

HBaseWireCompatibility20120221

HBase wire compatibility meeting notes

2/21/2012

Versioning:

0.94: just make it a performance release

keep wire compatibility with 92 for:

- 0.92 clients to 0.94 servers
- 0.94 clients to 0.92 servers (we'll check if this is possible, but we think so.. maybe one or two spots. Lars and Stack to investigate)
- rolling upgrade within the cluster for advanced users only (eg they have to turn off prefix compression during upgrade, etc)

0.96: "the singularity"

- aim to have stuff moved over to protobuf at this time, and probably remove old writable-based calls.
- try to remove ROOT, maybe HBASE-2600 if it got done?
- migration: perhaps offline, but has to be in-place, and preferably automatically done (eg master rewrites META)

ideas for compatibility of old clients to 0.96 (to help people move to singularity):

- maybe we backport a forward-compatible client onto 0.94 or 0.92 releases (rather than trying to make 0.96 back-compatible). eg. if we remove ROOT, add a forward-compatibility path which can use either ROOTless or with-ROOT, make people upgrade to 0.94.X client, then can upgrade server

Tasks:

Core types:

- [KeyValue](#), HRegionInfo, HRegionLocation, HTableDescriptor, [ServerName](#),
- some deprecated but still all over: HServerLoad?
- [ClusterStatus](#)

Actual conversion to protobuf RPCs:

- HRegionInterface - Jimmy
- HMasterInterface - Greg
- HMasterRegionInterface - Greg (while doing this, collapse RPCs where necessary)

Zookeeper data:

- Client-facing stuff: data in ROOT region location node, master address node (Stack)
- Cluster-internal stuff:
- /hbase/regionserver znodes
- distributed log splitting status, etc

Filters: (tricky because they're user-definable?)

- todd to investigate

Coprocessors:

- this is somewhat tricky
- Enis to investigate

Security:

salesforce folks to investigate (might come for free)

Replication:

- there's stuff in ZK, also uses client APIs (Lars/Chris/Jesse)

RPC payload + envelope

- Devaraj
- (move the headers, etc to Protobuf)

Management stuff for multi-version:

- add version string for each node in the cluster to ZK
- display this info on web UIs, JMX, etc

HBase shell: (Chris)

- has to know to deserialize the proto types instead, for example? we're not sure if it ever reads/deserializes stuff directly

On-disk formats:

- eg reference files use Writable
- HLog: HLogKey, WALEdits, etc (affects unclean upgrade, and replication)
- HFile metadata - currently it's a [MapWritable](#) - need to figure out how to keep this back-compatible against Writable
- /hbase/hbase.version file

How to integrate with Thrift? (facebook to investigate)