RingBFT: Resilient Consensus over Sharded Ring Topology

Sajjad Rahnama

Suyash Gupta

Rohan Sogani

Dhruv Krishnan

Mohammad Sadoghi

Exploratory Systems Lab

University of California Davis

March, 2022

ExpoLab Creativity Unfolded

ResilientDB

UCDAVIS UNIVERSITY OF CALIFORNIA

Chronology of Consensus

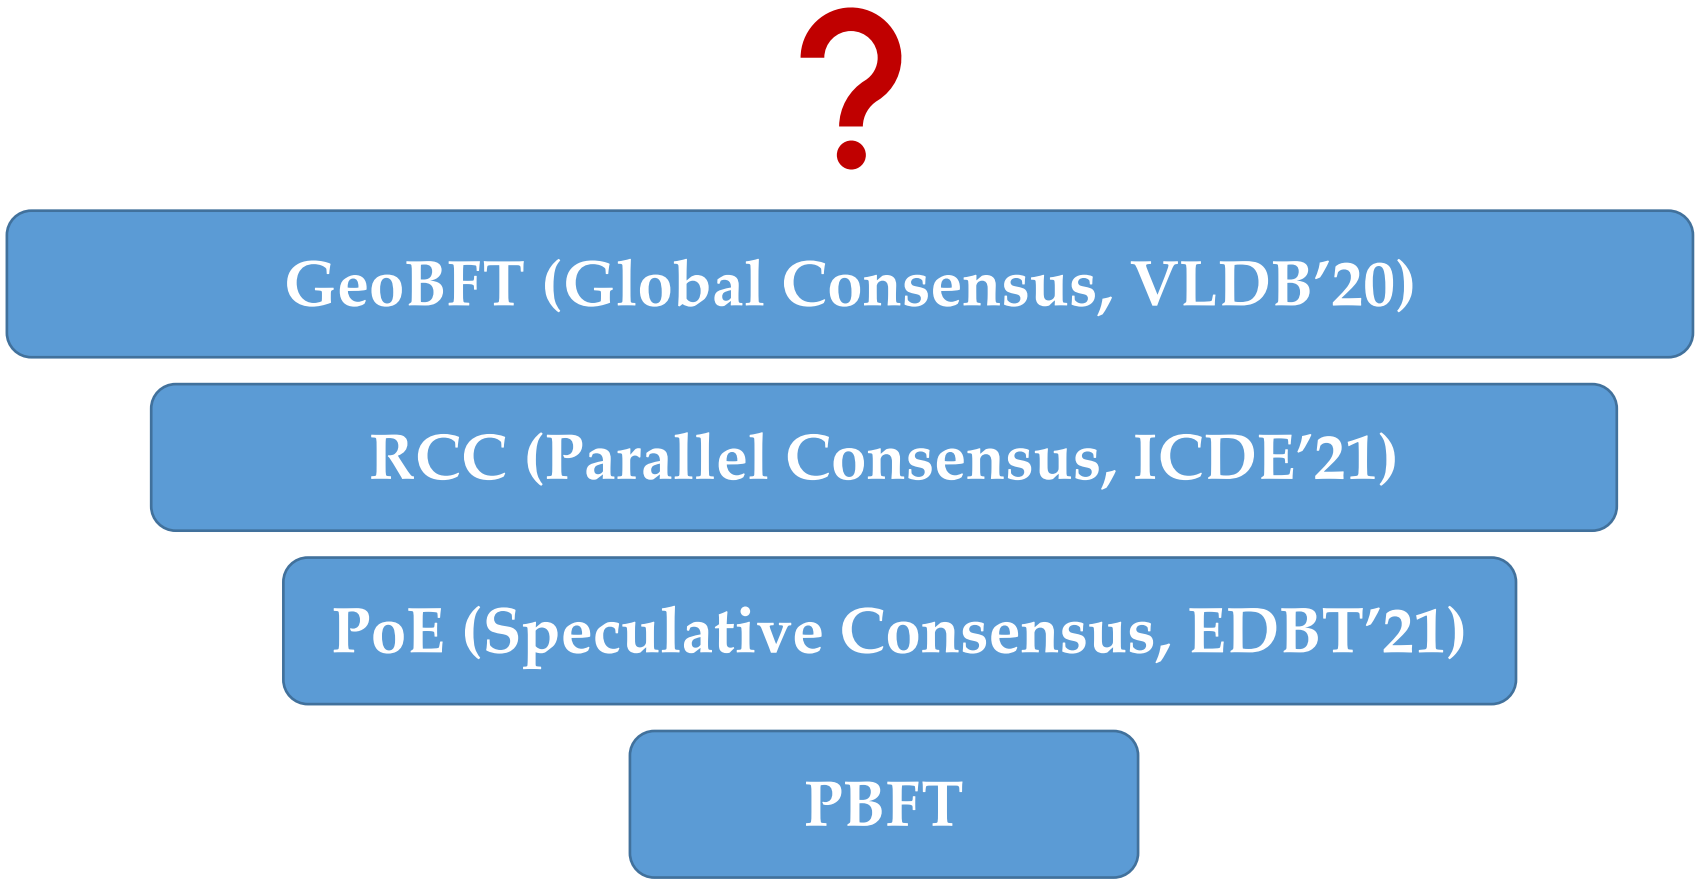

Sharding Solution

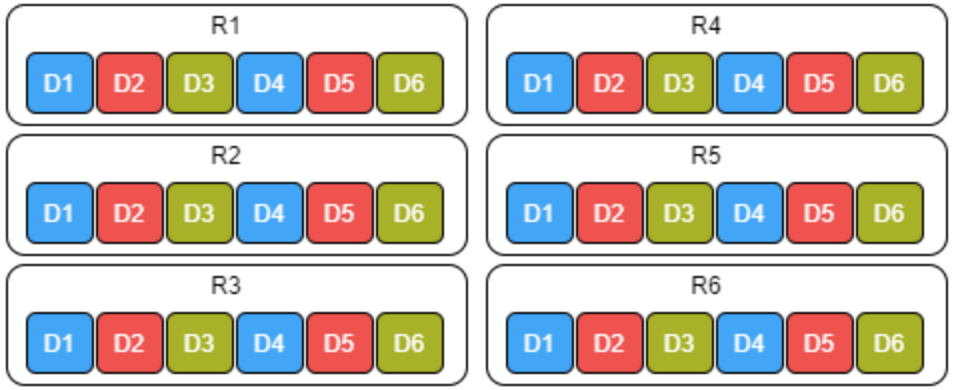

Fully Replicated System

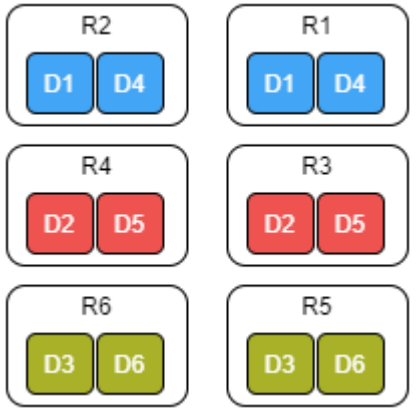

Sharded System

Sharded Consensus

- Single Shard Consensus: Cheap and Parallelized
- Multi Shard Consensus: Main Challenge, Expensive **Focus of This work**

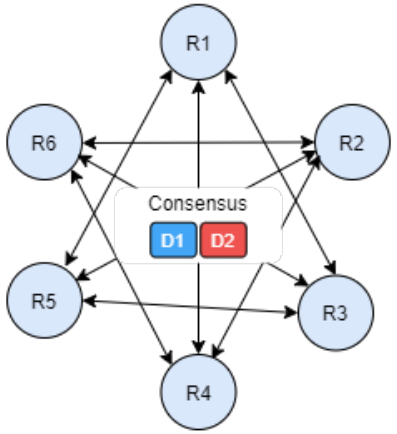

Replicated

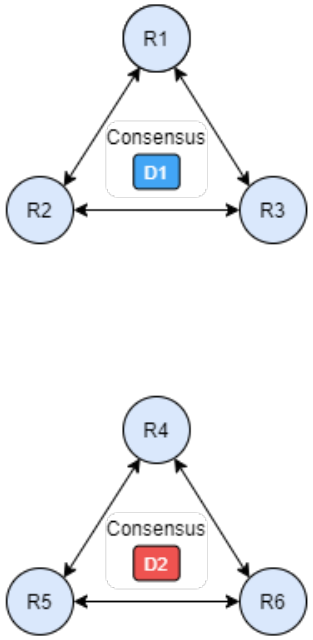

Single Shard

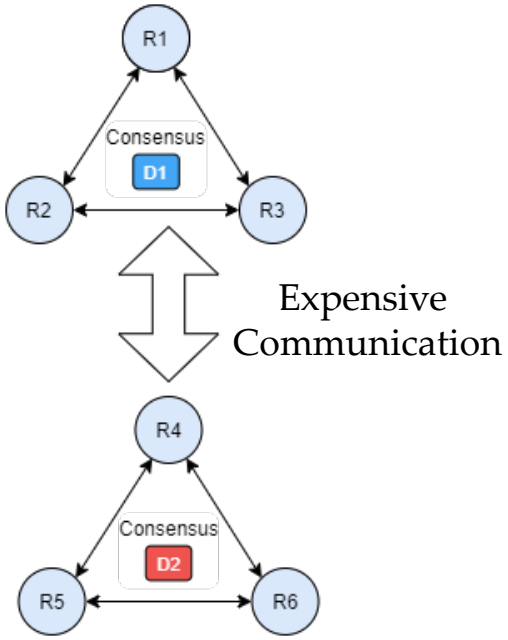

Multi Shard

AHL [Sigmod'19]

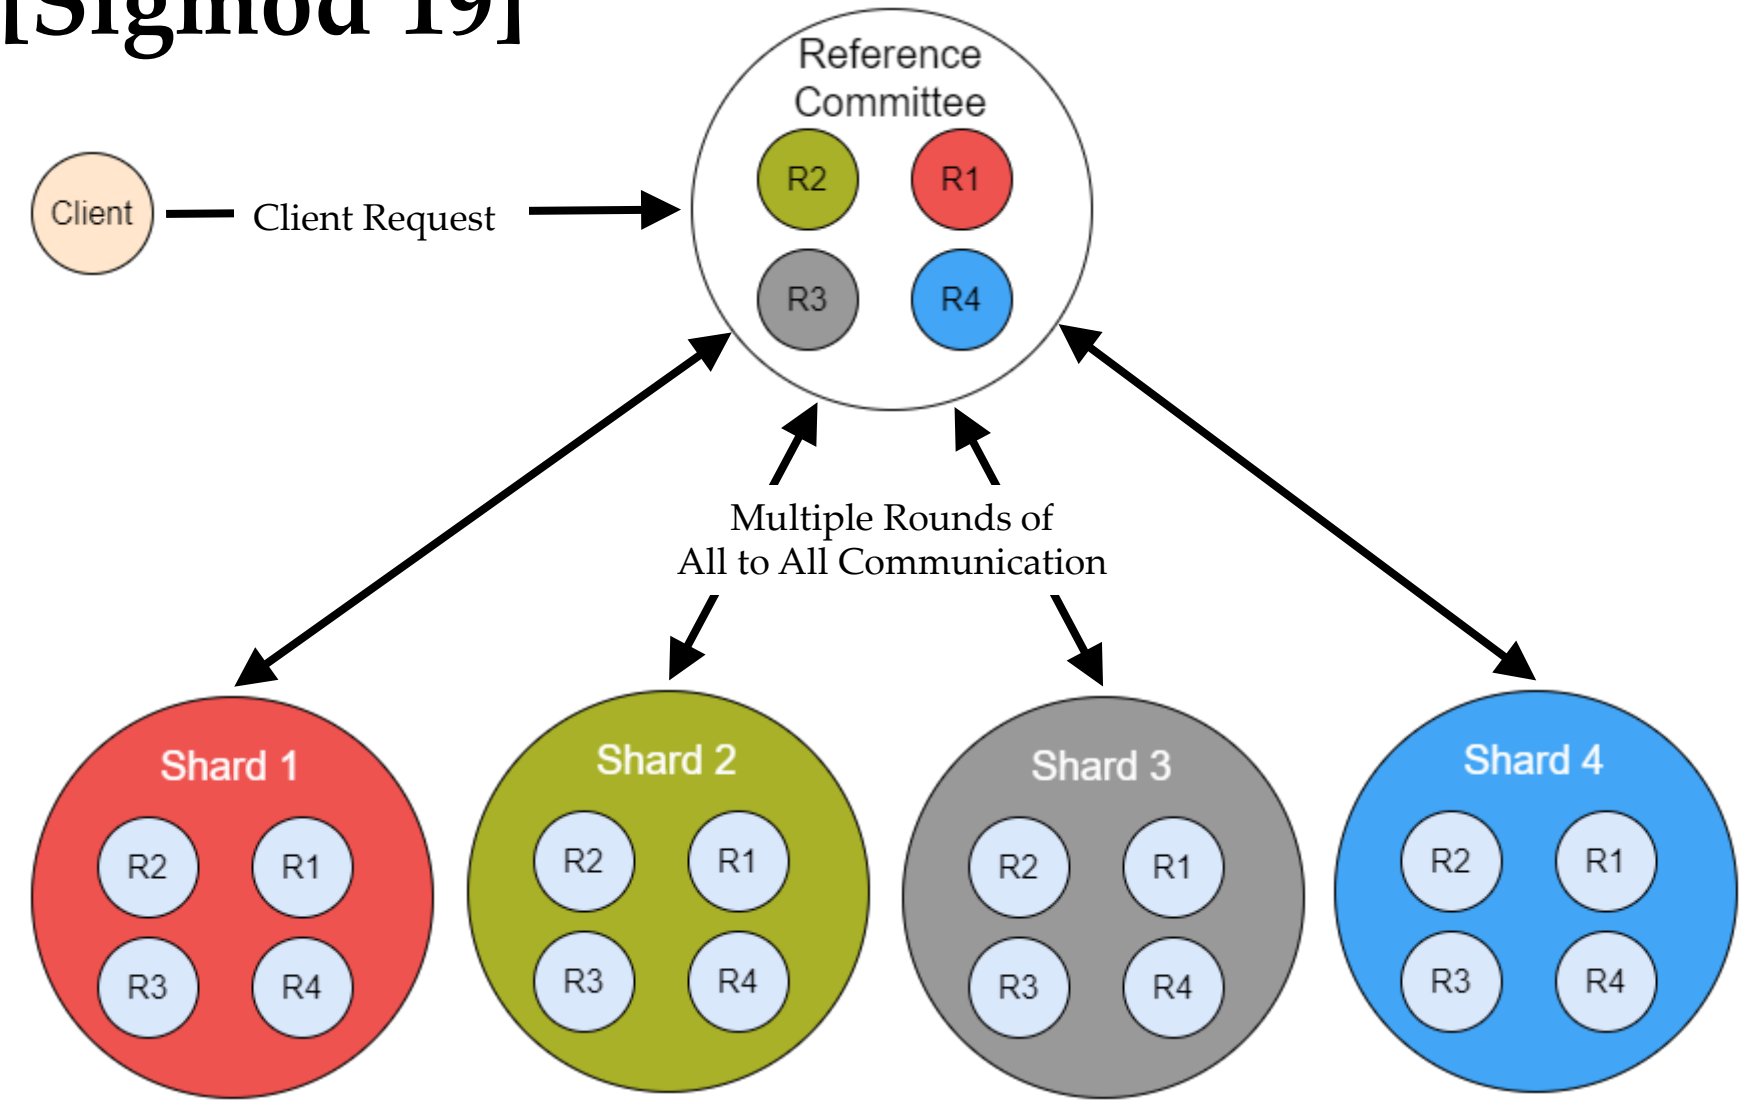

ResilientDB ExpoLab Creativity Unfolded

Sharper [Sigmod'21]

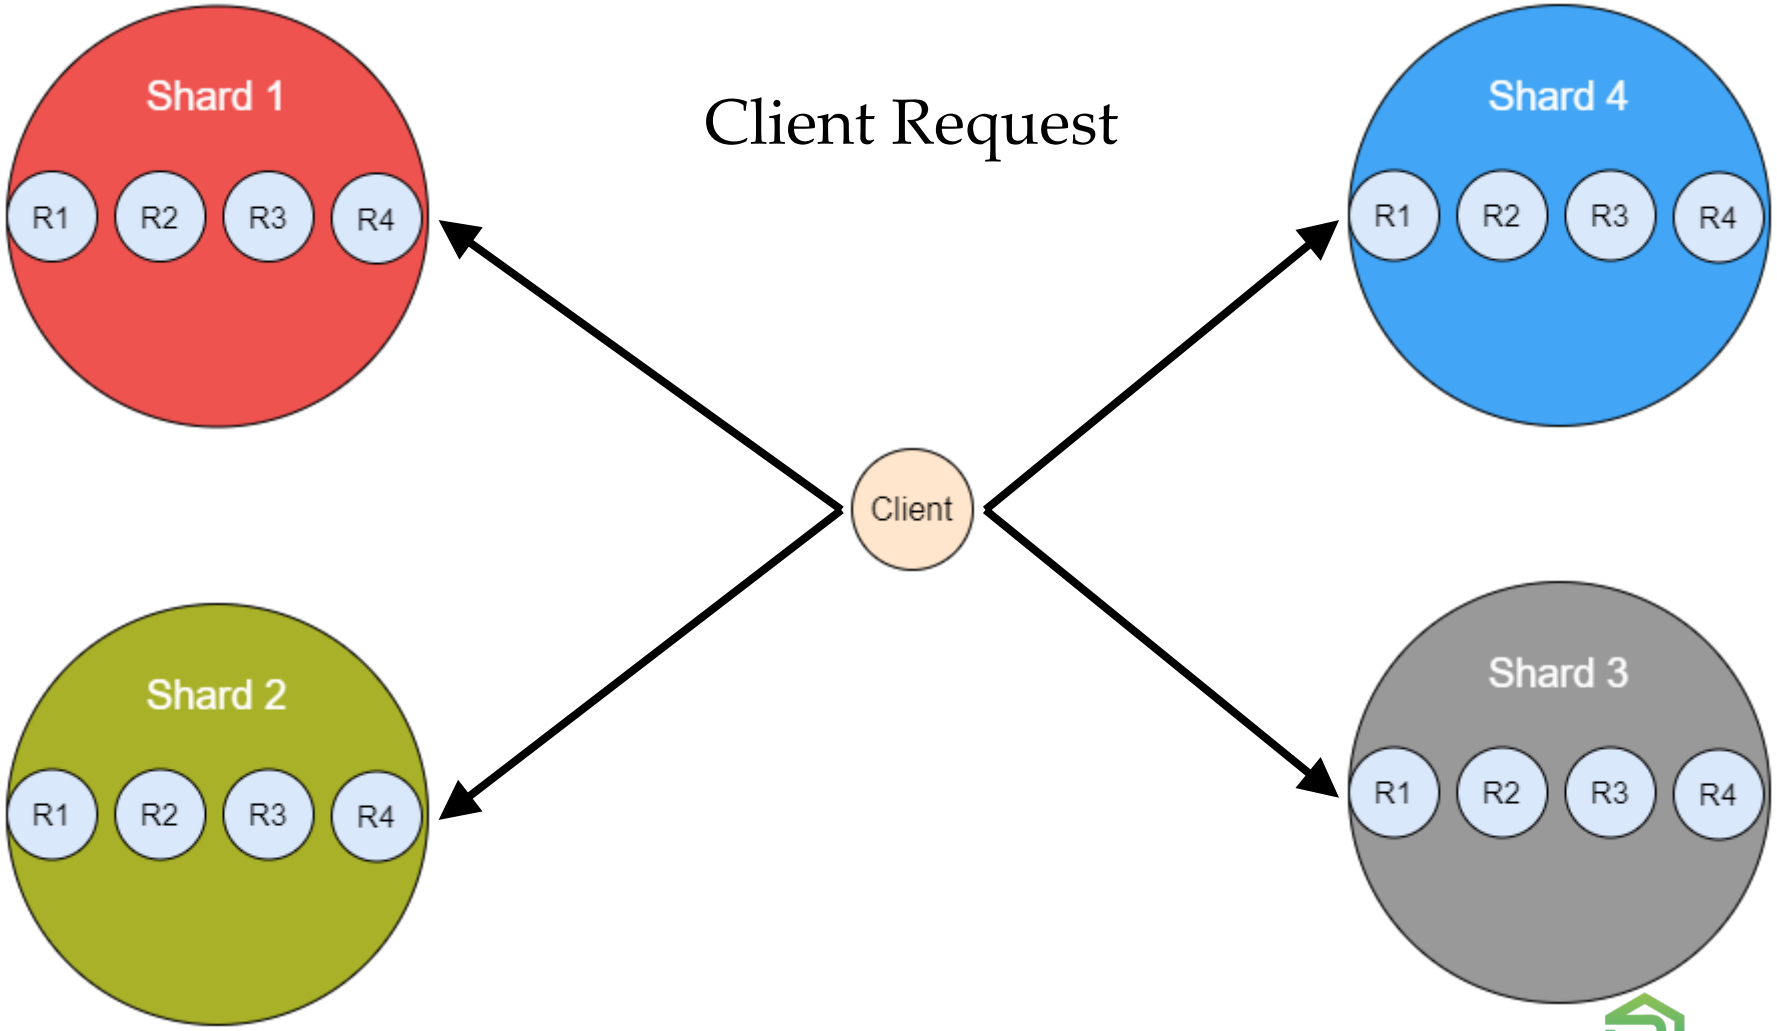

Sharper [Sigmod'21]

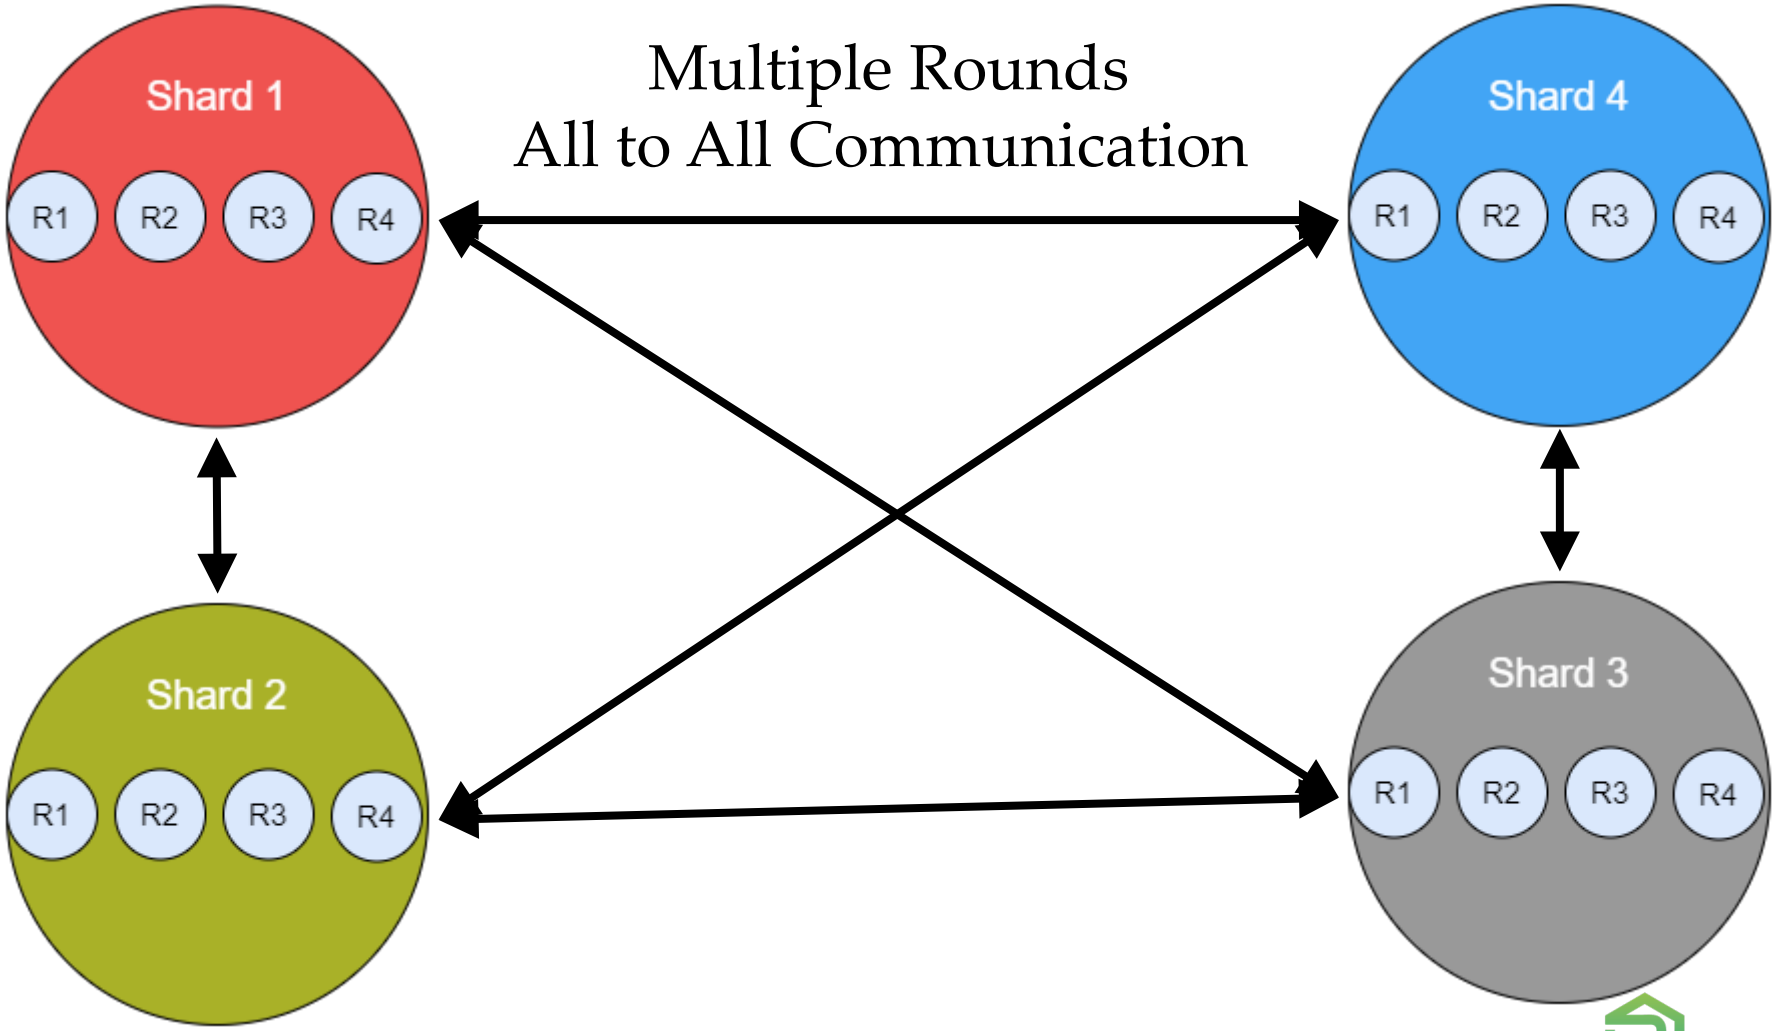

ResilientDB ExpoLab Creativity Unfolded

The Overlooked Problem

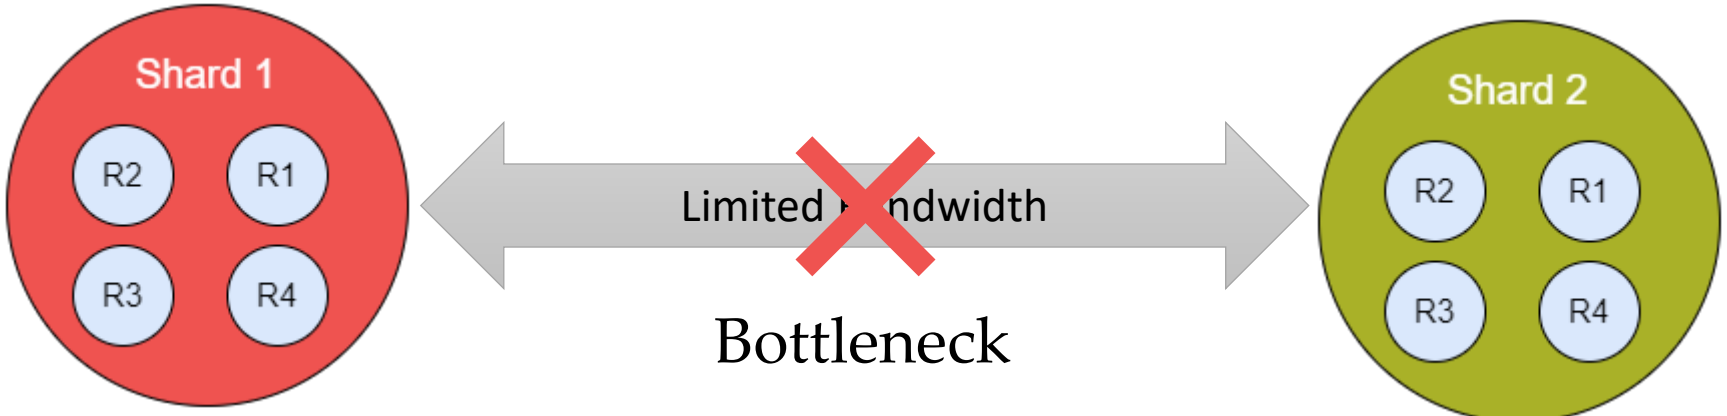

Region 2

Bottleneck

Message Processing

Queuing Delay

Region 2

Our Work: RingBFT

Process

Forward

Retransmit

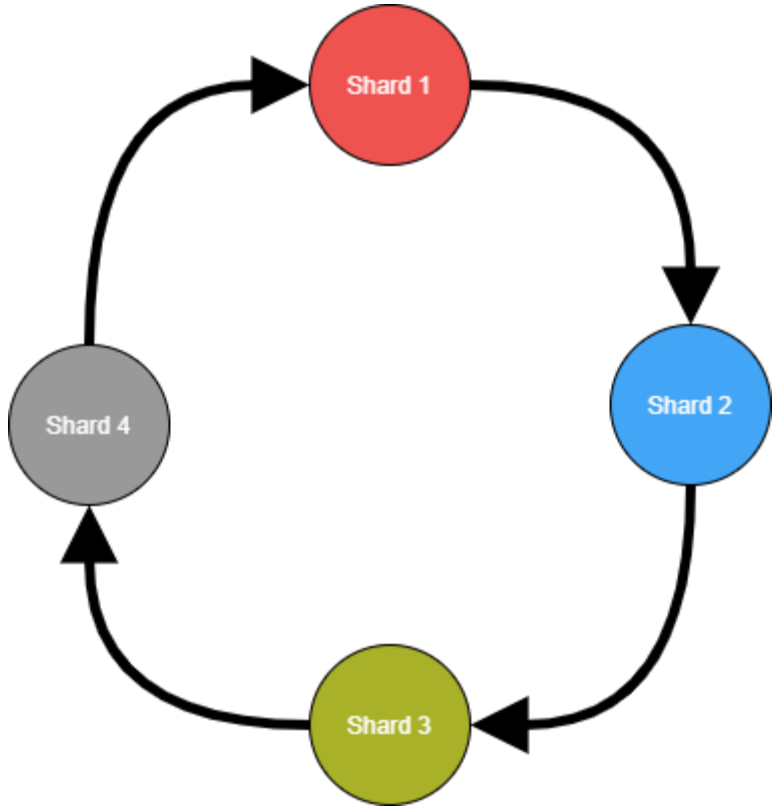

Go around the Ring: No All to All Communication

RingBFT Cross-Shard Consensus Protocol

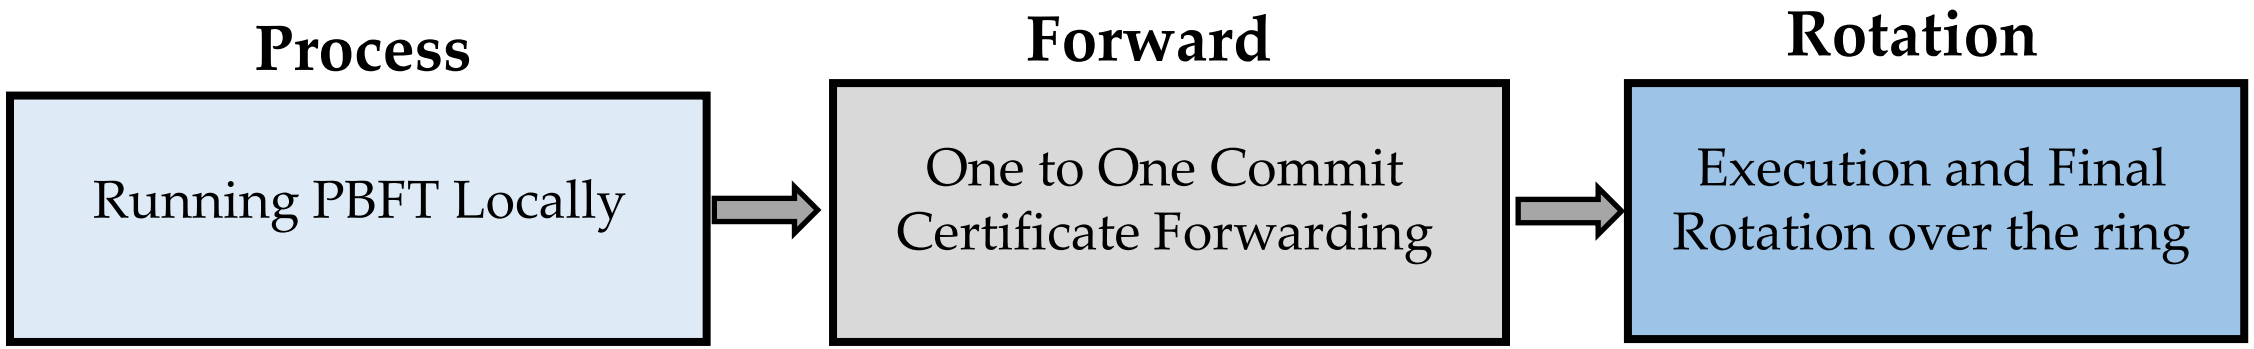

Diving into RingBFT

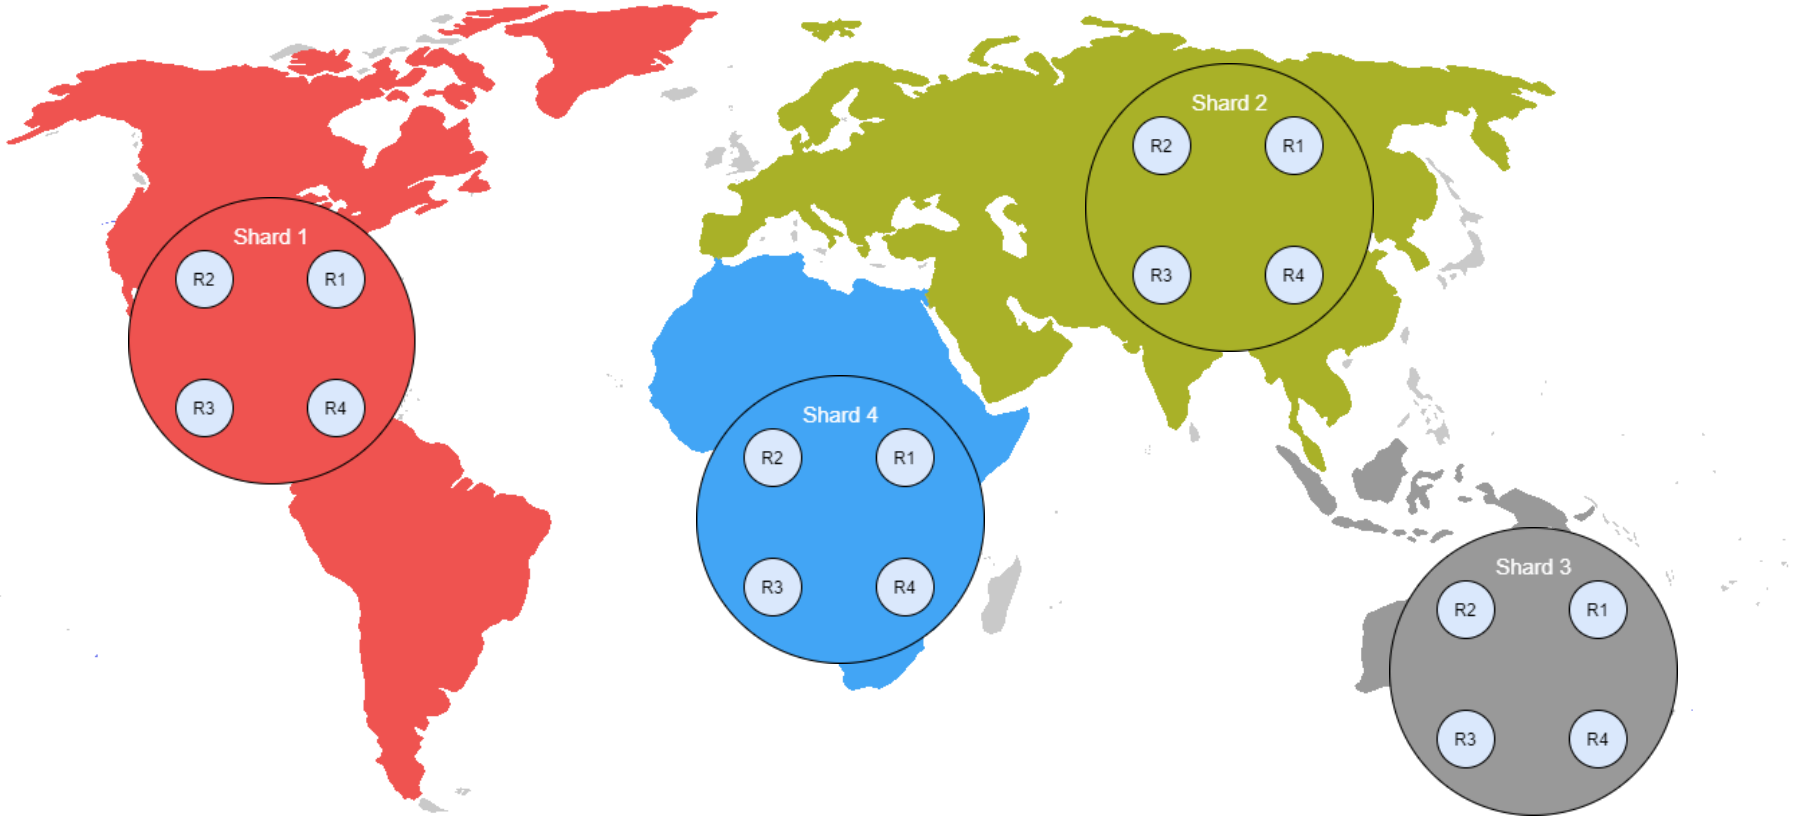

RingBFT Cross-shard Consensus Flow

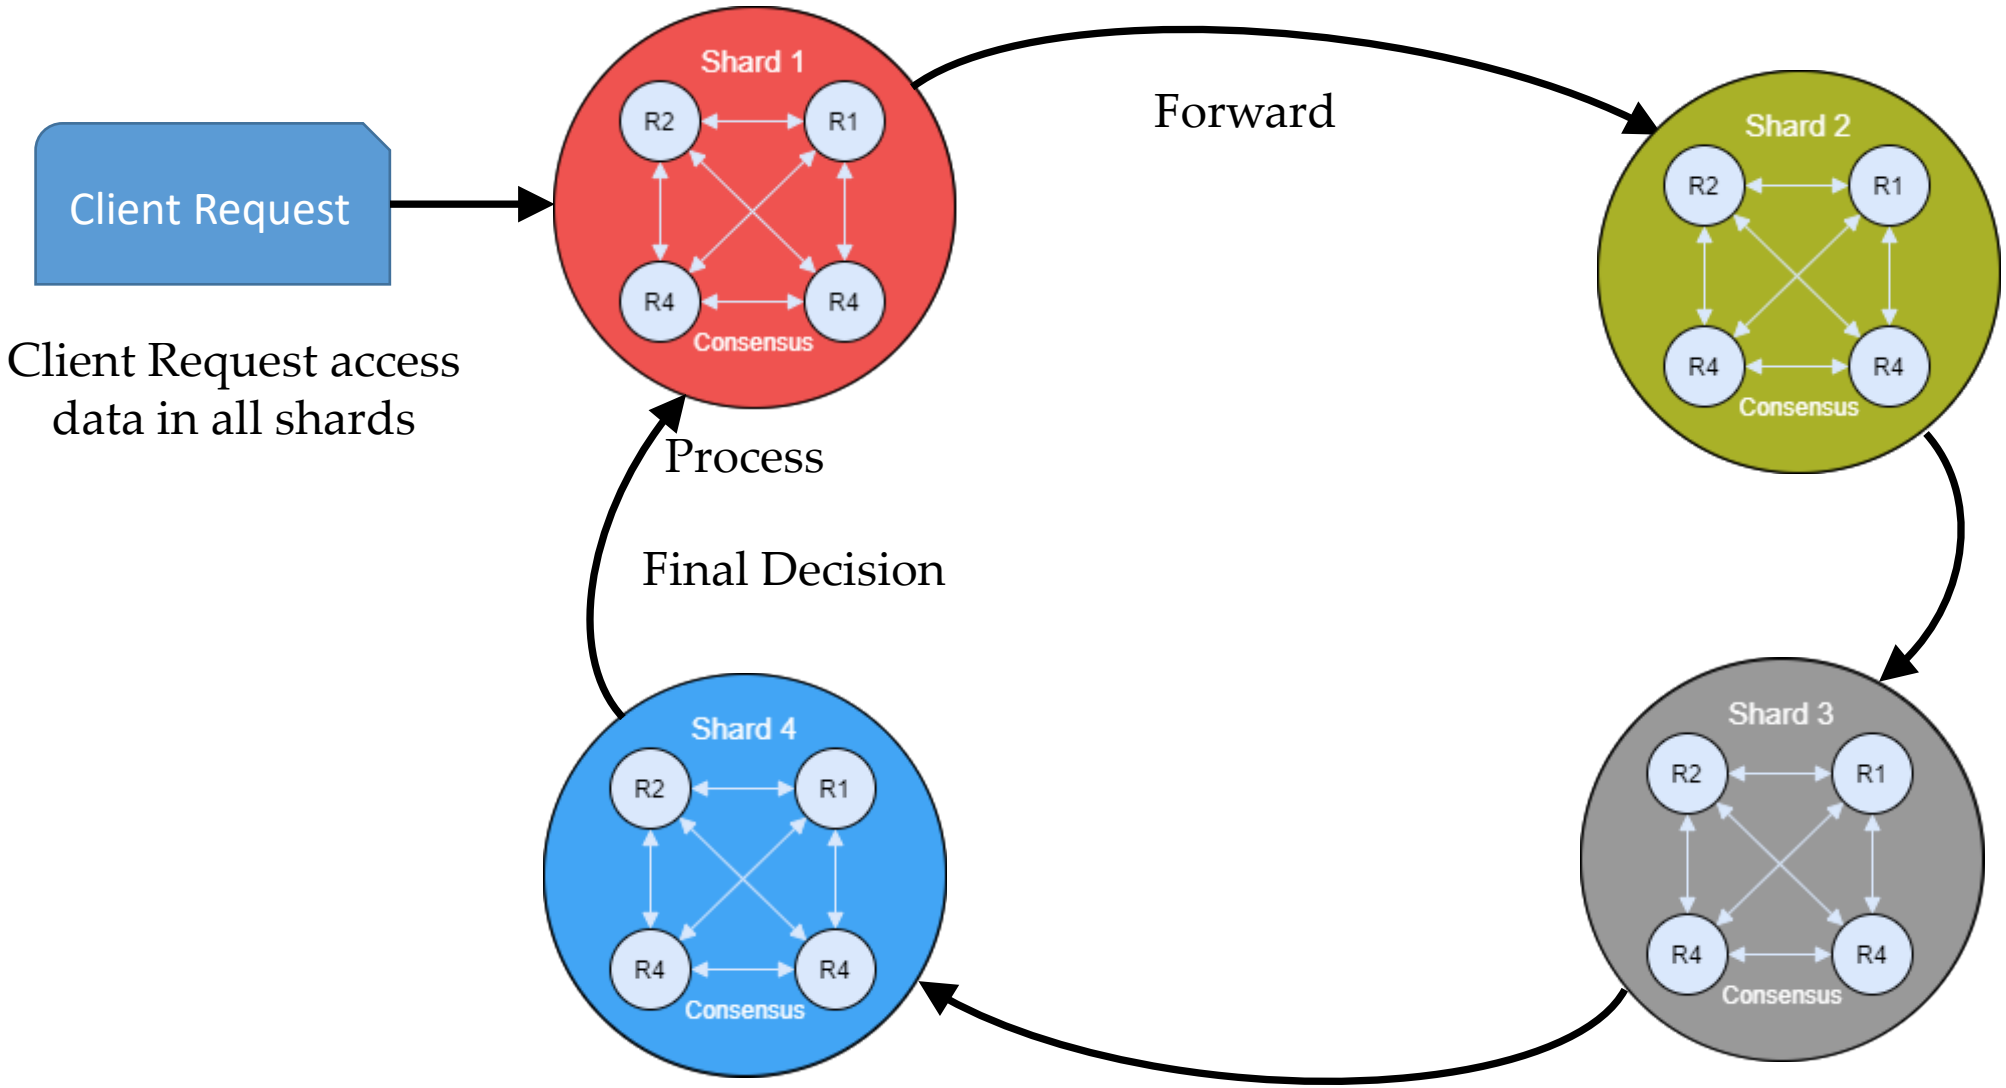

Linear Forward Communication

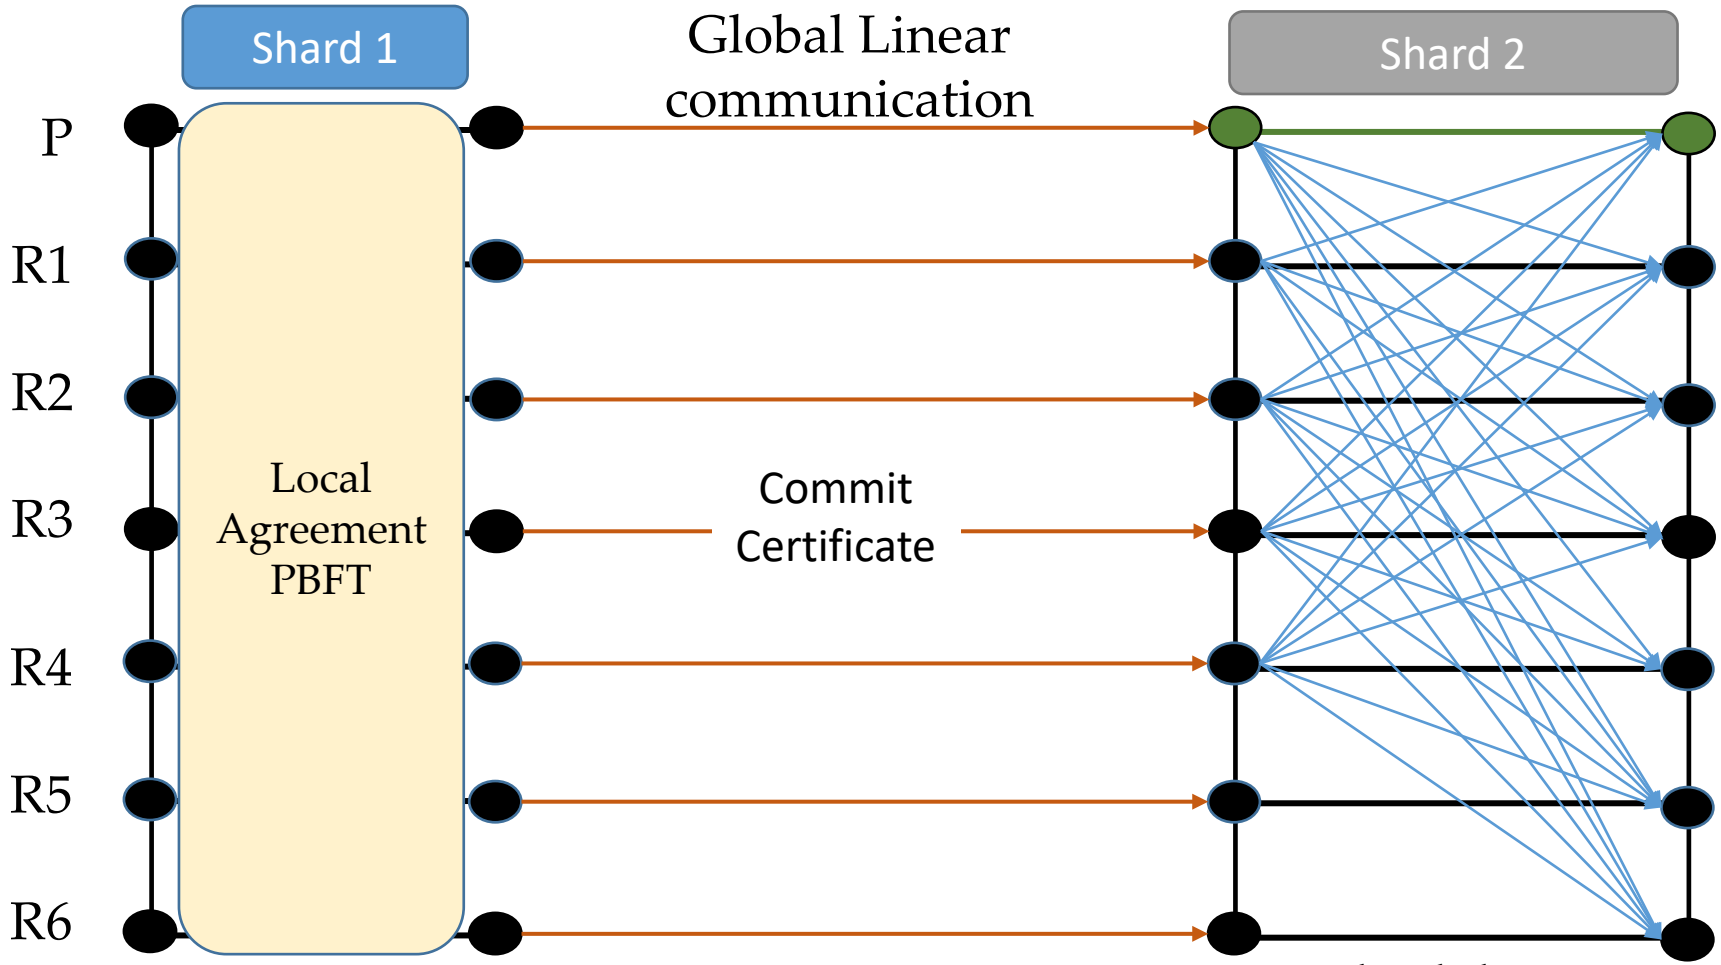

Recovery: Local Timer

Local Failure

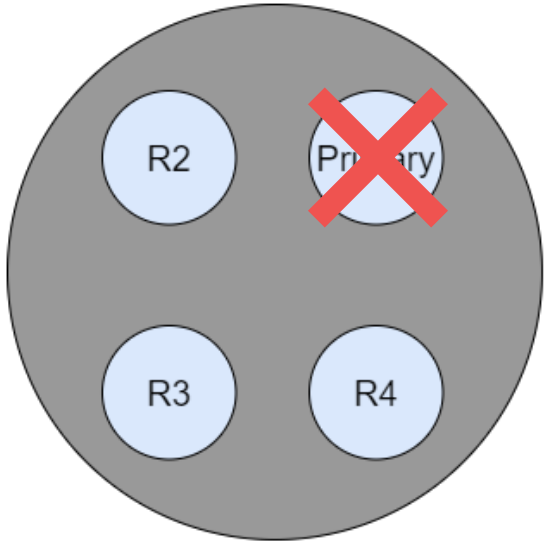

Local Timer Expires

View Change

Recovery: Transmit Timer

No Communication

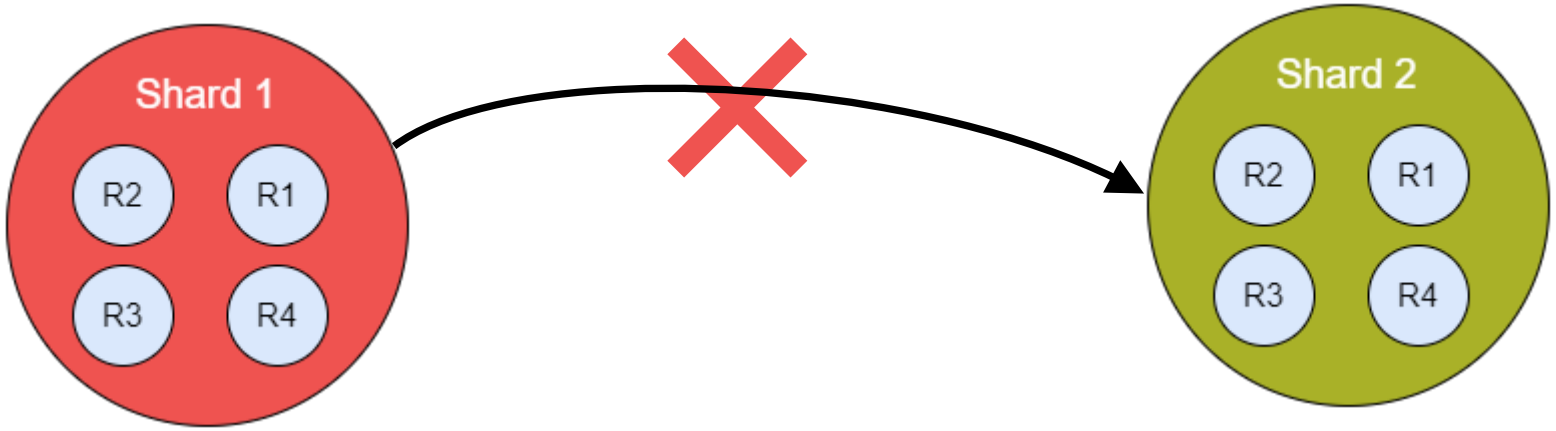

Transmit Timer Expires

Re-transmit Certificates

Recovery: Remote Timer

Partial Communication

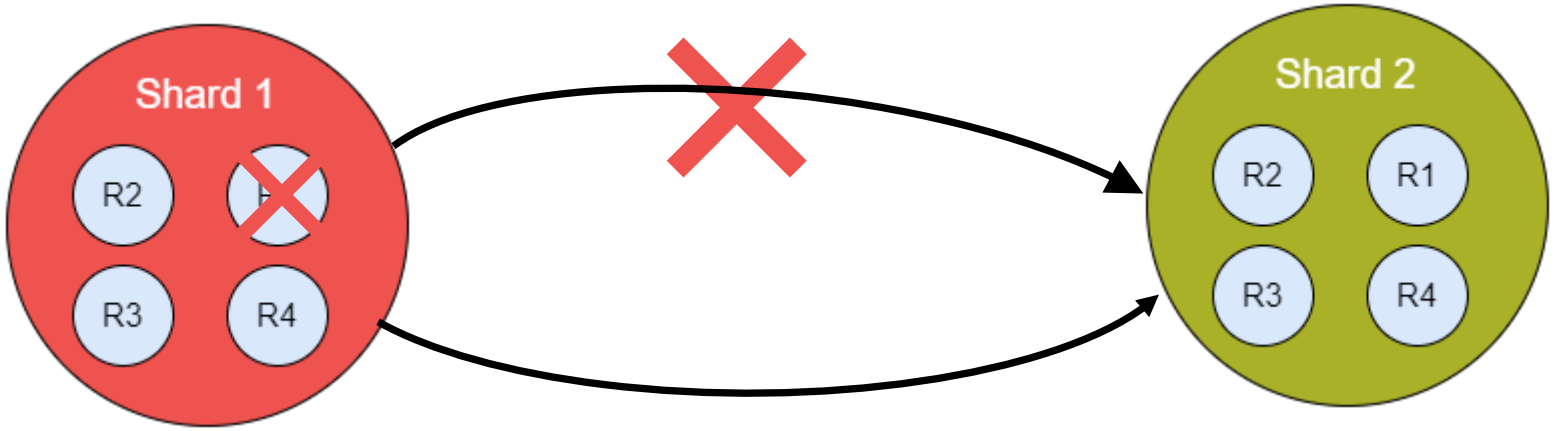

Not Enough Certificates

Remote Timer Expires

Remote View Change

Interdependent Transactions

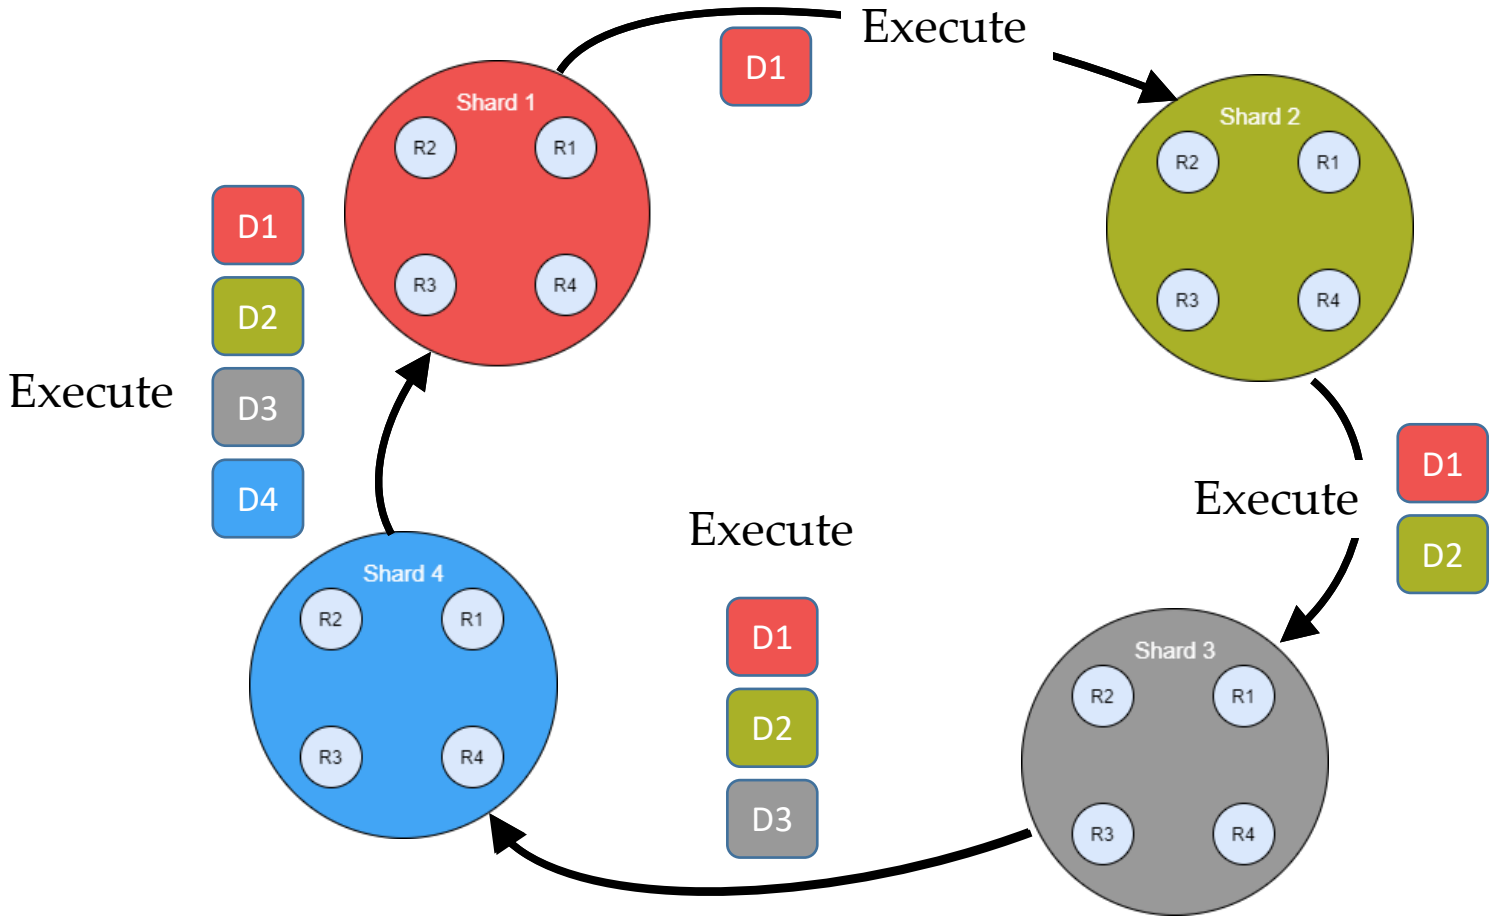

Open Sourced at: https://resilientdb.com

ResilientDB Architecture

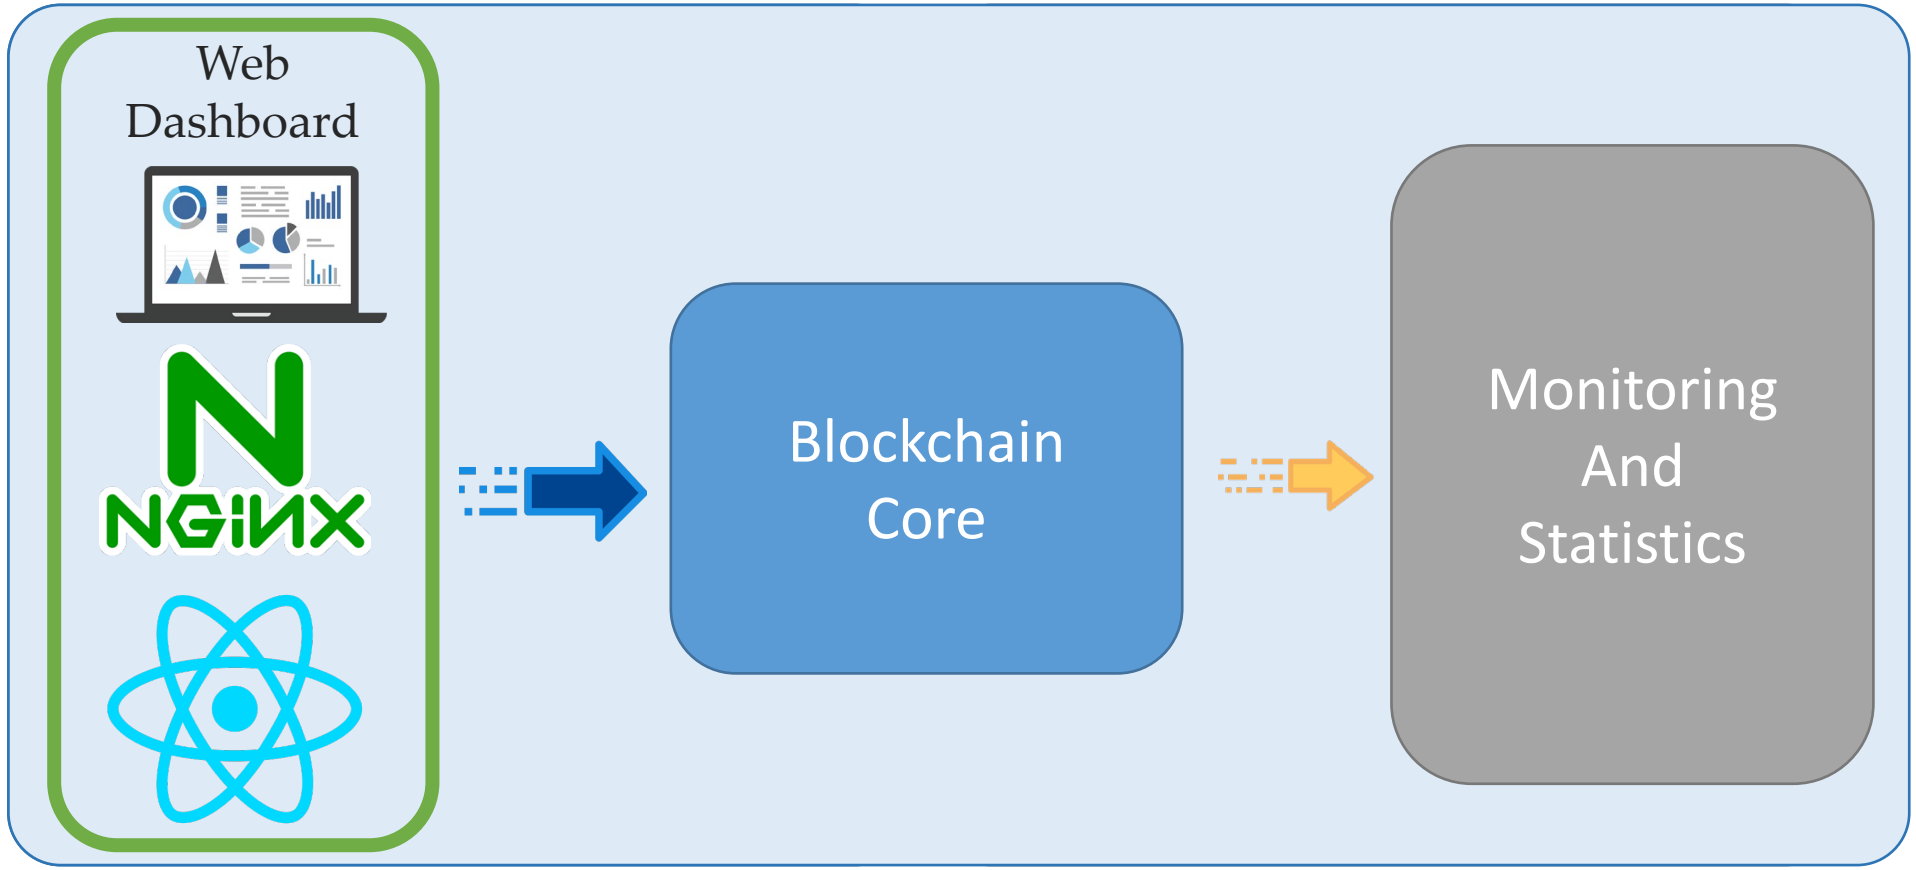

ResilientDB Runtime

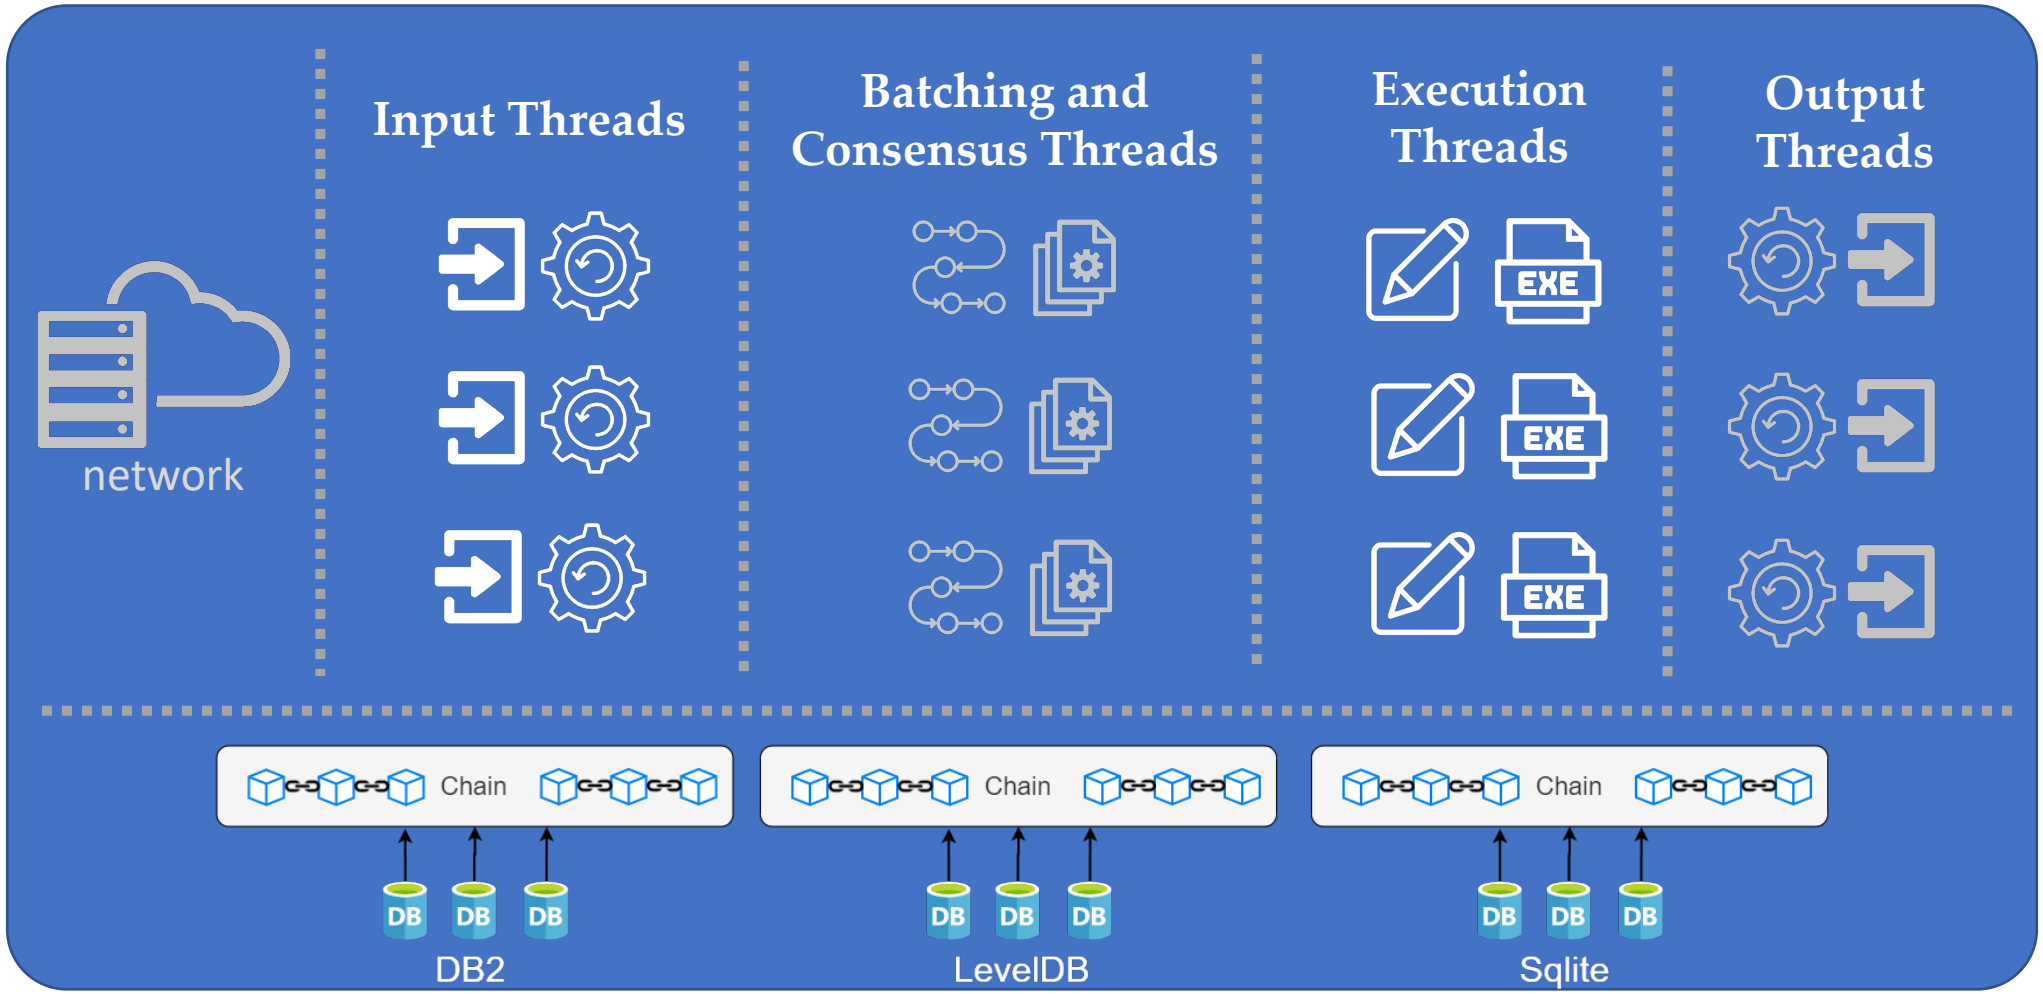

RingBFT Evaluation on ResilientDB

- Scale up to **540 replicas** on Google Cloud
- **16-core** machines with **16 GB memory.**
- Yahoo Cloud Serving Benchmark (**YCSB**).
- Up to **15 regions, 28 replicas per region.**
- **Shuffled** shards to prevent shards proximity.

	Country/City	Region
0	oregon	us-west1
1	iowa	us-central1
2	montreal	northamerica-northeast1
3	netherland	europe-west4
4	taiwan	asia-east1
5	sydney	australia-southeast1
6	singapore	asia-southeast1
7	south-carolina	us-east1
8	north-virginia	us-east4
9	los angeles	us-west2
10	las vegas	us-west4
11	london	europe-west2
12	belgium	europe-west1
13	tokyo	asia-northeast1
14	hong-kong	asia-east2

Scalability: Impact of shards

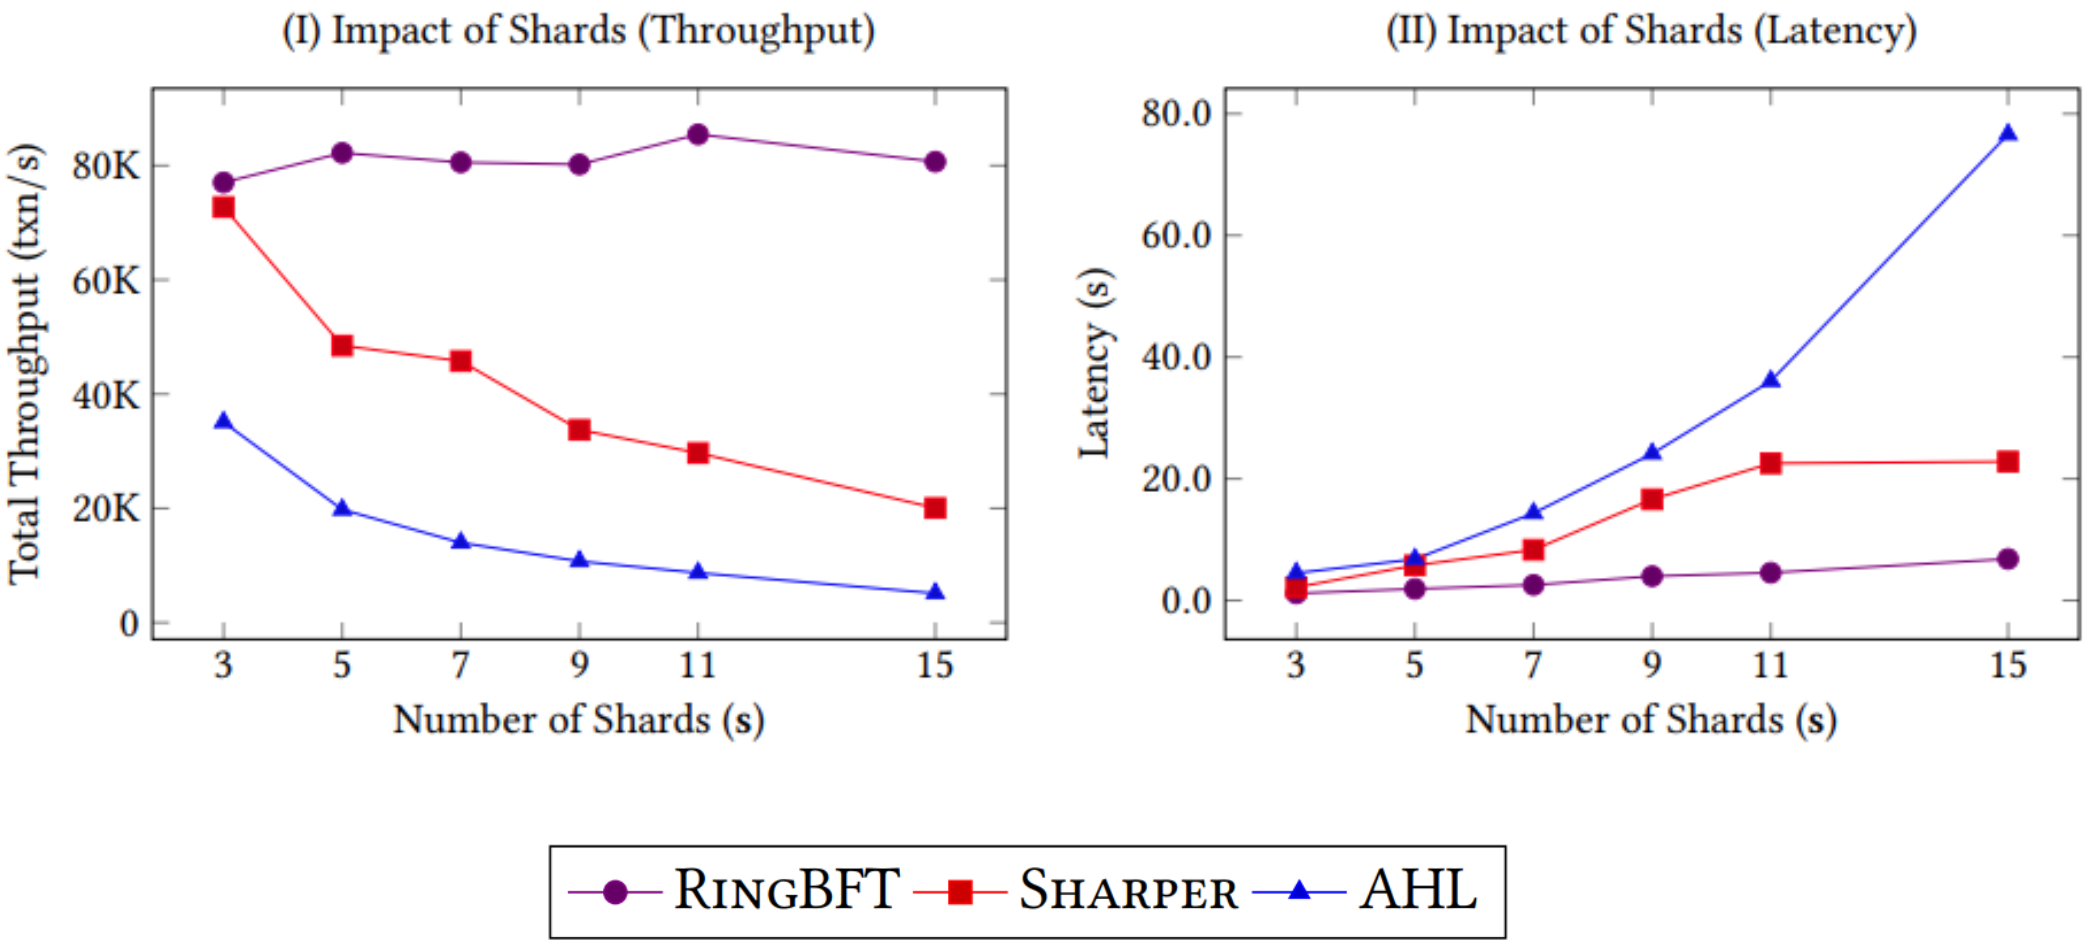

Scalability: Impact of X-shard workload rate

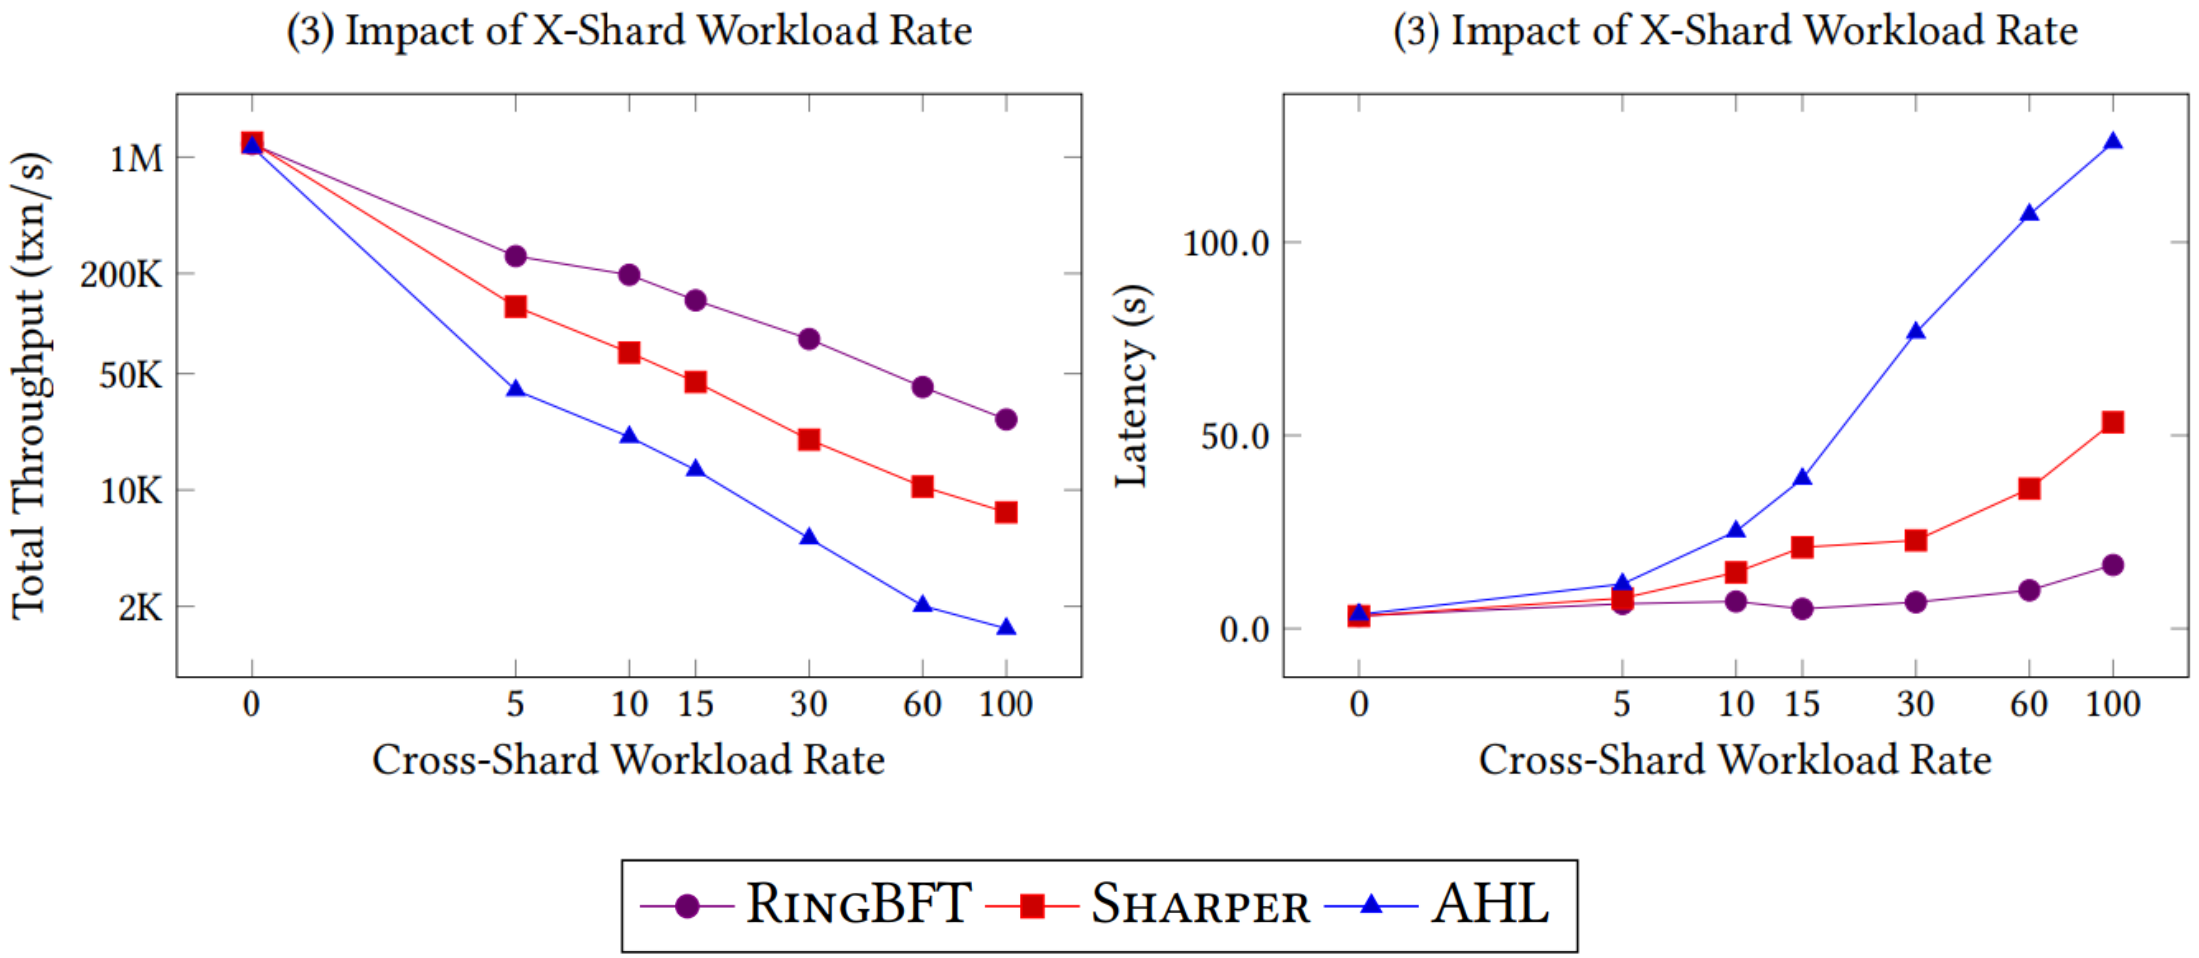

ResilientDB Roadmap

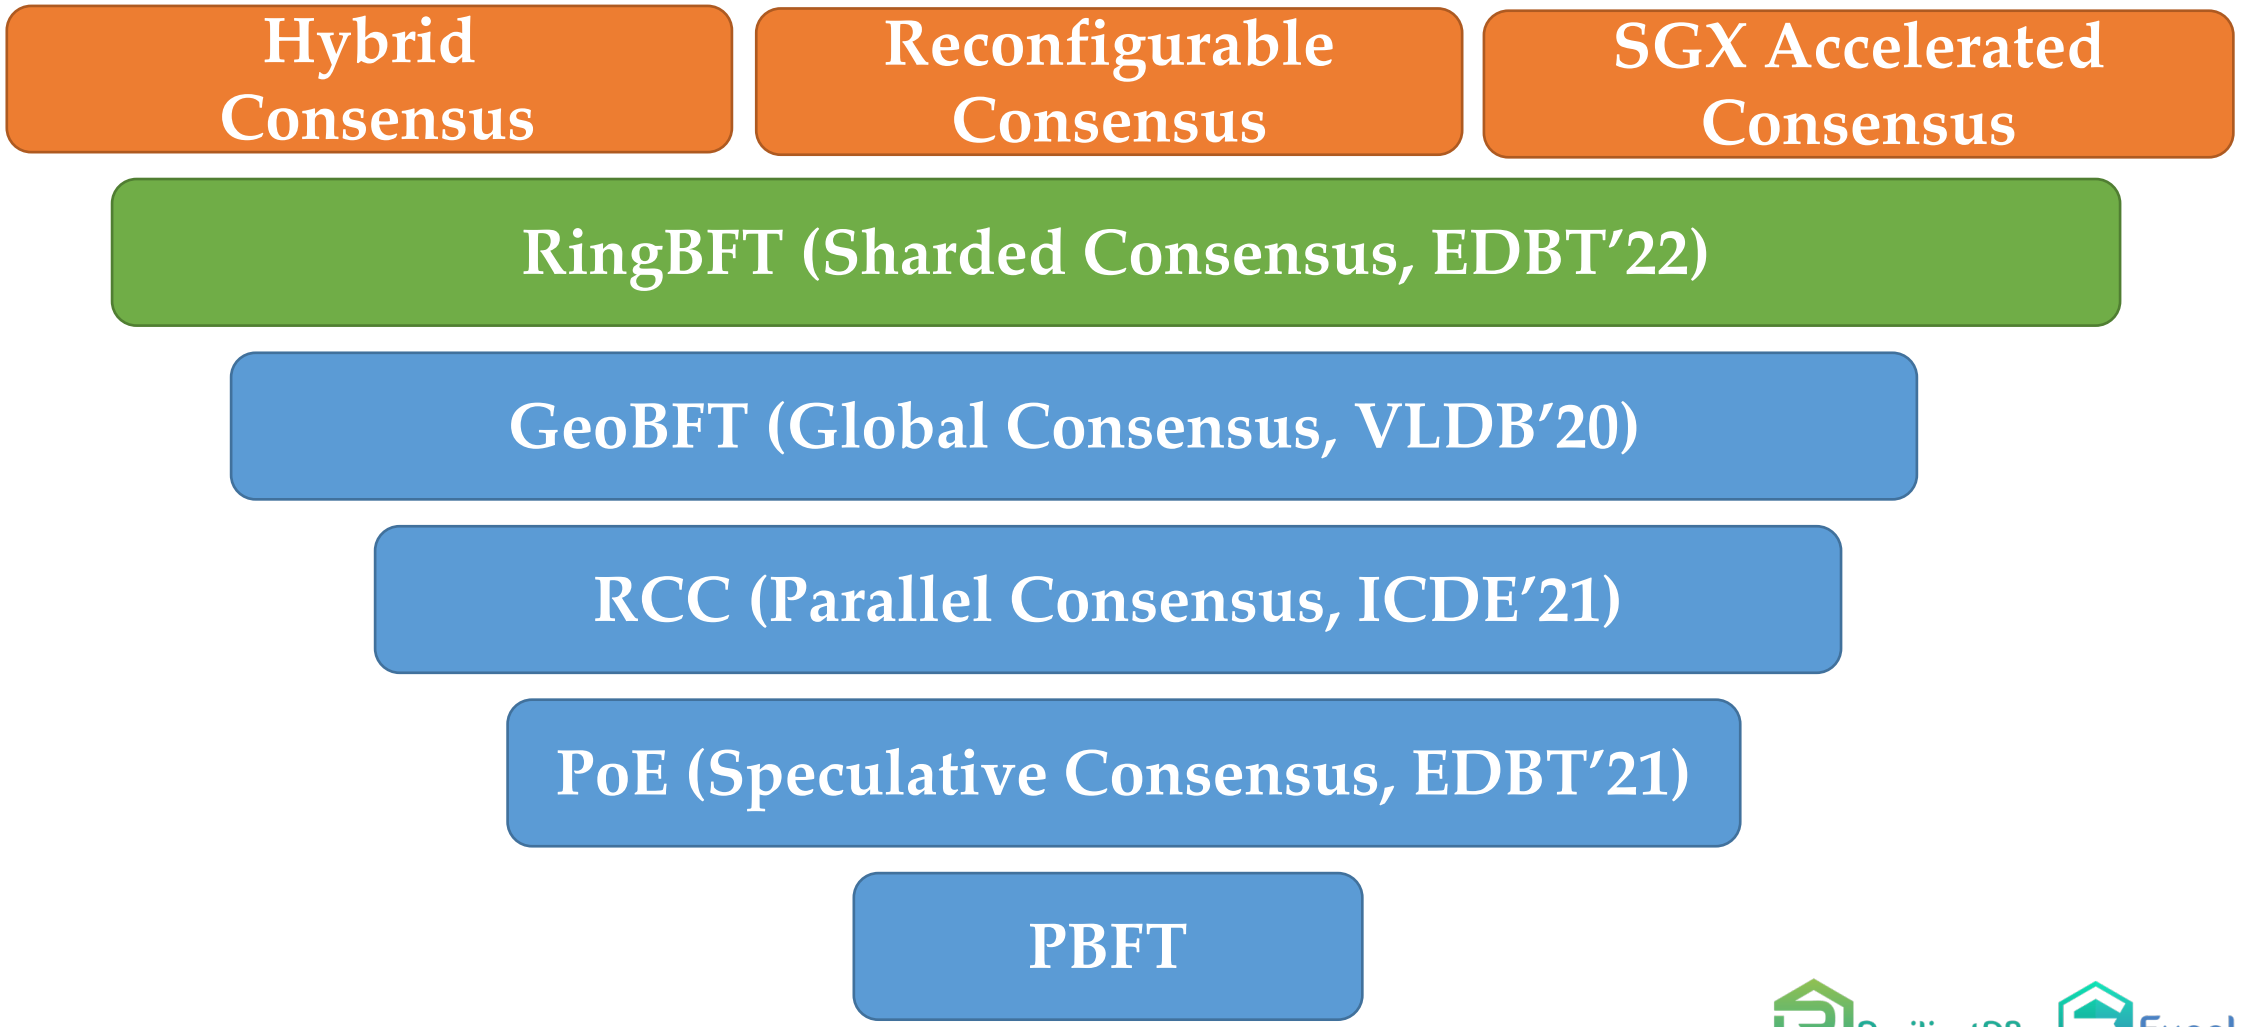

ResilientDB ExpoLab Creativity Unfolded

Thanks!