

Use of HDFS Snapshots for External Table Replication

At Source: Two Snapshots are maintained (initial & old) which are used to compute the diff.

At Target: One Snapshot is maintained which is the same as the older snapshot at the source.(old)

DistCp Copy Modes:

- **InitialCopy:** The mode used when the directory is being copied for the first time using snapshots. This is a normal distCp from the source snapshot directory to target. Since during first time only one snapshot is available, so diff can not be computed
- **DiffCopy:** The copy mode used when two snapshots are available and the copy is done using snapshot diff.
- **FallbackCopy:** The copy mode used when snapshot is not available or we are unable to create snapshots at source, like if the FileSystem doesn't support snapshots, or the client can not allow snapshot on the directory being copied.

Configurations:

The snapshots if enabled would work implicitly for the external warehouse location and tables/partitions outside the Warehouse location can be explicitly added as part of the configuration. The explicitly added paths, needs to be managed by the client itself.

Name of Config	Value to be Set	Description
hive.repl.externaltable.snapshotdiff.copy	true	Enables the snapshot diff copy
hive.repl.externaltable.snapshot.prefix	"Some string"	Unique string prefix to be used for creating snapshots
hive.repl.externatable.snapshot.paths	Comma separated paths	List of external paths to be used for snapshots, for tables outside the warehouse location
hive.repl.external.warehouse.single.copy.task	true	To enable single copy task for warehouse directory

***Process Flow:**

>> The DirCopyTasks are always created during dump time, execution of them during load or dump is dependent on the configuration, so the decision to use which mode of snapshot is done during DirCopyTask creation itself.

- ❖ Check for the configurations if snapshot copy is enabled, and get the snapshotprefix and the list of external table paths from the configuration.
- ❖ Identify whether this directory is being copied first time using snapshots or has already been copied using snapshots, GetFileStatus with the snapshot path will get that information.
- ❖ If first time, attempt allowing & creating snapshots for the paths to be processed using snapshots(The Wh path & the configured path) and create a task as InitialCopy
- ❖ If not the first time, rename the initial snapshot to old snapshot and create a new initial snapshot and create a task as DiffCopy

- ❖ If there is any error during the process, create a task as FallbackCopy.

***Fallback:** Provide a mechanism to fallback to regular distCp mode, in case snapshots can not be created due to some issues.

Allow and create snapshots are done in a safe way, if they fail, the exception is swallowed and the copy task is created with FallbackCopy mode, irrespective of snapshot copy enabled.

***Metrics:** The paths where snapshots can not be enabled during dump are stored in metrics, for further analysis and troubleshooting.

***CleanUp:** Delete and disallow snapshots during dump, if the database location is Altered to another location.

Discussed Cases :

- Creation of .snapshot directory explicitly: .snapshot is a reserved keyword in HDFS, so can't be created. Will throw : org.apache.hadoop.ipc.RemoteException(org.apache.hadoop.HadoopIllegalArgumentExpection): ".snapshot" is a reserved name.
- Deletion of Snapshot on Alter Location change : during dump the Alter database location attempts to disallow and drop the snapshots.
- Fallback, in case unable to create a snapshot: Checks present during copy task creation, which marks a Fallback case, Secondly checks present, if the distCp copy with diff isn't success for whichever reason, normal distCp would be attempted