

1. Root/Meta table should be up at first, and can recover themselves if corresponding region server dies;
2. Only after AssignmentManager (AM) has recovered its in-memory region states, it can take ZooKeeper (ZK) events, and server shutdown handler (SSH) can assign user regions, because AM is not ready yet. ZK event handling and SSH could modify the states AM is rebuilding. Sometimes, they even try to access regions not yet known to AM at this moment.
3. In memory region states, region assignments, and region to region server mapping is consistent because of synchronization.

