

HBase Connection Control

1.1 Overview

It is desirable to set the limit on the number of client connections permitted to the HBase server by controlling with certain system variables/parameters. Too many connections to the HBase server imply too many queries and MR jobs running on HBase. This can slow down the performance of the system and lead to denial of service. Hence such connections need to be controlled. Using too many connections may just cause thrashing rather than get more useful work done.

This is kind off inspired from <http://www.ebaytechblog.com/2014/08/21/quality-of-service-in-hadoop/#.VO2JXXyUe9y>

In case an operation is attempted and max configured connections are exceeding the global/user limit an error should be returned.

HBase will support controlling the maximum client connection to the server by the following checks:

- ✓ Enable/disable connection control check (**hbase.connectioncontrol.enable**)
- ✓ Global connection counter for connections does not exceed the global limit (**hbase.connectioncontrol.maxconnections**)
- ✓ Counter for the user does not exceed the per-user limit (**hbase.connectioncontrol.maxconnectionsperuser**)
- ✓ Counter for new coming connection does not exceed the connection limit (**hbase.connectioncontrol.maxconnectionsinperiod**) for the specified time-window frame (**hbase.connectioncontrol.limitperiod**) unit in **seconds**

If and only if all these checks are passed, request would be accepted and connection request will be granted.

1.2 Implementation Details

- By default connection control will be disabled.
- Connections to the server will be controlled only when HBase Kerberos security and connection control both are enabled.
- There will no limit/control on the number of connections requested by user having same name as Master principal or RegionServer principal or null.
- Default value for **hbase.connectioncontrol.maxconnections**, **hbase.connectioncontrol.maxconnectionsperuser** and **hbase.connectioncontrol.maxconnectionsperiod** is set to **Integer.MAX_VALUE** i.e., 2147483647.
- Default value for **hbase.connectioncontrol.limitperiod** is set 10 seconds.
- If value for **hbase.connectioncontrol.maxconnections**, **hbase.connectioncontrol.maxconnectionsperuser** and **hbase.connectioncontrol.maxconnectionsperiod** is found to be less than 0 then internally we will set these properties value as **Integer.MAX_VALUE** i.e., 2147483647.
- If value for **hbase.connectioncontrol.limitperiod** is found to be less than or equal to 0 then internally we will set this property value as 0.
- If **hbase.connectioncontrol.limitperiod** is 0 then no connection request will be checked against **hbase.connectioncontrol.maxconnectionsperiod** value.
- If a connection request exceeds the limit then server will throw **AccessDeniedException**.