

HBase Backup and Restore CLI Specification (HBase-7912)

Last updated: March 31, 2014

HBase Backup/Restore CLI Specification

Specification

\$ hbase backup help

Usage: hbase backup COMMAND

where COMMAND is one of:

create	create a new backup image
cancel	cancel an ongoing backup
delete	delete an existing backup image
describe	show the detailed information of a backup image
history	show history of all successful backups
status	show the status of the latest backup request
convert	convert incremental backup WAL files into HFiles
merge	merge backup images
stop	remove table(s) from backup table set
show	show table(s) in backup table set

Enter 'help COMMAND' to see help message for each command

\$ hbase backup help create

Usage: hbase backup create <type> <backup_root_path> [tables] [-s name] [-convert] [-silent]

type	"full" to create a full backup image; "incremental" to create an incremental backup image
backup_root_path	The full root path to store the backup image, the prefix can be gpfs, hdfs or webhdfs

Options:

tables	If no tables (") are specified, all tables are backed up. Otherwise it is a comma separated list of tables.
-s name	Use the specified snapshot for full backup
-convert	For an incremental backup, convert WAL files to HFiles
-silent	Exit silently w/o blocking to show progress

\$ hbase backup help cancel

Usage: hbase backup cancel <backup_id>

Cancel an ongoing backup operation

backup_id	The id identifying the backup operation
-----------	---

\$ hbase backup help delete

Usage: hbase backup delete <backup_id> [tables]
backup_id The token identifies the backup and decides which backup image to delete. The token can be obtained through "hbase backup history" command. The backup instance is removed from history list if successful.

Options:

tables Only delete backup image of selected tables from the backup path, table list is separated by comma.

\$ hbase backup help describe

Usage: hbase backup describe <backup_root_path> [backup_id]
Show the detailed information of the backup image, including its backup lineages

backup_root_path The parent location where the backup images are stored

Options:

backup_id The id identifying the backup image

Notes:

1. This command is used on the target cluster

\$ hbase backup help history

Usage: hbase backup history [tables]
Show the history of all successful backups, including start time, complete time, type, location, size and tables

Options:

tables Only show the history of successful backups containing the specified table(s)

Notes:

1. This command is used on the source cluster

\$ hbase backup help status

Usage: hbase backup status
Show the status of the latest backup request

\$ hbase backup help convert

Usage: hbase backup convert <backup_root_path> <backup_id> [-keep]

Convert incremental backup WAL files to HFiles

backup_root_path The parent location of backup images
backup_id The id identifying the backup image

Options:

-keep Keep original HLogs after conversion

\$ hbase backup help merge

Usage: hbase backup merge <backup_root_path> <end_backup_id>

[start_backup_id] [-t tables] [-keep]

Merge the backup images between "start_backup_id" and "end_backup_id". If "start_backup_id" is missing, merge all backup images up to the

"end_backup_id"

backup_root_path The parent location of backup images
end_backup_id The id identifying the last backup image to merge

Options:

start_backup_id The id identifying the first backup image to merge

-t tables A comma separated list of tables

-keep After merge, move the original files

to .mergeBackups directory

\$ hbase backup help stop

Usage: hbase backup stop <tables|-all>

Remove table(s) from the backup table set. The backup table set is used to track the lineage of the backup images

tables A comma separated list of table names.

-all Remove the backup table set

\$ hbase backup help show

Usage: hbase backup show

Show the table(s) in the backup table set. The backup table set is used to track the lineage of the backup images

\$ hbase restore help

Usage: hbase restore <backup_root_path> <backup_id> [<tables>
[tableMapping]] [-nodelete] [-overwrite] [-check] [-automatic]
backup_root_path The parent location where the backup images are
stored
backup_id The id identifying the backup image

Options:

tables If no tables ("") are specific, all tables in the
backup image are restored. Otherwise only selected tables from the
backup image are restored. Tables are separated by comma.
tableMapping A comma separated list of target tables. If
specified, each table in <tables> must have a mapping.
-nodelete With this option, backup image is not deleted after
restore. Otherwise the image is deleted if restore succeeds.
-overwrite With this option, restore overwrites to the existing
table if there's any in restore target. The existing table must be
online before restore.
-check With this option, restore sequence and dependencies
are checked and verified without executing the restore
-automatic With this option, all the dependencies are
automatically restored together with this backup image following the
correct order. The restore dependencies can be checked by using "-
check" option, or using "hbase backup describe" command. Without this
option, only this backup image is restored

Descriptions

[backup_root]

This denotes the parent directory of all backup images in a target cluster. This parameter is required in many of the commands. It helps the backup/restore utility to track the lineages of the backup images. The directory names of the backup images are automatically generated by the backup utility. Users use a backup_id together with backup_root_path to specify a backup image in the target cluster.

[tables]

This parameter is a comma delimited list of table names.

[backup_id]

Each backup image is uniquely identified by a internally generated backup_id. Merge operation will choose a backup_id from the backup_ids of the backup images being merged as the unique id of the new backup image.

Backup Table Set

As of today, HBase's transaction logs (HLogs) are still global (i.e., database level). In order to perform automatic conversion and automatic restore based on the backup image lineage, backup/restore utility internally tracks a global set of tables. This set of tables can be modified with "backup add_table", "backup remove_table" commands.