Hortonworks Data Platform

User Guides

(Feb 10, 2014)

docs.hortonworks.com

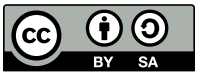
Hortonworks Data Platform : User Guides

Copyright © 2012, 2013, 2014 Hortonworks, Inc. Some rights reserved.

The Hortonworks Data Platform, powered by Apache Hadoop, is a massively scalable and 100% open source platform for storing, processing and analyzing large volumes of data. It is designed to deal with data from many sources and formats in a very quick, easy and cost-effective manner. The Hortonworks Data Platform consists of the essential set of Apache Hadoop projects including MapReduce, Hadoop Distributed File System (HDFS), HCatalog, Pig, Hive, HBase, Zookeeper and Ambari. Hortonworks is the major contributor of code and patches to many of these projects. These projects have been integrated and tested as part of the Hortonworks Data Platform release process and installation and configuration tools have also been included.

Unlike other providers of platforms built using Apache Hadoop, Hortonworks contributes 100% of our code back to the Apache Software Foundation. The Hortonworks Data Platform is Apache-licensed and completely open source. We sell only expert technical support, training and partner-enablement services. All of our technology is, and will remain free and open source.

Please visit the Hortonworks Data Platform page for more information on Hortonworks technology. For more information on Hortonworks services, please visit either the Support or Training page. Feel free to Contact Us directly to discuss your specific needs.



Except where otherwise noted, this document is licensed under
Creative Commons Attribution ShareAlike 3.0 License.
http://creativecommons.org/licenses/by-sa/3.0/legalcode

Hortonworks Data Platform (HDP) and any of its components are not anticipated to be combined with any hardware, software or data, except as expressly recommended in this documentation.

Table of Contents

1. HBase Snapshots 1
    1.1. Configuration 1
    1.2. Take a Snapshot 1
    1.3. Listing Snapshots 1
    1.4. Deleting Snapshots 2
    1.5. Clone a table from snapshot 2
    1.6. Restore a snapshot 2
    1.7. Snapshots operations and ACLs 2
    1.8. Export to another cluster 3
2. User Guide - HDFS NFS Gateway 4
    2.1. Prerequisites 4
    2.2. Instructions 4
3. User Guide - HDFS Snapshots 9
    3.1. Snapshottable Directories 9
    3.2. Snapshot Paths 9
    3.3. Snapshot Operations 10
        3.3.1. Administrator Operations 10
        3.3.2. User Operations 11
4. Add HDP Maven Repository to Existing Project 13

List of Tables

3.1. Administrator Operations - Allow Snapshots 10
3.2. Administrator Operations - Disallow Snapshots 10
3.3. User Operations - Create Snapshots 11
3.4. User Operations - Delete Snapshots 11
3.5. User Operations - Rename Snapshots 12
3.6. User Operations - Get Snapshottable Directory Listing 12
3.7. User Operations - Get Snapshots Difference Report 12

1. HBase Snapshots

HBase Snapshots allow you to take a snapshot of a table without too much impact on Region Servers. Snapshot, Clone and restore operations don't involve data copying. Also, Exporting the snapshot to another cluster doesn't have impact on the Region Servers.

Prior to version 0.94.6, the only way to backup or to clone a table is to use CopyTable/ExportTable, or to copy all the hfiles in HDFS after disabling the table. The disadvantages of these methods are that you can degrade region server performance (Copy/Export Table) or you need to disable the table, that means no reads or writes; and this is usually unacceptable. In this section:

- Configuration
- Take a Snapshot
- Listing Snapshots
- Deleting Snapshots
- Clone a table from snapshot
- Restore a snapshot
- Snapshots operations and ACLs
- Export to another cluster

1.1. Configuration

To turn on the snapshot support just set the `hbase.snapshot.enabled` property to true. (Snapshots are enabled by default in 0.95+ and off by default in 0.94.6+)

```
<property>
  <name>hbase.snapshot.enabled</name>
  <value>true</value>
</property>
```

1.2. Take a Snapshot

You can take a snapshot of a table regardless of whether it is enabled or disabled. The snapshot operation doesn't involve any data copying.

```
  $ hbase shell
  hbase> snapshot 'myTable', 'myTableSnapshot-122112'
```

1.3. Listing Snapshots

List all snapshots taken (by printing the names and relative information).

```
$ hbase shell
hbase> list_snapshots
```

1.4. Deleting Snapshots

You can remove a snapshot, and the files retained for that snapshot will be removed if no longer needed.

```
$ hbase shell
hbase> delete_snapshot 'myTableSnapshot-122112'
```

1.5. Clone a table from snapshot

From a snapshot you can create a new table (clone operation) with the same data that you had when the snapshot was taken. The clone operation, doesn't involve data copies, and a change to the cloned table doesn't impact the snapshot or the original table.

```
$ hbase shell
hbase> clone_snapshot 'myTableSnapshot-122112', 'myNewTestTable'
```

1.6. Restore a snapshot

The restore operation requires the table to be disabled, and the table will be restored to the state at the time when the snapshot was taken, changing both data and schema if required.

```
$ hbase shell
hbase> disable 'myTable'
hbase> restore_snapshot 'myTableSnapshot-122112'
```



Note

Since Replication works at log level and snapshots at file-system level, after a restore, the replicas will be in a different state from the master. If you want to use restore, you need to stop replication and redo the bootstrap.

In case of partial data-loss due to misbehaving client, instead of a full restore that requires the table to be disabled, you can clone the table from the snapshot and use a Map-Reduce job to copy the data that you need, from the clone to the main one.

1.7. Snapshots operations and ACLs

If you are using security with the AccessController Coprocessor, only a global administrator can take, clone, or restore a snapshot, and these actions do not capture the ACL rights. This

means that restoring a table preserves the ACL rights of the existing table, while cloning a table creates a new table that has no ACL rights until the administrator adds them.

1.8. Export to another cluster

The ExportSnapshot tool copies all the data related to a snapshot (hfiles, logs, snapshot metadata) to another cluster. The tool executes a Map-Reduce job, similar to distcp, to copy files between the two clusters, and since it works at file-system level the hbase cluster does not have to be online. The HBase Snapshot Export tool must be run as hbase user. The HBase Snapshot Export tool must have temp directory, specified by "hbase.tmp.dir" (ie / grid/0/var/log/hbase), created on HDFS with hbase user as the owner.

To copy a snapshot called MySnapshot to an HBase cluster srv2 (hdfs://srv2:8020/hbase) using 16 mappers:

```
$ hbase org.apache.hadoop.hbase.snapshot.tool.ExportSnapshot -snapshot
 MySnapshot -copy-to hdfs://srv2:8020/hbase -mappers 16
```

2. User Guide - HDFS NFS Gateway

The NFS Gateway for HDFS supports NFSv3 and lets you mount HDFS as part of the client's local file system.

This release of NFS Gateway supports and enables the following usage patterns:

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

2.1. Prerequisites

In this section:

- The NFS gateway machine must have everything to run an HDFS client such as a Hadoop core JAR file and a HADOOP_CONF directory.
- The NFS gateway can be either on any DataNode, NameNode, or any HDP client machine. Start the NFS server on that machine.

2.2. Instructions

To configure and use the HDFS NFS gateway, complete the following steps:

1. Configure settings for HDFS NFS gateway:

   NFS gateway uses the same configurations as used by the NameNode and DataNode. Configure the following properties based on your application's requirements:

   a. Edit the `hdfs-default.xml` file on your NFS gateway machine and modify the following property:

```
<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, make sure this property is not disabled in the configuration file. Only 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 following property to `hdfs-site.xml`:

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



Note

NFS client often reorders writes. Sequential writes can arrive at the NFS gateway at random order. This directory is used to temporarily save out-of-order writes before writing to HDFS. One needs to make sure 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 if a worst-case write reorder happens to every file.

c. Update the following property to `hdfs-site.xml`:

```
<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 specifiy an access privilege, the default machine access to exports is `readonly`. Separate machine dentries by `;`. For example, `192.168.0.0/22 rw ; host*.example.com ; host1.test.org ro;`.

Restart the NFS gateway after this property is updated.

d. Optional - Customize log settings.

Edit the `log4j.property` file to add the following:

To change trace level, add the following:

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

To get more details of ONCRPC requests, add the following:

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

2. Start the NFS gateway service.

Three daemons are required to provide NFS service: `rpcbind` (or `portmap`), `mountd` and `nfsd`. The NFS gateway process has both `nfsd` and `mountd`. It shares the HDFS root "/" as the only export. We recommend using the `portmap` included in NFS gateway package as shown below:

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

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

b. Start the included `portmap` package (needs root privileges):

```
hadoop portmap
```

OR

```
hadoop-daemon.sh start portmap
```

c. Start `mountd` and `nfsd`.

No root privileges are required for this command. However, verify that the user starting the Hadoop cluster and the user starting the NFS gateway are same.

```
hadoop nfs3
```

OR

```
hadoop-daemon.sh start nfs3
```



Note

If the `hadoop-daemon.sh` script starts the NFS gateway, its log can be found in the hadoop log folder.

d. Stop NFS gateway services.

```
hadoop-daemon.sh stop nfs3
hadoop-daemon.sh stop portmap
```

3. Verify validity of NFS-related services.

a. Execute the following command to verify that all the services are up and running:

```
rpcinfo -p $nfs_server_ip
```

You should see 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
```

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

```
showmount -e $nfs_server_ip
```

You should see output similar to the following:

```
Exports list on $nfs_server_ip :
/ (everyone)
```

4. Mount the export "/".

Currently NFS v3 is supported and uses TCP as the transportation protocol is TCP. The users can mount the HDFS namespace as shown below:

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

Then the users can access HDFS as part of the local file system except that hard/symbolic link and random write are not supported in this release.



Note

Because NLM is not supported, the mount option `nolock` is needed.

User authentication and mapping:

NFS gateway in this release uses AUTH_UNIX style authentication which means that the the login user on the client is the same user that NFS passes to the HDFS. For example, if the NFS client has current user as `admin`, when the user accesses the mounted directory, NFS gateway will access HDFS as 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.

5. Set up client machine users to interact with HDFS through NFS.

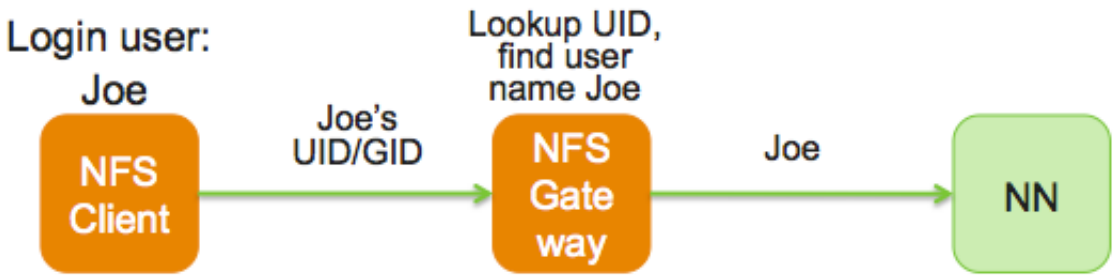
NFS gateway converts the UID to user name and HDFS uses username for checking permissions.

The system administrator must ensure that the user on NFS client machine has the same name and UID as that on the NFS gateway machine. This is usually not a problem if you use the same user management system such as LDAP/NIS to create and deploy users to HDP nodes and to client node.

If the user is created manually, you might need to modify the UID on either the client or NFS gateway host in order to make them the same:

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

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



3. User Guide - HDFS Snapshots

HDFS Snapshots are read-only point-in-time copies of the file system. Snapshots can be taken on a subtree of the file system or the entire file system. Some common use cases of snapshots are data backup, protection against user errors and disaster recovery.

The implementation of HDFS Snapshots is efficient:

1. Snapshot creation is instantaneous: the cost is **O(1)** excluding the inode lookup time.
2. Additional memory is used only when modifications are made relative to a snapshot: memory usage is **O(M)** where **M** is the number of modified files/directories.
3. Blocks in datanodes are not copied: the snapshot files record the block list and the file size. There is no data copying.
4. Snapshots do not adversely affect regular HDFS operations: modifications are recorded in reverse chronological order so that the current data can be accessed directly. The snapshot data is computed by subtracting the modifications from the current data.

In this document:

- Snapshottable Directories
- Snapshot Paths
- Snapshot Operations

3.1. Snapshottable Directories

Snapshots can be taken on any directory once the directory has been set as **snapshottable**. A snapshottable directory is able to accommodate 65,536 simultaneous snapshots. There is no limit on the number of snapshottable directories. Administrators may set any directory to be snapshottable. If there are snapshots in a snapshottable directory, the directory can be neither deleted nor renamed before all the snapshots are deleted.

3.2. Snapshot Paths

For a snapshottable directory, the path component **".snapshot"** is used for accessing its snapshots. Suppose `/foo` is a snapshottable directory, `/foo/bar` is a file/directory in `/foo`, and `/foo` has a snapshot `s0`. Then, the path `/foo/.snapshot/s0/bar` refers to the snapshot copy of `/foo/bar`. The usual API and CLI can work with the ".snapshot" paths. The following are some examples:

- Listing all the snapshots under a snapshottable directory: `hadoop dfs -ls /foo/.snapshot`
- Listing the files in snapshot `s0`: `hadoop dfs -ls /foo/.snapshot/s0`
- Copying a file from snapshot `s0`: `hadoop dfs -cp /foo/.snapshot/s0/bar /tmp`

The name ".snapshot" is now a reserved file name in HDFS so that users cannot create a file/directory with ".snapshot" as the name. If ".snapshot" is used in a previous version of HDFS, it must be renamed before upgrade; otherwise, upgrade will fail.

3.3. Snapshot Operations

Snapshot operations are grouped into the following two categories:

- Administrator Operations
- User Operations

3.3.1. Administrator Operations

The operations described in this section require superuser privileges.

- Allow Snapshots: Allowing snapshots of a directory to be created. If the operation completes successfully, the directory becomes snapshottable.
  - Command:

```
hadoop dfsadmin -allowSnapshot $path
```

  - Arguments:

Table 3.1. Administrator Operations - Allow Snapshots

Parameter name	Description
path	The path of the snapshottable directory.

See also the corresponding Java API `void allowSnapshot(Path path)` in `HdfsAdmin`.

- Disallow Snapshots: Disallowing snapshots of a directory to be created. All snapshots of the directory must be deleted before disallowing snapshots.
  - Command:

```
hadoop dfsadmin -disallowSnapshot $path
```

  - Arguments:

Table 3.2. Administrator Operations - Disallow Snapshots

Parameter name	Description
path	The path of the snapshottable directory.

See also the corresponding Java API `void disallowSnapshot(Path path)` in `HdfsAdmin`.

3.3.2. User Operations

The section describes user operations. Note that HDFS superuser can perform all the operations without satisfying the permission requirement in the individual operations.

- Create Snapshots: Create a snapshot of a snapshottable directory. This operation requires owner privilege to the snapshottable directory.
  - Command:

```
hadoop dfs -createSnapshot $path $snapshotName
```

  - Arguments:

Table 3.3. User Operations - Create Snapshots

Parameter name	Description
path	The path of the snapshottable directory.
snapshotName	The snapshot name, which is an optional argument. When it is omitted, a default name is generated using a timestamp with the format "'s'yyyyMMdd-HHmmss.SSS", e.g. "s20130412-151029.033".

See also the corresponding Java API `Path createSnapshot(Path path)` and `Path createSnapshot(Path path, String snapshotName)` in `FileSystem`. The snapshot path is returned in these methods.

- Delete Snapshots: Delete a snapshot of from a snapshottable directory. This operation requires owner privilege of the snapshottable directory.
  - Command:

```
hadoop dfs -deleteSnapshot $path $snapshotName
```

  - Arguments:

Table 3.4. User Operations - Delete Snapshots

Parameter name	Description
path	The path of the snapshottable directory.
snapshotName	The snapshot name.

See also the corresponding Java API `void deleteSnapshot(Path path, String snapshotName)` in `FileSystem`.

- Rename Snapshots: Rename a snapshot. This operation requires owner privilege of the snapshottable directory..
  - Command:

```
hadoop dfs -renameSnapshot $path $oldName $newName
```

  - Arguments:

Table 3.5. User Operations - Rename Snapshots

Parameter name	Description
path	The path of the snapshottable directory.
oldName	The old snapshot name.
newName	The new snapshot name.

See also the corresponding Java API `void renameSnapshot(Path path, String oldName, String newName)` in `FileSystem`.

- Get Snapshottable Directory Listing: Get all the snapshottable directories where the current user has permission to take snapshots.
  - Command:

```
hadoop lsSnapshottableDir
```

  - Arguments:

Table 3.6. User Operations - Get Snapshottable Directory Listing

Parameter name	Description
path	The path of the snapshottable directory.
snapshotName	The snapshot name.

See also the corresponding Java API `SnapshottableDirectoryStatus[] getSnapshottableDirectoryListing()` in `DistributedFileSystem`.

- Get Snapshots Difference Report: Get the differences between two snapshots. This operation requires read access privilege for all files/directories in both snapshots.
  - Command:

```
hadoop snapshotDiff $path $fromSnapshot $toSnapshot
```

  - Arguments:

Table 3.7. User Operations - Get Snapshots Difference Report

Parameter name	Description
path	The path of the snapshottable directory.
fromSnapshot	The name of the starting snapshot.
toSnapshot	The name of the ending snapshot.

4. Add HDP Maven Repository to Existing Project

Apache Maven is a software project management and comprehension tool. Based on the concept of a project object model (POM), Maven can manage a project's build, reporting and documentation from a central piece of information. Maven projects are defined by their **Project Object Model** or `pom`. This file is located in the base directory of a maven project and is called pom.xml.

Use one of the following options to add HDP Maven repository as a default repository in your existing project:

- **Option I: Add HDP Maven repository to existing Maven project**

  A repository in Maven is used to hold build artifacts and dependencies of varying types. There are strictly only two types of repositories: local and remote.

  The local repository refers to a copy on your own installation that is a cache of the remote downloads, and also contains the temporary build artifacts that you have not yet released. Remote repositories refer to any other type of repository, accessed by a variety of protocols such as `file://` and `http://` These repositories might be a truly remote repository set up by a third party to provide their artifacts for downloading (for example, `repo.maven.apache.org` hosts Maven's central repository). Other "remote" repositories may be internal repositories set up on a file or HTTP server within your company, used to share private artifacts between development teams and for releases.

  To add HDP Maven repository, add the following lines to your Maven project's pom.xml file:

```
<repositories>
    <repository>
      <releases>
        <enabled>true</enabled>
        <updatePolicy>always</updatePolicy>
        <checksumPolicy>warn</checksumPolicy>
      </releases>
      <snapshots>
        <enabled>false</enabled>
        <updatePolicy>never</updatePolicy>
        <checksumPolicy>fail</checksumPolicy>
      </snapshots>
      <id>HDPReleases</id>
      <name>HDP Releases</name>
      <url>http://repo.hortonworks.com/content/repositories/releases/</url>
      <layout>default</layout>
    </repository>
  </repositories>
```

- **Option II: Add HDP Maven repository to existing Ant/Ivy project**

  Apache Ivy repositories are configured inside the `<resolvers>` element of an `ivysettings.xml` file. Usually, the resolvers (where to get those required artifacts) are provided through a separate file,`ivysettings.xml` file.

The `ivysettings.xml` file holds a chain of Ivy resolvers used for both resolution and publishing (deployment). Resolvers exist for both regular artifacts and Ivy module files. Apache Ivy uses `chain` to define the preference order for the repositories. Inside the `<chain>` element, you will find a `<url>` element. The `<url>` element is a remote site that contains bundle dependencies.

To add HDP Maven repository to existing Ant/Ivy project, add a new resolver to the existing Ivy chain so that HDP versioned artifacts can be resolved.

- **Option III: Setup Maven proxy**

  It is often the case that users wish to set up a Maven proxy repository inside their corporate firewall and have developer instances resolve artifacts through such a proxy. Proxy repositories provide a single point of remote download for an organization. In addition to control and security concerns, Proxy repositories are primarily important for increased speed across a team. These scenarios can be realized by using internal Maven repositories and a Maven proxy.

  To setup maven proxy pointing to HDP Maven or Nexus repository, use the following URL (http://repo.hortonworks.com/content/repositories/releases/) for caching the HDP artifacts to your local or internal Maven, Nexus, or Archiva repositories respectively.