HarperDB

Clustering & Replication

Your Data Where You Need It, When You Need It

Clustering: Reap the benefits of parallel computing
Connect enterprise servers and/or devices together, across multiple geographic areas, to access the same database.

Replication: You specify where your data resides
Define residence of your tables to specify which servers and/or devices your data should reside. If the data doesn't live on a node, that node will know where to retrieve the data.

Time-to-Live: Edge data doesn't have to be big data
Shrink your edge data footprint by configuring time-to-live (TTL) on tables that reside on the edge. With TTL, data will replicate to your enterprise servers and be removed from your devices after a user-specified duration.

Every node is the master and the follower
HarperDB is configured to act as both the master and the follower. When a HarperDB node receives a transaction it immediately begins acting as the master for that transaction. Meanwhile, it could also be acting as a follower, receiving data from another HarperDB node.

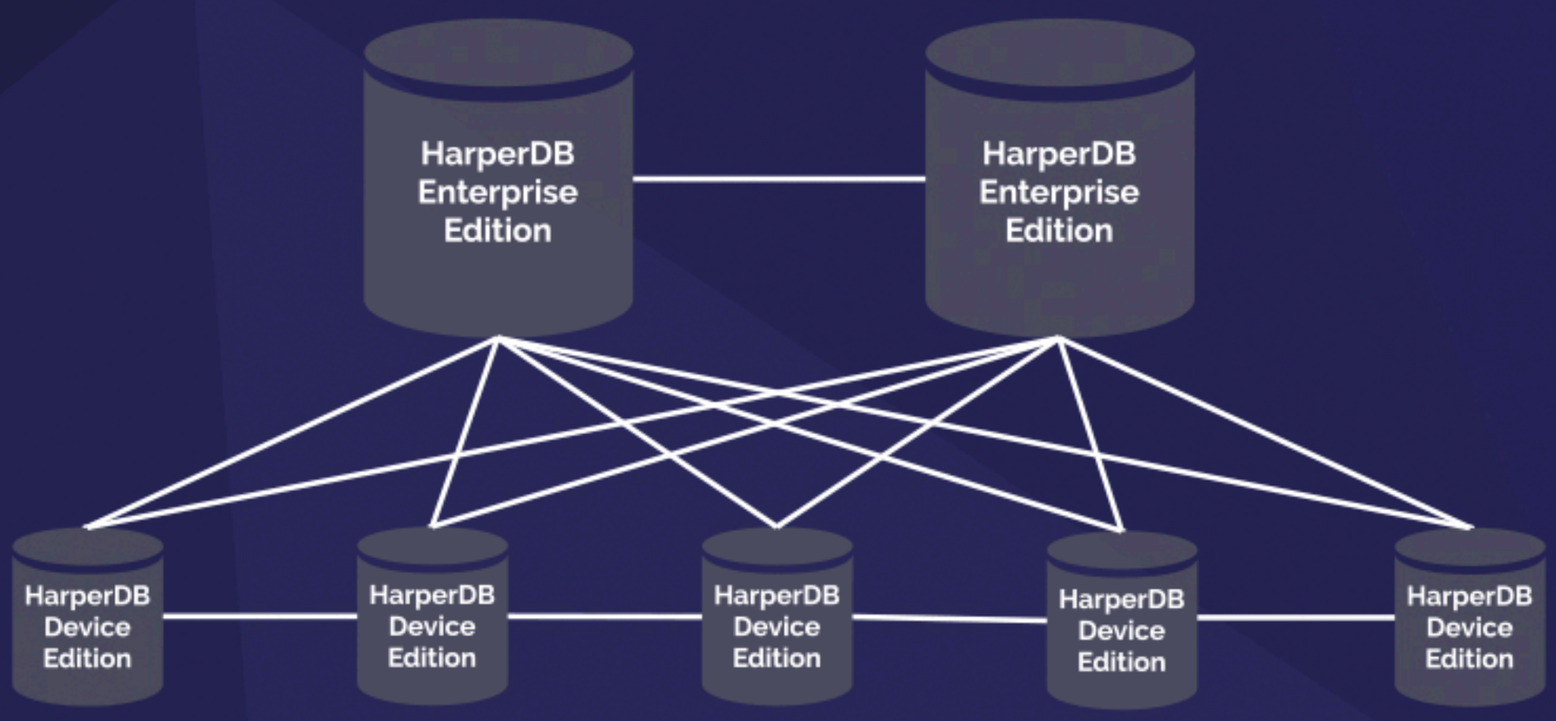

- Peer-to-Peer Configuration
- Zero Code Replication from the Edge to Cloud
- Increased Performance
- Reduced Hardware Cost
- Distributed Edge Querying

HarperDB's clustering and replication can be used for scalability, failover, data transport, and/or backups. Replication is robust, designed to function in environments with low or sporadic bandwidth and in cases where connections are not always available. Transactions are ACID and reliably queued to ensure consistency.

Using peer-to-peer clustering and replication, all nodes are aware of all schemas and tables, which enables distributed querying and true parallel computing. This allows for economical, redundant, and performant benefits of horizontal scale.

3000 Lawrence St.
Suite 145
Denver, CO 80205

www.harperdb.io
hello@harperdb.io
(650) 479-5641

HarperDB
Simplicity without Sacrifice