

HIVE-13966:- DbNotificationListener: can loose DDL operation notifications

Rahul Sharma - 08/25/2016

Problem:

The code for each API in HiveMetaStore.java is like this:

1. openTransaction()
2. – operation--
3. commit() or rollback() based on result of the operation.
4. add entry to notification log (unconditionally)

Two Issues:

1. If the operation is failed (in step 2), we still add entry to notification log. It is still ok as this is the case of false positive.
2. If the operation is successful and adding to notification log failed, the user will get an MetaException. It will not rollback the operation, as it is already committed. We need to handle this case so that we will not have false negatives.

Solution:

After following the discussion upstream the solution was to provide a configurable property in hive, where all the listeners which need to run within the metastore operation transaction can be added.

METASTORE_TRANSACTIONAL_EVENT_LISTENERS("hive.metastore.transactional.event.listeners", "", + "The Listeners added to this property will run within the actual metastore ops transaction"),

The attached patch(HIVE-13966.2.patch) solves the second issue.

Notes:

There is a tag interface *TransactionalMetaStoreEventListener* which has to be implemented by the metastore listeners which want to be executed within in the transaction.

As per the solution implemented, the user is responsible for not adding a listener to both *hive.metastore.event.listeners* and *hive.metastore.transactional.event.listeners* property in hive-site.xml which may result in duplicate events.

Positive test cases can be tested by running the DBNotificationListener as a transactional listener.

Negative test cases are handled running the DBNotificationListener as transactional listener and renaming the notification_log table, which prevents writing to the notification log and then verify that the entire metastore operation is rolled back.