection. Increasing the heap will cause garbage collection to take longer but happen less frequently.

Decide to use the BucketCache

The BucketCache manages areas of memory called *buckets* for holding the cached blocks. You can use BucketCache if any of the conditions listed in here are true.

- If the result of a Get or Scan typically fits completely in the heap, the default configuration, which uses the on-heap LruBlockCache, is the best choice, as the L2 cache will not provide much benefit. If the eviction rate is low, garbage collection can be 50% less than that of the BucketCache, and throughput can be at least 20% higher.
- Otherwise, if your cache is experiencing a consistently high eviction rate, use the BucketCache, which causes 30-50% of the garbage collection of LruBlockCache when the eviction rate is high.
- BucketCache using file mode on solid-state disks has a better garbage-collection profile but lower throughput than BucketCache using off-heap memory.

About the Off-heap BucketCache

If the BucketCache is enabled, it stores data blocks, leaving the on-heap cache free for storing indexes and Bloom filters.

The physical location of the BucketCache storage can be either in memory (off-heap) or in a file stored in a fast disk.

- Off-heap: This is the default configuration.
- File-based: You can use the file-based storage mode to store the BucketCache on an SSD or FusionIO device,

You can configure a column family to keep its data blocks in the L1 cache instead of the BucketCache, using the HColumnDescriptor.cacheDataInL1(true) method or by using the following syntax in HBase Shell:

```
hbase> alter 'myTable', CONFIGURATION => {CACHE_DATA_IN_L1 => 'true'}}
```

Off-heap BucketCache

If the BucketCache is enabled, it stores data blocks, leaving the on-heap cache free for storing indexes and Bloom filters. The physical location of the BucketCache storage can be either in memory (off-heap) or in a file stored in a fast disk.

- Off-heap: This is the default configuration.
- File-based: You can use the file-based storage mode to store the BucketCache on an SSD or FusionIO device,

You can configure a column family to keep its data blocks in the L1 cache instead of the BucketCache, using the `HColumnDescriptor.cacheDataInL1(true)` method or by using the following syntax in HBase Shell:


```
hbase> alter 'myTable', CONFIGURATION => {CACHE_DATA_IN_L1 => 'true'}}
```

This table summarizes the important configuration properties for the BucketCache. To configure the BucketCache. The table is followed by three diagrams that show the impacts of different blockcache settings.

Table 1: BucketCache Configuration Properties


Property	Default	Description
<code>hbase.bucketcache.combinedcache.enabled</code>	true	When BucketCache is enabled, use it as a L2 cache for LruBlockCache. If set to true, indexes and Bloom filters are kept in the LruBlockCache and the data blocks are kept in the BucketCache.
<code>hbase.bucketcache.ioengine</code>	none (BucketCache is disabled by default)	Where to store the contents of the BucketCache. Its value can be <code>offheap</code> , <code>file:PATH</code> , <code>mmap:PATH</code> or <code>pmem:PATH</code> where <code>PATH</code> is the path to the file that host the file-based cache.
<code>hfile.block.cache.size</code>	0.4	A float between 0.0 and 1.0. This factor multiplied by the Java heap size is the size of the L1 cache. In other words, the percentage of the Java heap to use for the L1 cache.
<code>hbase.bucketcache.size</code>	not set	When using BucketCache, this is a float that represents one of two different values, depending on whether it is a floating-point decimal less than 1.0 or an integer greater than 1.0. • If less than 1.0, it represents a percentage of total heap memory size to give to the cache. • If greater than 1.0, it represents the capacity of the cache in megabytes
<code>hbase.bucketcache.bucket.sizes</code>	4, 8, 16, 32, 40, 48, 56, 64, 96, 128, 192, 256, 384, 512 KB	A comma-separated list of sizes for buckets for the BucketCache if you prefer to use multiple sizes. The sizes should be multiples of the default blocksize, ordered from smallest to largest. The sizes you use will depend on your data patterns. This parameter is experimental.
<code>-XX:MaxDirectMemorySize</code>	<code>MaxDirectMemorySize = BucketCache + 1</code>	A JVM option to configure the maximum amount of direct memory available for the JVM. It is automatically calculated and configured based on the following formula: <code>MaxDirectMemorySize = BucketCache size + 1 GB</code> for other features using direct memory, such as DFSClient. For example, if the BucketCache size is 8 GB, it will be <code>-XX:MaxDirectMemorySize=9G</code> .

Figure 1: Default LRU Cache, L1 only block cache `hbase.bucketcache.ioengine=NULL`



1. 20% minimum reserved for operations and rpc call queues
2. `hbase.regionserver.global.memstore.size`: default is 0.4, which means 40%
3. `hbase.regionserver.global.memstore.size` + `hfile.block.cache.size` #0.80, which means 80%
4. `hfile.block.cache.size`: default is 0.4, which means 40%
5. slack reserved for HDFS SCR/NIO: number of open HFiles * `hbase.dfs.client.read.shortcircuit.buffer.size`, where `hbase.dfs.client.read.shortcircuit.buffer.size` is set to 128k.

Figure 2: Default LRU Cache, L1 only block cache `hbase.bucketcache.ioengine=offheap`



1. 20% minimum reserved for operations and rpc call queues
2. hbase.regionserver.global.memstore.size: default is 0.4, which means 40%
3. hbase.regionserver.global.memstore.size + hfile.block.cache.size #0.80, which means 80%
4. hfile.block.cache.size: default is 0.4 which means 40%
5. slack reserved for HDFS SCR/NIO: number of open HFiles * hbase.dfs.client.read.shortcircuit.buffer.size, where hbase.dfs.client.read.shortcircuit.buffer.size is set to 128k.
6. hbase.bucketcache.size: default is 0.0

If hbase.bucketcache.size is float <1, it represents the percentage of total heap size.

If hbase.bucketcache.size is #1, it represents the absolute value in MB. It must be < HBASE_OFFHEAPSIZE

BucketCache IO engine

Use the hbase.bucketcache.ioengine parameter to define where to store the content of the BucketCache. Its value can be offheap, file:PATH, mmap:PATH, pmem:PATH, or it can be empty. By default it is empty which means that BucketCache is disabled.

You can set the following values in the `hbase.bucketcache.ioengine` parameter to define where to store the BucketCache:

- `offheap`: When `hbase.bucketcache.ioengine` is set to `offheap` the content of the BucketCache is stored off-heap.
- `file:PATH`: When `hbase.bucketcache.ioengine` is set to `file:PATH`, the BucketCache uses file caching.
- `mmap:PATH`: When `hbase.bucketcache.ioengine` is set to `mmap:PATH`, the content of the BucketCache is stored and accessed through memory mapping to a file under the specified path.
- `pmem:PATH`: When `hbase.bucketcache.ioengine` is set to `pmem:PATH`, BucketCache uses direct memory access to and from a file on the specified path. The specified path must be under a volume that is mounted on a persistent memory device that supports direct access to its own address space.

The advantage of the `pmem` engine over the `mmap` engine is that it supports large cache size. That is because `pmem` allows for reads straight from the device address, which means in this mode no copy is created on DRAM. Therefore, swapping due to DRAM free memory exhaustion is not an issue when large cache size is specified. With devices currently available, the bucket cache size can be set to the order of hundreds of GBs or even a few TBs.

When bucket cache size is set to larger than 256GB, the OS limit must be increased, which can be configured by the `max_map_count` property. Make sure you have an extra 10% for other processes on the host that require the use of memory mapping. This additional overhead depends on the load of processes running on the RS hosts. To calculate the OS limit divide the block cache size in GB by 4 MB and then multiply it by 1.1: $(\text{block cache size in GB} / 4 \text{ MB}) * 1.1$.

Configure BucketCache IO engine

You must configure the BucketCache IO engine using Cloudera Manager.

Set the value `offheap` and `file:PATH`

1. In Cloudera Manager select the HBase service and go to Configuration.
2. Search for BucketCache IOEngine and set it to the required value.

Set the value `mmap:PATH` and `pmem:PATH`



Important: These values can only be set using safety valves.

1. In Cloudera Manager select the HBase service and go to Configuration.
2. Search for RegionServer Advanced Configuration Snippet (Safety Valve) for `hbase-site.xml`.
3. Click the plus icon.
4. Set the required value:
 - Name: Add `hbase.bucketcache.ioengine`.
 - Value: Add either `mmap:PATH` or `pmem:PATH`.

Configure the off-heap BucketCache using Cloudera Manager

You can configure the off-heap BucketCache engine using Cloudera Manager.

Procedure

1. Go to the HBase service.
2. Click the Configuration tab.
3. Select the RegionServer scope and do the following:
 - a) Ensure that Enable Combined BucketCache is selected.
 - b) Set BucketCache IOEngine to `offheap`.
 - c) Update the value of BucketCache Size according to the required BucketCache size.

4. In the Region Server Environment Advanced Configuration Snippet (Safety Valve), edit the HBASE_REGIONSERVER_OPTS parameter: Add the JVM option \$HBASE_REGIONSERVER_OPTS -XX:MaxDirectMemorySize=<size>G, replacing <size> with a value not smaller than the aggregated heap size expressed as a number of gigabytes + the off-heap BucketCache, expressed as a number of gigabytes + around 1GB used for HDFS short circuit read. For example, if the off-heap BucketCache is 16GB and the heap size is 15GB, the total value of MaxDirectMemorySize could be 32: -XX:MaxDirectMemorySize=32G

```
HBASE_REGIONSERVER_OPTS="$HBASE_REGIONSERVER_OPTS -XX:MaxDirectMemorySize=<size>G"
```

5. Optionally, when combined BucketCache is in use, you can decrease the heap size ratio allocated to the L1 BlockCache, and increase the Memstore size. The on-heap BlockCache only stores indexes and Bloom filters, the actual data resides in the off-heap BucketCache. A larger Memstore is able to accommodate more write request before flushing them to disks.
 - Decrease HFile Block Cache Size to 0.3 or 0.2.
 - Increase Maximum Size of All Memstores in RegionServer to 0.5 or 0.6 respectively.
6. Enter a Reason for change, and then click Save Changes to commit the changes.
7. Restart or rolling restart your RegionServers for the changes to take effect.

Configure the off-heap BucketCache using the command line

You can configure the off-heap BucketCache engine from the command-line interface.

About this task

Procedure

1. Configure the MaxDirectMemorySize option for the RegionServers JVMs. Add the JVM option \$HBASE_REGIONSERVER_OPTS -XX:MaxDirectMemorySize=<size>G, replacing <size> with a value not smaller than the aggregated heap size expressed as a number of gigabytes + the off-heap BucketCache, expressed as a number of gigabytes + around 1GB used for HDFS short circuit read. For example, if the off-heap BucketCache is 16GB and the heap size is 15GB, the total value of MaxDirectMemorySize could be 32: -XX:MaxDirectMemorySize=32G. This can be done adding the following line in hbase-env.sh:

```
HBASE_REGIONSERVER_OPTS="$HBASE_REGIONSERVER_OPTS -XX:MaxDirectMemorySize=<size>G"
```

2. Next, in the hbase-site.xml files on the RegionServers, configure the properties in BucketCache configuration properties as appropriate, using the example below as a model.

```
<property>
  <name>hbase.bucketcache.combinedcache.enabled</name>
  <value>true</value>
</property>
<property>
  <name>hbase.bucketcache.ioengine</name>
  <value>offheap</value>
</property>
<property>
  <name>hbase.bucketcache.size</name>
  <value>8388608</value>
</property>
<property>
  <name>hfile.block.cache.size</name>
  <value>0.2</value>
</property>
<property>
  <name>hbase.regionserver.global.memstore.size</name>
  <value>0.6</value>
```

```
</property>
```

Optionally, when combined BucketCache is in use, you can decrease the heap size ratio allocated to the L1 BlockCache, and increase the Memstore size as it is done in the above example. The on-heap BlockCache only stores indexes and Bloom filters, the actual data resides in the off-heap BucketCache. A larger Memstore is able to accommodate more write request before flushing them to disks.

- Decrease `hfile.block.cache.size` to 0.3 or 0.2.
- Increase `hbase.regionserver.global.memstore.size` to 0.5 or 0.6 respectively.

3. Restart each RegionServer for the changes to take effect.

Cache eviction priorities

You must decide on the cache eviction priorities to allow for scan-resistance and in-memory column families.

Both the on-heap cache and the off-heap BucketCache use the same cache priority mechanism to decide which cache objects to evict to make room for new objects. Three levels of block priority allow for scan-resistance and in-memory column families. Objects evicted from the cache are subject to garbage collection.

- Single access priority: The first time a block is loaded from HDFS, that block is given single access priority, which means that it will be part of the first group to be considered during evictions. Scanned blocks are more likely to be evicted than blocks that are used more frequently.
- Multi access priority: If a block in the single access priority group is accessed again, that block is assigned multi access priority, which moves it to the second group considered during evictions, and is therefore less likely to be evicted.
- In-memory access priority: If the block belongs to a column family which is configured with the in-memory configuration option, its priority is changed to in memory access priority, regardless of its access pattern. This group is the last group considered during evictions, but is not guaranteed not to be evicted. Catalog tables are configured with in-memory access priority.

To configure a column family for in-memory access, use the following syntax in HBase Shell:

```
hbase> alter 'myTable', 'myCF', CONFIGURATION => {IN_MEMORY => 'true'}
```

To use the Java API to configure a column family for in-memory access, use the `HColumnDescriptor.setInMemory(true)` method.

Bypass the BlockCache

You can bypass the BlockCache if the data needed for a specific but atypical operation does not all fit in memory.

For an atypical operation does not all fit in memory, using the BlockCache can be counter-productive because data that you are still using may be evicted, or even if other data is not evicted, excess garbage collection can adversely effect performance. For this type of operation, you may decide to bypass the BlockCache. To bypass the BlockCache for a given Scan or Get, use the `setCacheBlocks(false)` method.

In addition, you can prevent a specific column family's contents from being cached, by setting its BLOCKCACHE configuration to false. Use the following syntax in HBase Shell:

```
hbase> alter 'myTable', CONFIGURATION => {NAME => 'myCF', BLOCKCACHE => 'false'}
```

Monitor the BlockCache

Cloudera Manager provides metrics to monitor the performance of the BlockCache, to assist you in tuning your configuration.

You can view further detail and graphs using the RegionServer UI. To access the RegionServer UI in Cloudera Manager, go to the Cloudera Manager page for the host, click the RegionServer process, and click HBase RegionServer Web UI.

If you do not use Cloudera Manager, access the BlockCache reports at http://regionServer_host:22102/rs-status#memoryStats, replacing regionServer_host with the hostname or IP address of your RegionServer.

Using quota management

Two types of HBase quotas are well established: throttle quota and number-of tables-quota. These two quotas can regulate users and tables.

In a multitenant HBase environment, ensuring that each tenant can use only its allotted portion of the system is key in meeting SLAs.

Table 2: Quota Support Matrix

Quota Type	Resource Type	Purpose	Namespace applicable?	Table applicable?	User applicable?
Throttle	Network	Limit overall network throughput and number of RPC requests	Yes	Yes	Yes
New space	Storage	Limit amount of storage used for table or namespaces	Yes	Yes	No
Number of tables	Metadata	Limit number of tables for each namespace or user	Yes	No	Yes
Number of regions	Metadata	Limit number of regions for each namespace	Yes	No	No

Configuring quotas

HBase quotas are disabled by default. To enable quotas, the relevant `hbase-site.xml` property must be set to true and the limit of each quota specified on the command line.

Before you begin

hbase superuser privileges

Procedure

1. Set the `hbase.quota.enabled` property in the `hbase-site.xml` file to true.
2. Enter the command to set the limit of the quota, type of quota, and to which entity to apply the quota. The command and its syntax are:

```
$hbase_shell> set_quota TYPE =>
    quota_type,
    arguments
```

General Quota Syntax

The general quota syntax are `THROTTLE_TYPE`, Request sizes and space limit, Number of requests, Time limits and Number of tables or regions.

THROTTLE_TYPE

Can be expressed as `READ-only`, `WRITE-only`, or the default type (both `READ` and `WRITE` permissions)

Timeframes

Can be expressed in the following units of time:

- `sec` (second)
- `min` (minute)
- `hour`
- `day`

Request sizes and space limit

Can be expressed in the following units:

- `B`: bytes
- `K`: kilobytes
- `M`: megabytes
- `G`: gigabytes
- `P`: petabytes

When no size units is included, the default value is bytes.

Number of requests

Expressed as integer followed by the string request

Time limits

Expressed as requests per unit-of-time or size per unit-of-time

Examples: `10req/day` or `100P/hour`

Number of tables or regions

Expressed as integers

Space quotas

Space quotas, also known as filesystem space quotas, limit the amount of stored data. It can be applied at a table or namespace level where table-level quotas take priority over namespace-level quotas.

Space quotas are special in that they can trigger different policies when storage goes above thresholds. The following list describes the policies, and they are listed in order of least strict to most strict:

NO_INSERTS

Prohibits new data from being ingested (for example, data from `Put`, `Increment`, and `Append` operations are not ingested).

NO_WRITES

Performs the same function as `NO_INSERTS` but `Delete` operations are also prohibited.

NO_WRITES_COMPACTIONS

Performs the same function as NO_INSERTS but compactions are also prohibited.

DISABLE

Disables tables.

Examples of Adding Space Quotas

Add quota with the condition that Insert operations are rejected when table t1 reaches 1 GB of data:

```
hbase> set_quota TYPE => SPACE, TABLE => 't1', LIMIT => '1G', POLICY => NO_INSERTS
```

Add quota with the condition that table t2 is disabled when 50 GB of data is exceeded:

```
hbase> set_quota TYPE => SPACE, TABLE => 't2', LIMIT => '50G', POLICY => DISABLE
```

Add quota with the condition that Insert and Delete operations cannot be applied to namespace ns1 when it reaches 50 terabytes of data:

```
hbase> set_quota TYPE => SPACE, NAMESPACE => 'ns1', LIMIT => '50T', POLICY => NO_WRITES
```

Listing Space Quotas

See "Examples of Listing Throttle Quotas Commands" above for the supported syntax.

Examples of Updating and Deleting Space Quotas

A quota can be removed by setting LIMIT to NONE.

```
hbase> set_quota TYPE => SPACE, TABLE => 't1', LIMIT => NONE
```

```
hbase> set_quota TYPE => SPACE, NAMESPACE => 'ns1', LIMIT => NONE
```

Quota enforcement

When a quota limit is exceeded, the Master server instructs RegionServers to enable an enforcement policy for the namespace or table that violated the quota.

It is important to note the storage quota is not reported in real-time. There is a window when threshold is reached on RegionServers but the threshold accounted for on the Master server is not updated.



Note:

Set a storage limit lower than the amount of available disk space to provide extra buffer.

Quota violation policies

If quotas are set for the amount of space each HBase tenant can fill on HDFS, then a quota violation policy should be planned and implemented on the system.

When a quota violation policy is enabled, the table owner should not be allowed to remove the policy. The expectation is that the Master automatically removes the policy. However, the HBase superuser should still have permission.

Automatic removal of the quota violation policy after the violation is resolved can be accomplished via the same mechanisms that it was originally enforced. But the system should not immediately disable the violation policy when the violation is resolved.

The following describes quota violation policies that you might consider.

Disabling Tables

This is the “brute-force” policy, disabling any tables that violated the quota. This policy removes the risk that tables over quota affect your system. For most users, this is likely not a good choice as most sites want READ operations to still succeed.

One hypothetical situation when a disabling tables policy might be advisable is when there are multiple active clusters hosting the same data and, because of a quota violation, it is discovered that one copy of the data does not have all of the data it should have. By disabling tables, you can prevent further discrepancies until the administrator can correct the problem.

Rejecting All WRITE Operations, Bulk Imports, and Compactions

This policy rejects all WRITES and bulk imports to the region which the quota applies. Compactions for this region are also disabled to prevent the system from using more space because of the temporary space demand of a compaction. The only resolution in this case is administrator intervention to increase the quota that is being exceeded.

Rejecting All WRITE Operations and Bulk Imports

This is the same as the previous policy, except that compactions are still allowed. This allows users to set or alter a TTL on table and then perform a compaction to reduce the total used space. Inherently, using this violation policy means that you let used space to slightly rise before it is ultimately reduced.

Allowing DELETE Operations But Rejecting WRITE Operations and Bulk Imports

This is another variation of the two previously listed policies. This policy allows users to run processes to delete data in the system. Like the previous policy, using this violation policy means that you let used space slightly rises before it is ultimately reduced. In this case, the deletions are propagated to disk and a compaction actually removes data previously stored on disk. TTL configuration and compactions can also be used to remove data.

Using HBase scanner heartbeat

A scanner heartbeat check enforces a time limit on the execution of scan RPC requests. This helps prevent scans from taking too long and causing a timeout at the client.

When the server receives a scan RPC request, a time limit is calculated to be half of the smaller of two values: `hbase.client.scanner.timeout.period` and `hbase.rpc.timeout` (which both default to 60000 milliseconds, or one minute). When the time limit is reached, the server returns the results it has accumulated up to that point. This result set may be empty. If your usage pattern includes that scans will take longer than a minute, you can increase these values.

To make sure the timeout period is not too short, you can configure `hbase.cells.scanned.per.heartbeat.check` to a minimum number of cells that must be scanned before a timeout check occurs. The default value is 10000. A smaller value causes timeout checks to occur more often.

Configure the scanner heartbeat using Cloudera Manager

You can configure the HBase scanner heartbeat using Cloudera Manager.

Procedure

1. Go to the HBase service.
2. Click the Configuration tab.
3. Select HBase or HBase Service-Wide.

4. Locate the RPC Timeout property by typing its name in the Search box, and edit the property.
5. Locate the HBase RegionServer Lease Period property by typing its name in the Search box, and edit the property.
6. Enter a Reason for change, and then click Save Changes to commit the changes.
7. Restart the role.
8. Restart the service.

Limiting the speed of compactions

You can limit the speed at which HBase compactions run, by configuring `hbase.regionserver.throughput.controller` and its related settings.

The default controller is `org.apache.hadoop.hbase.regionserver.throttle.PressureAwareCompactionThroughputController`, which uses the following algorithm:

- If compaction pressure is greater than 1.0, there is no speed limitation.
- In off-peak hours, use a fixed throughput limitation, configured using `hbase.hstore.compaction.throughput.offpeak`, `hbase.offpeak.start.hour`, and `hbase.offpeak.end.hour`.
- In normal hours, the max throughput is tuned between `hbase.hstore.compaction.throughput.higher.bound` and `hbase.hstore.compaction.throughput.lower.bound` (which default to 20 MB/sec and 10 MB/sec respectively), using the following formula, where `compactionPressure` is between 0.0 and 1.0. The `compactionPressure` refers to the number of store files that require compaction.

$$\text{lower} + (\text{higher} - \text{lower}) * \text{compactionPressure}$$

To disable compaction speed limits, set `hbase.regionserver.throughput.controller` to `org.apache.hadoop.hbase.regionserver.throttle.NoLimitCompactionThroughputController`.

Configure the compaction speed using Cloudera Manager

Procedure

1. Go to the HBase service.
2. Click the Configuration tab.
3. Select HBase or HBase Service-Wide.
4. Search for HBase Service Advanced Configuration Snippet (Safety Valve) for `hbase-site.xml`. Paste the relevant properties from the following example into the field and modify the values as needed:

```
<property>
  <name>hbase.hstore.compaction.throughput.higher.bound</name>
  <value>20971520</value>
  <description>The default is 20 MB/sec</description>
</property>
<property>
  <name>hbase.hstore.compaction.throughput.lower.bound</name>
  <value>10485760</value>
  <description>The default is 10 MB/sec</description>
</property>
<property>
  <name>hbase.hstore.compaction.throughput.offpeak</name>
  <value>9223372036854775807</value>
  <description>The default is Long.MAX_VALUE, which effectively means no limitation</description>
</property>
```

```
<property>
  <name>hbase.offpeak.start.hour</name>
  <value>20</value>
  <description>When to begin using off-peak compaction settings, expressed
as an integer between 0 and 23.</description>
</property>
<property>
  <name>hbase.offpeak.end.hour</name>
  <value>6</value>
  <description>When to stop using off-peak compaction settings, expressed
as an integer between 0 and 23.</description>
</property>
<property>
  <name>hbase.hstore.compaction.throughput.tune.period</name>
  <value>60000</value>
</property>
```

5. Enter a Reason for change, and then click Save Changes to commit the changes.
6. Restart the service.

Enable HBase indexing

You can enable HBase indexing using Cloudera Manager.

Before you begin

HBase indexing is dependent on the Key-Value Store Indexer service.

Procedure

1. Go to the HBase service.
2. Select ScopeHBASE-1 (Service Wide)
3. Select CategoryBackup.
4. Select the Enable Replication and Enable Indexing properties.
5. Click Save Changes.

Using HBase coprocessors

You can configure HBase coprocessors to run your own custom code. The HBase coprocessor framework provides a way to extend HBase with custom functionality. Coprocessors provide a way to run server-level code against locally-stored data.

Coprocessors are not designed to be used by end users of HBase, but by HBase developers who need to add specialized functionality to HBase. One example of the use of coprocessors is pluggable compaction and scan policies.

Related Information

[Apache HBase blog: Coprocessor Introduction](#)

Add a custom coprocessor

You can add a custom coprocessor to to extend HBase with custom functionality using Cloudera Manager.

About this task

To configure these properties in Cloudera Manager:

Procedure

1. Select the HBase service.
2. Click the Configuration tab.
3. Select Scope All .
4. Select Category All .
5. Type HBase Coprocessor in the Search box.
6. You can configure the values of the following properties:
 - HBase Coprocessor Abort on Error (Service-Wide)
 - HBase Coprocessor Master Classes (Master Default Group)
 - HBase Coprocessor Region Classes (RegionServer Default Group)
7. Enter a Reason for change, and then click Save Changes to commit the changes.

Disable loading of coprocessors

You can disable loading of coprocessors using Cloudera Manager.

About this task

Cloudera recommends against disabling loading of system coprocessors, because HBase security functionality is implemented using system coprocessors. However, disabling loading of user coprocessors may be appropriate.

Procedure

1. Select the HBase service.
2. Click the Configuration tab.
3. Search for HBase Service Advanced Configuration Snippet (Safety Valve) for hbase-site.xml.
4. To disable loading of all coprocessors, add a new property with the name `hbase.coprocessor.enabled` and set its value to false. Cloudera does not recommend this setting.
5. To disable loading of user coprocessors, add a new property with the name `hbase.coprocessor.user.enabled` and set its value to false.
6. Enter a Reason for change, and then click Save Changes to commit the changes.

Configuring HBase MultiWAL

You can configure multiple write-ahead logs (MultiWAL) for HBase. If you do not configure MultiWAL, each region on a RegionServer writes to the same WAL.

A busy RegionServer might host several regions, and each write to the WAL is serial because HDFS only supports sequentially written files. This causes the WAL to negatively impact performance.

MultiWAL allows a RegionServer to write multiple WAL streams in parallel by using multiple pipelines in the underlying HDFS instance, which increases total throughput during writes.



Note: In the current implementation of MultiWAL, incoming edits are partitioned by Region. Therefore, throughput to a single Region is not increased.

To configure MultiWAL for a RegionServer, set the value of the property `hbase.wal.provider` to `multiwal` and restart the RegionServer. To disable MultiWAL for a RegionServer, unset the property and restart the RegionServer.

RegionServers using the original WAL implementation and those using the MultiWAL implementation can each handle recovery of either set of WALs, so a zero-downtime configuration update is possible through a rolling restart.

Configuring MultiWAL support using Cloudera Manager

You can configure MultiWAL using Cloudera Manager.

Procedure

1. Go to the HBase service.
2. Click the Configuration tab.
3. Select Scope RegionServer .
4. Select Category Main .
5. Set WAL Provider to MultiWAL.
6. Set the Per-RegionServer Number of WAL Pipelines to a value greater than 1.
7. Enter a Reason for change, and then click Save Changes to commit the changes.
8. Restart the RegionServer roles.

Configuring the storage policy for the Write-Ahead Log (WAL)

You can configure the preferred HDFS storage policy for HBase's write-ahead log (WAL) replicas. This feature allows you to tune HBase's use of SSDs to your available resources and the demands of your workload.

These instructions assume that you have followed the instructions to configure storage directories for DataNodes, and that your cluster has SSD storage available to HBase. If HDFS is not configured to use SSDs, these configuration changes will have no effect on HBase. The following policies are available:

- NONE: no preference about where the replicas are written.
- ONE_SSD: place one replica on SSD storage and the remaining replicas in default storage. This allows you to derive some benefit from SSD storage even if it is a scarce resource in your cluster.



Warning: ONE_SSD mode has not been thoroughly tested with HBase and is not recommended.

- ALL_SSD: place all replicas on SSD storage.

Configure the storage policy for WALs using Cloudera Manager

You can configure the preferred HDFS storage policy for HBase's write-ahead log (WAL) replicas using Cloudera Manager.

Procedure

1. Go to the HBase service.
2. Click the Configuration tab.
3. Search for the property WAL HSM Storage Policy.
4. Select your desired storage policy.
5. Save your changes. Restart all HBase roles.

Results

Changes will take effect after the next major compaction.

Configure the storage policy for WALs using the Command Line

You can configure the preferred HDFS storage policy for HBase's write-ahead log (WAL) replicas using the command line.

About this task



Important: Follow these command-line instructions on systems that do not use Cloudera Manager.

Procedure

- Paste the following XML into hbase-site.xml. Uncomment the <value> line that corresponds to your desired storage policy.

```
<property>
  <name>hbase.wal.storage.policy</name>
  <value>NONE</value>
  <!--<value>ONE_SSD</value>-->
  <!--<value>ALL_SSD</value>-->
</property>
```



Warning: ONE_SSD mode has not been thoroughly tested with HBase and is not recommended.

- Restart HBase. Changes will take effect for a given region during its next major compaction.

Using RegionServer grouping

You can use RegionServer Grouping (rsgroup) to impose strict isolation between RegionServers by partitioning RegionServers into distinct groups. You can use HBase Shell commands to define and manage RegionServer Grouping.

You must first create an rsgroup before you can add RegionServers to it. Once you have created an rsgroup, you can move your HBase tables into this rsgroup so that only the RegionServers in the same rsgroup can host the regions of the table.



Note: RegionServers and tables can only belong to one rsgroup at a time. By default, all the tables and RegionServers belong to the default rsgroup.

A custom balancer implementation tracks assignments per rsgroup and moves regions to the relevant RegionServers in that rsgroup. The rsgroup information is stored in a regular HBase table, and a ZooKeeper-based read-only cache is used at cluster bootstrap time.

Enable RegionServer grouping using Cloudera Manager

You must use Cloudera Manager to enable RegionServer Grouping before you can define and manage rsgroups.

Procedure

1. Go to the HBase service.
2. Click the Configuration tab.

3. Select Scope Master .
4. Locate the HBase Coprocessor Master Classes property or search for it by typing its name in the Search box.
5. Add the following property value: `org.apache.hadoop.hbase.rsgroup.RSGroupAdminEndpoint`.
6. Locate the Master Advanced Configuration Snippet (Safety Valve) for `hbase-site.xml` property or search for it by typing its name in the Search box.
7. Click View as XML and add the following property:

```
<property>
  <name>hbase.master.loadbalancer.class</name>
  <value>org.apache.hadoop.hbase.rsgroup.RSGroupBasedLoadBalancer</value>
</property>
```

8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Restart the role.
10. Restart the service.

Configure RegionServer grouping

When you add a new rsgroup, you are creating an rsgroup other than the default group.

About this task

To configure a rsgroup, in the HBase shell:

Procedure

1. Add an rsgroup: `$hbase> add_rsgroup 'MYGROUP'`.
2. Add RegionServers and tables to this rsgroup: `$hbase> move_servers_tables_rsgroup 'MYGROUP', ['SERVER1:PORT','SERVER2:PORT'],['TABLE1','TABLE2']`.
3. Run the `balance_rsgroup` command if the tables are slow to migrate to the group's dedicated server.



Note: The term *rsgroup* refers to servers in a cluster with only the hostname and port. It does not make use of the HBase `ServerName` type identifying RegionServers (hostname + port + start time) to distinguish RegionServer instances.

Monitor RegionServer grouping

You can monitor the status of the commands using the Tables tab on the HBase Master UI home page.

You can monitor the status of the commands using the Tables tab on the HBase Master UI home page. If you click on a table name, you can see the RegionServers that are deployed.

You must manually align the RegionServers referenced in rsgroups with the actual state of nodes in the cluster that is active and running.

Remove a RegionServer from RegionServer grouping

You can remove a RegionServer by moving it to the default rsgroup. Edits made using shell commands to all rsgroups, except the default rsgroup, are persisted to the system `hbase:rsgroup` table. If an rsgroup references a decommissioned RegionServer, then the rsgroup should be updated to undo the reference.

Procedure

1. Move the RegionServer to the default rsgroup using the command: `$hbase> move_servers_rsgroup 'default', ['SERVER1:PORT']`.

2. Check the list of RegionServers in your rsgroup to ensure that the RegionServer is successfully removed using the command: `$hbase> get_rsgroup 'MYGROUP'`

The default rsgroup's RegionServer list mirrors the current state of the cluster. If you shut down a RegionServer that was part of the default rsgroup, and then run the `get_rsgroup 'default'` command to list its content in the shell, the server is no longer listed. If you move the offline server from the non-default rsgroup to default, it will not show in the default list; the server will just be removed from the list.

Enable ACL for RegionServer grouping

You need to be a Global Admin to manage rsgroups if authorization is enabled.

About this task

To enable ACL:

Procedure

1. Add the following to the `hbase-site.xml` file:

```
<property>
  <name>hbase.security.authorization</name>
  <value>true</value>
</property>
```

2. Restart your HBase Master server.

Best practices when using RegionServer grouping

You must keep in mind the following best practices when using rsgroups.

- Isolate system tables: You can either have a system rsgroup where all the system tables are present or just leave the system tables in default rsgroup and have all user-space tables in non-default rsgroups.
- Handle dead nodes: You can have a special rsgroup of dead or questionable nodes to help you keep them without running until the nodes are repaired. Be careful when replacing dead nodes in an rsgroup, and ensure there are enough live nodes before you start moving out the dead nodes. You can move the good live nodes first before moving out the dead nodes.

If you have configured a table to be in a rsgroup, but all the RegionServers in that rsgroup die, the tables become unavailable and you can no longer access those tables.

Disable RegionServer grouping

When you no longer require rsgroups, you can disable it for your cluster.

About this task

Removing RegionServer Grouping for a cluster on which it was enabled involves more steps in addition to removing the relevant properties from `hbase-site.xml`. You must ensure that you clean the RegionServer grouping-related metadata so that if the feature is re-enabled in the future, the old metadata will not affect the functioning of the cluster.

To disable RegionServer Grouping:

Procedure

1. Move all the tables in non-default rsgroups to default RegionServer group.

```
#Reassigning table t1 from the non-default group - hbase shell
hbase> move_tables_rsgroup 'default', ['T1']
```

2. Move all RegionServers in non-default rsgroups to default regionserver group.

```
#Reassigning all the servers in the non-default rsgroup to default - hbase
shell
hbase> move_servers_rsgroup 'default',
['REGIONSERVER1:PORT', 'REGIONSERVER2:PORT', 'REGIONSERVER3:PORT']
```

3. Remove all non-default rsgroups. default rsgroup created implicitly does not have to be removed.

```
#removing non-default rsgroup - hbase shell
hbase> remove_rsgroup 'MYGROUP'
```

4. Remove the changes made in hbase-site.xml and restart the cluster.

5. Drop the table hbase:rsgroup from HBase.

```
#Through hbase shell drop table hbase:rsgroup
hbase> disable 'hbase:rsgroup'
0 row(s) in 2.6270 seconds
hbase> drop 'hbase:rsgroup'
0 row(s) in 1.2730 seconds
```

6. Remove the znode rsgroup from the cluster ZooKeeper using zkCli.sh.

```
#From ZK remove the node /hbase/rsgroup through zkCli.sh
rmr /hbase/rsgroup
```

Optimizing HBase I/O

You can optimize HBase I/O using several ways. Two HBase key concepts that helps you in the process are BlockCache and MemStore tuning.

The information in this section is oriented toward basic BlockCache and MemStore tuning. As such, it describes only a subset of cache configuration options. HDP supports additional BlockCache and MemStore properties, as well as other configurable performance optimizations such as remote procedure calls (RPCs), HFile block size settings, and HFile compaction. For a complete list of configurable properties, see the hbase-default.xml source file in GitHub.

HBase I/O components

The concepts related to HBase file operations and memory (RAM) caching are HFile, Block, BlockCache, MemStore and Write Ahead Log (WAL).

HBase Component	Description
HFile	An HFile contains table data, indexes over that data, and metadata about the data.
Block	An HBase block is the smallest unit of data that can be read from an HFile. Each HFile consists of a series of blocks. (Note: an HBase block is different from an HDFS block or other underlying file system blocks.)


HBase Component	Description
BlockCache	BlockCache is the main HBase mechanism for low-latency random read operations. BlockCache is one of two memory cache structures maintained by HBase. When a block is read from HDFS, it is cached in BlockCache. Frequent access to rows in a block cause the block to be kept in cache, improving read performance.
MemStore	MemStore ("memory store") is in-memory storage for a RegionServer. MemStore is the second of two cache structures maintained by HBase. MemStore improves write performance. It accumulates data until it is full, and then writes ("flushes") the data to a new HFile on disk. MemStore serves two purposes: it increases the total amount of data written to disk in a single operation, and it retains recently written data in memory for subsequent low-latency reads.
Write Ahead Log (WAL)	The WAL is a log file that records all changes to data until the data is successfully written to disk (MemStore is flushed). This protects against data loss in the event of a failure before MemStore contents are written to disk.

HBase Read/Write Operations

BlockCache and MemStore reside in random-access memory (RAM). HFiles and the Write Ahead Log are persisted to HDFS.

The following figure shows these simplified write and read paths:


- During write operations, HBase writes to WAL and MemStore. Data is flushed from MemStore to disk according to size limits and flush interval.
- During read operations, HBase reads the block from BlockCache or MemStore if it is available in those caches. Otherwise, it reads from disk and stores a copy in BlockCache.



By default, BlockCache resides in an area of RAM that is managed by the Java Virtual Machine (JVM) garbage collector; this area of memory is known as on-heap memory or the JVM heap. The BlockCache implementation that manages the on-heap cache is called LruBlockCache.

If you have stringent read latency requirements and you have more than 20 GB of RAM available on your servers for use by HBase RegionServers, consider configuring BlockCache to use both on-heap and off-heap memory. BucketCache is the off-heap memory equivalent to LruBlockCache in on-heap memory. Read latencies for BucketCache tend to be less erratic than LruBlockCache for large cache loads because BucketCache (not JVM garbage collection) manages block cache allocation. The MemStore always resides in the on-heap memory.

Figure 3: Relationship among Different BlockCache Implementations and MemStore



- Additional notes:
- BlockCache is enabled by default for all HBase tables.
- BlockCache is beneficial for both random and sequential read operations although it is of primary consideration for random reads.
- All regions hosted by a RegionServer share the same BlockCache.
- You can turn BlockCache caching on or off per column family.

Advanced configuration for write-heavy workloads

HBase includes several advanced configuration parameters for adjusting the number of threads available to service flushes and compactions in the presence of write-heavy workloads. Tuning these parameters incorrectly can severely degrade performance and is not necessary for most HBase clusters. If you use Cloudera Manager, configure these options using the HBase Service Advanced Configuration Snippet (Safety Valve) for `hbase-site.xml`

`hbase.hstore.flusher.count`

The number of threads available to flush writes from memory to disk. Never increase `hbase.hstore.flusher.count` to more of 50% of the number of disks available to HBase. For example, if you have 8 solid-state drives (SSDs), `hbase.hstore.flusher.count` should never exceed 4. This allows scanners and compactions to proceed even in the presence of very high writes.

`hbase.regionserver.thread.compaction.large` and `hbase.regionserver.thread.compaction.small`

The number of threads available to handle small and large compactions, respectively. Never increase either of these options to more than 50% of the number of disks available to HBase.

Ideally, `hbase.regionserver.thread.compaction.small` should be greater than or equal to `hbase.regionserver.thread.compaction.large`, since the large compaction threads do more intense work and will be in use longer for a given operation.

In addition to the above, if you use compression on some column families, more CPU will be used when flushing these column families to disk during flushes or compaction. The impact on CPU usage depends on the size of the flush or the amount of data to be decompressed and compressed during compactions.