

# Permission Inheritance in Hive

This document describes how attributes (permission, group, extended ACL's) of files representing Hive data are determined.

## HDFS Background

- When a file or directory is created, its owner is the user identity of the client process, and its group is inherited from parent (the BSD rule). Permissions are taken from default umask. Extended Acl's are taken from parent unless they are set explicitly.

## Goals

To reduce need to set fine-grain file security props after every operation, users may want the following Hive warehouse file/dir to auto-inherit security properties from their directory parents:

- Directories created by new database/table/partition/bucket
- Files added to tables via load/insert
- Table directories exported/imported (open question of whether exported table inheriting perm from new parent needs another flag)

What is inherited:

- Basic file permission
- Groups (already done by HDFS for new directories)
- Extended ACL's (already done by HDFS for new directories)

This inheritance of extended ACL's is literal, all extended ACL's are copied to children as is, including ACL's for the defaultGroup.



One room for improvement may be to follow HDFS semantics for the defaultGroup, which is as follows:

"When a new file or sub-directory is created, it automatically copies the default ACL of its parent into its own access ACL. A new sub-directory also copies it to its own default ACL. In this way, the default ACL will be copied down through arbitrarily deep levels of the file system tree as new sub-directories get created." ([https://hadoop.apache.org/docs/current/hadoop-project-dist/hadoop-hdfs/HdfsPermissionsGuide.html#ACLS\\_Access\\_Control\\_Lists](https://hadoop.apache.org/docs/current/hadoop-project-dist/hadoop-hdfs/HdfsPermissionsGuide.html#ACLS_Access_Control_Lists))

See [HIVE-11481](#).

## Behavior

- When "[hive.warehouse.subdir.inherit.perms](#)" flag is enabled in Hive, Hive will try to do all the following inheritances.
  - Database directory inherits from warehouse directory.
  - Table directory inherits from database directory, or from warehouse directory if it is part of the default database.
  - External table directory inherits from parent directory.
  - Partition directory inherits from table directory. (As of [Hive 1.1.0](#).)
  - Data files inherit from table or partition directory.
- Failure by Hive to inherit will not cause operation to fail. Rule of thumb of when security-prop inheritance will happen is the following:
  - To run chmod, a user must be the owner of the file, or else a super-user.
  - To run chgrp, a user must be the owner of files, or else a super-user.
  - Hence, user that hive runs as (either 'hive' or the logged-in user in case of impersonation), must be super-user or owner of the file whose security properties are going to be changed

## Version Information

Most of this functionality was added as of Hive 0.14.

See umbrella JIRA [HIVE-6892](#) for details.

[hive.warehouse.subdir.inherit.perms](#) was removed in Hive 3.0.0. The feature is no longer needed in Hive as the traditional permission model has largely been replaced by external security systems such as Ranger and Sentry. A user may choose SQLStdAuth which ships with Hive if user doesn't want to use an external security system.