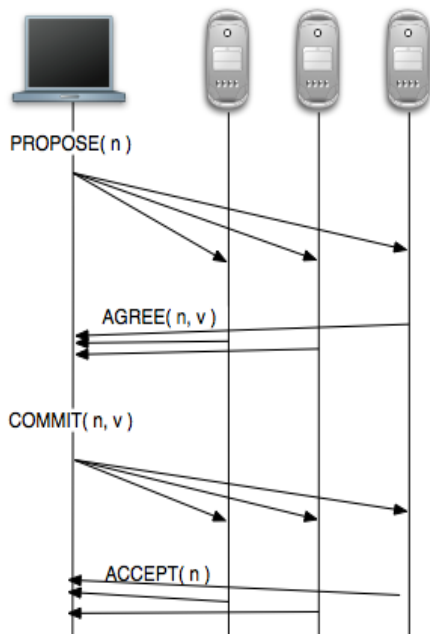


Big Data Blockchain Databases For the Enterprise


Trent McConaghy
@trentmc0

BIGCHAIN^{DB}

Everyone uses databases. How do they scale to big data? Answer: Distribute storage across many machines (sharding)



A “consensus” algorithm keeps distributed nodes in sync.



Example: Cassandra scaling.

More nodes = more throughput, more capacity

Banks & Other Enterprises Have Operationalized “Big Data” Distributed Databases



Example: MongoDB select customers



Along came blockchains, With new “blue ocean” benefits

Decentralized / shared control

Immutability / audit trail

Tokens / exchanges

A blockchain caveat or two

Completely new code bases

Reinventing consensus

No sharding = no scaling

No querying // single-node querying

But they're cool... right?

The BigchainDB approach: big data + blockchain.



1. Start with an Enterprise-Grade Distributed DB (MongoDB)
2. Engineer in Blockchain Characteristics

Decentralized /
Shared Control

- Each DB node is a federation node

Immutable /
Audit Trails

- Hash Previous Blocks
- Append-only

Native assets

- “Own” = have private key
- Asset lives on the database

Sample Engagements



Energy Deregulation – RWE

HR Credentials – Recruit Technologies

Financial network interoperability – Interledger / Ripple

Loyalty and Reward System – CapGemini

Financial Infrastructure – <private>

Medical Supply Chain – Tangent90

Voting – SettleMint

Streaming Music – Resonate

Database Partner – MongoDB

On Deploying



No need to re-do the whole stack,
Just add one more (special) database.



mongoDB



BigchainDB: The best of big data + blockchain = The Big Data Blockchain Database



	Traditional blockchains	Big Data Databases	BigchainDB
Immutability	✓		✓
Decentralized Control	✓		✓
Native Assets	✓		✓
Scalable		✓	✓
Queryability		✓	✓
Operationalized		✓	✓