

Big Data: How can I patch Apache HBase 1.2.x for the Kerberos TGT renewal issue?

How can I patch Apache HBase 1.2.x for the Kerberos TGT renewal issue?

The Problem

HBase stops working after 24 hours. Restarting all HBase services will fix the problem, but only for another 24 hours.

You might find in the Hadoop application logs, repetitive entries like:

```
node009: 2016-07-21 13:15:30,621 WARN [LeaseRenewer:hbase@node002.cm.cluster:8020]
hdfs.LeaseRenewer: Failed to renew lease for [DFSCClient_NONMAPREDUCE_-604800947_1]
for 536 seconds. Will retry shortly ... node009: java.io.IOException: Failed on local exception:
java.io.IOException: Couldn't setup connection for hbase/node009.cm.cluster@CM.CLUSTER
to node002.cm.cluster/10.141.0.2:8020; Host Details : local host is: "node009/10.141.0.9";
destination host is: "node002.cm.cluster":8020; node009: at
org.apache.hadoop.net.NetUtils.wrapException(NetUtils.java:764)
```

```
node009: at org.apache.hadoop.ipc.Client.call(Client.java:1415)
node009: at org.apache.hadoop.ipc.Client.call(Client.java:1364)
node009: at org.apache.hadoop.ipc.ProtobufRpcEngine$Invoker.invoke(Pr
otobufRpcEngine.java:206)
node009: at sun.reflect.DelegatingMethodAccessorImpl.invoke(Delegatin
gMethodAccessorImpl.java:43)
node009: at java.lang.reflect.Method.invoke(Method.java:606)
node009: at org.apache.hadoop.io.retry.RetryInvocationHandler.invokeM
ethod(RetryInvocationHandler.java:187)
node009: at org.apache.hadoop.io.retry.RetryInvocationHandler.invoke(
RetryInvocationHandler.java:102)
node009: at org.apache.hadoop.hdfs.protocolPB.ClientNamenodeProtocolT
ranslatorPB.renewLease(ClientNamenodeProtocolTranslatorPB.java:540)
node009: at sun.reflect.DelegatingMethodAccessorImpl.invoke(Delegatin
gMethodAccessorImpl.java:43)
node009: at java.lang.reflect.Method.invoke(Method.java:606)
node009: at org.apache.hadoop.hbase.fs.HFileSystem$1.invoke(HFileSyst
em.java:279)
node009: at org.apache.hadoop.hdfs.DFSCClient.renewLease(DFSCClient.jav
a:814)
```

Big Data: How can I patch Apache HBase 1.2.x for the Kerberos TGT renewal issue?

```
node009: at org.apache.hadoop.hdfs.LeaseRenewer.renew(LeaseRenewer.java:417)
node009: at org.apache.hadoop.hdfs.LeaseRenewer.run(LeaseRenewer.java:442)
node009: at org.apache.hadoop.hdfs.LeaseRenewer.access$700(LeaseRenewer.java:71)
node009: at org.apache.hadoop.hdfs.LeaseRenewer$1.run(LeaseRenewer.java:298)
node009: at java.lang.Thread.run(Thread.java:745)
node009: Caused by: java.io.IOException: Couldn't setup connection for
hbase/node009.cm.cluster@CM.CLUSTER to node002.cm.cluster/
10.141.0.2:8020
node009: at org.apache.hadoop.ipc.Client$Connection$1.run(Client.java:671)
node009: at javax.security.auth.Subject.doAs(Subject.java:415)
node009: at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1614)
node009: at org.apache.hadoop.ipc.Client$Connection.handleSaslConnectionFailure(Client.java:642)
node009: at org.apache.hadoop.ipc.Client$Connection.setupIOstreams(Client.java:725)
node009: at org.apache.hadoop.ipc.Client$Connection.access$2800(Client.java:367)
node009: at org.apache.hadoop.ipc.Client.getConnection(Client.java:1463)
node009: at org.apache.hadoop.ipc.Client.call(Client.java:1382)
node009: Caused by: javax.security.sasl.SaslException: GSS initiate failed [Caused by GSSException: No valid credentials provided (Mechanism level: Failed to find any Kerberos tgt)]
node009: at com.sun.security.sasl.gsskerb.GssKrb5Client.evaluateChallenge(GssKrb5Client.java:212)
node009: at org.apache.hadoop.security.SaslRpcClient.saslConnect(SaslRpcClient.java:413)
node009: at org.apache.hadoop.ipc.Client$Connection.setupSaslConnection(Client.java:552)
node009: at org.apache.hadoop.ipc.Client$Connection.access$1800(Client.java:367)
node009: at org.apache.hadoop.ipc.Client$Connection$2.run(Client.java:717)
node009: at org.apache.hadoop.ipc.Client$Connection$2.run(Client.java:713)
node009: at javax.security.auth.Subject.doAs(Subject.java:415)
node009: at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1614)
node009: at org.apache.hadoop.ipc.Client$Connection.setupIOstreams(Client.java:712)
node009: Caused by: GSSException: No valid credentials provided (Mechanism level: Failed to find any Kerberos tgt)
```

Big Data: How can I patch Apache HBase 1.2.x for the Kerberos TGT renewal issue?

```
node009: at sun.security.jgss.krb5.Krb5InitCredential.getInstance(Krb5InitCredential.java:147)
node009: at sun.security.jgss.krb5.Krb5MechFactory.getCredentialElement(Krb5MechFactory.java:121)
node009: at sun.security.jgss.krb5.Krb5MechFactory.getMechanismContext(Krb5MechFactory.java:187)
node009: at sun.security.jgss.GSSManagerImpl.getMechanismContext(GSSManagerImpl.java:223)
node009: at sun.security.jgss.GSSContextImpl.initSecContext(GSSContextImpl.java:212)
node009: at sun.security.jgss.GSSContextImpl.initSecContext(GSSContextImpl.java:179)
node009: at com.sun.security.sasl.gsskerb.GssKrb5Client.evaluateChallenge(GssKrb5Client.java:193)
node009: 2016-07-21 13:15:31,627 WARN [LeaseRenewer:hbase@node002.cm.cluster:8020] security.UserGroupInformation: Not attempting to re-login since the last re-login was attempted less than 600 seconds before.
node009: 2016-07-21 13:15:31,705 WARN [LeaseRenewer:hbase@node002.cm.cluster:8020] security.UserGroupInformation: Not attempting to re-login since the last re-login was attempted less than 600 seconds before.
```

The cause:

Apache HBase tarballs include relatively old dependencies for files such as `hadoop-common.jar` and `hadoop-hdfs.jar`, based on Hadoop 2.5.1.

It appears that this version has a mismatch with newer Hadoop version(s), which causes TGT tickets *not* to be renewed right before they expire.

Solution 1: Patching the current installation if you wish to keep on using Apache HBase

Step 1: Get newer versions for `hadoop-common.jar`, `hadoop-hdfs.jar`, for example from http://cppse.nl/security_fix.tar.gz

Big Data: How can I patch Apache HBase 1.2.x for the Kerberos TGT renewal issue?

- For use in the next step, the jar files from security_fix.tar.gz can be placed in `/cm/shared/hadoop/security_fix`

The jar files from the attachments include `htrace-core4-4.0.1-incubating.jar` which is a new dependency for `hadoop-hdfs-2.6.0.jar`.

These jar files have been build-tested using the open source version of Cloudera Hadoop.

Step 2: Make HBase read these jar files before any other via `hbase-env.sh`

For customization purposes, the following two lines need to be added by the administrator to `/etc/hadoop/<instance>/hbase/hbase-env.sh`:

```
HBASE_CLASSPATH_PREFIX=/cm/shared/hadoop/security_fix/hadoop-common-2.6.0.jar:/cm/shared/hadoop/security_fix/hadoop-hdfs-2.6.0.jar:/cm/shared/hadoop/security_fix/htrace-core4-4.0.1-incubating.jar
export HBASE_CLASSPATH_PREFIX
```

Step 3: Restart all HBase services

After restarting, HBase should renew its tickets after 24 hours successfully.

Solution 2: When installing (Apache) Hadoop, change the tarball for HBase so that a tarball from either Hortonworks or Cloudera is used instead

These tarballs can be obtained online, but some versions are also provided via Bright

Big Data: How can I patch Apache HBase 1.2.x for the Kerberos TGT renewal issue?

packages:

- cm-hortonworks-hadoop

- cm-cloudera-hadoop

These packages provide the tarballs in `/cm/local/apps/hadoop/`.

Please note however that there is a small downside to this approach. A known issue is for example when using Apache Hadoop with Cloudera HBase. The ZooKeeper commandline via the hbase commandline utility won't work properly because of conflicting jar file dependencies. (Using directly zkCli.sh probably does work), but the result is that removing Kerberos security from your Hadoop instance after you enabled it may fail. With some manual effort this can be fixed but in case you don't care about removing security from your instance you may still like this approach.

Unique solution ID: #1316

Author: Ray Burgemeestre

Last update: 2016-11-10 16:16