



# HBase Disaster Recovery Solution at Huawei

Ashish Singhi

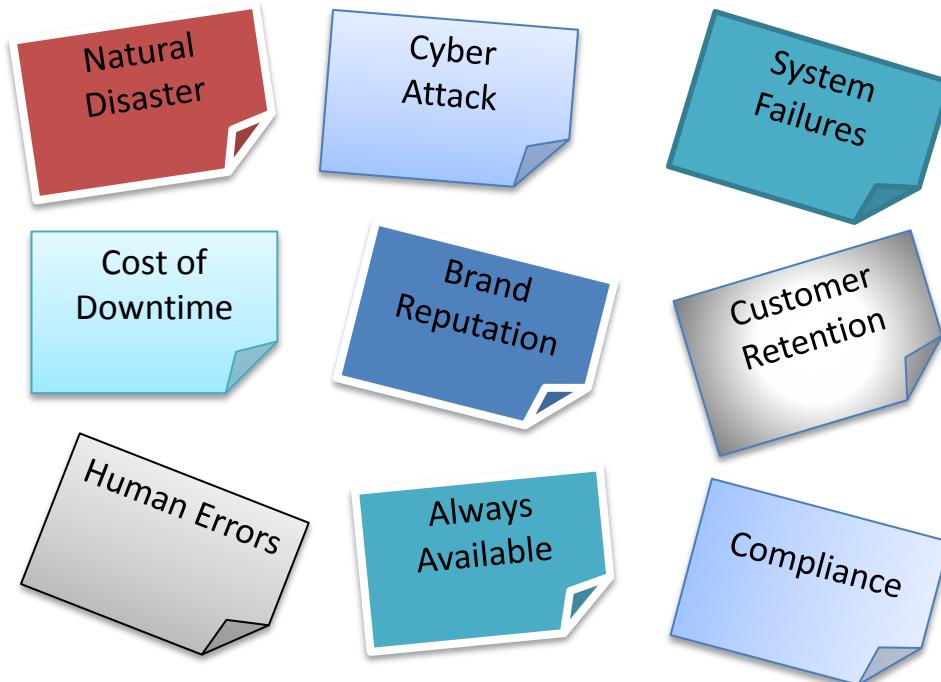
# About.html

- Senior Technical Leader at Huawei
- Around 6 years of experience in Big Data related projects
- Apache HBase Committer

# Agenda

- Why Disaster Recovery ?
- Backup Vs Disaster Recovery
- HBase Disaster Recovery
- Solution
- Miscellaneous
- Future Work

# Why Disaster Recovery ?



# Agenda

- Why Disaster Recovery ?
- **Backup Vs Disaster Recovery**
- HBase Disaster Recovery
- Solution
- Miscellaneous
- Future Work

# Backup Vs Disaster Recovery

Two different problems and solutions

	<b>Backup</b>	<b>Disaster Recovery</b>
Process	Archive items to cold media	Replicate to secondary site
Infrastructure	Medium level	Duplicate of active cluster (high level)
Cost	Affordable	Expensive
Restore process	One to few at a time	One to everything
Restore time	Slow	Fast
Production usage	Common	Rare

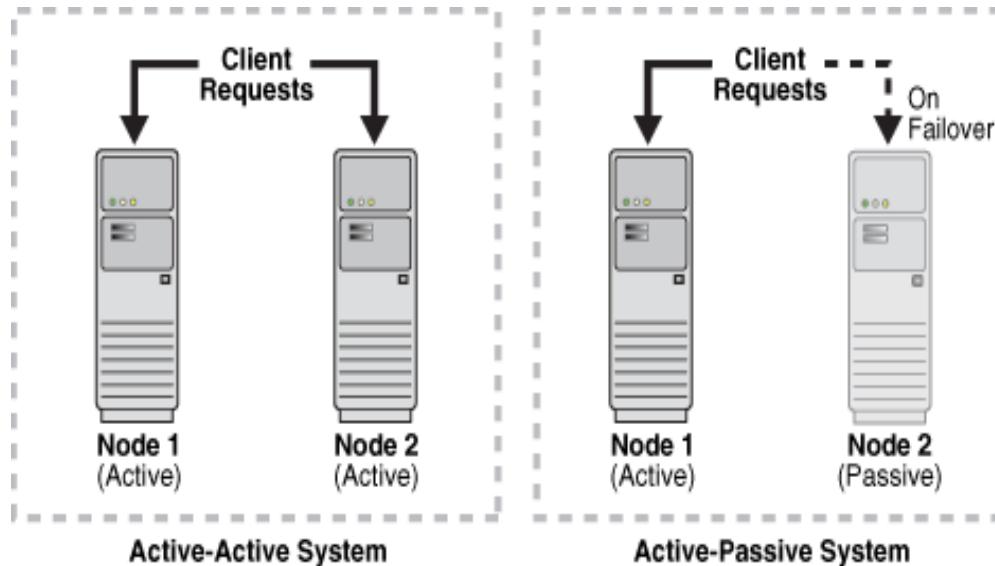
# Agenda

- Why Disaster Recovery ?
- Backup Vs Disaster Recovery
- **HBase Disaster Recovery**
- Solution
- Miscellaneous
- Future Work

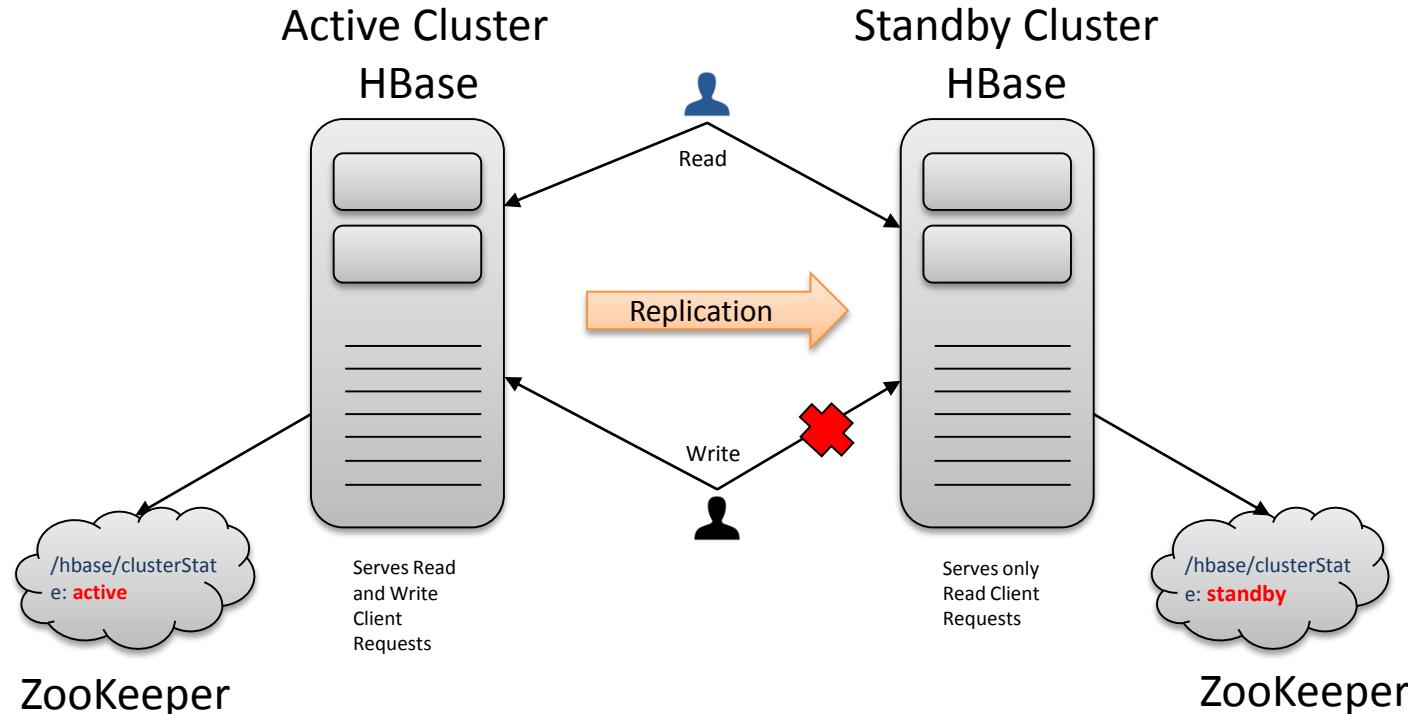
# HBase Disaster Recovery

- HBase Disaster recovery is based on replication, which mirrors data across a network in real time.
- The technology is used to move data from a local source location to one or more target locations.
- Replication over WAN has become an ideal technology for disaster recovery to prevent data loss in the event of failure.

# Deployment Strategies



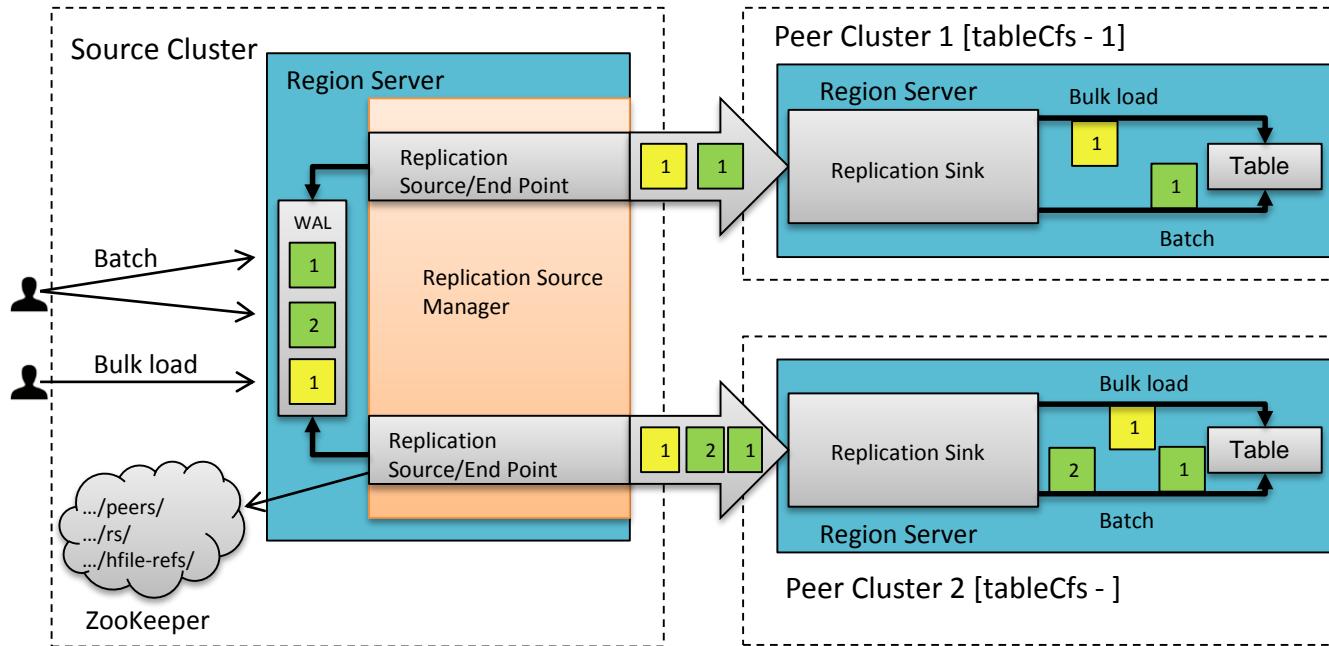
# Active – Standby Cluster



# Agenda

- Why Disaster Recovery ?
- Backup Vs Disaster Recovery
- HBase Disaster Recovery
- **Solution**
- Miscellaneous
- Future Work

# Replication



# Sync DDL Operations

- Synchronize the table properties across clusters
  - Any change in the source cluster, reflects immediately in the peer clusters.
  - Does not break the replication.
- An additional option with DDL command to sync
  - Internally sync those changes to peer clusters.

```
You can sync alter table operation in peer clusters also:  
hbase> alter 't1', NAME => 'f1', VERSIONS => 5, SYNC_PEER => true  
hbase> alter 't1', 'f1', {NAME => 'f2', VERSIONS => 10}, SYNC_PEER => true  
hbase> alter 't1', MAX_FILESIZE => '134217728', SYNC_PEER => true  
hbase> alter 'ns1:t1', NAME => 'f1', METHOD => 'delete', SYNC_PEER => true
```

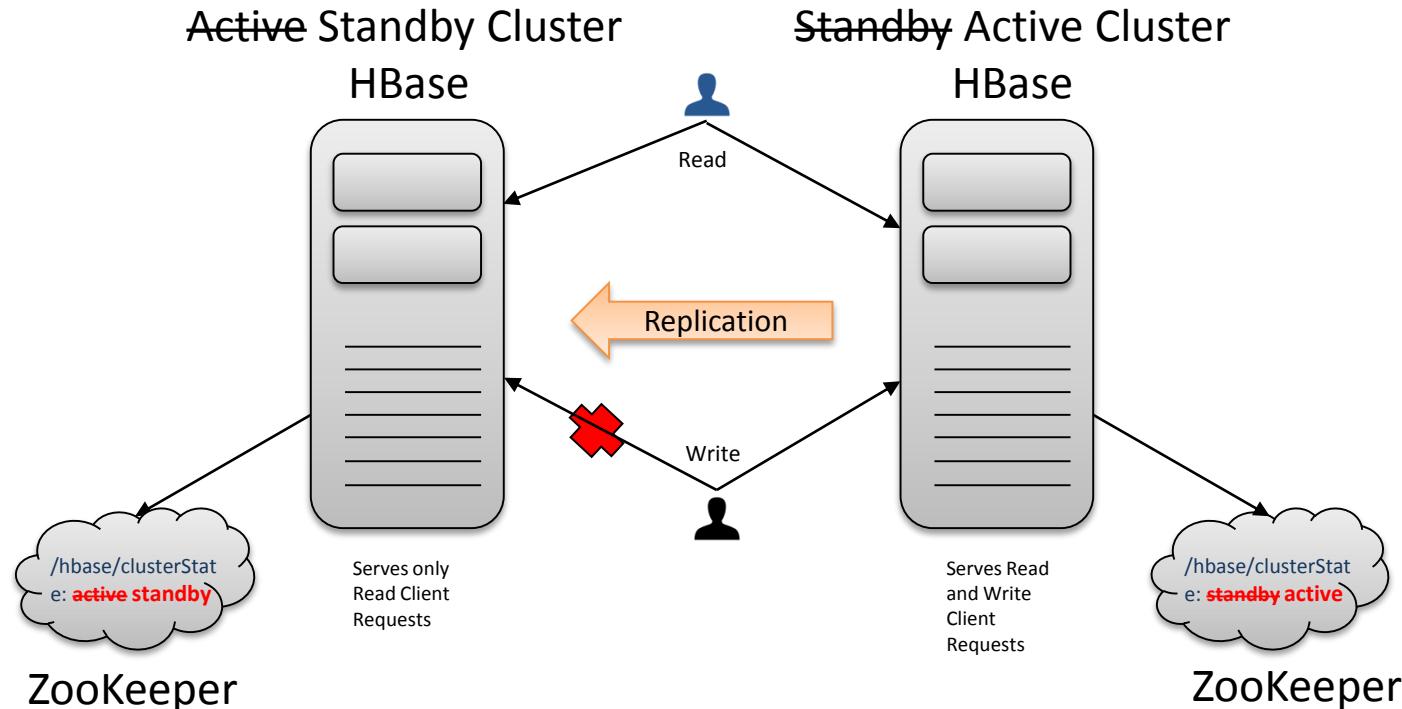
# Sync Security related Data

- Synchronize security related HBase data across the clusters
  - Any update in the source cluster ACL, Quota or Visibility Labels table, reflects immediately in peer clusters.
  - A custom WAL entry filter is added in replication for this.
  - Does not break the security for HBase data access.

# Read Only Cluster

- Enable a cluster to serve only read requests
  - A coprocessor based solution
  - Standby cluster will serve all the read requests
  - Standby cluster will serve write requests only if the requests is coming from a,
    - Super user
    - From a list of accepted IPs

# Cluster Recovery



# Agenda

- Why Disaster Recovery ?
- Backup Vs Disaster Recovery
- HBase Disaster Recovery
- Solution
- **Miscellaneous**
- Future Work

# Miscellaneous

- Increased the default *replication.source.ratio* to 0.5
- Adaptive *hbase.replication.rpc.timeout*
- Active cluster HDFS server configurations are maintained in Standby cluster ZooKeeper for bulk loaded data replication.

# Agenda

- Why Disaster Recovery ?
- Backup Vs Disaster Recovery
- HBase Disaster Recovery
- Solution
- Miscellaneous
- **Future Work**

# Future work

- Move HBase Replication tracking from ZooKeeper to HBase table (HBASE-15867)
- Copy bulk loaded data to peer with data locality
- Replication data network bandwidth throttling.

# Thank You !

<mailto:ashishsinghi@apache.org>

Twitter: ashishsinghi89