

Cross-Datacenter Replication

- Single “Master” cluster
 - Other “Slave” clusters replicate the same tables
 - Master cluster synchronously sends HLogs over to slave clusters
 - Coordination among clusters is via Zookeeper
 - Zookeeper can be used like a file system to store control information
1. */hbase/replication/state*
 2. */hbase/replication/peers/<peer cluster number>*
 3. */hbase/replication/rs/<hlog>*