

Managing Data Storage

Date published: 2020-02-20

Date modified: 2021-08-05



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

Optimizing data storage.....	6
Balancing data across disks of a DataNode.....	6
Plan the data movement across disks.....	6
Run the Disk Balancer plan.....	7
Disk Balancer commands.....	8
Erasure coding overview.....	10
Understanding erasure coding policies.....	11
Comparing replication and erasure coding.....	12
Best practices for rack and node setup for EC.....	14
Prerequisites for enabling erasure coding.....	14
Limitations of erasure coding.....	14
Using erasure coding for existing data.....	14
Using erasure coding for new data.....	15
Advanced erasure coding configuration.....	15
Erasure coding CLI command.....	16
Erasure coding examples.....	17
Increasing storage capacity with HDFS compression.....	19
Enable GZipCodec as the default compression codec.....	19
Use GZipCodec with a one-time job.....	20
Set HDFS quotas.....	20
Setting HDFS quotas in Cloudera Manager.....	21
Configuring heterogeneous storage in HDFS.....	21
HDFS storage types.....	22
HDFS storage policies.....	22
Commands for configuring storage policies.....	22
Set up a storage policy for HDFS.....	23
Set up SSD storage using Cloudera Manager.....	23
Configure archival storage.....	24
The HDFS mover command.....	25
Balancing data across an HDFS cluster.....	25
Why HDFS data becomes unbalanced.....	25
Configurations and CLI options for the HDFS Balancer.....	26
Configuring and running the HDFS balancer using Cloudera Manager.....	29
Cluster balancing algorithm.....	31
Exit statuses for the HDFS Balancer.....	33
HDFS heap sizing.....	33
 Optimizing performance.....	 34
Improving performance with centralized cache management.....	34
Benefits of centralized cache management in HDFS.....	34
Use cases for centralized cache management.....	35
Centralized cache management architecture.....	35
Caching terminology.....	36
Properties for configuring centralized caching.....	37
Commands for using cache pools and directives.....	39
Specifying racks for hosts.....	42
Viewing racks assigned to cluster hosts.....	42
Editing rack assignments for hosts.....	42

Customizing HDFS.....	43
Customize the HDFS home directory.....	43
Properties to set the size of the NameNode edits directory.....	43
Optimizing NameNode disk space with Hadoop archives.....	43
Overview of Hadoop archives.....	43
Hadoop archive components.....	44
Create a Hadoop archive.....	45
List files in Hadoop archives.....	45
Format for using Hadoop archives with MapReduce.....	46
Detecting slow DataNodes.....	46
Enable disk IO statistics.....	46
Enable detection of slow DataNodes.....	47
Allocating DataNode memory as storage.....	48
HDFS storage types.....	48
LAZY_PERSIST memory storage policy.....	48
Configure DataNode memory as storage.....	48
Improving performance with short-circuit local reads.....	49
Prerequisites for configuring short-circuit local reads.....	49
Properties for configuring short-circuit local reads on HDFS.....	49
Configure mountable HDFS.....	52
Add HDFS system mount.....	52
Optimize mountable HDFS.....	53
 Configuring Proxy Users to Access HDFS.....	 53
 Using DistCp to copy files.....	 55
Using DistCp.....	55
Distcp syntax and examples.....	56
Using DistCp with Highly Available remote clusters.....	57
Using DistCp with Amazon S3.....	59
Using a credential provider to secure S3 credentials.....	60
Examples of DistCp commands using the S3 protocol and hidden credentials.....	61
Kerberos setup guidelines for Distcp between secure clusters.....	62
Distcp between secure clusters in different Kerberos realms.....	63
Configure source and destination realms in krb5.conf.....	63
Configure HDFS RPC protection.....	64
Specify truststore properties.....	64
Set HADOOP_CONF to the destination cluster.....	64
Launch distcp.....	65
Copying data between a secure and an insecure cluster using DistCp and WebHDFS.....	65
Post-migration verification.....	66
Using DistCp between HA clusters using Cloudera Manager.....	66
 Using the NFS Gateway for accessing HDFS.....	 68
Install the NFS Gateway.....	68
Configure the NFS Gateway.....	68
Start and stop the NFS Gateway services.....	70
Start the NFS Gateway services.....	71
Stop the NFS Gateway services.....	71
Verify validity of the NFS services.....	71
Access HDFS from the NFS Gateway.....	72
How NFS Gateway authenticates and maps users.....	72

APIs for accessing HDFS.....	73
Set up WebHDFS on a secure cluster.....	73
Using HttpFS to provide access to HDFS.....	74
Add the HttpFS role.....	74
Using Load Balancer with HttpFS.....	74
HttpFS authentication.....	75
Use curl to access a URL protected by Kerberos HTTP SPNEGO.....	75
Data storage metrics.....	76
Using JMX for accessing HDFS metrics.....	76
HDFS Metrics.....	77
Using HdfsFindTool to find files.....	78
Downloading Hdfsfindtool from the CDH archives.....	78

Optimizing data storage

You can consider various options to optimize data storage, such as increasing storage space through erasure coding, using codecs for compressing data, balancing data across an HDFS cluster, and so on.

To know more about HDFS heap calculations, see [HDFS heap sizing](#).

Balancing data across disks of a DataNode

The HDFS Disk Balancer is a command-line tool that evenly distributes data across all the disks of a DataNode to ensure that the disks are effectively utilized. Unlike the HDFS Balancer that balances data across DataNodes of a cluster, the Disk Balancer balances data at the level of individual DataNodes.

Disks on a DataNode can have an uneven spread of data because of reasons such as large amount of writes and deletes or disk replacement. To ensure uniformity in the distribution of data across disks of a specified DataNode, the Disk Balancer operates against the DataNode and moves data blocks from one disk to another.

For a specified DataNode, the Disk Balancer determines the ideal amount of data to store per disk based on its total capacity and used space, and computes the amount of data to redistribute between disks. The Disk Balancer captures this information about data movement across the disks of a DataNode in a plan. On executing the plan, the Disk Balancer moves the data as specified.

Plan the data movement across disks

A Disk Balancer plan identifies the amount of data that should move between the disks of a specified DataNode. The plan contains move steps, where each move step specifies the source and destination disks for data movement and the number of bytes to move.

Before you begin

You must set the `dfs.disk.balancer.enabled` property in `hdfs-site.xml` to `true`.

Procedure

Run the `hdfs diskbalancer -plan` command by specifying the path to the DataNode.

For example, `hdfs diskbalancer -plan node1.mycluster.com`.

The specified command generates the following two JSON documents as output in the HDFS namespace:

`<nodename>.before.json` that indicates the state of the cluster before running Disk Balancer, and `<nodename>.plan.json` that details the data movement plan for the DataNode.

By default, the JSON documents are placed in the following location: `/system/diskbalancer/<Creation_Timestamp>`. Here, `<Creation_Timestamp>` indicates the folder named according to the date and time of creation of the JSON documents.

What to do next

Run the generated plan.

Related Information

[Run the Disk Balancer plan](#)

Parameters to configure the Disk Balancer

You must configure various parameters through HDFS Service Advanced Configuration Snippet (Safety Valve) in `hdfs-site.xml` for effectively planning and executing the Disk Balancer.

Parameter	Description
dfs.disk.balancer.enabled	Controls whether the Disk Balancer is enabled for the cluster. The Disk Balancer executes only if this parameter is set to <i>True</i> . The default value is <i>False</i> .
dfs.disk.balancer.max.disk.throughputInMBperSec	The maximum disk bandwidth that Disk Balancer consumes while transferring data between disks. The default value is 10 MB/s.
dfs.disk.balancer.max.disk.errors	<p>The maximum number of errors to ignore for a move operation between two disks before abandoning the move. The default value of the maximum errors to ignore is 5.</p> <p>For example, if a plan specifies data move operation between three pairs of disks, and if the move between the first pair encounters more than five errors, that move is abandoned and the next move between the second pair of disks starts.</p>
dfs.disk.balancer.block.tolerance.percent	<p>Specifies a threshold value in percentage to consider a move operation successful, and stop moving further data.</p> <p>For example, setting this value to 20% for an operation requiring 10GB data movement indicates that the movement will be considered successful only after 8GB of data is moved.</p>
dfs.disk.balancer.plan.threshold.percent	<p>The ideal storage value for a set of disks in a DataNode indicates the amount of data each disk should have for achieving perfect data distribution across those disks. The threshold percentage defines the value at which disks start participating in data redistribution or balancing operations. Minor imbalances are ignored because normal operations automatically correct some of these imbalances.</p> <p>The default threshold percentage for a disk is 10%; indicating that a disk is used in balancing operations only if the disk contains 10% more or less data than the ideal storage value.</p>

Run the Disk Balancer plan

For a DataNode, the HDFS Disk Balancer moves data according to the details specified in the JSON document generated as the Disk Balancer plan.

Before you begin

You must have generated the data movement plan for the DataNode.

Procedure

Run the `hdfs diskbalancer -execute` command by specifying the path to the JSON plan document.

```
hdfs diskbalancer -execute /system/diskbalancer/<Creation_Timestamp>/nodename.plan.json
```

The command reads the DataNode's address from the JSON file and executes the plan.

Related Information

[Plan the data movement across disks](#)

Disk Balancer commands

In addition to planning for data movement across disks and executing the plan, you can use `hdfs diskbalancer` sub-commands to query the status of the plan, cancel the plan, identify at a cluster level the DataNodes that require balancing, or generate a detailed report on a specific DataNode that can benefit from running the Disk Balancer.

Planning the data movement for a DataNode

Command: `hdfs diskbalancer -plan <datanode>`

Argument	Description
<code><datanode></code>	Fully qualified name of the DataNode for which you want to generate the plan.

`hdfs diskbalancer -plan node1.mycluster.com`

The following table lists the additional options that you can use with the `hdfs diskbalancer -plan` command:

Option	Description
<code>-out</code>	Specify the location within the HDFS namespace where you want to save the output JSON documents that contain the generated plans.
<code>-bandwidth</code>	Specify the maximum bandwidth to use for running the Disk Balancer. This option helps in minimizing the amount of data moved by the Disk Balancer on an operational DataNode. Disk Balancer uses the default bandwidth of 10 MB/s if you do not specify this value.
<code>-thresholdPercentage</code>	The ideal storage value for a set of disks in a DataNode indicates the amount of data each disk should have for achieving perfect data distribution across those disks. The threshold percentage defines the value at which disks start participating in data redistribution or balancing operations. Minor imbalances are ignored because normal operations automatically correct some of these imbalances. The default value of <code>-thresholdPercentage</code> for a disk is 10%; indicating that a disk is used in balancing operations only if the disk contains 10% more or less data than the ideal storage value.
<code>-maxerror</code>	Specify the number of errors to ignore for a move operation between two disks before abandoning the move. Disk Balancer uses the default if you do not specify this value.
<code>-v</code>	Verbose mode. Specify this option for Disk Balancer to display a summary of the plan as output.
<code>-fs</code>	Specify the NameNode to use. Disk Balancer uses the default NameNode from the configuration if you do not specify this value.

Executing the plan

Command: `hdfs diskbalancer -execute <JSON file path>`

Argument	Description
<i><JSON file path></i>	Path to the JSON document that contains the generated plan (nodename.plan.json).

`hdfs diskbalancer -execute /system/diskbalancer/nodename.plan.json`

Querying the current status of execution

Command: `hdfs diskbalancer -query <datanode>`

Argument	Description
<i><datanode></i>	Fully qualified name of the DataNode for which the plan is running.

`hdfs diskbalancer -query nodename.mycluster.com`

Cancelling a running plan

Commands:

`hdfs diskbalancer -cancel <JSON file path>`

Argument	Description
<i><JSON file path></i>	Path to the JSON document that contains the generated plan (nodename.plan.json).

`hdfs diskbalancer -cancel /system/diskbalancer/nodename.plan.json`

OR

`hdfs diskbalancer -cancel <planID> -node <nodename>`

Argument	Description
<i>planID</i>	ID of the plan to cancel. You can get this value from the output of the <code>hdfs diskbalancer -query</code> command.
<i>nodename</i>	The fully qualified name of the DataNode on which the plan is running.

Viewing detailed report of DataNodes that require Disk Balancer

Commands:

`hdfs diskbalancer -fs http://namenode.uri -report -node <file://>`

Argument	Description
<i><file://></i>	Hosts file listing the DataNodes for which you want to generate the reports.

OR

`hdfs diskbalancer -fs http://namenode.uri -report -node [<DataNodeID|IP|Hostname>,...]`

Argument	Description
[<DataNodeID/IP/Hostname>,...]	Specify the DataNode ID, IP address, and the host name of the DataNode for which you want to generate the report. For multiple DataNodes, provide the details using comma-separated values.

Viewing details of the top DataNodes in a cluster that require Disk Balancer

Command: `hdfs diskbalancer -fs http://namenode.uri -report-node -top <topnum>`

Argument	Description
<topnum>	The number of the top DataNodes that require Disk Balancer to be run.

Erasure coding overview

Data durability describes how resilient data is to loss. When data is stored in HDFS, CDP provides two options for data durability. You can use replication, which HDFS was originally built on, or Erasure Coding (EC).



Note: The comparisons between EC and replication use a replication factor of 3 (three copies of data are maintained) since that is the default.

Replication

HDFS creates two copies of data, resulting in three total instances of data. These copies are stored on separate DataNodes to guard against data loss when a node is unreachable. When the data stored on a node is lost or inaccessible, it is replicated from one of the other nodes to a new node so that there are always multiple copies. The number of replications is configurable, but the default is three. Cloudera recommends keeping the replication factor to at least three when you have three or more DataNodes. A lower replication factor leads to a situation where the data is more vulnerable to DataNode failures since there are fewer copies of data spread out across fewer DataNodes..

When data is written to an HDFS cluster that uses replication, additional copies of the data are automatically created. No additional steps are required.

Replication supports all data processing engines that CDP supports.

Erasure Coding (EC)

EC is an alternative to replication. When an HDFS cluster uses EC, no additional direct copies of the data are generated. Instead, data is striped into blocks and encoded to generate parity blocks. If there are any missing or corrupt blocks, HDFS uses the remaining data and parity blocks to reconstruct the missing pieces in the background. This process provides a similar level of data durability to 3x replication but at a lower storage cost.

Additionally, EC is applied when data is written. This means that to use EC, you must first create a directory and configure it for EC. Then, you can either replicate existing data or write new data into this directory.

EC supports the following data processing engines:

- Hive
- MapReduce
- Spark

With both data durability schemes, replication and EC, recovery happens in the background and requires no direct input from a user.

HDFS clusters can be configured with a single data durability scheme (3x replication or EC), or with a hybrid data durability scheme where EC enabled directories co-exist on a cluster with other directories that are protected with the

traditional 3x replication model. This decision should be based on the temperature of the data (how often the data is accessed) stored in HDFS. Hot data, data that is accessed frequently, should use replication. Cold data, data that is accessed less frequently, can take advantage of EC's storage savings.

Understanding erasure coding policies

The EC policy determines how data is encoded and decoded. An EC policy is made up of the following parts: codec-number of data blocks-number of parity blocks-cell size.

- **Codec:** The erasure codec that the policy uses. CDP currently supports Reed-Solomon (RS).
- **Number of Data Blocks:** The number of data blocks per stripe. The higher this number, the more nodes that need to be read when accessing data because HDFS attempts to distribute the blocks evenly across DataNodes.
- **Number of Parity Blocks:** The number of parity blocks per stripe. Even if a file does not use up all the data blocks available to it, the number of parity blocks will always be the total number listed in the policy.
- **Cell Size:** The size of one basic unit of striped data.

For example, a RS-6-3-1024k policy has the following attributes:

- **Codec:** Reed-Solomon
- **Number of Data Blocks:** 6
- **Number of Parity Blocks:** 3
- **Cell Size:** 1024k

The sum of the number of data blocks and parity blocks is the data-stripe width. When you make hardware plans for your cluster, the number of racks should at least equal the stripe width in order for the data to be resistant to rack failures.

The following image compares the data durability and storage efficiency of different RS codecs and replication:

	Data Durability	Storage Efficiency
Single copy of data	0	100%
3x replication	2	33%
RS(6,3)	3	67%
RS(3,2)	2	60%

Storage efficiency is the ratio of data blocks to total blocks as represented by the following formula: data blocks / (data blocks + parity blocks)

Comparing replication and erasure coding

You must consider factors such as data temperature, i/o cost, storage cost, and file size when comparing replication and erasure coding.

Data Temperature

Data temperature refers to how often data is accessed. EC works best with cold data that is accessed and modified infrequently. There is no data locality, and all reads are remote. Replication is more suitable for hot data, data that is accessed and modified frequently because data locality is a part of replication.

I/O Cost

EC has higher I/O costs than replication for the following reasons:

- EC spreads data across nodes and racks, which means reading and writing data comes at a higher network cost.
- A parity block is generated when data is written, thus impacting write speed. This can be slower than writing to a file when the replication factor is one but is faster than writing two or more replicas.
- If data is missing or corrupt, a DataNode reads the remaining data and parity blocks in order to reconstruct the data. This process requires CPU and network resources.

Cloudera recommends at least a 10 GB network connection if you want to use EC.

Storage Cost

EC has a lower storage overhead than replication because multiple copies of data are not maintained. Instead, a number of parity blocks are generated based on the EC policy. For the same amount of data, EC will store fewer blocks than 3x replication in most cases. For example with a Reed-Solomon (6,3), HDFS stores three parity blocks for each set of 6 data blocks. With replication, HDFS stores 12 replica blocks for every six data blocks, the original block and three replicas. The case where 3x replication requires fewer blocks is when data is stored in small files.


File Size

Erasure coding works best with larger files. The total number of blocks is determined by data blocks + parity blocks, which is the data-stripe width discussed earlier.


128 MB is the default block size. With RS (6,3), each block group can hold up to $(128 \text{ MB} * 6) = 768 \text{ MB}$ of data. Inside each block group, there will be 9 total blocks, 6 data blocks, each holding up to 128 MB, and 3 parity blocks. This is why EC works best with larger files. For a chunk of data less than the block size, replication writes one data block to three DataNodes; EC, on the other hand, still needs to stripe the data to data blocks and calculate parity blocks. This leads to a situation where erasure coded files will generate more blocks than replication because of the parity blocks required. The bytes/blocks ratio is worse for small files which increases the memory usage of the NameNode.

The following two figures show how replication (with a replication factor of 3) compares to EC based on the number of blocks relative to file size.

For very small files of sizes smaller than the value of data blocks * cell size (in case of RS(6, 3) it is $(6 * 1) \text{ MB}$), the number of actual data blocks is less than the data blocks defined by the erasure coding policy, though the number of parity blocks is always the same. For example, in the case of RS(6,3) with a cell size of 1 MB, a 1MB file consists of one data block, rather than six, but it still has three parity blocks. A 1MB file would therefore require four block objects in total. If 3x replication were used, the same file would require only three block objects.



As shown in the following figure, the RS(6,3) EC policy ensures that the number of data blocks continues to remain less than that used by 3x replication for larger file sizes.



Supported Engines

Replication supports all data processing engines that CDP supports.

EC supports the following data processing engines: Hive, MapReduce, and Spark.

Unsupported Features

The XOR codec for EC is not supported. Additionally, certain HDFS functions are not supported with EC: hflush, hsync, concat, setReplication, truncate and append. For more information, see [Erasure Coding Limitations](#), and [HDFS Erasure Coding Limitations](#).

Best practices for rack and node setup for EC

In a CDP Private Cloud Base deployment, when setting up a cluster to take advantage of EC, consider the number of racks and nodes in your setup.

Ideally, the number of racks exceed the data-stripe width to account for downtime and outages. If there are fewer racks than the data-stripe width, HDFS spreads data blocks across multiple nodes to maintain fault tolerance at the node level. When distributing blocks to racks, HDFS attempts to distribute the blocks evenly across all racks. Because of this behavior, Cloudera recommends setting up each rack with a similar number of DataNodes. Otherwise, racks with fewer DataNodes may be filled up faster than racks with more DataNodes.

To achieve node-level fault tolerance, the number of nodes needs to equal the data-stripe width. For example, in order for a RS-6-3-1024k policy to be node failure tolerant, you need at least nine nodes. Note that the write will fail if the number of DataNodes is less than the policy's number of data blocks. The write will succeed but show a warning message if the number of DataNodes is less than the policy's number of data blocks + parity blocks. For example, with RS(6,3), if there are six to eight DataNodes, the write will succeed but show a warning message. If there are less than six DataNodes, the write fails.

For rack-level fault tolerance, there must be enough number of racks such that each rack contains at most the same number of blocks as the erasure coding parity blocks. The number of racks is calculated as (data stripe width) / (number of parity blocks). For example, with RS(6,3), the minimum number of racks must be $(6+3)/3 = 3$. Cloudera recommends at least nine racks, which leads to one data or parity block on each rack. Ideally, the number of racks exceeds the sum of the number of data and parity blocks.

If you decide to use three racks, spread the nine nodes evenly across three racks. The data and parity blocks, when distributed evenly, lead to the following placement on the racks:

- Rack 1 with three nodes: Three blocks
- Rack 2 with three nodes: Three blocks
- Rack 3 with three nodes: Three blocks

A policy with a wide data-stripe width like RS-6-3-1024k comes with a tradeoff though. Data must be read from six blocks, increasing the overall network load. Choose your EC policy based on your network settings and expected storage efficiency. Note, the larger the cluster and colder the data, the more appropriate it is to use EC policies with large data-stripe widths. Larger data-stripe widths have the benefit of a better storage efficiency.

Prerequisites for enabling erasure coding

Before enabling erasure coding on your data, you must consider various factors such as the type of policy to use, the type of data, and the rack or node requirements.

- Note the limitations for EC.
- Determine which EC policy you want to use:
- Determine if you want to use EC for existing data or new data. If you want to use EC for existing data, you need to replicate that data with distcp or BDR.
- Verify that your cluster setup meets the rack and node requirements.

Limitations of erasure coding

The limitations of erasure coding include non-support of XOR codecs and certain HDFS functions.

EC does not support:

- XOR codecs
- certain HDFS functions such as hflush, hsync, concat, setReplication, and truncate. For more information, see [Erasure Coding Limitations](#).

Using erasure coding for existing data

You must set a supported EC policy for a directory and copy the existing data to the directory.

Procedure

1. Create a new directory or choose an existing directory.
2. View the supported EC policies.

```
hdfs ec -listPolicies
```

3. Enable a supported EC policy.

```
hdfs ec -enablePolicy -policy <policy>
```

4. Set the EC policy for the directory you want to use.

```
hdfs ec -setPolicy -path <directory> [-policy <policyName>]
```

- path. Required. Specify the HDFS directory you want to apply the EC policy to.
- policy. Optional. The EC policy you want to use for the directory you specified. If you do not provide this parameter, the EC policy you specified in the Default Policy when Setting Erasure Coding setting from Cloudera Manager is used.

5. Copy the data to the directory you set an EC policy for.

You can use the `distcp` tool or Cloudera Manager's Backup and Disaster Recovery process.

Using erasure coding for new data

You must create a new directory and then set a supported EC policy for the directory.

Procedure

1. Create a new directory or choose an existing directory.
2. View the supported EC policies.

```
hdfs ec -listPolicies
```

3. Enable a supported EC policy.

```
hdfs ec -enablePolicy -policy <policy>
```

4. Set the EC policy for the directory you want to use.

```
hdfs ec -setPolicy -path <directory> [-policy <policyName>]
```

- path. Required. Specify the HDFS directory you want to apply the EC policy to.
- policy. Optional. The EC policy you want to use for the directory you specified. If you do not provide this parameter, the EC policy you specified in the Fallback Erasure Coding Policy setting from Cloudera Manager is used.

5. Set the destination for the data to the directory you enabled EC for. No action beyond that is required.

When data is written to the directory, it will be erasure coded based on the policy you set.

Advanced erasure coding configuration

You can customize the behavior of EC through a combination of the `hdfs ec` subcommand and the Cloudera Manager Admin Console.

About this task

This procedure provides information about configuring certain properties for EC with the Cloudera Manager Admin Console.

For information about the `hdfs ec` see [Erasure coding CLI command](#) on page 16.

Procedure

1. Select Clusters and choose the HDFS cluster you want to configure.
2. Navigate to the Configuration tab and select the Erasure Coding category.
3. Configure the advanced EC properties.
 - DataNode Striped Read Timeout: The timeout for striped reads during background data reconstruction in milliseconds.
 - DataNode Striped Read Threads: The number of threads that a DataNode can use during background data reconstruction.
 - Erasure Coding Reconstruction Weight: The relative weight of resources used by EC for data recovery. The number of blocks that must be read is based on the EC policy used. For example, RS-6-3-1024k requires six blocks to be read. Replication only requires one block to be read. Higher values result in fewer reconstruction tasks being able to run concurrently. The number of blocks required to be read to recover data is multiplied by this weight to determine the total weight of the recovery task. The total weight of the recovery task counts against the limit set with the `dfs.namenode.replication.max-streams` property.
 - Fallback Erasure Coding Policy: The fallback Erasure Coding policy that HDFS uses if no policy is specified when you run the `-setPolicy` command.



Note: You must disable erasure coding in small clusters to prevent potential data loss. If the cluster falls back to an Erasure Coding policy value that requires a greater number of DataNodes or racks than available, this situation could result in a potential data loss. As a result, the Erasure Coding Policy Verification Test returns "Concerning" health. To prevent the issue, set the value of the fallback Erasure Coding policy to No Default Erasure Coding Policy. In addition, you must disable all the erasure coding policies that are enabled on the cluster.

Erasure coding CLI command

Use the `hdfs ec` command to set erasure coding policies on directories.

```
hdfs ec [generic options]
  [-setPolicy -path <path> [-policy <policyName>] [-replicate]]
  [-getPolicy -path <path>]
  [-unsetPolicy -path <path>]
  [-listPolicies]
  [-addPolicies -policyFile <file>]
  [-listCodecs]
  [-removePolicy -policy <policyName>]
  [-enablePolicy -policy <policyName>]
  [-disablePolicy -policy <policyName>]
  [-help [cmd ...]]
```

Options:

- `[-setPolicy [-p <policyName>] <path>]`: Sets an EC policy on a directory at the specified path. The following EC policies are supported: RS-3-2-1024k (Reed-Solomon with 3 data blocks, 2 parity blocks and 1024 KB cell size), RS-6-3-1024k, RS-LEGACY-6-3-1024k, and XOR-2-1-1024k.

`<path>`: A directory in HDFS. This is a mandatory parameter. Setting a policy only affects newly created files, and does not affect existing files.

`<policyName>`: The EC policy to be used for files under the specified directory. This is an optional parameter, specified using the `-p` flag. If no policy is specified, the system default Erasure Coding policy is used. The default policy is RS-6-3-1024k.

- `-replicate`: Forces a directory to use the default 3x replication scheme.



Note: You cannot specify `-replicate` and `-policy <policyName>` at the same time. Both the arguments are optional.

- `-getPolicy -path <path>`: Gets details of the EC policy of a file or directory for the specified path.

- `[-unsetPolicy -path <path>]`: Removes an EC policy already set by a `setPolicy` on a directory. This option does not work on a directory that inherits the EC policy from a parent directory. If you run this option on a directory that does not have an explicit policy set, no error is returned.
- `[-addPolicies -policyFile <file>]`: Adds a list of EC policies. HDFS allows you to add 64 policies in total, with the policy ID in the range of 64 to 127. Adding policies fails if there are already 64 policies.
- `[-listCodecs]`: Lists all the supported EC erasure coding codecs and coders in the system.
- `[-removePolicy -policy <policyName>]`: Removes an EC policy.
- `[-enablePolicy -policy <policyName>]`: Enables an EC policy.
- `[-disablePolicy -policy <policyName>]`: Disables an EC policy.
- `[-help]`: Displays help for a given command, or for all commands if none is specified.

Erasure coding background recovery work on the DataNodes can be tuned using the following configuration parameters in `hdfs-site.xml`.

- `dfs.datanode.ec.reconstruction.stripedread.timeout.millis`: Timeout for striped reads. Default value is 5000 ms.
- `dfs.datanode.ec.reconstruction.threads`: Number of threads used by the DataNode for the background recovery task. The default value is 8 threads.
- `dfs.datanode.ec.reconstruction.stripedread.buffer.size`: Buffer size for reader service. Default value is 64 KB.
- `dfs.datanode.ec.reconstruction.xmits.weight`: The relative weight of xmits used by the EC background recovery task when compared to replicated block recovery. The default value is 0.5.

If the parameter is set to 0, the EC task always has one xmit. The xmits of an erasure coding recovery task are calculated as the maximum value between the number of read streams and the number of write streams. For example, if an EC recovery task needs to read from six nodes and write to two nodes, the xmit value is $\max(6, 2) * 0.5 = 3$.

Erasure coding examples

You can use the `hdfs ec` command with its various options to set erasure coding policies on directories.

Viewing the list of erasure coding policies

The following example shows how you can view the list of available erasure coding policies:

```
hdfs ec -listPolicies
Erasure Coding Policies:
ErasureCodingPolicy=[Name=RS-10-4-1024k, Schema=[ECSchema=[Codec=rs,
numDataUnits=10, numParityUnits=4]], CellSize=1048576, Id=5], State=DISABLED
ErasureCodingPolicy=[Name=RS-3-2-1024k, Schema=[ECSchema=[Codec=rs,
numDataUnits=3, numParityUnits=2]], CellSize=1048576, Id=2], State=DISABLED
ErasureCodingPolicy=[Name=RS-6-3-1024k, Schema=[ECSchema=[Codec=rs,
numDataUnits=6, numParityUnits=3]], CellSize=1048576, Id=1], State=ENABLED
ErasureCodingPolicy=[Name=RS-LEGACY-6-3-1024k, Schema=[ECSchema=[Codec=rs-
legacy, numDataUnits=6, numParityUnits=3]], CellSize=1048576, Id=3],
State=DISABLED
ErasureCodingPolicy=[Name=XOR-2-1-1024k, Schema=[ECSchema=[Codec=xor,
numDataUnits=2, numParityUnits=1]], CellSize=1048576, Id=4], State=DISABLED
```

Enabling an erasure coding policy

In the previous example, the list of erasure coding policies indicates that RS-6-3-1024k is already enabled. If required, you can enable additional policies as mentioned in the following example:

```
hdfs ec -enablePolicy -policy RS-3-2-1024k
Erasure coding policy RS-3-2-1024k is enabled
```

Setting an erasure coding policy

The following example shows how you can set the erasure coding policy RS-6-3-1024k on a particular directory:

```
hdfs ec -setPolicy -path /data/dir1 -policy RS-6-3-1024k
Set erasure coding policy RS-6-3-1024k on /data/dir1
```

To confirm whether the specified directory has the erasure coding policy applied, run the `hdfs ec -getPolicy` command:

```
hdfs ec -getPolicy -path /data/dir1
RS-6-3-1024k
```

Checking the block status on an erasure-coded directory

After enabling erasure coding on a directory, you can check the block status by running the `hdfs fsck` command. The following example output shows the status of the erasure-coded blocks:

```
hdfs fsck /data/dir1
.
.
.
Erasure Coded Block Groups:
Total size: 434424 B
Total files: 1
Total block groups (validated): 1 (avg. block group size 434424 B)
Minimally erasure-coded block groups: 1 (100.0 %)
Over-erasure-coded block groups: 0 (0.0 %)
Under-erasure-coded block groups: 0 (0.0 %)
Unsatisfactory placement block groups: 0 (0.0 %)
Average block group size: 4.0
Missing block groups: 0
Corrupt block groups: 0
Missing internal blocks: 0 (0.0 %)
FSCK ended at Fri Mar 21 19:39:11 UTC 2018 in 1 milliseconds

The filesystem under path '/data/dir1' is HEALTHY
```

Changing the erasure coding policy

You can use the `hdfs ec setPolicy` command to change the erasure coding policy applied on a particular directory.

```
hdfs ec -setPolicy -path /data/dir1 -policy RS-3-2-1024k
Set erasure coding policy RS-3-2-1024k on /data/dir1
```

You can check the check the block status after applying the new policy. The following example output shows the status of the erasure-coded blocks for a directory that has the RS-3-2-1024k policy:

```
hdfs fsck /data/dir1
.
.
.
Erasure Coded Block Groups:
Total size: 68644 B
Total files: 2
Total block groups (validated): 2 (avg. block group size 34322 B)
Minimally erasure-coded block groups: 2 (100.0 %)
Over-erasure-coded block groups: 0 (0.0 %)
Under-erasure-coded block groups: 0 (0.0 %)
Unsatisfactory placement block groups: 0 (0.0 %)
Average block group size: 2.5
Missing block groups: 0
Corrupt block groups: 0
Missing internal blocks: 0 (0.0 %)
FSCK ended at Mon Apr 09 10:11:06 UTC 2018 in 3 milliseconds
```

```
The filesystem under path '/data/dir1' is HEALTHY
```

You can apply the default 3x replication policy and check the block status as specified in the following examples:

```
hdfs ec -setPolicy -path /data/dir1 -replicate
Set erasure coding policy replication on /tmp/data1/
Warning: setting erasure coding policy on a non-empty directory will not
automatically convert existing files to replication
```

```
hdfs fsck /data/dir1
.
.
.
Erasure Coded Block Groups:
Total size: 34322 B
Total files: 1
Total block groups (validated): 1 (avg. block group size 34322 B)
Minimally erasure-coded block groups: 1 (100.0 %)
Over-erasure-coded block groups: 0 (0.0 %)
Under-erasure-coded block groups: 0 (0.0 %)
Unsatisfactory placement block groups: 0 (0.0 %)
Average block group size: 2.0
Missing block groups: 0
Corrupt block groups: 0
Missing internal blocks: 0 (0.0 %)
FSCK ended at Tue Apr 10 04:34:14 UTC 2018 in 2 milliseconds

The filesystem under path '/data/dir1' is HEALTHY
```

Increasing storage capacity with HDFS compression

Linux supports GzipCodec, DefaultCodec, BZip2Codec, LzoCodec, and SnappyCodec. Typically, GzipCodec is used for HDFS compression.

To configure data compression, you can either enable a data compression codec, for example, GZipCodec, as the default or use the codec from the command line with a one-time job.

Enable GZipCodec as the default compression codec

For the MapReduce framework, update relevant properties in `core-site.xml` and `mapred-site.xml` to enable GZipCodec as the default compression codec.

Procedure

1. Edit the `core-site.xml` file on the NameNode host machine.

```
<property>
  <name>io.compression.codecs</name>
  <value>org.apache.hadoop.io.compress.GzipCodec,
    org.apache.hadoop.io.compress.DefaultCodec,com.hadoop.compression.lzo.
    LzoCodec,org.apache.hadoop.io.compress.SnappyCodec</value>
  <description>A list of the compression codec classes that can be used
    for compression/decompression.</description>
</property>
```

2. Edit the `mapred-site.xml` file on the JobTracker host machine.

```
<property>
  <name>mapreduce.map.output.compress</name>
```

```

    <value>true</value>
  </property>

  <property>
    <name>mapreduce.map.output.compress.codec</name>
    <value>org.apache.hadoop.io.compress.GzipCodec</value>
  </property>

  <property>
    <name>mapreduce.output.fileoutputformat.compress.type</name>
    <value>BLOCK</value>
  </property>

```

3. Enable the following two configuration parameters to enable job output compression. Edit the `mapred-site.xml` file on the Resource Manager host machine.

```

<property>
  <name>mapreduce.output.fileoutputformat.compress</name>
  <value>true</value>
</property>

<property>
  <name>mapreduce.output.fileoutputformat.compress.codec</name>
  <value>org.apache.hadoop.io.compress.GzipCodec</value>
</property>

```

4. Restart the cluster.

Use GZipCodec with a one-time job

You can configure GZipcodec to compress the output of a MapReduce job.

Procedure

To use GzipCodec with a one-time only job, add the options to configure compression for the MapReduce job and configure GZipCodec for the output of the job.

```

hadoop jar hadoop-examples-1.1.0-SNAPSHOT.jar sort sbr "-Dmapred.compress.map
.output=true"
sbr "-Dmapred.map.output.compression.codec=org.apache.hadoop.io.compress.G
zipCodec"
sbr "-Dmapred.output.compress=true"
sbr "-Dmapred.output.compression.codec=org.apache.hadoop.io.compress.GzipC
odec" sbr -outKey org.apache.hadoop.io.Text sbr
-outValue org.apache.hadoop.io.Text input output

```

Set HDFS quotas

As an administrator, you can set HDFS quotas for the number of file and directory names used and the amount of space used by directories after the Reports Manager service is installed as part of Cloudera Management Service. Before you can set a quota, the Reports Manager must index the file or directory.

Considerations for working with HDFS quotas

- The quotas for names and the quotas for space are independent of each other.
- You cannot create more files and directories if their creation would cause the quotas to exceed.
- Block allocation fails if the quota prevents a full block from being written.
- If you are using replication, each replica of a block counts against the quota.

File count limits

- The file count quota is a limit on the number of file and directory names in the directory configured.
- A directory counts against its own quota. Therefore, a quota value of 1 forces the directory to remain empty.
- File counts are based on the intended replication factor for the files. Changing the replication factor for a file will increase or decrease the corresponding quota values.

Disk space limits

- The space quota is a hard limit on the number of bytes used by files in the tree rooted at the directory being configured.
- Each replica of a block counts against the quota.
- The disk space quota calculation takes replication into account. Therefore, the calculation uses the replicated size of each file and not the user-facing size.
- The disk space quota calculation includes open files (that are being written) and files that already written.
- Block allocations for files being written fail if the quota does not allow a full block to be written.

Related Information

[HDFS Quotas Guide](#)

[Set up a storage policy for HDFS](#)

Setting HDFS quotas in Cloudera Manager

You can set the file count or space quotas for a directory using Cloudera Manager. Ensure that you have administrator privileges before you set the quotas.

Procedure

1. Go to the Cloudera Manager Clusters *HDFS service* File Browser tab.
2. Browse the file system to find the directory for which you want to set quotas.
3. Click the directory name so that it appears in the gray panel above the listing of its contents and in the detail section to the right of the File Browser table.
4. Click the Edit Quota option for the directory.



Tip: The Edit Quota option appears if the Reports Manager service is installed as part of Cloudera Management Service.

5. Set the file count or disk space limits for the directory in the Manage Quota dialog box.



Tip: Before you can set a quota, the Reports Manager must index the file or directory.

6. Click OK.

Configuring heterogeneous storage in HDFS

A variety of storage types are supported with HDFS. You can choose which storage type to assign to each DataNode Data Directory. Specifying a storage type allows you to optimize your data usage and lower your costs, based on your data usage frequency.

Each DataNode in a cluster is configured with a set of data directories. You can configure each data directory with a storage type. The storage policy dictates which storage types to use when storing the file or directory.

Some reasons to consider using different types of storage are as follows:

- You have datasets with temporal locality (for example, time-series data). The latest data can be loaded initially into SSD for improved performance, then migrated out to disk as it ages.

- You need to move cold data to denser archival storage because the data will rarely be accessed and archival storage is much cheaper. This could be done with simple age-out policies: for example, moving data older than six months to archival storage.

HDFS storage types

The storage type associated with HDFS identifies the underlying storage media.

HDFS supports the following storage types:

- ARCHIVE** - Archival storage is for very dense storage and is useful for rarely accessed data. This storage type is typically cheaper per TB than normal hard disks.
- DISK** - Hard disk drives are relatively inexpensive and provide sequential I/O performance. This is the default storage type.
- SSD** - Solid state drives are useful for storing hot data and I/O-intensive applications.
- RAM_DISK** - This special in-memory storage type is used to accelerate low-durability, single-replica writes.

When you add the DataNode Data Directory, you can specify which type of storage it uses, by prefixing the path with the storage type, in brackets. If you do not specify a storage type, it is assumed to be DISK.

HDFS storage policies

A storage policy contains information that describes the type of storage to use. This policy also defines the fallback storage type if the primary type is out of space or out of quota. If a target storage type is not available, HDFS attempts to place replicas on the default storage type.

Each storage policy consists of a policy ID, a policy name, a list of storage types, a list of fallback storage types for file creation, and a list of fallback storage types for replication.

HDFS has six preconfigured storage policies.

- Hot** - All replicas are stored on DISK.
- Cold** - All replicas are stored ARCHIVE.
- Warm** - One replica is stored on DISK and the others are stored on ARCHIVE.
- All_SSD** - All replicas are stored on SSD.
- One_SSD** - One replica is stored on SSD and the others are stored on DISK.
- Lazy_Persist** - The replica is written to RAM_DISK and then lazily persisted to DISK.



Note: You cannot create your own storage policy. You must use one of the six pre-configured policies. HDFS clients such as HBase may support different storage policies.

Commands for configuring storage policies

Depending on your requirements, use the `hdfs storagepolicy` sub-commands to set storage policies or list the storage policies applied on a file or a directory.

Setting storage policies

Command: `hdfs storagepolicies -setStoragePolicy -path <path> -policy <policyName>`

Argument	Description
<path>	The path to a directory or file.
<policyName>	The name of the storage policy.

`hdfs storagepolicies -setStoragePolicy -path /cold1 -policy COLD`

Listing storage policies

Command: `hdfs storagepolicies -getStoragePolicy -path <path>`

Argument	Description
<path>	The path to a directory or file.

```
hdfs storagepolicies -getStoragePolicy -path /cold1
```

Unsetting/clearing storage policies

Command: `hdfs storagepolicies -unsetStoragePolicy -path <path>`

Argument	Description
<path>	The path to a directory or file.

```
hdfs storagepolicies -unsetStoragePolicy -path /cold1
```

Set up a storage policy for HDFS

You can use Cloudera Manager to set up storage policy on a DataNode data directory.

Procedure

1. Check the HDFS Service Advanced Configuration Snippet (Safety Valve) for `hdfs-site.xml` to be sure that `dfs.storage.policy.enabled` has not been changed from its default value of `true`.
2. Specify the storage types for each DataNode Data Directory that is not a standard disk, by adding the storage type in brackets at the beginning of the directory path.

For example:

```
[SSD]/dfs/dn1
[DISK]/dfs/dn2
[ARCHIVE]/dfs/dn3
```

3. Open a terminal session on any HDFS host and run the `hdfs storagepolicies -setStoragePolicy` command for each path on which you want to set a storage policy.

```
hdfs storagepolicies -setStoragePolicy -path <path> -policy <policy>
path_to_file_or_directory -policy policy_name
```

4. To move the data to the appropriate storage based on the current storage policy, use the mover utility, from any HDFS host.

Use `mover -h` to get a list of available options. To migrate all data at once (this may take a long time), you can set the path to `/`.

```
hdfs mover -p <path>
```



Note: Quotas are enforced at the time you set the storage policy or when writing the file, not when quotas are changed. The Mover tool does not recognize quota violations. It only verifies that a file is stored on the storage types specified in its policy.

Related Information

[Set HDFS quotas](#)

Set up SSD storage using Cloudera Manager

You can use Cloudera Manager to set up SSD storage for your data directories.

Procedure

1. Set up your cluster normally, but customize your DataNodes with the `[ssd]` prefix for data directories. Adding `[ssd]` can also be done after initial setup (which requires an extra HDFS restart).

2. Stop HBase.
3. Using the HDFS client, move /hbase to /hbase_backup.
4. Re-create /hbase using the Cloudera Manager command in the HBase service (this ensures that proper permissions are used).
5. Using the HDFS client, set the storage policy for /hbase to be SSD only.
6. Use the DistCp to copy /hbase_backup to /hbase.

```
hadoop distcp /hbase_backup /hbase
```

7. Start HBase.

Configure archival storage

To configure archival storage for a DataNode, you must assign the ARCHIVE storage type to the DataNode, set storage policies, and move blocks that violate the storage policy to the appropriate storage type.

Procedure

1. Shut down the DataNode.
2. Assign the ARCHIVE Storage Type to the DataNode.

You can use the `dfs.datanode.data.dir` property in the `/etc/hadoop/conf/hdfs-site.xml` file to assign the ARCHIVE storage type to a DataNode.

The `dfs.datanode.data.dir` property determines where on the local filesystem a DataNode should store its blocks.

If you specify a comma-delimited list of directories, data will be stored in all named directories, typically on different devices. Directories that do not exist are ignored. You can specify that each directory resides on a different type of storage: DISK, SSD, ARCHIVE, or RAM_DISK.

To specify a DataNode as DISK storage, specify [DISK] and a local file system path. For example:

```
<property>
  <name>dfs.datanode.data.dir</name>
  <value>[DISK]/grid/1/tmp/data_trunk</value>
</property>
```

To specify a DataNode as ARCHIVE storage, insert [ARCHIVE] at the beginning of the local file system path. For example:

```
<property>
  <name>dfs.datanode.data.dir</name>
  <value>[ARCHIVE]/grid/1/tmp/data_trunk</value>
</property>
```

3. Depending on your requirements, either set a storage policy or list the already applied storage policies on a specified file or directory.

If you want to...	Use this command...
Set a storage policy	<code>hdfs storagepolicies -setStoragePolicy <path> <policyname></code>
List storage policies	<code>hdfs storagepolicies -getStoragePolicy <path></code>



Note: When you update a storage policy setting on a file or directory, the new policy is *not* automatically enforced. You must use the HDFS mover data migration tool to actually move blocks as specified by the new storage policy.

4. Start the DataNode.
5. Use the HDFS mover tool according to the specified storage policy.

The HDFS mover data migration tool scans the specified files in HDFS and verifies if the block placement satisfies the storage policy. For the blocks that violate the storage policy, the tool moves the replicas to a different storage type in order to fulfill the storage policy requirements.

The HDFS mover command

You can use the `hdfs mover` command to move replicas of data blocks that violate the storage policy set on a file or a directory to a storage type that fulfills the policy requirements.

Command: `hdfs mover [-p <files/dirs> | -f <local file name>]`

Argument	Description
<code>-f <local file></code>	Specify a local file containing a list of HDFS files or directories to migrate.
<code>-p <files/dirs></code>	Specify a space-separated list of HDFS files or directories to migrate.



Note: When both `-p` and `-f` options are omitted, the default path is the root directory.

Balancing data across an HDFS cluster

The HDFS Balancer is a tool for balancing the data across the storage devices of a HDFS cluster.

You can also specify the source DataNodes, to free up the spaces in particular DataNodes. You can use a block distribution application to pin its block replicas to particular DataNodes so that the pinned replicas are not moved for cluster balancing.

Why HDFS data becomes unbalanced

Factors such as addition of DataNodes, block allocation in HDFS, and behavior of the client application can lead to the data stored in HDFS clusters becoming unbalanced.

Addition of DataNodes

When new DataNodes are added to a cluster, newly created blocks are written to these DataNodes from time to time. The existing blocks are not moved to them without using the HDFS Balancer.

Behavior of the client application

In some cases, a client application might not write data uniformly across the DataNode machines. A client application might be skewed in writing data, and might always write to some particular machines but not others. HBase is an example of such a client application. In other cases, the client application is not skewed by design, for example, MapReduce or YARN jobs.

The data is skewed so that some of the jobs write significantly more data than others. When a Datanode receives the data directly from the client, it stores a copy to its local storage for preserving data locality. The DataNodes receiving more data generally have higher storage utilization.

Block allocation in HDFS

HDFS uses a constraint satisfaction algorithm to allocate file blocks. Once the constraints are satisfied, HDFS allocates a block by randomly selecting a storage device from the candidate set uniformly. For large clusters, the blocks are essentially allocated randomly in a uniform distribution, provided that the client applications write data

to HDFS uniformly across the DataNode machines. Uniform random allocation might not result in a uniform data distribution because of randomness. This is generally not a problem when the cluster has sufficient space. The problem becomes serious when the cluster is nearly full.

Configurations and CLI options for the HDFS Balancer

You can configure the HDFS Balancer by changing various configuration options or by using the command line.

Properties for configuring the Balancer

Depending on your requirements, you can configure various properties for the HDFS Balancer.

dfs.datanode.balance.max.concurrent.moves

Limits the maximum number of concurrent block moves that a DataNode is allowed for balancing the cluster. If you set this configuration in a DataNode, the DataNode throws an exception when the limit is exceeded. If you set this configuration in the HDFS Balancer, the HDFS Balancer schedules concurrent block movements within the specified limit. The DataNode setting and the HDFS Balancer setting can be different. As both settings impose a restriction, an effective setting is the minimum of them.

It is recommended that you set this to the highest possible value in DataNodes and adjust the runtime value in the HDFS Balancer to gain the flexibility. The default value is 100.

You can reconfigure without DataNode restart. Follow these steps to reconfigure a DataNode:

1. Change the value of *dfs.datanode.balance.max.concurrent.moves* from the Configuration tab of the HDFS service from Cloudera Manager.
2. Refresh the cluster.

You can use the default value of 100 as the maximum number of concurrent block moves in most of the situations. If you want to set it to a lower value, you can consider a value between 25 and 50. The recommended maximum value for this parameter is 200.

dfs.datanode.balance.bandwidthPerSec

Limits the bandwidth in each DataNode using for balancing the cluster. Changing this configuration requires restarting DataNodes.

The default is 100 MB/s.

To dynamically change balancer bandwidth, use the following command:

```
hdfs dfsadmin -setBalancerBandwidth <bandwidth in bytes per second>
```

dfs.balancer.moverThreads

Limits the number of total concurrent moves for balancing in the entire cluster. Set this property to the number of threads in the HDFS Balancer for moving blocks. Each block move requires a thread.

The default is 1000.

dfs.balancer.max-size-to-move

With each iteration, the HDFS Balancer chooses DataNodes in pairs and moves data between the DataNode pairs. Limits the maximum size of data that the HDFS Balancer moves between a chosen DataNode pair. If you increase this configuration when the network and disk are not saturated, increases the data transfer between the DataNode pair in each iteration while the duration of an iteration remains about the same.

The default is 10GB.

dfs.balancer.getBlockSize

Specifies the total data size of the block list returned by a `getBlocks(..)`.

When the HDFS Balancer moves a certain amount of data between source and destination DataNodes, it repeatedly invokes the `getBlocks(..)` rpc to the Namenode to get lists of blocks from the source DataNode until the required amount of data is scheduled.

The default is 2GB.

dfs.balancer.getBlocks.min-block-size

Specifies the minimum block size for the blocks used to balance the cluster.

The default is 10MB.



Note: If the majority of files in the cluster have block size smaller than 10MB, then set a lower value for this parameter, because the HDFS Balancer might not be able to find enough blocks greater than this size to move.

dfs.datanode.block-pinning.enabled

Specifies if block-pinning is enabled. When you create a file, a user application can specify a list of favorable DataNodes by way of the file creation API in DistributedFileSystem. The NameNode uses its best effort, allocating blocks to the favorable DataNodes. If `dfs.datanode.block-pinning.enabled` is set to true, if a block replica is written to a favorable DataNode, it is “pinned” to that DataNode. The pinned replicas are not moved for cluster balancing to keep them stored in the specified favorable DataNodes. This feature is useful for block distribution aware user applications such as HBase.

The default is false.

Balancer commands

You can use various command line options with the `hdfs balancer` command to work with the HDFS Balancer.

Balancing policy, threshold, and blockpools

[-policy <policy>]

Specifies which policy to use to determine if a cluster is balanced.

The two supported policies are `blockpool` and `datanode`. Setting the policy to `blockpool` means that the cluster is balanced if each pool in each node is balanced while `datanode` means that a cluster is balanced if each DataNode is balanced.

The default policy is `datanode`.

[-threshold <threshold>]

Specifies a number in [1.0, 100.0] representing the acceptable threshold of the percentage of storage capacity so that storage utilization outside the average +/- the threshold is considered as over/under utilized.

The default threshold is 10.0.

[-blockpools <comma-separated list of blockpool ids>]

Specifies a list of block pools on which the HDFS Balancer runs. If the list is empty, the HDFS Balancer runs on all existing block pools.

The default value is an empty list.

Include and exclude lists

[-include [-f <hosts-file> | <comma-separated list of hosts>]]

When the include list is non-empty, only the DataNodes specified in the list are balanced by the HDFS Balancer. An empty include list means including all the DataNodes in the cluster. The default value is an empty list.

[-exclude [-f <hosts-file> | <comma-separated list of hosts>]]

The DataNodes specified in the exclude list are excluded so that the HDFS Balancer does not balance those DataNodes. An empty exclude list means that no DataNodes are excluded. When a DataNode is specified in both in the include list and the exclude list, the DataNode is excluded. The default value is an empty list.

Idle-iterations and run during upgrade

[-idleiterations <idleiterations>]

Specifies the number of consecutive iterations in which no blocks have been moved before the HDFS Balancer terminates with the NO_MOVE_PROGRESS exit status.

Specify -1 for infinite iterations. The default is 5.

[-runDuringUpgrade]

If specified, the HDFS Balancer runs even if there is an ongoing HDFS upgrade. If not specified, the HDFS Balancer terminates with the UNFINALIZED_UPGRADE exit status.

When there is no ongoing upgrade, this option has no effect. It is usually not desirable to run HDFS Balancer during upgrade. To support rollback, blocks being deleted from HDFS are moved to the internal trash directory in DataNodes and not actually deleted. Running the HDFS Balancer during upgrading cannot reduce the usage of any DataNode storage.

Source DataNodes

[-source [-f <hosts-file> | <comma-separated list of hosts>]]

Specifies the source DataNode list. The HDFS Balancer selects blocks to move from only the specified DataNodes. When the list is empty, all the DataNodes are chosen as a source. The option can be used to free up the space of some particular DataNodes in the cluster. Without the -source option, the HDFS Balancer can be inefficient in some cases.

The default value is an empty list.

The following table shows an example, where the average utilization is 25% so that D2 is within the 10% threshold. It is unnecessary to move any blocks from or to D2. Without specifying the source nodes, HDFS Balancer first moves blocks from D2 to D3, D4 and D5, since they are under the same rack, and then moves blocks from D1 to D2, D3, D4 and D5. By specifying D1 as the source node, HDFS Balancer directly moves blocks from D1 to D3, D4 and D5.

Table 1: Example of utilization movement

Datanodes (with the same capacity)	Utilization	Rack
D1	95%	A
D2	30%	B
D3, D4, and D5	0%	B

Recommended configurations for the Balancer

The HDFS Balancer can run in either Background or Fast modes. Depending on the mode in which you want the Balancer to run, you can set various properties to recommended values.

Background and fast modes

HDFS Balancer runs as a background process. The cluster serves other jobs and applications at the same time.

Fast Mode

HDFS Balancer runs at maximum (fast) speed.

Table 2: DataNode configuration properties

Property	Default	Background Mode	Fast Mode
dfs.datanode.balance.max.concurrent.moves	50	4 x (# of disks)	4 x (# of disks)
dfs.datanode.balance.bandwidthPerSec	10485760 (10 MB)	use default	10737418240 (10 GB)

Table 3: Balancer configuration properties

Property	Default	Background Mode	Fast Mode
dfs.datanode.balance.max.concurrent.moves	50	# of disks	4 x (# of disks)
dfs.balancer.moverThreads	1000	use default	20,000
dfs.balancer.max-size-to-move	10737418240 (10 GB)	1073741824 (1GB)	107374182400 (100 GB)
dfs.balancer.getBlockSize.min-block-size	10485760 (10 MB)	use default	104857600 (100 MB)

Troubleshooting tip

When the block report length exceeds the configured limit, you might receive the following exception in the NameNode logs:

```
java.io.IOException: Requested data length 141557760 is longer than maximum configured RPC length 134217728
```

Causes

When there are too many blocks in one or more Datanodes, for example in the range of several millions or 10M+, it leads to a large size of block report from those particular Datanodes, and the block report length exceeds the default value of 128M. In that case the NameNode only receives a part of block report instead of the entire block report as the block report gets truncated at the default configured length of 128M. Since the HDFS Balancer receives block information from NameNode, with wrong information of block number the balancer might not move the blocks from the Datanode with the above issue.

Solutions

Increase the value of `ipc.maximum.data.length` to 256M or above the reported data length from the above exception, so that the NameNode can receive a large size of block report. The balancer would then be able to progress further.



Note: If the above mentioned exception is sustained in the NameNode logs for a prolonged period of time, then the property modification would be required regardless of whether the balancer needs to be run or not and even for normal and correct HDFS operations.

Configuring and running the HDFS balancer using Cloudera Manager

Learn how to run HDFS Balancer by using Cloudera Manager. You configure the balancer threshold, concurrent moves, and block size. Then you run the HDFS Balancer. You also learn about recommended values for the HDFS Balancer.

In Cloudera Manager, the HDFS balancer utility is implemented by the Balancer role. The Balancer role usually shows a health of None on the HDFS Instances tab because it does not run continuously.

The Balancer role is normally added (by default) when the HDFS service is installed. If it has not been added, you must add a Balancer role to rebalance HDFS and to see the Rebalance action.

Configuring the balancer threshold

Learn how to configure the Balancer Threshold property, before using HDFS Balancer feature.

About this task

The Balancer has a default threshold of 10%, which ensures that disk usage on each DataNode differs from the overall usage in the cluster by no more than 10%. For example, if overall usage across all the DataNodes in the cluster is 40% of the cluster's total disk-storage capacity, the script ensures that DataNode disk usage is between 30% and 50% of the DataNode disk-storage capacity. To change the threshold, perform the following steps:

Procedure

1. Got to the HDFS service in Cloudera Manager.
2. Select the Configuration tab.
3. Select Scope Balancer .
4. Select Category Main .
5. Set the Rebalancing Threshold property.

To apply this configuration property to other role groups as needed, edit the value for the appropriate role group. For more information, see [Modifying Configuration Properties Using Cloudera Manager](#).

6. Enter a reason for change, and then click Save Changes.

Configuring concurrent moves

Learn how to configure concurrent moves for the DataNode and Balancer to prevent the balancer from taking too many resources from the DataNode and interfering with normal cluster operations.

About this task

The property `dfs.datanode.balance.max.concurrent.moves` sets the maximum number of threads used by the DataNode balancer for pending moves. It is a throttling mechanism to prevent the balancer from taking too many resources from the DataNode and interfering with normal cluster operations. Increasing the value allows the balancing process to complete more quickly. Decreasing the value allows rebalancing to complete more slowly, but is less likely to compete for resources with other tasks on the DataNode. To use this property, you need to set the value on both the DataNode and the Balancer.

Procedure

1. Configure concurrent moves on the DataNode.
 - a) Go to the HDFS service.
 - b) Select the Configuration tab.
 - c) Search for DataNode Advanced Configuration Snippet (Safety Valve) for `hdfs-site.xml`.
 - d) Add the following code to the configuration field, for example, setting the value to 50.

```
<property>
  <name>dfs.datanode.balance.max.concurrent.moves</name>
  <value>50</value>
</property>
```

- e) Click Save Changes.
 - f) Restart the DataNode.
2. Configure concurrent moves on the Balancer.
 - a) Go to the HDFS service.
 - b) Select the Configuration tab.
 - c) Search for Balancer Advanced Configuration Snippet (Safety Valve) for `hdfs-site.xml`.
 - d) Add the following code to the configuration field, for example, setting the value to 50.

```
<property>
  <name>dfs.datanode.balance.max.concurrent.moves</name>
  <value>50</value>
```

```
</property>
```

- e) Click Save Changes.

Recommended configurations for the balancer

Depending on whether you want the HDFS Balancer to run in the background or at maximum speed, you can configure the values of certain properties through safety valves.

Property	Values for Running the Balancer in the Background	Value for Running the Balancer at Maximum Speed
DataNode		
dfs.datanode.balance.bandwidthPerSec	10 MB	10 GB
Balancer		
dfs.balancer.moverThreads	1000	20000
dfs.balancer.max-size-to-move	10 GB	100 GB
dfs.balancer.getBlocks.min-block-size	10 MB	100 MB

Running the balancer

Learn how to run the HDFS Balancer.

Procedure

1. Go to the HDFS service.
 2. Ensure that the service has a Balancer role.
 3. Select Actions Rebalance .
 4. Click Rebalance to confirm.
- If you see a Finished status, the Balancer ran successfully.

Configuring block size

Learn how to configure the block metadata batch size and minimum block size for HDFS.

About this task

The Block Metadata Batch Size property configures the amount of block metadata that gets retrieved. The Minimum Block Size property configures the smallest block to consider for moving. Tuning these properties can improve performance during balancing.

Procedure

1. In the Cloudera Manager Admin Console, select Clusters <HDFS cluster> .
2. Select the Configuration tab.
3. Search for the following properties:
 - Block Metadata Batch Size (dfs.balancer.getBlocks.size)
 - Minimum Block Size (dfs.balancer.getBlocks.min-block-size)
4. Set the values for each, and click Save Changes.

Cluster balancing algorithm

The HDFS Balancer runs in iterations. Each iteration contains the following four steps: storage group classification, storage group pairing, block move scheduling, and block move execution.

Storage group classification

The HDFS Balancer first invokes the `getLiveDatanodeStorageReport` rpc to the Namenode to the storage report for all the storages in all Datanodes. The storage report contains storage utilization information such as capacity, dfs used space, remaining space, and so forth, for each storage in each DataNode.

A Datanode can contain multiple storages and the storages can have different storage types. A storage group $G_{i,T}$ is defined to be the group of all the storages with the same storage type T in Datanode i . For example, $G_{i,DISK}$ is the storage group of all the DISK storages in Datanode i . For each storage type T in each DataNode i , HDFS Balancer computes Storage Group Utilization (%)

$U_{i,T} = 100\% \text{ (storage group used space) / (storage group capacity),}$

and Average Utilization (%)

$U_{avg,T} = 100\% * \text{(sum of all used spaces) / (sum of all capacities).}$

Let $\#$ be the threshold parameter (default is 10%) and $G_{i,T}$ be the storage group with storage type T in DataNode I .

Average + Threshold	Over-Utilized:	$\{G_{i,T} : U_{avg,T} + \# < U_{i,T}\},$

Average	Above-Average:	$\{G_{i,T} : U_{avg,T} < U_{i,T} \leq U_{avg,T} + \#\},$

Average - Threshold	Below-Average:	$\{G_{i,T} : U_{avg,T} - \# \leq U_{i,T} \leq U_{avg,T}\},$

	Under-Utilized:	$\{G_{i,T} : U_{i,T} < U_{avg,T} - \#\}.$

A storage group is over-utilized or under-utilized if its utilization is larger or smaller than the difference between the average and the threshold. A storage group is above-average or below-average if its utilization is larger or smaller than average but within the threshold.

If there are no over-utilized storages and no under-utilized storages, the cluster is said to be balanced. The HDFS Balancer terminates with a SUCCESS state. Otherwise, it continues with storage group pairing.

Storage group pairing

The HDFS Balancer selects over-utilized or above-average storage as source storage, and under-utilized or below-average storage as target storage. It pairs a source storage group with a target storage group (source # target) in a priority order depending on whether or not the source and the target storage reside in the same rack.

The Balancer uses the following priority orders for pairing storage groups from the source and the target.

- Same-Rack (where the source and the target storage reside in the same rack)
 - Over-Utilized # Under-Utilize
 - Over-Utilized # Below-Average
 - Above-Average # Under-Utilized
- Any (where the source and target storage do not reside in the same rack)
 - Over-Utilized # Under-Utilized
 - Over-Utilized # Below-Average
 - Above-Average # Under-Utilized

Block move scheduling

For each source-target pair, the HDFS Balancer chooses block replicas from the source storage groups and schedules block moves.

A block replica in a source DataNode is a good candidate if it satisfies all of the following conditions:

- The storage type of the block replica in the source DataNode is the same as the target storage type.
- The storage type of the block replica is not already scheduled.
- The target does not already have the same block replica.
- The number of racks of the block is not reduced after the move.

Logically, the HDFS Balancer schedules a block replica to be “moved” from a source storage group to a target storage group. In practice, a block usually has multiple replicas. The block move can be done by first copying the replica from a proxy, which can be any storage group containing one of the replicas of the block, to the target storage group, and then deleting the replica in the source storage group.

After a candidate block in the source DataNode is specified, the HDFS Balancer selects a storage group containing the same replica as the proxy. The HDFS Balancer selects the closest storage group as the proxy in order to minimize the network traffic.

When it is impossible to schedule a move, the HDFS Balancer terminates with a `NO_MOVE_BLOCK` exit status.

Block move execution

The HDFS Balancer dispatches a scheduled block move by invoking the `DataTransferProtocol.replaceBlock(..)` method to the target DataNode.

The Balancer specifies the proxy, and the source as delete-hint in the method call. The target DataNode copies the replica directly from the proxy to its local storage. When the copying process has been completed, the target DataNode reports the new replica to the NameNode with the delete-hint. NameNode uses delete-hint to delete the extra replica, that is, delete the replica stored in the source.

After all block moves are dispatched, the HDFS Balancer waits until all the moves are completed. Then, the HDFS Balancer continues running a new iteration and repeats all of the steps. If the scheduled moves fail for 5 consecutive iterations, the HDFS Balancer terminates with a `NO_MOVE_PROGRESS` exit status.

Exit statuses for the HDFS Balancer

The HDFS Balancer concludes a cluster balancing operation with a specific exit status that indicates whether the operation succeeded or failed, with supporting reasons.

Table 4: Exit statuses for the HDFS balancer

Status	Value	Description
SUCCESS	0	The cluster is balanced. There are no over or under-utilized storages, with regard to the specified threshold.
ALREADY_RUNNING	-1	Another HDFS Balancer is running.
NO_MOVE_BLOCK	-2	The HDFS Balancer is not able to schedule a move.
NO_MOVE_PROGRESS	-3	All of the scheduled moves have failed for 5 consecutive iterations.
IO_EXCEPTION	-4	An IOException occurred.
ILLEGAL_ARGUMENTS	-5	An invalid argument in the command or configuration occurred.
INTERRUPTED	-6	The HDFS Balancer process was interrupted.
UNFINALIZED_UPGRADE	-7	The cluster is being upgraded.

HDFS heap sizing

You can provision an HDFS cluster for optimal performance based on the desired storage capacity.

Component	Memory	CPU	Disk
JournalNode	1 GB (default) Set this value using the Java Heap Size of JournalNode in Bytes HDFS configuration property.	1 core minimum	1 dedicated disk

Component	Memory	CPU	Disk
NameNode	<ul style="list-style-type: none"> Minimum: 1 GB (for proof-of-concept deployments) Add an additional 1 GB for each additional 1,000,000 blocks <p>Snapshots and encryption can increase the required heap memory.</p> <p>See <i>Sizing NameNode Heap Memory</i>.</p> <p>Set this value using the Java Heap Size of NameNode in Bytes HDFS configuration property.</p>	Minimum of 4 dedicated cores; more may be required for larger clusters	<ul style="list-style-type: none"> Minimum of 2 dedicated disks for metadata 1 dedicated disk for log files (This disk may be shared with the operating system.) Maximum disks: 4
DataNode	<p>Minimum: 4 GB</p> <p>Maximum: 8 GB</p> <p>Increase the memory for higher replica counts or a higher number of blocks per DataNode. When increasing the memory, Cloudera recommends an additional 1 GB of memory for every 1 million replicas above 4 million on the DataNodes. For example, 5 million replicas require 5 GB of memory.</p> <p>Set this value using the Java Heap Size of DataNode in Bytes HDFS configuration property.</p>	Minimum: 4 cores. Add more cores for highly active clusters.	<p>Minimum: 4</p> <p>Maximum: 24</p> <p>The maximum acceptable size will vary depending upon how large average block size is. The DN's scalability limits are mostly a function of the number of replicas per DN, not the overall number of bytes stored. That said, having ultra-dense DNs will affect recovery times in the event of machine or rack failure. Cloudera does not support exceeding 100 TB per data node. You could use 12 x 8 TB spindles or 24 x 4TB spindles. Cloudera does not support drives larger than 8 TB.</p>



Warning: Running Runtime on storage platforms other than direct-attached physical disks can provide suboptimal performance. Cloudera Enterprise and the majority of the Hadoop platform are optimized to provide high performance by distributing work across a cluster that can utilize data locality and fast local I/O.

Optimizing performance

You can consider the following options to optimize the performance of an HDFS cluster: swapping disk drives on a DataNode, caching data, specifying racks for hosts, customizing HDFS, optimizing NameNode disk space with Hadoop archives, identifying slow DataNodes and improving them, optimizing small write operations by using DataNode memory as storage, and implementing short-circuit reads.

Improving performance with centralized cache management

Centralized cache management enables you to specify paths to directories that are cached by HDFS, thereby improving performance for applications that repeatedly access the same data.

Centralized cache management in HDFS is an explicit caching mechanism. The NameNode communicates with DataNodes that have the required data blocks available on disk, and instructs the DataNodes to cache the blocks in off-heap caches.

Benefits of centralized cache management in HDFS

Centralized cache management in HDFS offers many significant advantages such as explicit pinning, querying cached blocks for task placement, and improving cluster memory utilization.

- Explicit pinning prevents frequently used data from being evicted from memory. This is particularly important when the size of the working set exceeds the size of main memory, which is common for many HDFS workloads.
- Because DataNode caches are managed by the NameNode, applications can query the set of cached block locations when making task placement decisions. Co-locating a task with a cached block replica improves read performance.
- When a block has been cached by a DataNode, clients can use a more efficient zero-copy read API. Since checksum verification of cached data is done once by the DataNode, clients can incur essentially zero overhead when using this new API.
- Centralized caching can improve overall cluster memory utilization. When relying on the operating system buffer cache on each DataNode, repeated reads of a block will result in all n replicas of the block being pulled into the buffer cache. With centralized cache management, you can explicitly pin only m of the n replicas, thereby saving $n-m$ memory.

Use cases for centralized cache management


Centralized cache management is useful for files that are accessed repeatedly and for mixed workloads that have performance SLAs.

- Files that are accessed repeatedly: For example, a small fact table in Hive that is often used for joins is a good candidate for caching. Conversely, caching the input of a once-yearly reporting query is probably less useful, since the historical data might only be read once.
- Mixed workloads with performance SLAs: Caching the working set of a high priority workload ensures that it does not compete with low priority workloads for disk I/O.

Centralized cache management architecture

In a centralized cache management, the NameNode is responsible for coordinating all of the DataNode off-heap caches in the cluster. The NameNode periodically receives a cache report from each DataNode. The cache report describes all of the blocks cached on the DataNode. The NameNode manages DataNode caches by piggy-backing cache and uncache commands on the DataNode heartbeat.

The following figure illustrates the centralized cached management architecture.



The NameNode queries its set of cache directives to determine which paths should be cached. Cache directives are persistently stored in the fsimage and edit logs, and can be added, removed, and modified through Java and command-line APIs. The NameNode also stores a set of cache pools, which are administrative entities used to group cache directives together for resource management, and to enforce permissions.

The NameNode periodically re-scans the namespace and active cache directives to determine which blocks need to be cached or uncached, and assigns caching work to DataNodes. Re-scans can also be triggered by user actions such as adding or removing a cache directive or removing a cache pool.

Cache blocks that are under construction, corrupt, or otherwise incomplete are not cached. If a Cache directive covers a symlink, the symlink target is not cached.

Caching can only be applied to directories and files.

Related Information

[Caching terminology](#)

Caching terminology

A cache directive defines the path to cache while a cache pool manages groups of cache directives.

Cache directive

Defines the path to be cached. Paths can point either directories or files. Directories are cached non-recursively, meaning only files in the first-level listing of the directory will be cached.

Cache directives also specify additional parameters, such as the cache replication factor and expiration time. The replication factor specifies the number of block replicas to cache. If multiple cache directives refer to the same file, the maximum cache replication factor is applied.

The expiration time is specified on the command line as a time-to-live (TTL), which represents a relative expiration time in the future. After a cache directive expires, it is no longer taken into consideration by the NameNode when making caching decisions.

Cache pool

An administrative entity that manages groups of cache directives. Cache pools have UNIX-like permissions that restrict which users and groups have access to the pool. Write permissions allow users to add and remove cache directives to the pool. Read permissions allow users to list the Cache Directives in a pool, as well as additional metadata. Execute permissions are unused.

Cache pools are also used for resource management. Cache pools can enforce a maximum memory limit, which restricts the aggregate number of bytes that can be cached by directives in the pool. Normally, the sum of the pool limits will approximately equal the amount of aggregate memory reserved for HDFS caching on the cluster. Cache pools also track a number of statistics to help cluster users track what is currently cached, and to determine what else should be cached.

Cache pools can also enforce a maximum time-to-live. This restricts the maximum expiration time of directives being added to the pool.

Related Information

[Centralized cache management architecture](#)

Properties for configuring centralized caching

You must enable JNI to use centralized caching. In addition, you must configure various properties and consider the locked memory limit for configuring centralized caching.

Native libraries

In order to lock block files into memory, the DataNode relies on native JNI code found in libhadoop.so. Be sure to enable JNI if you are using HDFS centralized cache management.

Configuration properties

Configuration properties for centralized caching are specified in the `hdfs-site.xml` file.

Required properties

Only the following property is required:

- `dfs.datanode.max.locked.memory` This property determines the maximum amount of memory (in bytes) that a DataNode will use for caching. The "locked-in-memory size" `ulimit (ulimit -l)` of the DataNode user also needs to be increased to exceed this parameter (for more details, see the following section on). When setting this value, remember that you will need space in memory for other things as well, such as the DataNode and application JVM heaps, and the operating system page cache. Example:

```
<property>
  <name>dfs.datanode.max.locked.memory</name>
  <value>268435456</value>
</property>
```

Optional properties

The following properties are not required, but can be specified for tuning.

- `dfs.namenode.path.based.cache.refresh.interval.ms` The NameNode will use this value as the number of milliseconds between subsequent cache path re-scans. By default, this parameter is set to 300000, which is five minutes. Example:

```
<property>
  <name>dfs.namenode.path.based.cache.refresh.interval.ms</name>
  <value>300000</value>
</property>
```

- `dfs.time.between.resending.caching.directives.ms` The NameNode will use this value as the number of milliseconds between resending caching directives. Example:

```
<property>
  <name>dfs.time.between.resending.caching.directives.ms</name>
  <value>300000</value>
</property>
```

- `dfs.datanode.fsdatasetcache.max.threads.per.volume` The DataNode will use this value as the maximum number of threads per volume to use for caching new data. By default, this parameter is set to 4. Example:

```
<property>
  <name>dfs.datanode.fsdatasetcache.max.threads.per.volume</name>
  <value>4</value>
</property>
```

- `dfs.cachereport.intervalMsec` The DataNode will use this value as the number of milliseconds between sending a full report of its cache state to the NameNode. By default, this parameter is set to 10000, which is 10 seconds. Example:

```
<property>
  <name>dfs.cachereport.intervalMsec</name>
  <value>10000</value>
</property>
```

- `dfs.namenode.path.based.cache.block.map.allocation.percent` The percentage of the Java heap that will be allocated to the cached blocks map. The cached blocks map is a hash map that uses chained hashing. Smaller maps may be accessed more slowly if the number of cached blocks is large. Larger maps will consume more memory. The default value is 0.25 percent. Example:

```
<property>
  <name>dfs.namenode.path.based.cache.block.map.allocation.percent</name>
  <value>0.25</value>
</property>
```

OS limits

If you get the error "Cannot start datanode because the configured max locked memory size...is more than the datanode's available RLIMIT_MEMLOCK ulimit," this means that the operating system is imposing a lower limit on the amount of memory that you can lock than what you have configured. To fix this, you must adjust the `ulimit -l` value that the DataNode runs with. This value is usually configured in `/etc/security/limits.conf`, but this may vary depending on what operating system and distribution you are using.

You have correctly configured this value when you can run `ulimit -l` from the shell and get back either a higher value than what you have configured or the string "unlimited", which indicates that there is no limit. Typically, `ulimit -l` returns the memory lock limit in kilobytes (KB), but `dfs.datanode.max.locked.memory` must be specified in bytes.

For example, if the value of `dfs.datanode.max.locked.memory` is set to 128000 bytes:

```
<property>
  <name>dfs.datanode.max.locked.memory</name>
  <value>128000</value>
</property>
```

Set the memlock (max locked-in-memory address space) to a slightly higher value. For example, to set memlock to 130 KB (130,000 bytes) for the `hdfs` user, you would add the following line to `/etc/security/limits.conf`.

```
hdfs - memlock 130
```

Commands for using cache pools and directives

You can use the Command-Line Interface (CLI) to create, modify, and list cache pools and cache directives using the `hdfs cacheadmin` subcommand.

Cache Directives are identified by a unique, non-repeating, 64-bit integer ID. IDs will not be reused even if a Cache Directive is removed.

Cache Pools are identified by a unique string name.

You must first create a Cache Pool, and then add Cache Directives to the Cache Pool.

Cache Pool Commands

- `addPool` -- Adds a new Cache Pool.

Usage:

```
hdfs cacheadmin -addPool <name> [-owner <owner>] [-group <group>]
[-mode <mode>] [-limit <limit>] [-maxTtl <maxTtl>]
```

Options:

Table 5: Cache Pool Add Options

Option	Description
<name>	The name of the pool.
<owner>	The user name of the owner of the pool. Defaults to the current user.
<group>	The group that the pool is assigned to. Defaults to the primary group name of the current user.
<mode>	The UNIX-style permissions assigned to the pool. Permissions are specified in octal (e.g. 0755). Pool permissions are set to 0755 by default.
<limit>	The maximum number of bytes that can be cached by directives in the pool, in aggregate. By default, no limit is set.
<maxTtl>	The maximum allowed time-to-live for directives being added to the pool. This can be specified in seconds, minutes, hours, and days (e.g. 120s, 30m, 4h, 2d). Valid units are [smhd]. By default, no maximum is set. A value of "never" specifies that there is no limit.

- `modifyPool` -- Modifies the metadata of an existing Cache Pool.

Usage:

```
hdfs cacheadmin -modifyPool <name> [-owner <owner>] [-group <group>]
[-mode <mode>] [-limit <limit>] [-maxTtl <maxTtl>]
```

Options:

Table 6: Cache Pool Modify Options

Option	Description
<name>	The name of the pool to modify.
<owner>	The user name of the owner of the pool.
<group>	The group that the pool is assigned to.
<mode>	The UNIX-style permissions assigned to the pool. Permissions are specified in octal (e.g. 0755).
<limit>	The maximum number of bytes that can be cached by directives in the pool, in aggregate.

Option	Description
<maxTtl>	The maximum allowed time-to-live for directives being added to the pool. This can be specified in seconds, minutes, hours, and days (e.g. 120s, 30m, 4h, 2d). Valid units are [smdh]. By default, no maximum is set. A value of "never" specifies that there is no limit.

- **removePool --** Removes a Cache Pool. This command also "un-caches" paths that are associated with the pool.

Usage:

```
hdfs cacheadmin -removePool <name>
```

Options:

Table 7: Cache Pool Remove Options

Option	Description
<name>	The name of the Cache Pool to remove.

- **listPools --** Displays information about one or more Cache Pools, such as name, owner, group, permissions, and so on.

Usage:

```
hdfs cacheadmin -listPools [-stats] [<name>]
```

Options:

Table 8: Cache Pools List Options

Option	Description
-stats	Displays additional Cache Pool statistics.
<name>	If specified, lists only the named Cache Pool.

- **help --** Displays detailed information about a command.

Usage:

```
hdfs cacheadmin -help <command-name>
```

Options:

Table 9: Cache Pool Help Options

Option	Description
<command-name>	Displays detailed information for the specified command name. If no command name is specified, detailed help is displayed for all commands.

Cache Directive Commands

- **addDirective --** Adds a new Cache Directive.

Usage:

```
hdfs cacheadmin -addDirective -path <path> -pool <pool-name> [-force]
```



```
[-replication <replication>] [-ttl <time-to-live>]
```

Options:

Table 10: Cache Pool Add Directive Options

Option	Description
<path>	The path to the cache directory or file.
<pool-name>	The Cache Pool to which the Cache Directive will be added. You must have Write permission for the Cache Pool in order to add new directives.
-force	Skips checking of the Cache Pool resource limits.
<replication>	The cache replication factor to use. Default setting is 1.
<time-to-live>	How long the directive is valid. This can be specified in minutes, hours and days (e.g. 30m, 4h, 2d). Valid units are [smdh]. A value of "never" indicates a directive that never expires. If unspecified, the directive never expires.

- `removeDirective` -- Removes a Cache Directive.

Usage:

```
hdfs cacheadmin -removeDirective <id>
```

Options:

Table 11: Cache Pool Remove Directive Options

Option	Description
<id>	The ID of the Cache Directive to remove. You must have Write permission for the pool that the directive belongs to in order to remove it. You can use the <code>-listDirectives</code> command to display a list of Cache Directive IDs.

- `removeDirectives` -- Removes all of the Cache Directives in a specified path.

Usage:

```
hdfs cacheadmin -removeDirectives <path>
```

Options:

Table 12: Cache Pool Remove Directives Options

Option	Description
<path>	The path of the Cache Directives to remove. You must have Write permission for the pool that the directives belong to in order to remove them. You can use the <code>-listDirectives</code> command to display a list of Cache Directives.

- `listDirectives` -- Returns a list of Cache Directives.

Usage:

```
hdfs cacheadmin -listDirectives [-stats] [-path <path>] [-pool <pool>]
```

Options:

Table 13: Cache Pools List Directives Options

Option	Description
<path>	Lists only the Cache Directives in the specified path. If there is a Cache Directive in the <path> that belongs to a Cache Pool for which you do not have Read access, it will not be listed.
<pool>	Lists on the Cache Directives in the specified Cache Pool.
-stats	Lists path-based Cache Directive statistics.

Specifying racks for hosts

HDFS uses the network location of your deployment, such as hosts and racks, to place block replicas more intelligently and trade off performance and resilience. When placing jobs on hosts, CDP prefers within-rack transfers, where there is more bandwidth available, to off-rack transfers. These computations are performed with the assistance of rack awareness scripts. Cloudera Manager includes internal rack awareness scripts, but you must specify the racks where the hosts in your cluster are located.

Procedure

1. In Cloudera Manager, go to **Hosts All Hosts**.
2. Select the checkboxes corresponding to the hosts for which you want to specify the racks.
3. Select **Actions for Selected >Assign Rack**.
4. Enter a rack name or ID that begins with a slash (/); for example, /rack123 or /aisle1/rack123, and click **Confirm**.
5. Restart the affected services.

Viewing racks assigned to cluster hosts

You can view details of the racks assigned to the cluster hosts from the **All Hosts** page of Cloudera Manager.

Procedure

1. In Cloudera Manager, go to **Hosts All Hosts**.
2. Expand the **Columns** drop-down list, and select **Rack**.

The table displaying the host details now includes a **Rack** column next to **Name**, and lists the racks assigned to every host.

Editing rack assignments for hosts

You can edit individual host rack assignments from Cloudera Manager.

Procedure

1. In Cloudera Manager, go to **Hosts All Hosts**.
2. Select the checkboxes corresponding to the hosts for which you want to specify the racks.
3. Select **Actions for Selected Assign Rack**.
4. Enter a rack name different from that assigned to the selected hosts, and click **Confirm**.
5. Restart the affected services.

Customizing HDFS

You can use the `dfs.user.home.base.dir` property to customize the HDFS home directory. In addition, you can configure properties to control the size of the directory that holds the NameNode edits directory.

Customize the HDFS home directory

By default, the HDFS home directory is set to `/user/<user_name>`. Use the `dfs.user.home.dir.prefix` property to customize the HDFS home directory.

Procedure

In `hdfs-site.xml` file, set the value of the `dfs.user.home.dir.prefix` property.

```
<property>
  <name>dfs.user.home.dir.prefix</name>
  <value>/user</value>
  <description>Base directory of user home.</description>
</property>
```

In the example, `<value>` is the path to the new home directory.

Properties to set the size of the NameNode edits directory

You can configure the `dfs.namenode.num.checkpoints.retained` and `dfs.namenode.num.extra.edits.retained` properties to control the size of the directory that holds the NameNode edits directory.

- `dfs.namenode.num.checkpoints.retained`: The number of image checkpoint files that are retained in storage directories. All edit logs necessary to recover an up-to-date namespace from the oldest retained checkpoint are also retained.
- `dfs.namenode.num.extra.edits.retained`: The number of extra transactions that should be retained beyond what is minimally necessary for a NameNode restart. This can be useful for audit purposes, or for an HA setup where a remote Standby Node may have been offline for some time and require a longer backlog of retained edits in order to start again.

Optimizing NameNode disk space with Hadoop archives

Hadoop Archives (HAR) are special format archives that efficiently pack small files into HDFS blocks.

The Hadoop Distributed File System (HDFS) is designed to store and process large data sets, but HDFS can be less efficient when storing a large number of small files. When there are many small files stored in HDFS, these small files occupy a large portion of the namespace. As a result, disk space is under-utilized because of the namespace limitation.

Hadoop Archives (HAR) can be used to address the namespace limitations associated with storing many small files. A Hadoop Archive packs small files into HDFS blocks more efficiently, thereby reducing NameNode memory usage while still allowing transparent access to files. Hadoop Archives are also compatible with MapReduce, allowing transparent access to the original files by MapReduce jobs.

Overview of Hadoop archives

Storing a large number of small files in HDFS leads to inefficient utilization of space – the namespace is overutilized while the disk space might be underutilized. Hadoop Archives (HAR) address this limitation by efficiently packing small files into large files without impacting the file access.

The Hadoop Distributed File System (HDFS) is designed to store and process large (terabytes) data sets. For example, a large production cluster may have 14 PB of disk space and store 60 million files.

However, storing a large number of small files in HDFS is inefficient. A file is generally considered to be "small" when its size is substantially less than the HDFS block size. Files and blocks are name objects in HDFS, meaning that they occupy namespace (space on the NameNode). The namespace capacity of the system is therefore limited by the physical memory of the NameNode.

When there are many small files stored in the system, these small files occupy a large portion of the namespace. As a consequence, the disk space is underutilized because of the namespace limitation. In one real-world example, a production cluster had 57 million files less than 256 MB in size, with each of these files taking up one block on the NameNode. These small files used up 95% of the namespace but occupied only 30% of the cluster disk space.

Hadoop Archives (HAR) can be used to address the namespace limitations associated with storing many small files. HAR packs a number of small files into large files so that the original files can be accessed transparently (without expanding the files).

HAR increases the scalability of the system by reducing the namespace usage and decreasing the operation load in the NameNode. This improvement is orthogonal to memory optimization in the NameNode and distributing namespace management across multiple NameNodes.

Hadoop Archive is also compatible with MapReduce — it allows parallel access to the original files by MapReduce jobs.

Hadoop archive components

You can use the Hadoop archiving tool to create Hadoop Archives (HAR). The Hadoop Archive is integrated with the Hadoop file system interface. Files in a HAR are exposed transparently to users. File data in a HAR is stored in multipart files, which are indexed to retain the original separation of data.

Hadoop archiving tool

Hadoop Archives can be created using the Hadoop archiving tool. The archiving tool uses MapReduce to efficiently create Hadoop Archives in parallel. The tool can be invoked using the command:

```
hadoop archive -archiveName name -p <parent> <src>* <dest>
```

A list of files is generated by traversing the source directories recursively, and then the list is split into map task inputs. Each map task creates a part file (about 2 GB, configurable) from a subset of the source files and outputs the metadata. Finally, a reduce task collects metadata and generates the index files.

HAR file system

Most archival systems, such as tar, are tools for archiving and de-archiving. Generally, they do not fit into the actual file system layer and hence are not transparent to the application writer in that the archives must be expanded before use.

The Hadoop Archive is integrated with the Hadoop file system interface. The `HarFileSystem` implements the `FileSystem` interface and provides access via the `har://` scheme. This exposes the archived files and directory tree structures transparently to users. Files in a HAR can be accessed directly without expanding them.

For example, if we have the following command to copy an HDFS file to a local directory:

```
hdfs dfs -get hdfs://namenode/foo/file-1 localdir
```

Suppose a Hadoop Archive `bar.har` is created from the `foo` directory. With the HAR, the command to copy the original file becomes:

```
hdfs dfs -get har://namenode/bar.har/foo/file-1 localdir
```

Users only need to change the URI paths. Alternatively, users may choose to create a symbolic link (from `hdfs://namenode/foo` to `har://namenode/bar.har/foo` in the example above), and then even the URIs do not need to be changed. In either case, `HarFileSystem` will be invoked automatically to provide access to the files in the HAR.

Because of this transparent layer, HAR is compatible with the Hadoop APIs, MapReduce, the FS shell command-line interface, and higher-level applications such as Pig, Zebra, Streaming, Pipes, and DistCp.

HAR format data model

The Hadoop Archive data format has the following layout:

```
foo.har/_masterindex //stores hashes and offsets
foo.har/_index //stores file statuses
foo.har/part-[1..n] //stores actual file data
```

The file data is stored in multipart files, which are indexed in order to retain the original separation of data. Moreover, the file parts can be accessed in parallel by MapReduce programs. The index files also record the original directory tree structures and file status.

Create a Hadoop archive

Use the `hadoop archive` command to invoke the Hadoop archiving tool.

Procedure

Run the `hadoop archive` command by specifying the archive name to create, the parent directory relative to the archive location, the source files to archive, and the destination archive location.

```
hadoop archive -archiveName name -p <parent> <src>* <dest>
```

The archive name must have a `.har` extension



Note:

- Archiving does not delete the source files. If you want to delete the input files after creating an archive to reduce namespace, you must manually delete the source files.
- Although the `hadoop archive` command can be run from the host file system, the archive file is created in the HDFS file system from directories that exist in HDFS. If you reference a directory on the host file system and not HDFS, the system displays the following error:

```
The resolved paths set is empty. Please check whether the srcPaths exist, where srcPaths
= [</directory/path>]
```

Example

Consider the following example of archiving two files:

```
hadoop archive -archiveName foo.har -p /user/hadoop dir1 dir2 /user/zoo
```

This example creates an archive using `/user/hadoop` as the relative archive directory. The directories `/user/hadoop/dir1` and `/user/hadoop/dir2` will be archived in the `/user/zoo/foo.har` archive.

List files in Hadoop archives

Use the `hdfs dfs -ls` command to list files in Hadoop archives.

Procedure

Run the `hdfs dfs -ls` command by specifying the archive directory location.

To specify the directories in an archive directory `foo.har` located in `/usr/zoo`, run the following command:

```
hdfs dfs -ls har:///user/zoo/foo.har/
```

Assuming the archive directory `foo.har` contains two directories `dir1` and `dir2`, the command returns the following

```
har:///user/zoo/foo.har/dir1
har:///user/zoo/foo.har/dir2
```

**Note:**

Consider an archive created using the following command:

```
hadoop archive -archiveName foo.har -p /user/ hadoop/dir1 hadoop/dir2 /
user/zoo
```

If you list the files of the archive created in the preceding command, the command returns the following:

```
har:///user/zoo/foo.har/hadoop
har:///user/zoo/foo.har/hadoop/dir1
har:///user/zoo/foo.har/hadoop/dir2
```

Note that the modified parent argument causes the files to be archived relative to `/user/`.

Format for using Hadoop archives with MapReduce

To use Hadoop Archives with MapReduce, you must reference files differently than you would with the default file system. If you have a Hadoop Archive stored in HDFS in `/user/zoo/foo.har`, you must specify the input directory as `har:///user/zoo/foo.har` to use it as a MapReduce input.

Because Hadoop Archives are exposed as a file system, MapReduce can use all of the logical input files in Hadoop Archives as input.

Detecting slow DataNodes

Slow DataNodes in a CDP Private Cloud Base cluster can negatively impact the cluster performance. Therefore, HDFS provides a mechanism to detect and report slow DataNodes that have a negative impact on the performance of the cluster.

HDFS is designed to detect and recover from complete failure of DataNodes:

- There is no single point of failure.
- Automatic NameNode failover takes only a few seconds.
- Because data replication can be massively parallelized in large clusters, recovery from DataNode loss occurs within minutes.
- Most jobs are not affected by DataNode failures.

However, partial failures can negatively affect the performance of running DataNodes:

- Slow network connection due to a failing or misconfigured adapter.
- Bad OS or JVM settings that affect service performance.
- Slow hard disk.
- Bad disk controller.

Slow DataNodes can have a significant impact on cluster performance. A slow DataNode may continue sending heartbeats successfully, and the NameNode will keep redirecting clients to slow DataNodes. HDFS DataNode monitoring provides detection and reporting of slow DataNodes that negatively affect cluster performance.

Enable disk IO statistics

Disk IO statistics are disabled by default. To enable disk IO statistics on a CDP Private Cloud Base cluster, you must set the file IO sampling percentage to a non-zero value in the `hdfs-site.xml` file.

Procedure

1. Set the `dfs.datanode.fileio.profiling.sampling.percentage` property to a non-zero value in `hdfs-site.xml`.

```
<property>
  <name>dfs.datanode.fileio.profiling.sampling.percentage</name>
  <value>25</value>
</property>
```



Note: Sampling disk IO might have a minor impact on cluster performance.

2. Access the disk IO statistics from the NameNode JMX page at `http://<namenode_host>:50070/jmx`. In the following JMX output example, the time unit is milliseconds, and the disk is healthy because the IO latencies are low:

```
"name" : "Hadoop:service=DataNode,name=DataNodeVolume-/data/disk2/dfs/
data/",
"modelerType" : "DataNodeVolume-/data/disk2/dfs/data/",
"tag.Context" : "dfs",
"tag.Hostname" : "n001.hdfs.example.com",
"TotalMetadataOperations" : 67,
"MetadataOperationRateAvgTime" : 0.08955223880597014,
...
"WriteIoRateNumOps" : 7321,
"WriteIoRateAvgTime" : 0.050812730501297636
```

Enable detection of slow DataNodes

When slow DataNode detection is enabled, DataNodes collect latency statistics on their peers during write pipelines, and use periodic outlier detection to determine slow peers. The NameNode aggregates reports from all DataNodes and flags potentially slow nodes. Slow DataNode detection is disabled by default.

Procedure

1. To enable slow DataNode detection, set the value of the `dfs.datanode.peer.stats.enabled` property to `true` in `hdfs-site.xml`.

```
<property>
  <name>dfs.datanode.peer.stats.enabled</name>
  <value>true</value>
</property>
```

2. Access the slow DataNode statistics either from the NameNode JMX page at `http://<namenode_host>:50070/jmx` or from the DataNode JMX page at `http://<datanode_host>:50075/jmx`. In the following JMX output example, the time unit is milliseconds, and the peer DataNodes are healthy because the latencies are in milliseconds:

```
"name" : "Hadoop:service=DataNode,name=DataNodeInfo",
"modelerType" : "org.apache.hadoop.hdfs.server.datanode.DataNode",
"SendPacketDownstreamAvgInfo" : "{
  \"[192.168.7.202:50075]RollingAvgTime\" : 1.4476967370441458,
  \"[192.168.7.201:50075]RollingAvgTime\" : 1.5569170444798432
}"
```

Allocating DataNode memory as storage

HDFS supports efficient writes of large data sets to durable storage, and also provides reliable access to the data. This works well for batch jobs that write large amounts of persistent data. Emerging classes of applications are driving use cases for writing smaller amounts of temporary data. Using DataNode memory as storage addresses the use case of applications that want to write relatively small amounts of intermediate data sets with low latency.

Writing block data to memory reduces durability, as data can be lost due to process restart before it is saved to disk. HDFS attempts to save replica data to disk in a timely manner to reduce the window of possible data loss.

DataNode memory is referenced using the `RAM_DISK` storage type and the `LAZY_PERSIST` storage policy.

HDFS storage types

The storage type associated with HDFS identifies the underlying storage media.

HDFS supports the following storage types:

- **ARCHIVE** - Archival storage is for very dense storage and is useful for rarely accessed data. This storage type is typically cheaper per TB than normal hard disks.
- **DISK** - Hard disk drives are relatively inexpensive and provide sequential I/O performance. This is the default storage type.
- **SSD** - Solid state drives are useful for storing hot data and I/O-intensive applications.
- **RAM_DISK** - This special in-memory storage type is used to accelerate low-durability, single-replica writes.

When you add the DataNode Data Directory, you can specify which type of storage it uses, by prefixing the path with the storage type, in brackets. If you do not specify a storage type, it is assumed to be `DISK`.

LAZY_PERSIST memory storage policy

Use the `LAZY_PERSIST` storage policy to store data blocks on the configured DataNode memory.

For `LAZY_PERSIST`, the first replica is stored on `RAM_DISK` (DataNode memory), and the remaining replicas are stored on `DISK`. The fallback storage for both creation and replication is `DISK`.

The following table summarizes these replication policies:

Policy ID	Policy Name	Block Placement (for n replicas)	Fallback storage for creation	Fallback storage for replication
15	LAZY_PERSIST	RAM_DISK: 1, DISK:n-1	DISK	DISK

Configure DataNode memory as storage

Configuring memory on a DataNode as storage requires you to shut down the particular DataNode, set `RAM_DISK` as the storage type, set the `LAZY_PERSIST` storage policy to store data, and then start the DataNode.

Procedure

1. Shut down the DataNode.
2. Use required mount commands to allocate a certain portion of the DataNode memory as storage. The following example shows how you can allocate 2GB memory for use by HDFS.

```
sudo mkdir -p /mnt/hdfsramdisk
sudo mount -t tmpfs -o size=2048m tmpfs /mnt/hdfsramdisk
sudo mkdir -p /usr/lib/hadoop-hdfs
```


3. Assign the RAM_DISK storage type to ensure that HDFS can assign data to the DataNode memory configured as storage.

To specify the DataNode as RAM_DISK storage, insert [RAM_DISK] at the beginning of the local file system mount path and add it to the `dfs.name.dir` property in `hdfs-default.xml`.

The following example shows the updated mount path values for `dfs.datanode.data.dir`

```
<property>
  <name>dfs.datanode.data.dir</name>
  <value>file:///grid/3/aa/hdfs/data/,[RAM_DISK]file:///mnt/hdfsramdisk/</value>
</property>
```

4. Set the LAZY_PERSIST storage policy to store data on the configured DataNode memory.

The following example shows how you can use the `hdfs dfsadmin -getStoragePolicy` command to configure the LAZY_PERSIST storage policy:

```
hdfs dfsadmin -getStoragePolicy /memory1 LAZY_PERSIST
```



Note: When you update a storage policy setting on a file or directory, the new policy is *not* automatically enforced. You must use the HDFS mover data migration tool to actually move blocks as specified by the new storage policy.

5. Start the DataNode.
6. Use the HDFS mover tool to move data blocks according to the specified storage policy.

The HDFS mover data migration tool scans the specified files in HDFS and verifies if the block placement satisfies the storage policy. For the blocks that violate the storage policy, the tool moves the replicas to a different storage type in order to fulfill the storage policy requirements.

Improving performance with short-circuit local reads

In HDFS, reads normally go through the DataNode. Thus, when a client asks the DataNode to read a file, the DataNode reads that file off of the disk and sends the data to the client over a TCP socket. "Short-circuit" reads bypass the DataNode, allowing the client to read the file directly. This is only possible in cases where the client is co-located with the data. Short-circuit reads provide a substantial performance boost to many applications.

Prerequisites for configuring short-circuit local reads

To configure short-circuit local reads, you must enable `libhadoop.so`.

See the Native Libraries Guide for details on enabling this library.

Related Information

[Native Libraries Guide](#)

Properties for configuring short-circuit local reads on HDFS

To configure short-circuit local reads, you must add various properties to the `hdfs-site.xml` file. Short-circuit local reads must be configured on both the DataNode and the client.

Property Name	Property Value	Description
<code>dfs.client.read.shortcircuit</code>	<code>true</code>	Set this to true to enable short-circuit local reads.

Property Name	Property Value	Description
dfs.domain.socket.path	/var/lib/hadoop-hdfs/ dn_socket	<p>The path to the domain socket. Short-circuit reads make use of a UNIX domain socket. This is a special path in the file system that allows the client and the DataNodes to communicate. You will need to set a path to this socket. The DataNode needs to be able to create this path. On the other hand, it should not be possible for any user except the hdfs user or root to create this path. For this reason, paths under /var/run or /var/lib are often used.</p> <p>In the file system that allows the client and the DataNodes to communicate. You will need to set a path to this socket. The DataNode needs to be able to create this path. On the other hand, it should not be possible for any user except the hdfs user or root to create this path. For this reason, paths under /var/run or /var/lib are often used.</p>
dfs.client.domain.socket.data.traffic	false	<p>This property controls whether or not normal data traffic will be passed through the UNIX domain socket. It is recommended that you set the value of this property to false.</p> <p>Unusual data traffic will be passed through the UNIX domain socket.</p>
dfs.client.use.legacy.blockreader.local	false	Setting this value to false specifies that the new version (based on HDFS-347) of the short-circuit reader is used. Setting this value to true would mean that the legacy short-circuit reader would be used.
dfs.datanode.hdfs-blocks-metadata.enabled	true	Boolean which enables back-end DataNode-side support for the experimental DistributedFileSystem#getFile VBlockStorageLocationsAPI.
dfs.client.file-block-storage-locations.timeout	60	<p>Timeout (in seconds) for the parallel RPCs made in DistributedFileSystem #getFileBlockStorageLocations().</p> <p>This property is deprecated but is still supported for backward compatibility</p>
dfs.client.file-block-storage-locations.timeout.millis	60000	<p>Timeout (in milliseconds) for the parallel RPCs made in DistributedFileSystem #getFileBlockStorageLocations().</p> <p>This property replaces dfs.client.file-block-storage-locations.timeout, and offers a finer level of granularity.</p>
dfs.client.read.shortcircuit.skip.checksum	false	If this configuration parameter is set, short-circuit local reads will skip checksums. This is normally not recommended, but it may be useful for special setups. You might consider using this if you are doing your own checksumming outside of HDFS.

Property Name	Property Value	Description
dfs.client.read.shortcircuit.streams.cache.size	256	The DFSClient maintains a cache of recently opened file descriptors. This parameter controls the size of that cache. Setting this higher will use more file descriptors, but potentially provide better performance on workloads involving many seeks.
dfs.client.read.shortcircuit.streams.cache.expiry.ms	300000	This controls the minimum amount of time (in milliseconds) file descriptors need to sit in the client cache context before they can be closed for being inactive for too long.

The XML for these entries:

```
<configuration>
<property>
  <name>dfs.client.read.shortcircuit</name>
  <value>true</value>
</property>

<property>
  <name>dfs.domain.socket.path</name>
  <value>/var/lib/hadoop-hdfs/dn_socket</value>
</property>

<property>
  <name>dfs.client.domain.socket.data.traffic</name>
  <value>>false</value>
</property>

<property>
  <name>dfs.client.use.legacy.blockreader.local</name>
  <value>>false</value>
</property>

<property>
  <name>dfs.datanode.hdfs-blocks-metadata.enabled</name>
  <value>true</value>
</property>

<property>
  <name>dfs.client.file-block-storage-locations.timeout.millis</name>
  <value>60000</value>
</property>

<property>
  <name>dfs.client.read.shortcircuit.skip.checksum</name>
  <value>>false</value>
</property>

<property>
  <name>dfs.client.read.shortcircuit.streams.cache.size</name>
  <value>256</value>
</property>

<property>
  <name>dfs.client.read.shortcircuit.streams.cache.expiry.ms</name>
  <value>300000</value>
</property>
</configuration>
```

Configure mountable HDFS

CDP includes a FUSE (Filesystem in UserSpace) interface into HDFS. The `hadoop-hdfs-fuse` command enables you to use your HDFS cluster as if it were a traditional filesystem on Linux.

Before you begin

You must have a working HDFS cluster and know the hostname and port that your NameNode exposes.

Procedure

1. Set up and test your mount point depending on whether you have an HA or a non-HA configuration.

- To set up and test your mount point in a non-HA installation:

```
$ mkdir -p <mount_point>
$ hadoop-fuse-dfs dfs://<name_node_hostname>:<namenode_port> <mount_point>
```

where *namenode_port* is the NameNode's RPC port, `dfs.namenode.servicerpc-address`.

- To set up and test your mount point in an HA installation:

```
mkdir -p <mount_point>
hadoop-fuse-dfs dfs://<nameservice_id> <mount_point>
```

- where *nameservice_id* is the value of `fs.defaultFS`. In this case the port defined for `dfs.namenode.rpc-address`. `[nameservice ID].[name node ID]` is used automatically.

You can now run operations as if they are on your mount point.



Note:

To find its configuration directory, `hadoop-fuse-dfs` uses the `HADOOP_CONF_DIR` configured at the time the mount command is invoked.

2. Clean up your test mount by using the `umount` command.

```
umount <mount_point>
```

What to do next

You can now add a permanent HDFS mount which persists through reboots.

Add HDFS system mount

After installing `hadoop-hdfs-fuse` and configuring the test mount, you can add a persistent HDFS system mount.

About this task

Procedure

1. Open `/etc/fstab` and append information about the hostname, the NameNode port, and the mount point in the specified format.

The format is as follows:

```
hadoop-fuse-dfs#dfs://<name_node_hostname>:<namenode_port> <mount_point>  
fuse allow_other,usetrash,rw 2 0
```

For example:

```
hadoop-fuse-dfs#dfs://localhost:8020 /mnt/hdfs fuse allow_other,usetrash  
,rw 2 0
```



Note:

In an HA deployment, use the HDFS nameservice instead of the NameNode URI; that is, use the value of `dfs.nameservices` in `hdfs-site.xml`.

2. Test to make sure the mount point is working properly

```
mount <mount_point>
```

Results

Your system is now configured to allow you to use the `ls` command and use that mount point as if it were a normal system disk.

For more information, see the help for `hadoop-fuse-dfs`:

```
hadoop-fuse-dfs --help
```

Optimize mountable HDFS

You can consider certain recommendations to optimize the mountable HDFS.

- Cloudera recommends that you use the `-obig_writes` option on kernels later than 2.6.26. This option allows for better performance of writes.
- By default, the package installation creates the `/etc/default/hadoop-fuse` file with a maximum heap size of 128 MB. You might need to change the JVM minimum and maximum heap size for better performance. For example:

```
export LIBHDFS_OPTS="-Xms64m -Xmx256m"
```

Be careful not to set the minimum to a higher value than the maximum.

Configuring Proxy Users to Access HDFS

Hadoop allows you to configure proxy users to submit jobs or access HDFS on behalf of other users; this is called *impersonation*. When you enable impersonation, any jobs submitted using a proxy are executed with the impersonated user's existing privilege levels rather than those of a superuser (such as `hdfs`).

All proxy users are configured in one location, `core-site.xml`, for Hadoop administrators to implement centralized access control.

To configure proxy users, set the `hadoop.proxyuser.<proxy_user>.hosts`, `hadoop.proxyuser.<proxy_group>.groups`, and `hadoop.proxyuser.<proxy_user>.users` in `core-site.xml` properties.

For example, to allow user alice to impersonate a user belonging to group_a and group_b, set `hadoop.proxyuser.<proxy_group>.groups` as follows:

```
<property>
  <name>hadoop.proxyuser.alice.groups</name>
  <value>group_a,group_b</value>
</property>
```

To limit the hosts from which impersonated connections are allowed, use `hadoop.proxyuser.<proxy_user>.hosts`. For example, to allow user alice impersonated connections only from host_a and host_b:

```
<property>
  <name>hadoop.proxyuser.alice.hosts</name>
  <value>host_a,host_b</value>
</property>
```

If the configuration properties described are not present, impersonation is not allowed and connections will fail.

For looser restrictions, use a wildcard (*) to allow impersonation from any host and of any user. For example, to allow user bob to impersonate any user belonging to any group, and from any host, set the properties as follows:

```
<property>
  <name>hadoop.proxyuser.bob.hosts</name>
  <value>*</value>
</property>
<property>
  <name>hadoop.proxyuser.bob.groups</name>
  <value>*</value>
</property>
```

The `hadoop.proxyuser.<proxy_user>.hosts` property also accepts comma-separated lists of IP addresses, IP address ranges in CIDR format, or host names. For example, to allow user kate access from hosts in the range 10.222.0.0-15 and 10.113.221.221, to impersonate user_a and user_b, set the proxy user properties as follows:

```
<property>
  <name>hadoop.proxyuser.kate.hosts</name>
  <value>10.222.0.0/16,10.113.221.221</value>
</property>
<property>
  <name>hadoop.proxyuser.kate.users</name>
  <value>user1,user2</value>
</property>
```



Note: If you require proxy user to be setup on HTTPFS, you also need to configure them through the safety valve for the httpfs-site.xml. The properties required are the same as for the general proxy user setup, only they are prefixed with httpfs instead of hadoop. For example,

```
<property>
  <name>httpfs.proxyuser.#USER#.hosts</name>
  <value>*</value>
  <description>
    List of hosts the '#USER#' user is allowed to perform 'doAs'
    operations.
    The '#USER#' must be replaced with the username of the user who is
    allowed to perform 'doAs' operations.

    The value can be the '*' wildcard or a list of hostnames.

    For multiple users copy this property and replace the user name
    in the property name.
  </description>
</property>

<property>
  <name>httpfs.proxyuser.#USER#.groups</name>
  <value>*</value>
  <description>
    List of groups the '#USER#' user is allowed to impersonate users
    from to perform 'doAs' operations.

    The '#USER#' must be replaced with the username of the user who is
    allowed to perform 'doAs' operations.
    The value can be the '*' wildcard or a list of groups.

    For multiple users copy this property and replace the user name
    in the property name.
  </description>
</property>
```

The content of the properties should be the same as for the hadoop equivalent. For HTTPFS, you need to go into HttpFS Advanced Configuration Snippet (Safety Valve) for httpfs-site.xml.

Using DistCp to copy files

Hadoop DistCp (distributed copy) can be used to copy data between CDP clusters (and also within a CDP cluster).

DistCp uses MapReduce to implement its distribution, error handling, and reporting. It expands a list of files and directories into map tasks, each of which copies a partition of the files specified in the source list.

Using DistCp

The distributed copy command, `distcp`, is a general utility for copying large data sets between distributed filesystems within and across clusters. You can also use `distcp` to copy data to and from an Amazon S3 bucket. The `distcp` command submits a regular MapReduce job that performs a file-by-file copy.

To see the `distcp` command options, run the built-in help:

```
hadoop distcp
```

**Important:**

- Do not run distcp as the hdfs user which is disallowed for MapReduce jobs by default.
- Do not use [Hadoop shell commands](#) (such as cp, copyfromlocal, put, get) for large copying jobs or you may experience I/O bottlenecks.

Distcp syntax and examples

You can use distcp for copying data between CDP clusters. In addition, you can also use it to copy data between a CDP cluster and Amazon S3 or Azure Data Lake Storage Gen 2.

Common use of distcp

The most common use of distcp is an inter-cluster copy:

```
hadoop distcp hdfs://nn1:8020/source hdfs://nn2:8020/destination
```

Where hdfs://nn1:8020/source is the data source, and hdfs://nn2:8020/destination is the destination. This will expand the name space under /source on NameNode "nn1" into a temporary file, partition its contents among a set of map tasks, and start copying from "nn1" to "nn2". Note that DistCp requires absolute paths.

You can also specify multiple source directories:

```
hadoop distcp hdfs://nn1:8020/source/a hdfs://nn1:8020/source/b hdfs://nn2:8020/destination
```

Or specify multiple source directories from a file with the -f option:

```
hadoop distcp -f hdfs://nn1:8020/srclist hdfs://nn2:8020/destination
```

Where srclist contains:

```
hdfs://nn1:8020/source/a
hdfs://nn1:8020/source/b
```

Copying between major versions

Run the distcp command on the cluster that runs the higher version of CDP, which should be the destination cluster. Use the following syntax:

```
hadoop distcp webhdfs://<namenode>:<port> hdfs://<namenode>
```

Note the webhdfs prefix for the remote cluster, which should be your source cluster. You must use webhdfs when the clusters run different major versions. When clusters run the same version, you can use the hdfs protocol for better performance.

For example, the following command copies data from a CDP source cluster named example-source to another CDP version destination cluster named example-dest:

```
hadoop distcp webhdfs://example-source.cloudera.com:8020 hdfs://example-dest.cloudera.com
```

Copying to/from Amazon S3

The following syntax for distcp shows how to copy data to/from S3:

```
#Copying from S3
hadoop distcp s3a://<bucket>/<data> hdfs://<namenode>/<directory>/
```



```
#Copying to S3
hadoop distcp hdfs://<namenode>/<directory> s3a://<bucket>/<data>
```

This is a basic example of using distcp with S3.

Copying to/from ADLS Gen 2

The following syntax for distcp shows how to copy data to/from ADLS Gen 2:

```
#Copying from ABFS
hadoop distcp abfs[s]://<file_system>@<account_name>.dfs.core.windows.net/<path>/<file_name> hdfs://<hdfs_destination_path>
#Copying to ADLS Gen2
hadoop distcp hdfs://<hdfs_destination_path> abfs[s]://<file_system>@<account_name>.dfs.core.windows.net/<path>/<file_name>
```

Using DistCp with Highly Available remote clusters

You can use distcp to copy files between highly available clusters by configuring access to the remote cluster with the nameservice ID.

Procedure

1. Create a new directory and copy the contents of the /etc/hadoop/conf directory on the local cluster to this directory.

The local cluster is the cluster where you plan to run the distcp command.

Specify the new directory for the --config parameter when you run the distcp command in step 5.



Note: The following steps use distcpConf as the directory name. Substitute the name of the directory you created for distcpConf.

2. In the hdfs-site.xml file in the distcpConf directory, add the nameservice ID for the remote cluster to the dfs.nameservices property.



Note: If the remote cluster has the same nameservice ID as the local cluster, change the remote cluster's nameservice ID. Nameservice names must be unique.

For example, if the nameservice name for both clusters is nameservice1, change the nameservice ID of the remote cluster to a different ID, such as externalnameservice:

```
<property>
  <name>dfs.nameservices</name>
  <value>nameservice1,externalnameservice</value>
</property>
```

3. On the remote cluster, find the `hdfs-site.xml` file and copy the properties that refer to the nameservice ID to the end of the `hdfs-site.xml` file in the `distcpConf` directory you created in step 1.
 - `dfs.ha.namenodes.<nameserviceID>`
 - `dfs.client.failover.proxy.provider.<remote nameserviceID>`
 - `dfs.ha.automatic-failover.enabled.<remote nameserviceID>`
 - `dfs.namenode.rpc-address.<nameserviceID>.<namenode1>`
 - `dfs.namenode.servicerpc-address.<nameserviceID>.<namenode1>`
 - `dfs.namenode.http-address.<nameserviceID>.<namenode1>`
 - `dfs.namenode.https-address.<nameserviceID>.<namenode1>`
 - `dfs.namenode.rpc-address.<nameserviceID>.<namenode2>`
 - `dfs.namenode.servicerpc-address.<nameserviceID>.<namenode2>`
 - `dfs.namenode.http-address.<nameserviceID>.<namenode2>`
 - `dfs.namenode.https-address.<nameserviceID>.<namenode2>`

By default, you can find the `hdfs-site.xml` file in the `/etc/hadoop/conf` directory on a node of the remote cluster.

4. If you changed the nameservice ID for the remote cluster in step 2, update the nameservice ID used in the properties you copied in step 3 with the new nameservice ID, accordingly.

The following example shows the properties copied from the remote cluster with the following values:

- A remote nameservice called `externalnameservice`
- NameNodes called `namenode1` and `namenode2`
- A host named `remotecluster.com`

```
<property>
<name>dfs.ha.namenodes.externalnameservice</name>
<value>namenode1,namenode2</value>
</property>

<property>
<name>dfs.client.failover.proxy.provider.externalnameservice</name>
<value>org.apache.hadoop.hdfs.server.namenode.ha.ConfiguredFailoverProxyProvider</value>
</property>

<property>
<name>dfs.ha.automatic-failover.enabled.externalnameservice</name>
<value>true</value>
</property>

<property>
<name>dfs.namenode.rpc-address.externalnameservice.namenode1</name>
<value>remotecluster.com:8020</value>
</property>

<property>
<name>dfs.namenode.servicerpc-address.externalnameservice.namenode1</name>
<value>remotecluster.com:8022</value>
</property>
<property>
<name>dfs.namenode.http-address.externalnameservice.namenode1</name>
<value>remotecluster.com:20101</value>
</property>
<property>
<name>dfs.namenode.https-address.externalnameservice.namenode1</name>
<value>remotecluster.com:20102</value>
</property>

<property>
<name>dfs.namenode.rpc-address.externalnameservice.namenode2</name>
```

```
<value>remotecluster.com:8020</value>
</property>

<property>
<name>dfs.namenode.servicerpc-address.externalnameservice.namenode2</name>
<value>remotecluster.com:8022</value>
</property>
<property>
<name>dfs.namenode.http-address.externalnameservice.namenode2</name>
<value>remotecluster.com:20101</value>
</property>
<property>
<name>dfs.namenode.https-address.externalnameservice.namenode2</name>
<value>remotecluster.com:20102</value>
</property>
```

At this point, the `hdfs-site.xml` file in the `distcpConf` directory should have both the clusters and four NameNode IDs.

- Depending on the use case, the options specified when you run the `distcp` may differ.



Note: The remote cluster can be either the source or the target. The examples provided specify the remote cluster as the source.

To copy data from an insecure cluster, run the following command:

```
hadoop --config distcpConf distcp hdfs://<nameservice>/<source_directory>
<target_directory>
```

To copy data from a secure cluster, run the following command:

```
hadoop --config distcpConf distcp -Dmapreduce.job.hdfs-servers.token-renewal.exclude=<nameservice> hdfs://<nameservice>/<source_directory> <target_directory>
```

For example:

```
hadoop --config distcpConf distcp -Dmapreduce.job.hdfs-servers.token-renewal.exclude=ns1 hdfs://ns1/xyz /tmp/test
```

If the `distcp` source or target are in encryption zones, include the following `distcp` options: `-skipcrccheck` `-update`. The `distcp` command may fail if you do not include these options when the source or target are in encryption zones because the CRC for the files may differ.

For `distcp` between clusters that both use HDFS Transparent Encryption, you must include the `exclude` parameter.

Using DistCp with Amazon S3

You can copy HDFS files to and from an Amazon S3 instance. You must provision an S3 bucket using Amazon Web Services and obtain the access key and secret key.

You can pass these credentials on the `distcp` command line, or you can reference a credential store to "hide" sensitive credentials so that they do not appear in the console output, configuration files, or log files.

Amazon S3 block and native filesystems are supported with the `s3a://` protocol.

Example of an Amazon S3 Block Filesystem URI: `s3a://bucket_name/path/to/file`

S3 credentials can be provided in a configuration file (for example, `core-site.xml`):

```
<property>
  <name>fs.s3a.access.key</name>
```

```
<value>...</value>
</property>
<property>
  <name>fs.s3a.secret.key</name>
  <value>...</value>
</property>
```

You can also enter the configurations in the Advanced Configuration Snippet for `core-site.xml`, which allows Cloudera Manager to manage this configuration.

You can also provide the credentials on the command line:

```
hadoop distcp -Dfs.s3a.access.key=... -Dfs.s3a.secret.key=... s3a://
```

For example:

```
hadoop distcp -Dfs.s3a.access.key=myAccessKey -Dfs.s3a.secret.key=mySecretKey /user/hdfs/mydata s3a://myBucket/mydata_backup
```



Important: Entering secrets on the command line is inherently insecure. These secrets may be accessed in log files and other artifacts. Cloudera recommends that you use a credential provider to store secrets.



Note: Using the `-diff` option with the `distcp` command requires a Distributed FileSystem on both the source and destination and is not supported when using `distcp` to copy data to or from Amazon S3.

Using a credential provider to secure S3 credentials

You can run the `distcp` command without having to enter the access key and secret key on the command line. This prevents these credentials from being exposed in console output, log files, configuration files, and other artifacts.

About this task

Running the `distcp` command in this way requires that you provision a credential store to securely store the access key and secret key. The credential store file is saved in HDFS.



Note: Using a Credential Provider does not work with MapReduce v1 (MRV1).

Procedure

1. Provision the credentials by running the following commands:

```
hadoop credential create fs.s3a.access.key -value access_key -provider jceks://hdfs/path_to_credential_store_file
hadoop credential create fs.s3a.secret.key -value secret_key -provider jceks://hdfs/path_to_credential_store_file
```

For example:

```
hadoop credential create fs.s3a.access.key -value foobar -provider jceks://hdfs/user/alice/home/keystores/aws.jceks
hadoop credential create fs.s3a.secret.key -value barfoo -provider jceks://hdfs/user/alice/home/keystores/aws.jceks
```

You can omit the `-value` option and its value and the command will prompt the user to enter the value.

For more details on the `hadoop credential` command, see [Credential Management \(Apache Software Foundation\)](#).

2. Copy the contents of the `/etc/hadoop/conf` directory to a working directory.

3. Add the following to the core-site.xml file in the working directory:

```
<property>
<name>hadoop.security.credential.provider.path</name>
<value>jceks://hdfs/path_to_credential_store_file</value>
</property>
```

4. Set the HADOOP_CONF_DIR environment variable to the location of the working directory:

```
export HADOOP_CONF_DIR=path_to_working_directory
```

What to do next

After completing these steps, you can run the distcp command using the following syntax:

```
hadoop distcp source_path s3a://destination_path
```

You can also reference the credential store on the command line, without having to enter it in a copy of the core-site.xml file. You also do not need to set a value for HADOOP_CONF_DIR. Use the following syntax:

```
hadoop distcp source_path s3a://bucket_name/destination_path
-Dhadoop.security.credential.provider.path=jceks://
hdfspath_to_credential_store_file
```

Examples of DistCp commands using the S3 protocol and hidden credentials

You can use various distcp command options to copy files between your CDP clusters and Amazon S3.

Copying files to Amazon S3

```
hadoop distcp /user/hdfs/mydata s3a://myBucket/mydata_backup
```

Copying files from Amazon S3

```
hadoop distcp s3a://myBucket/mydata_backup //user/hdfs/mydata
```

Copying files to Amazon S3 using the -filters option to exclude specified source files

You specify a file name with the -filters option. The referenced file contains regular expressions, one per line, that define file name patterns to exclude from the distcp job. The pattern specified in the regular expression should match the fully-qualified path of the intended files, including the scheme (hdfs, webhdfs, s3a, etc.). For example, the following are valid expressions for excluding files:

```
hdfs://x.y.z:8020/a/b/c
webhdfs://x.y.z:50070/a/b/c
s3a://bucket/a/b/c
```

Reference the file containing the filter expressions using -filters option. For example:

```
hadoop distcp -filters /user/joe/myFilters /user/hdfs/mydata s
3a://myBucket/mydata_backup
```

Contents of the sample myFilters file:

```
.*foo.*
.*bar/.
hdfs://x.y.z:8020/tmp/.
hdfs://x.y.z:8020/tmp1/file1
```

The regular expressions in the `myFilters` exclude the following files:

- `.*foo.*` – excludes paths that contain the string "foo".
- `.*bar/.*` – excludes paths that include a directory named bar.
- `hdfs://x.y.z:8020/tmp/.*` – excludes all files in the /tmp directory.
- `hdfs://x.y.z:8020/tmp1/file1` – excludes the file /tmp1/file1.

Copying files to Amazon S3 with the `-overwrite` option.

The `-overwrite` option overwrites destination files that already exist.

```
hadoop distcp -overwrite /user/hdfs/mydata s3a://user/mydata_backup
```

For more information about the `-filters`, `-overwrite`, and other options, see [DistCp Guide: Command Line Options \(Apache Software Foundation\)](#).

Kerberos setup guidelines for Distcp between secure clusters

There are specific guidelines to consider while setting up Kerberos on secure CDP clusters for successfully performing distcp between them.

The guidelines mentioned here are only applicable for the following example deployment:

- You have two clusters, each in a different Kerberos realm (SOURCE and DESTINATION in this example)
- You have data that needs to be copied from SOURCE to DESTINATION
- A Kerberos realm trust exists, either between SOURCE and DESTINATION (in either direction), or between both SOURCE and DESTINATION and a common third realm (such as an Active Directory domain).

If your environment matches the one described above, use the following table to configure Kerberos delegation tokens on your cluster so that you can successfully distcp across two secure clusters. Based on the direction of the trust between the SOURCE and DESTINATION clusters, you can use the `mapreduce.job.hdfs-servers.token-renewal.exclude` property to instruct ResourceManagers on either cluster to skip or perform delegation token renewal for NameNode hosts.



Note: You must use the `mapreduce.job.hdfs-servers.token-renewal.exclude` parameter if both clusters use the HDFS Transparent Encryption feature.

Environment Type		Kerberos Delegation Token Setting
SOURCE trusts DESTINATION	Distcp job runs on the DESTINATION cluster	You do not need to set the <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property.
	Distcp job runs on the SOURCE cluster	Set the <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property to a comma-separated list of the hostnames of the NameNodes of the DESTINATION cluster.
DESTINATION trusts SOURCE	Distcp job runs on the DESTINATION cluster	Set the <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property to a comma-separated list of the hostnames of the NameNodes of the SOURCE cluster.
	Distcp job runs on the SOURCE cluster	You do not need to set the <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property.
Both SOURCE and DESTINATION trust each other		Set the <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property to a comma-separated list of the hostnames of the NameNodes of the DESTINATION cluster.

Environment Type	Kerberos Delegation Token Setting
Neither SOURCE nor DESTINATION trusts the other	<p>If a common realm is usable (such as Active Directory), set the <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property to a comma-separated list of hostnames of the NameNodes of the cluster that is not running the distcp job. For example, if you are running the job on the DESTINATION cluster:</p> <ol style="list-style-type: none"> 1. kinit on any DESTINATION YARN Gateway host using an AD account that can be used on both SOURCE and DESTINATION. 2. Run the distcp job as the hadoop user: <pre>\$ hadoop distcp -Ddfs.namenode.kerberos.principal.pattern=* \ -Dmapreduce.job.hdfs-servers.token-renewal.exclude=SOURCE-nn-host1,SOURCE-nn-host2 \ hdfs://source-nn-nameservice/source/path \ /destination/path</pre> <p>By default, the YARN ResourceManager renews tokens for applications. The <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property instructs ResourceManagers on either cluster to skip delegation token renewal for NameNode hosts.</p>

Distcp between secure clusters in different Kerberos realms

To use distcp between two secure clusters in different Kerberos realms, you must use a single Kerberos principal that can authenticate to both realms. In other words, a Kerberos realm trust relationship must exist between the source and destination realms. This can be a one-way trust (in either direction), a bi-directional trust, or even multiple one-way trusts where both the source and destination realms trust a third realm (such as an Active Directory domain).

Before you begin

Ensure that both the secure clusters run a supported version of JDK. For more information about the supported JDK versions, see [Java Requirements](#).

About this task

If there is no trust relationship between the source and destination realms, you cannot use distcp to copy data between the clusters, but you can use Replication Manager. For details, see [Enabling Replication Between Clusters with Kerberos Authentication](#).

Configure source and destination realms in krb5.conf

You must ensure that the `krb5.conf` file on the client from which the distcp job is submitted includes realm definitions and mappings for both source and destination realms.

Procedure

Include the realm definitions and mappings for both source and destination realms in the `krb5.conf` file, as specified in the following example.

```
[realms]
QA.EXAMPLE.COM = {
  kdc = kdc01.qa.example.com:88
  admin_server = kdc01.qa.example.com:749
}

DEV.EXAMPLE.COM = {
  kdc = kdc01.dev.example.com:88
  admin_server = kdc01.dev.example.com:749
}

[domain_realm]
```

```
.qa.example.com = QA.EXAMPLE.COM  
qa.example.com = QA.EXAMPLE.COM  
.dev.example.com = DEV.EXAMPLE.COM  
dev.example.com = DEV.EXAMPLE.COM
```

Configure HDFS RPC protection

Set the `hadoop.rpc.protection` property to authentication in both the clusters between which you want to run `distcp`. You can modify the property either in `hdfs-site.xml`, or using Cloudera Manager.

About this task

This procedure explains how you can configure the `hadoop.rpc.protection` property using Cloudera Manager. You must perform the procedure on both the clusters.

Procedure

1. Open the Cloudera Manager Admin Console.
2. Go to the HDFS service.
3. Click the Configuration tab.
4. Select Scope > HDFS-1 (Service-Wide).
5. Select Category > Security.
6. Locate the Hadoop RPC Protection property and select authentication.
7. Enter a Reason for Change and then click Save Changes.

Specify truststore properties

If TLS/SSL is enabled, then you must configure specific properties in the `ssl-client.xml` file on the client submitting the `distcp` job to establish trust between the target and destination clusters.

Procedure

Configure the `ssl.client.truststore.location`, `ssl.client.truststore.password`, and `ssl.client.truststore.type` properties in `ssl-client.xml` as specified.

```
<property>  
<name>ssl.client.truststore.location</name>  
<value>path_to_truststore</value>  
</property>  
  
<property>  
<name>ssl.client.truststore.password</name>  
<value>XXXXXX</value>  
</property>  
  
<property>  
<name>ssl.client.truststore.type</name>  
<value>jks</value>  
</property>
```

Set HADOOP_CONF to the destination cluster

Set the `HADOOP_CONF` path to be the destination environment. If you are not using HFTP, set the `HADOOP_CONF` path to the source environment instead.

Alternatively, you can point the `hadoop distcp` client to a configuration file that includes the parameters that allow it to point to the destination cluster. This can be done on the command line with the `--config` option. That option must

be pointed to a new copy of the contents of the `/etc/hadoop/conf` directory with the following parameters from the destination `/remote` cluster added to the `hdfs-site.xml` file:

```
dfs.ha.namenodes.<nameserviceID>
dfs.client.failover.proxy.provider.<remote nameserviceID>
dfs.ha.automatic-failover.enabled.<remote nameserviceID>
dfs.namenode.rpc-address.<nameserviceID>.<namenode1>
dfs.namenode.servicerpc-address.<nameserviceID>.<namenode1>
dfs.namenode.http-address.<nameserviceID>.<namenode1>
dfs.namenode.https-address.<nameserviceID>.<namenode1>
dfs.namenode.rpc-address.<nameserviceID>.<namenode2>
dfs.namenode.servicerpc-address.<nameserviceID>.<namenode2>
dfs.namenode.http-address.<nameserviceID>.<namenode2>
dfs.namenode.https-address.<nameserviceID>.<namenode2>
```

These parameters can be found in the `/etc/hadoop/conf/hdfs-site.xml` path on the remote cluster.

For example, the command would look something like the following:

```
hadoop --config distcpConf distcp -Dmapreduce.job.hdfs-servers.token-renewal
.exclude=<nameservice> hdfs://<nameservice>/<source_directory> <target dire
ctory>
```

For more details, see [Using DistCp with Highly Available remote clusters](#).

Launch distcp

Authenticate using kinit on the client and then launch the distcp job.

Consider the following example:

```
hadoop distcp hdfs://xyz01.dev.example.com:8020/user/alice hdfs://abc01.qa.e
xample.com:8020/user/alice
```

If launching distcp fails, force Kerberos to use TCP instead of UDP by adding the following parameter to the `krb5.conf` file on the client:

```
[libdefaults]
udp_preference_limit = 1
```

Copying data between a secure and an insecure cluster using DistCp and WebHDFS

You can use distcp and WebHDFS to copy data between a secure cluster and an insecure cluster.

About this task

When copying data, ensure that you run the distcp commands from the secure cluster.

Procedure

1. On the secure cluster, set `ipc.client.fallback-to-simple-auth-allowed` to `true` in `core-site.xml`.

```
<property>
  <name>ipc.client.fallback-to-simple-auth-allowed</name>
  <value>true</value>
</property>
```

Alternatively, you can also pass this as a parameter when you run the distcp command. If you want to do that, move onto step 2.

2. On the insecure cluster, add the secured cluster's realm name to the insecure cluster's configuration.
 - a) In the Cloudera Manager Admin Console for the insecure cluster, navigate to **Clusters <HDFS cluster>**.
 - b) On the Configuration tab, search for **Trusted Kerberos Realms** and add the secured cluster's realm name.



Note: This does not require Kerberos to be enabled but is a necessary step to allow the simple auth fallback to happen in the `hdfs://` protocol.

- c) Save the change.
3. Use commands such as the following only from the secure cluster side.

```
#This example uses the insecure cluster as the source and the secure cluster as the destination
distcp webhdfs://<insecure_namenode>:9870 webhdfs://<secure_namenode>:9871

#This example uses the secure cluster as the source and the insecure cluster as the destination
distcp webhdfs://<secure_namenode>:9871 webhdfs://<insecure_namenode>:9870
```

If TLS is enabled, replace `webhdfs` with `swebhdfs`.

If you did not configure `ipc.client.fallback-to-simple-auth-allowed` and want to pass it as a parameter, run commands such as the following from the secure cluster:

```
#This example uses the insecure cluster as the source and the secure cluster (with TLS enabled) as the destination cluster. swebhdfs is used instead of webhdfs when TLS is enabled.
hadoop distcp -D ipc.client.fallback-to-simple-auth-allowed=true webhdfs://<insecure_namenode>:9870 swebhdfs://<secure_namenode>:9871

#This example uses the secure cluster (with TLS enabled) as the source cluster and the insecure cluster as the destination. swebhdfs is used instead of webhdfs when TLS is enabled.
hadoop distcp -D ipc.client.fallback-to-simple-auth-allowed=true swebhdfs://<secure_namenode>:9871 webhdfs://<insecure_namenode>:9870
```

Post-migration verification

After migrating data between the two clusters, you can use `hadoop fs -ls /basePath` to verify the permissions, ownership and other aspects of your files, and correct any problems before using the files in your new cluster.

Using DistCp between HA clusters using Cloudera Manager

To copy data between HA clusters using `distcp`, you must configure specific name service properties to ensure that the HDFS clients in a cluster can access a remote HA cluster. You can configure these values through advanced configuration snippets in Cloudera Manager.

About this task

This procedure explains how you can configure the name service properties from Cloudera Manager to enable copying of data between two example clusters A and B. Here, A is the source cluster while B is the remote cluster.

Procedure

1. Select **Clusters** and choose the source HDFS cluster where you want to configure the properties.

- Navigate to the Configuration tab and search for HDFS Client Advanced Configuration Snippet (Safety Valve) for hdfs-site.xml.

Cluster 1 > Gateway Default Group [Undo](#) [View as XML](#)

HDFS Client Advanced Configuration Snippet (Safety Valve) for hdfs-site.xml

[hdfs_client_config_safety_valve](#)

Name	dfs.nfs.keytab.file	🗑️ 🔍
Value	{{CMF_CONF_DIR}}/hdfs.keytab	
Description		
	<input type="checkbox"/> Final	
Name		🗑️ 🔍
Value		
Description		
	<input type="checkbox"/> Final	

- Add the various properties as specified:

- Add the dfs.internal.nameservices property as the name service of the source cluster.
dfs.internal.nameservices = HAA
- Add the dfs.nameservices property as the name services of both the clusters.
dfs.nameservices = HAA,HAB
- Add the dfs.client.failover.proxy.provider.<nameservice> property for the remote cluster.
dfs.client.failover.proxy.provider.HAB = org.apache.hadoop.hdfs.server.namenode.ha.ConfiguredFailoverProxyProvider
- Add the dfs.ha.namenodes.<nameservice> property to reference the NameNodes of the remote cluster.
dfs.ha.namenodes.HAB = nn1,nn2
- Add the dfs.namenode.rpc-address and the dfs.namenode.servicerpc-address properties as the fully qualified RPC addresses on which each NameNode of the remote cluster can listen.
dfs.namenode.rpc-address.HAB.nn1 = <NN1_fqdn>:8020
dfs.namenode.rpc-address.HAB.nn2 = <NN2_fqdn>:8020
dfs.namenode.servicerpc-address.HAB.nn1 = <NN1_fqdn>:8022
dfs.namenode.servicerpc-address.HAB.nn2 = <NN2_fqdn>:8022
- Add the dfs.namenode.http-address and the dfs.namenode.https-address properties as the fully qualified HTTP and HTTPS addresses on which each NameNode of the remote cluster can listen.
dfs.namenode.http-address.HAB.nn1 = <NN1_fqdn>:9870
dfs.namenode.http-address.HAB.nn2 = <NN2_fqdn>:9870
dfs.namenode.https-address.HAB.nn1 = <NN1_fqdn>:9871
dfs.namenode.https-address.HAB.nn2 = <NN2_fqdn>:9871

- Specify the Reason for Change, and click Save Changes.

This saves the property values as client-side configurations.

- Restart the HDFS service, then run the distcp command using the NameService. For example:

```
hadoop distcp hdfs://HAA/tmp/testDistcp hdfs://HAB/tmp/
```



Note: When you try to connect a NameNode from one cluster to another cluster which resides in the same REALM and same KDC, you must add the above HDFS HA configurations on the client side to connect to the remote cluster.

Using the NFS Gateway for accessing HDFS

Available in a CDP Private Cloud Base deployment, the NFS Gateway for HDFS allows clients to mount HDFS and interact with it through NFS, as if it were part of their local file system. The Gateway supports NFSv3.

After mounting HDFS, a client user can perform the following tasks:

- Browse the HDFS file system through their local file system on NFSv3 client-compatible operating systems.
- Upload and download files between the HDFS file system and their local file system.
- Stream data directly to HDFS through the mount point. File append is supported, but random write is not supported.

Install the NFS Gateway

The NFS Gateway is part of HDFS service. The NFS Gateway can be installed on any host managed by Cloudera Manager, including DataNode, NameNode, or CDP client machine. Start the NFS server on that machine.

About this task

Perform the following steps to install NFS Gateway role:

Procedure

1. In the Cloudera Manager Admin Console, go to Clusters HDFS service .
2. Navigate to the Instances tab and click Add Role Instances.
3. Click NFS Gateway to choose the hosts to install the NFS Gateway role.
The NFS Gateway is added to the host.

Configure the NFS Gateway

You must ensure that the proxy user for the NFS Gateway can proxy all the users accessing the NFS mounts. In addition, you must configure settings specific to the Gateway.

Procedure

1. Ensure that the proxy user for the NFS Gateway can proxy all the users accessing the NFS mounts.

In non-secure mode, the user running the Gateway is the proxy user, while in secure mode the user in Kerberos keytab is the proxy user.

If a user `nfsserver` is running the Gateway and there are users belonging to groups `nfs-users1` and `nfs-users2`, then set the following values in `core-site.xml` on the NameNode.



Note:

Replace `nfsserver` with the user account that starts the Gateway in your cluster.

```
<property>
  <name>hadoop.proxyuser.nfsserver.groups</name>
  <value>nfs-users1,nfs-users2</value>
  <description>
    The 'nfsserver' user is allowed to proxy all members of the
    'nfs-users1' and 'nfs-users2' groups. Set this to '*' to allow
    nfsserver user to proxy any group.
  </description>
</property>
```

```
<property>
  <name>hadoop.proxyuser.nfsserver.hosts</name>
  <value>nfs-client-host1.com</value>
  <description>
    This is the host where the nfs gateway is running. Set this to
    '*' to allow requests from any hosts to be proxied.
  </description>
</property>
```

For a Kerberized cluster, set the following properties in `hdfs-site.xml`:

```
For a Kerberized cluster, set the following properties in
NFS Gateway Advanced Configuration Snippet (Safety Valve) for hdfs-site.x
ml
Name: nfs.keytab.file
Value: {{CMF_CONF_DIR}}/nfsserver.keytab

Name: nfs.kerberos.principal
Value: nfsserver/_HOST@YOUR-REALM.COM
```

2. Configure settings for the NFS Gateway.

The NFS Gateway uses the same settings that are used by the NameNode and DataNode. Configure various properties based on your application's requirements:

- a) Edit the `hdfs-site.xml` file on your NFS Gateway machine.

```
<property>
  <name>dfs.namenode.accesstime.precision</name>
  <value>3600000</value>
  <description>
    The access time for HDFS file is precise up to this value.
    The default value is 1 hour. Setting a value of 0 disables
    access times for HDFS.
  </description>
</property>
```



Note: If the export is mounted with access time update allowed, ensure that this property is not disabled in the configuration file. Only the NameNode needs to restart after this property is changed. If you have disabled access time update by mounting with `noatime`, you do NOT have to change this property nor restart your NameNode.

- b) Add the value of the `dfs.nfs3.dump.dir` property in `hdfs-site.xml`.

```
<property>
  <name>dfs.nfs3.dump.dir</name>
  <value>/tmp/.hdfs-nfs</value>
</property>
```



Note: The NFS client often reorders writes. Sequential writes can arrive at the NFS Gateway in a random order. This directory is used to temporarily save out-of-order writes before writing to HDFS. Ensure that the directory has enough space. For example, if the application uploads 10 files with each having 100MB, it is recommended for this directory to have 1GB space in case a write reorder happens to every file.

- c) Update the value of the `dfs.nfs.exports.allowed.hosts` property in `hdfs-site.xml` as specified.

```
<property>
  <name>dfs.nfs.exports.allowed.hosts</name>
  <value>* rw</value>
```

```
</property>
```



Note: By default, the export can be mounted by any client. You must update this property to control access. The value string contains the machine name and access privilege, separated by whitespace characters. The machine name can be in single host, wildcard, or IPv4 network format. The access privilege uses `rw` or `ro` to specify readwrite or readonly access to exports. If you do not specify an access privilege, the default machine access to exports is readonly. Separate machine entries by `;`. For example, `192.168.0.0/22 rw ; host*.example.com ; host1.test.org ro;`

- d) Restart the NFS Gateway.
- e) Customize log settings by modifying the `log4j.properties` file.

To change the trace level, add the following:

```
log4j.logger.org.apache.hadoop.hdfs.nfs=DEBUG
```

To view more information about ONCRPC requests, add the following:

```
log4j.logger.org.apache.hadoop.oncrpc=DEBUG
```

3. Specify JVM heap space (`HADOOP_NFS3_OPTS`) for the NFS Gateway.

You can increase the JVM heap allocation for the NFS Gateway using this option. To set this option, specify the following in `hadoop-env.sh`:

```
export HADOOP_NFS3_OPTS=<memory-setting(s)>
```

The following example specifies a 2GB process heap (2GB starting size and 2GB maximum):

```
export HADOOP_NFS3_OPTS="-Xms2048m -Xmx2048m"
```

4. To improve the performance of large file transfers, you can increase the values of the `dfs.nfs.rtmax` and `dfs.nfs.wtmax` properties.

These properties are configuration settings on the NFS Gateway server that change the maximum read and write request size supported by the Gateway. The default value for both settings is 1MB.

Start and stop the NFS Gateway services

You must start the following daemons to run the NFS services on the Gateway: `rpcbind` (or `portmap`), `mountd`, and `nfds`. The NFS Gateway process includes both `nfds` and `mountd`. Although NFS Gateway works with `portmap` included with most Linux distributions, you must use the `portmap` included in the NFS Gateway package on some Linux systems such as SLES 11 and RHEL 6.2.

About this task

Perform the following steps to use the bundled `portmap`:

Procedure

1. Stop the `nfs/rpcbind/portmap` services provided by the platform.

```
service nfs stop
service rpcbind stop
```

2. Start the included `portmap` service using one of the following commands:

```
hdfs portmap
OR
hdfs --daemon start portmap
```



Note: You must have root privileges to run these commands.

Start the NFS Gateway services

Learn how to start the NFS Gateway services.

About this task

Perform the following steps to start the NFS Gateway:

Procedure

1. In the Cloudera Manager Admin Console, go to Clusters HDFS service .
2. Navigate to the Instances tab and search for the NFS Gateway.
3. Select Start from the Actions for Selected dropdown menu.



Note: The NFS Gateway log file directory is configured by the property NFS Gateway Log Directory. For example, if NFS Gateway Log Directory is configured as /var/log/hadoop-hdfs and the NFS Gateway runs on host host1.example.com, the log file would be /var/log/hadoop-hdfs/hadoop-hdfs-hdfs-nfs3-host1.example.com.out.

What to do next

Verify the validity of the NFS services.

Stop the NFS Gateway services

Learn how to stop the NFS Gateway services. Before stopping the NFS Gateway services, all NFS mount points must be unmounted.

Procedure

1. Bring down all NFS mount points.

```
sudo umount {mountPoint}
```

2. Perform the following steps to stop the NFS Gateway services:
 - a) In the Cloudera Manager Admin Console, go to Clusters HDFS service .
 - b) Navigate to the Instances tab and search for NFS Gateway.
 - c) Select Stop from the drop down menu.

Verify validity of the NFS services

After starting the NFS services on the Gateway, you must verify whether all the services are running. Additionally, you must ensure that the HDFS namespace is exported and can be mounted.

Procedure

1. Run the `rpcinfo` command to verify whether all the NFS services are running.

```
rpcinfo -p $nfs_server_ip
```

The command returns an output similar to the following:

```
program vers proto  port
100005  1    tcp    4242  mountd
100005  2    udp    4242  mountd
100005  2    tcp    4242  mountd
100000  2    tcp    111   portmapper
100000  2    udp    111   portmapper
100005  3    udp    4242  mountd
100005  1    udp    4242  mountd
100003  3    tcp    2049  nfs
100005  3    tcp    4242  mountd
```

2. Verify that the HDFS namespace is exported and can be mounted.

```
showmount -e $nfs_server_ip
```

The command returns an exports list similar to the following:

Exports list on \$nfs_server_ip : / (everyone)

Access HDFS from the NFS Gateway

To access HDFS, you must first mount the namespace and then set up the NFS client hosts to interact with HDFS through the Gateway.

Procedure

1. Mount the HDFS namespace.

```
mount -t nfs -o vers=3,proto=tcp,nolock,rsync,rsync,rsync,rsync $server:/ $mount_point
```

Access HDFS as part of the local file system, except that hard link or symbolic link and random write are not supported.



Note: Because NLM is not supported, the mount option nolock is required.

You can use the sync option for improving the performance and reliability when writing large files. If the sync option is specified, the NFS client flushes write operations to the NFS Gateway before returning control to the client application. Additionally, the client does not issue reordered writes. This reduces buffering requirements on the NFS gateway.

2. Set up the NFS client users to interact with HDFS through the NFS Gateway.

Related Information

[How NFS Gateway authenticates and maps users](#)

How NFS Gateway authenticates and maps users

The manner in which NFS Gateway authenticates and maps users determines how NFS client users access HDFS through the NFS Gateway.

NFS Gateway uses AUTH_UNIX-style authentication, which means that the user logged on the client is the same that NFS passes to HDFS.

For example, suppose that the NFS client has current user as admin. When the user accesses the mounted HDFS directory, NFS gateway accesses HDFS as the user admin. To access HDFS as hdfs user, you must first switch the current user to hdfs on the client system before accessing the mounted directory.

For the NFS client users accessing HDFS, the Gateway converts the User Identifier (UID) to a username, which is used by HDFS to check permissions. Therefore, you must ensure that the user on the NFS client host has the same name and UID as that on the NFS Gateway host.




Note: If you use the same user management system such as LDAP/NIS to create and deploy users to the CDP nodes and to the client host, then the user names and UIDs are usually identical on both the client and the Gateway.

If the user is created manually on different hosts, you might need to modify the UID, as specified in the following example, on either the client or the NFS Gateway to make them the same:

```
usermod -u 123 $myusername
```

The following diagram illustrates how the UID and name are communicated between the NFS client, NFS gateway, and NameNode.



Related Information

[Access HDFS from the NFS Gateway](#)

APIs for accessing HDFS

Use the WebHDFS REST API to access an HDFS cluster from applications external to the cluster. WebHDFS supports all HDFS user operations including reading files, writing to files, making directories, changing permissions and renaming. In addition, WebHDFS uses Kerberos and delegation tokens for authenticating users.

Set up WebHDFS on a secure cluster

Setting up WebHDFS on a secure cluster requires you to update certain properties on `hdfs-site.xml`, create an HTTP service user principal, create a keytab for the principal, and verify the association of the principal and keytab with the correct HTTP service.

Procedure

1. Create an HTTP service user principal.

```
kadmin: addprinc -randkey HTTP/$<Fully_Qualified_Domain_Name>@$<Realm_Name>.COM
```

where:

- Fully_Qualified_Domain_Name: Host where the NameNode is deployed.
- Realm_Name: Name of your Kerberos realm.

2. Create a keytab file for the HTTP principal.

```
kadmin: xst -norandkey -k /etc/security/spnego.service.keytab HTTP/$<Fully_Qualified_Domain_Name>
```

3. Verify that the keytab file and the principal are associated with the correct service.

```
klist -k -t /etc/security/spnego.service.keytab
```

4. Add the `dfs.web.authentication.kerberos.principal` and `dfs.web.authentication.kerberos.keytab` properties to `hdfs-site.xml`.

```
<property>
  <name>dfs.web.authentication.kerberos.principal</name>
  <value>HTTP/$<Fully_Qualified_Domain_Name>@$<Realm_Name>.COM</value>
```

```
</property>
<property>
  <name>dfs.web.authentication.kerberos.keytab</name>
  <value>/etc/security/spnego.service.keytab</value>
</property>
```

5. Restart the NameNode and the DataNodes.

Using HttpFS to provide access to HDFS

Apache Hadoop HttpFS is a service that provides HTTP access to HDFS. HttpFS has a REST HTTP API supporting all HDFS filesystem operations (both read and write).

Common HttpFS use cases are:

- Read and write data in HDFS using HTTP utilities (such as curl or wget) and HTTP libraries from languages other than Java (such as Perl).
- Transfer data between HDFS clusters running different versions of Hadoop (overcoming RPC versioning issues), for example using Hadoop DistCp.
- Accessing WebHDFS using the NameNode Web UI port (default port 9870). Access to all data hosts in the cluster is required, because WebHDFS redirects clients to the DataNode port (default port 9864). If the cluster is behind a firewall, and you use WebHDFS to read and write data to HDFS, then Cloudera recommends you use the HttpFS server. The HttpFS server acts as a gateway. It is the only system that is allowed to send and receive data through the firewall.

HttpFS supports Hadoop pseudo-authentication, HTTP SPNEGO Kerberos, and additional authentication mechanisms using a plugin API. HttpFS also supports Hadoop proxy user functionality.

The webhdfs client file system implementation can access HttpFS using the Hadoop filesystem command (hadoop fs), by using Hadoop DistCp, and from Java applications using the Hadoop file system Java API.

The HttpFS HTTP REST API is interoperable with the WebHDFS REST HTTP API.

Add the HttpFS role

You can use Cloudera Manager to add the HttpFS role.

1. Go to the HDFS service.
2. Click the Instances tab.
3. Click Add Role Instances.
4. Click the text box below the HttpFS field. The Select Hosts dialog box displays.
5. Select the host on which to run the role and click OK.



Note: You can select any node other than the master nodes.

6. Click Continue.
7. Check the checkbox next to the HttpFS role and select Actions for Selected Start .

Using Load Balancer with HttpFS

Configure the HttpFS Service to work with the load balancer that you configured for the service.

1. In the Cloudera Manager Admin Console, navigate to Cluster <HDFS service> .

2. On the Configuration tab, search for the following property:

```
HttpFS Load Balancer
```

3. Enter the hostname and port for the load balancer in the following format:

```
<hostname>:<port>
```

4. Save the changes.



Note: When you set this property, Cloudera Manager regenerates the keytabs for HttpFS roles. The principal in these keytabs contains the load balancer hostname. If there is a Hue service that depends on this HDFS service, the Hue service has the option to use the load balancer as its HDFS Web Interface Role.

HttpFS authentication

To enable HttpFS to work with Kerberos security on your CDP cluster, ensure that you have configured Kerberos for the cluster.



Important: If the NameNode, Secondary NameNode, DataNode, JobTracker, TaskTrackers, ResourceManager, NodeManagers, HttpFS, or Oozie services are configured to use Kerberos HTTP SPNEGO authentication, and two or more of these services are running on the same host, then all of the running services must use the same HTTP principal and keytab file used for their HTTP endpoints.

Use curl to access a URL protected by Kerberos HTTP SPNEGO

To access a URL protected by Kerberos HTTP SPNEGO, ensure that your version of curl supports GSS and is capable of running `curl -V`.

Procedure

1. Run `curl -V`.

```
$ curl -V
curl 7.19.7 (universal-apple-darwin10.0) libcurl/7.19.7 OpenSSL/0.9.8l
zlib/1.2.3
Protocols: tftp ftp telnet dict ldap http file https ftps
Features: GSS-Negotiate IPv6 Largefile NTLM SSL libz
```

2. Log on to the KDC using `kinit`.

```
$ kinit
Please enter the password for username@LOCALHOST:
```

3. Use `curl` to fetch the protected URL.

```
$ curl --cacert
```

```
/path/to/truststore.pem --negotiate -u : -b ~/cookiejar.txt
-c ~/cookiejar.txt https://localhost:14000/webhdfs/v1/?op=liststatus
```

where:

- The `--cacert` option is required if you are using TLS/SSL certificates that curl does not recognize by default.
- The `--negotiate` option enables SPNEGO in curl.
- The `-u:` option is required but the username is ignored (the principal that has been specified for kinit is used).
- The `-b` and `-c` options are used to store and send HTTP cookies.
- Cloudera does not recommend using the `-k` or `--insecure` option as it turns off curl's ability to verify the certificate.

Data storage metrics

Use Java Management Extensions (JMX) APIs to collect the metrics exposed by the various HDFS daemons.

You can use various HDFS metrics to understand the state of your cluster.

Using JMX for accessing HDFS metrics

You can access HDFS metrics over Java Management Extensions (JMX) through either the web interface of an HDFS daemon or by directly accessing the JMX remote agent.

Using the HDFS daemon web interface

You can access JMX metrics through the web interface of an HDFS daemon. This is the recommended method.

For example, use the following command format to access the NameNode JMX:

```
curl -i http://localhost:50070/jmx
```

You can use the `qry` parameter to fetch only a particular key:

```
curl -i http://localhost:50070/jmx?qry=Hadoop:service=NameNode,name=NameNode
Info
```

Directly accessing the JMX remote agent

This method requires that the JMX remote agent is enabled with a JVM option when starting HDFS services.

For example, the following JVM options in `hadoop-env.sh` are used to enable the JMX remote agent for the NameNode. It listens on port 8004 with SSL disabled. The user name and password are saved in the `mxremote.password` file.

```
export HADOOP_NAMENODE_OPTS="-Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.password.file=$HADOOP_CONF_DIR/jmxremote.pass
word
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.sun.management.jmxremote.port=8004 $HADOOP_NAMENODE_OPTS"
```

See the Oracle Java SE documentation for more information about the related settings.

You can also use the `jmxquery` tool to retrieve information through JMX.

Hadoop has a built-in JMX query tool, `jmxget`. For example:

```
hdfs jmxget -server localhost -port 8004 -service NameNode
```



Note: jmxget requires that authentication be disabled, as it does not accept a user name and password.

Using JMX can be challenging for operations personnel who are not familiar with JMX setup, especially JMX with SSL and firewall tunnelling. Therefore, we recommend that you collect JMX information through the web interface of HDFS daemons rather than directly accessing the JMX remote agent.

Related Information

[JMX Query](#)

[Monitoring and Management Using JMX Query](#)

HDFS Metrics

Many aggregate metrics are available in addition to base metrics. If an entity type has parents defined, you can formulate all possible aggregate metrics using the formula `base_metric_across_parents`.

In addition, metrics for aggregate totals can be formed by adding the prefix `total_` to the front of the metric name.

Use the type-ahead feature in the Cloudera Manager chart browser to find the exact aggregate metric name, in case the plural form does not end in "s".

For example, the following metric names may be valid for HDFS:

- `alerts_rate_across_clusters`
- `total_alerts_rate_across_clusters`

Some metrics, such as `alerts_rate`, apply to nearly every metric context. Others only apply to a certain service or role.

Metric Name	Description	Unit	Parents
<code>alerts_rate</code>	The number of alerts.	events per second	cluster
<code>block_capacity</code>	The block capacity of the NameNode	blocks	cluster
<code>blocks_total</code>	Blocks total	blocks	cluster
<code>blocks_with_corrupt_replicas</code>	Blocks with corrupt replicas	blocks	cluster
<code>canary_duration</code>	Duration of the last or currently running canary job	ms	cluster
<code>cm_time_since_last_fsimage_fetch</code>	Time since last FsImage was fetched by Cloudera Reports Manager	seconds	cluster
<code>cm_time_since_last_fsimage_index</code>	Time since last FsImage was indexed By Cloudera Reports Manager	seconds	cluster
<code>dfs_capacity</code>	Total configured HDFS storage capacity	bytes	cluster
<code>dfs_capacity_used</code>	Storage space used by HDFS files	bytes	cluster
<code>dfs_capacity_used_non_hdfs</code>	Storage space used by non-HDFS files	bytes	cluster
<code>events_critical_rate</code>	The number of critical events.	events per second	cluster
<code>events_important_rate</code>	The number of important events.	events per second	cluster
<code>events_informational_rate</code>	The number of informational events.	events per second	cluster
<code>excess_blocks</code>	The total number of excess blocks	blocks	cluster
<code>expired_heartbeats</code>	The number of expired heartbeats	heartbeats	cluster
<code>files_total</code>	The number of files and directories in the HDFS	files	cluster
<code>health_bad_rate</code>	Percentage of Time with Bad Health	seconds per second	cluster

Metric Name	Description	Unit	Parents
health_concerning_rate	Percentage of Time with Concerning Health	seconds per second	cluster
health_disabled_rate	Percentage of Time with Disabled Health	seconds per second	cluster
health_good_rate	Percentage of Time with Good Health	seconds per second	cluster
health_unknown_rate	Percentage of Time with Unknown Health	seconds per second	cluster
missing_blocks	Missing blocks	blocks	cluster
pending_deletion_blocks	The number of replicas pending deletion.	replicas	cluster
pending_replication_blocks	The number of blocks with replication pending	blocks	cluster
scheduled_replication_blocks	The number of blocks with replication currently scheduled	blocks	cluster
under_replicated_blocks	Under-replicated blocks	blocks	cluster
xceivers	Transceivers	transceivers	cluster

Using HdfsFindTool to find files

HDFSFindTool is essentially the HDFS version of the Linux filesystem find command. You can use the tool to search one or more HDFS directory trees, find all HDFS files that match the specified expression, and apply selected actions to them.



Note: HDFSFindtool is *not* available for download with CDP. If you have the required payroll credentials to access the CDH archives, you can download the JAR file for HDFSFindtool and use it.

Downloading Hdfsfindtool from the CDH archives

If you have access to CDH credentials and want to use HDFSFindtool with your CDP deployment, then you can download the tool from the archives.

Before you begin

You must have the required payroll credentials to access the CDH archives.

About this task

You can download the CDH5 jar to run HDFSFindtool commands from the CDH archives.

Procedure

1. Log into the CDH archives with the help of your payroll credentials and download the search tarball.

```
$ wget --user=<your_paywall_username> --ask-password https://archive.cloudera.com/p/cdh5/cdh/5/search-1.0.0-cdh5.16.2.tar.gz
```

Here, <your_paywall_username> indicates your Paywall username. In addition, enter your password when prompted.

2. Extract the downloaded tarball file.

```
$ tar -xvf search-1.0.0-cdh5.16.2.tar.gz
```

3. Run HDFSFindtool.

```
$ hadoop jar cloudera-search-1.0.0-cdh5.16.2/dist/search-mr-1.0.0-cdh5.16.2.jar org.apache.solr.hadoop.HdfsFindTool -find /tmp/ -type f
```

Example

The following example shows how you can use HDFSFindtool:

```
$ hadoop jar cloudera-search-1.0.0-cdh5.16.2/dist/search-mr-1.0.0-cdh5.16.2.jar org.apache.solr.hadoop.HdfsFindTool -find /tmp/ -type f
WARNING: Use "yarn jar" to launch YARN applications.
hdfs://ns1/tmp/logs/systest/logs/application_1575879043308_0001/file1
hdfs://ns1/tmp/my_table.avsc
```