June2011DeveloperMeeting

This is the proposed agenda/meeting notes for the ZooKeeper Developer Meeting at Yahoo! on Jun 30, 2011 following the Hadoop Summit. This meeting is a chance for the core ZooKeeper developers (committers and contributors) to meet and kick around some ideas.(Please feel free to add items.)

Agenda items

- Code cleanup
- Move to maven
- Move to textile❓
- ZOOKEEPER-22
- Managing recipe implementations
- Move to standard serialization API (which one?)
- Change data tree to reflect future
- Hack/Bug fix/Commit athon

Agenda item suggestions

- phunt – should we start 4.x and when? If so, what should that entail:
    - API changes (e.g. long version fields rather than ints)
    - rip out nio entirely and replace with netty
    - simplify the concurrency model and implementation
    - we might consider doing as a sub-project, which would allow us to try out changes to our dev model (perhaps CTR rather than RTC, ease committership, etc...)
    - re-write everything, perhaps in scala 🍔

Meeting minutes

- Kerberos: we reviewed the patch for kerberos authentication. the approach looks good, but we suggested that the kerberos handshake should be done at the transport layer rather than all the way to FinalRequestProcessor. it would be nice to get into the next release.
- Code Cleanup: we should start standardizing on patterns to use. a place to start would be to look into using the executor api rather than threads.
- Move to Maven: pat already has a patch ready
- Move to Textile: yes!
- ZOOKEEPER-22: it is something we should work on 🍔
- Managing Recipes: it would be nice to manage recipes outside of the core api. the core committers don't have the bandwidth and it would be nice to decouple the releases. pat and ted to propose an approach to recruit recipe committers and handle releases.
- Serialization API: should we move to a more standardized serialization api. the general feeling seemed to be "if it aint broke, don't fix it".
- Reconfiguration: alex talked about the work for online cluster reconfiguration. (addition and removal of servers.) there was some skepticism that the solution was too complicated. alex will continue work and document the approach.
- Data Tree Changes: we talked about ZOOKEEPER-1092. one thing that emerged from the discussion is that what we really needed was a versioned data tree structure that was only used in PrepRequestProcessor but kept track of lastCommitted as well as future transactions. this would allow the simplified pipeline as proposed in ZOOKEEPER-1092 as well as allowing for read short cutting (servicing reads in PrepRP and responding directly to client) if there isn't any writes from that client ahead of the read.
- ZooKeeper at Facebook: Vishal gave an excellent review of large scale use of zookeeper at facebook. Please see the attached presentation.
- Load Balancing/Reresolution of DNS: the current method of random load balancing degrades with server failures and long lived client connections. we should do periodic rebalances. we also talk about the need to periodically reresolve DNS information for servers.
- DNS Bridge: it would be useful to have a DNS interface to zookeeper data.