Replicating Non-Volatile Main Memory

Yiying Zhang,

Jian Yang, Amirsaman Memaripour, Steven Swanson

Non-Volatile Main Memory (NVMM)

- Next generation non-volatile memory
- NVMM: NVM as (persistent) main memory

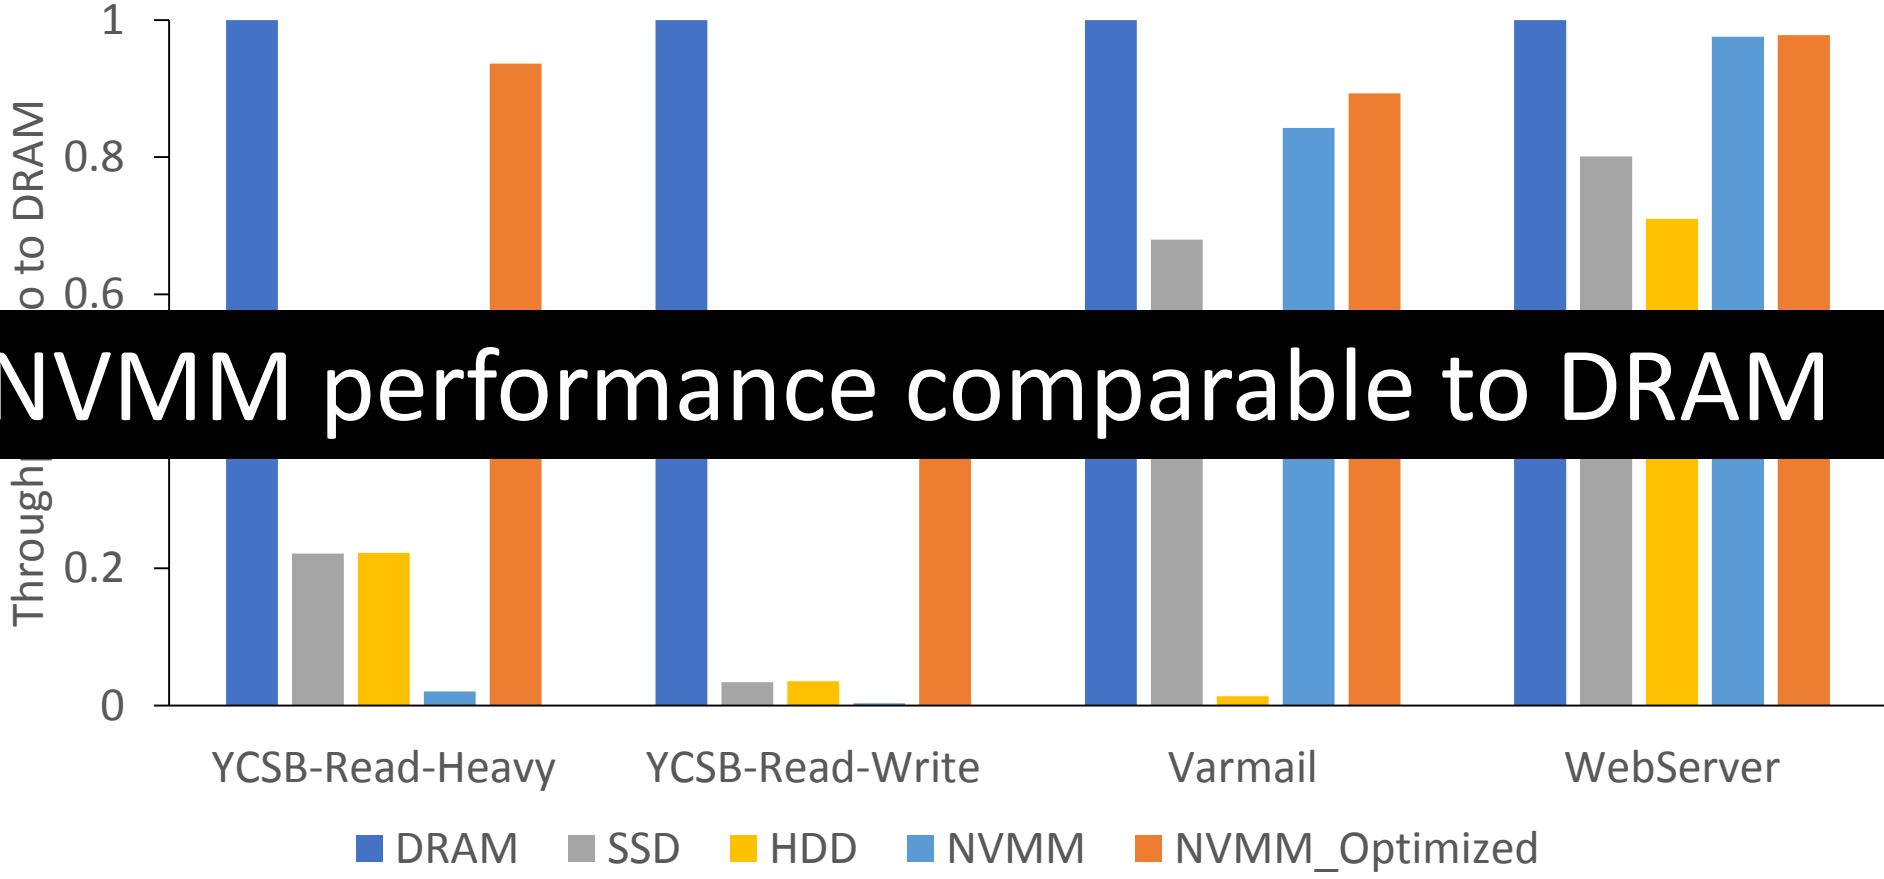

NVMM performance comparable to DRAM

NVMM in Data Center

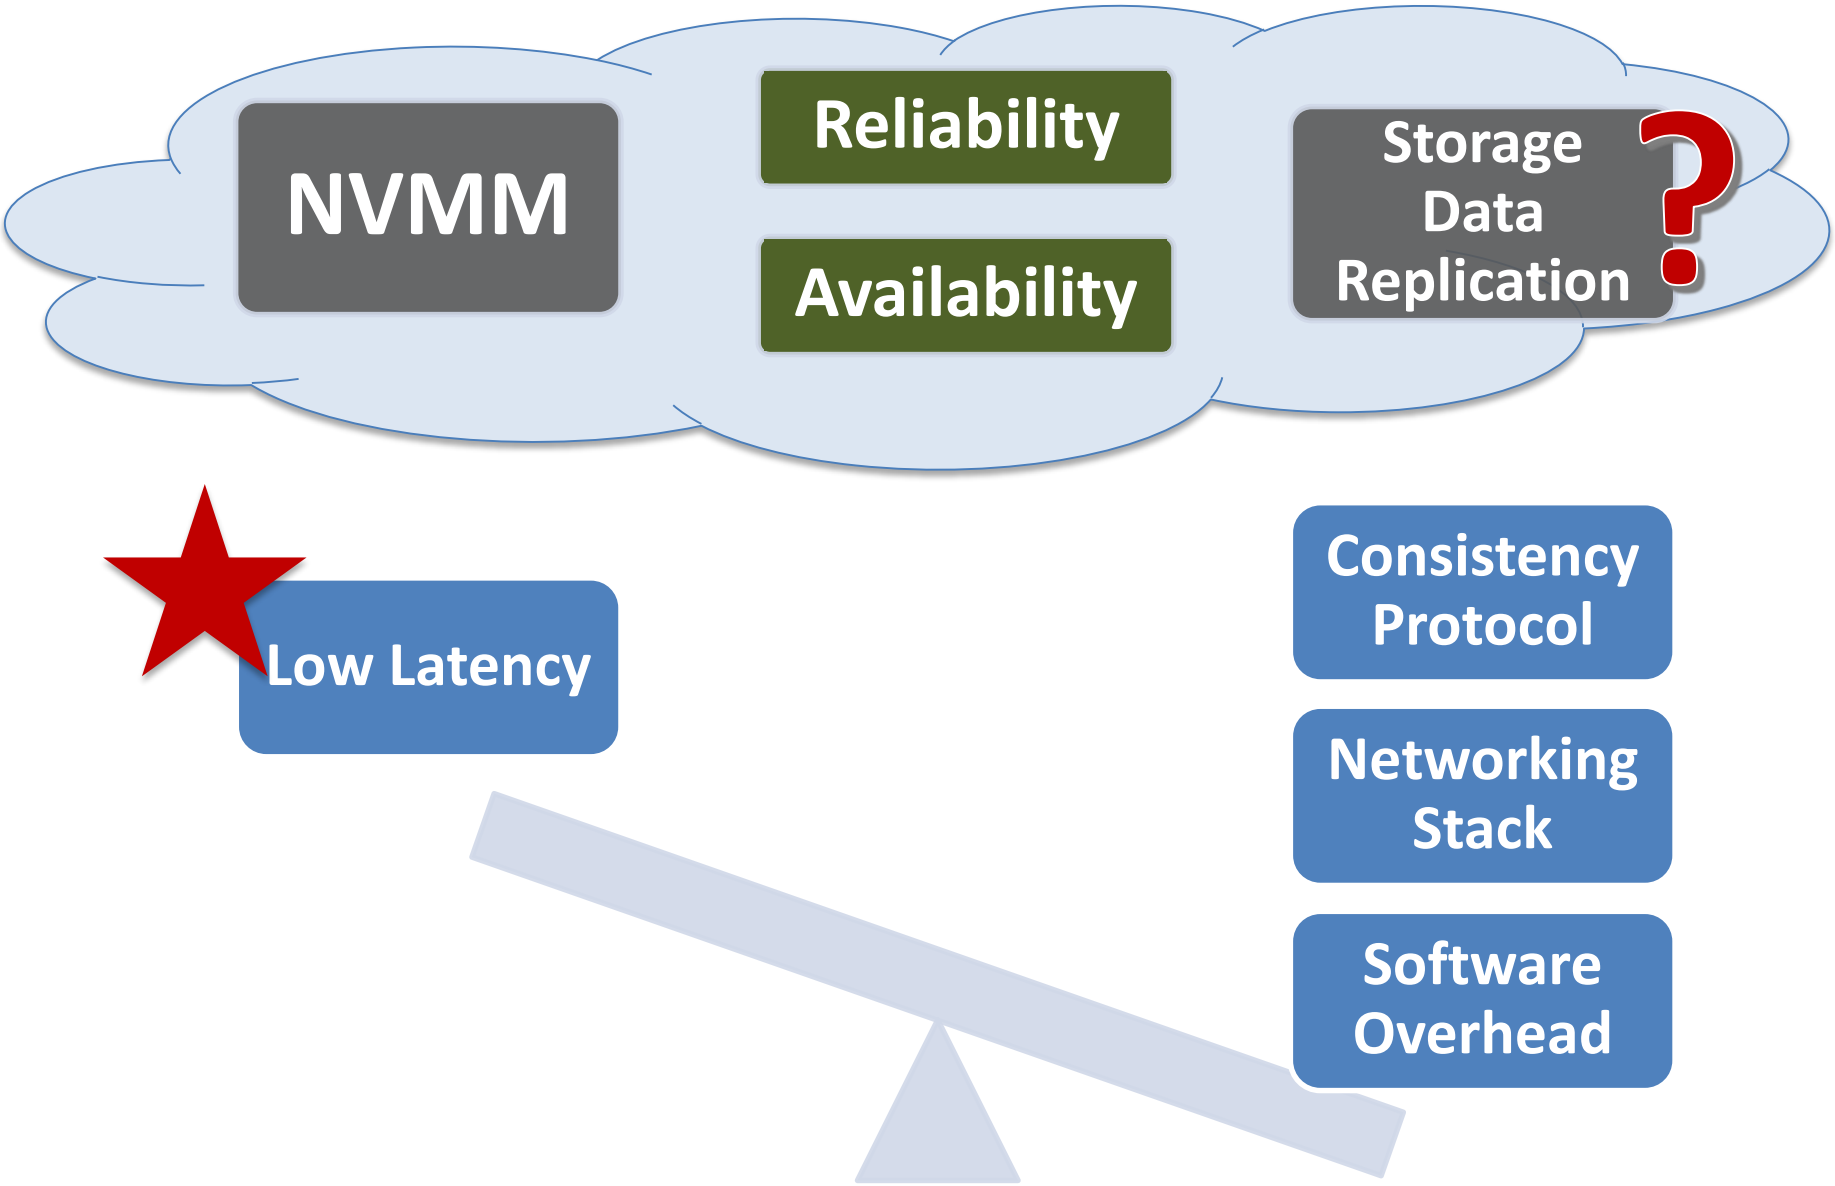

Mojim: Reliable and Highly-Available NVMM

- NVMM-to-NVMM fine-grained replication
- RDMA-based replication optimized for NVMM
- Two-tier architecture
- Flexible consistency, reliability, availability, costs

Mojim Results Highlight

- **29% – 73%** average latency of un-replicated
- **0.5x – 3.5x** bandwidth of un-replicated
- **3.4x – 4x** faster than MongoDB replication
- **And stronger consistency & reliability!**
- Instant fail over, 1.9 sec reconstruction

Outline

- Introduction
- Mojim Design and Architecture
- Mojim Implementation
- Evaluation Results
- Conclusion

Un-replicated NVMM

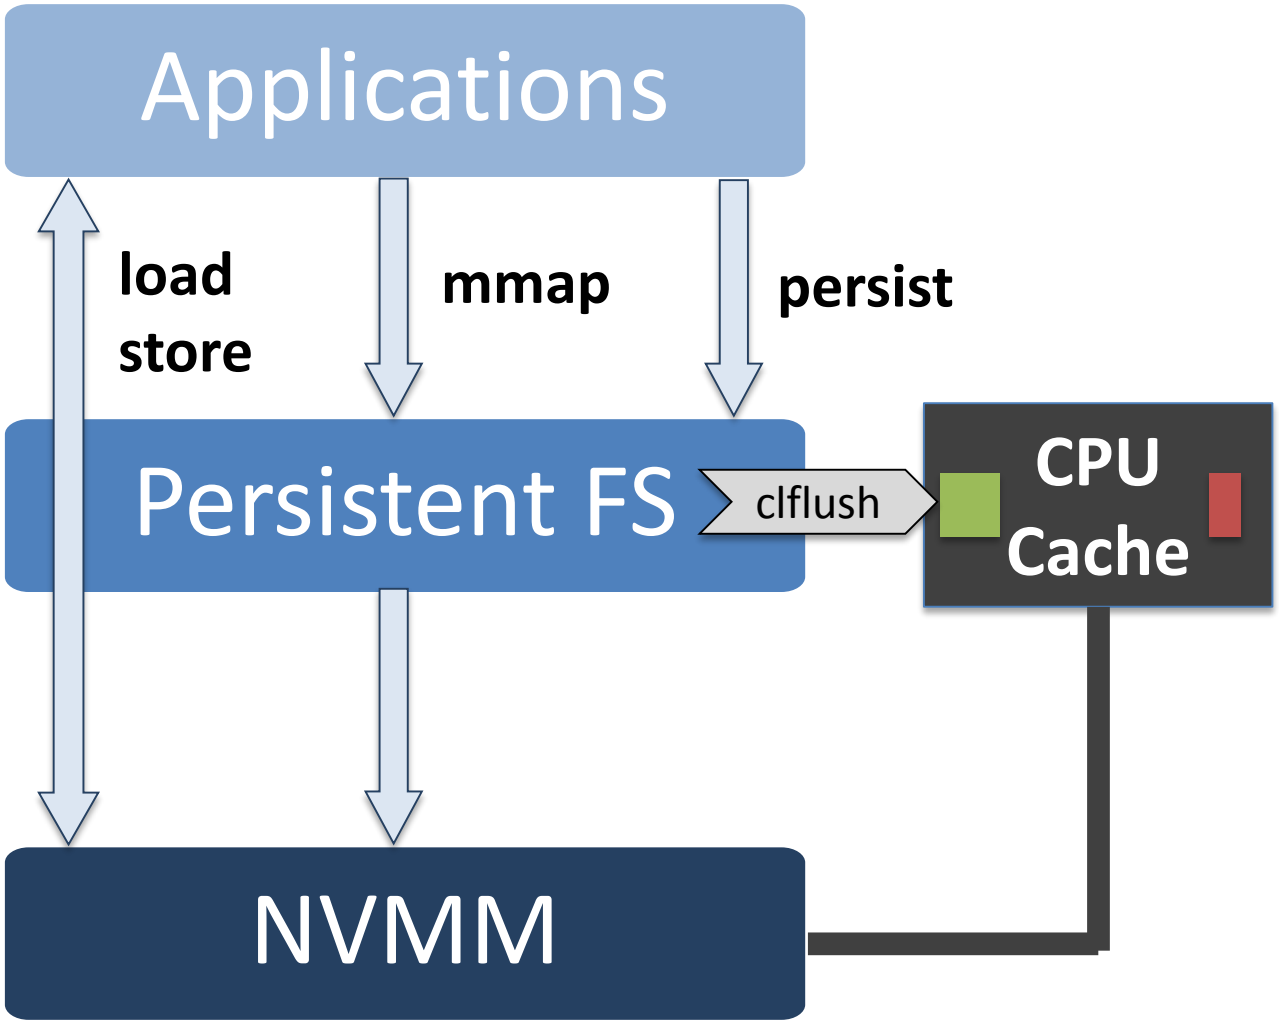

Mojim Architecture

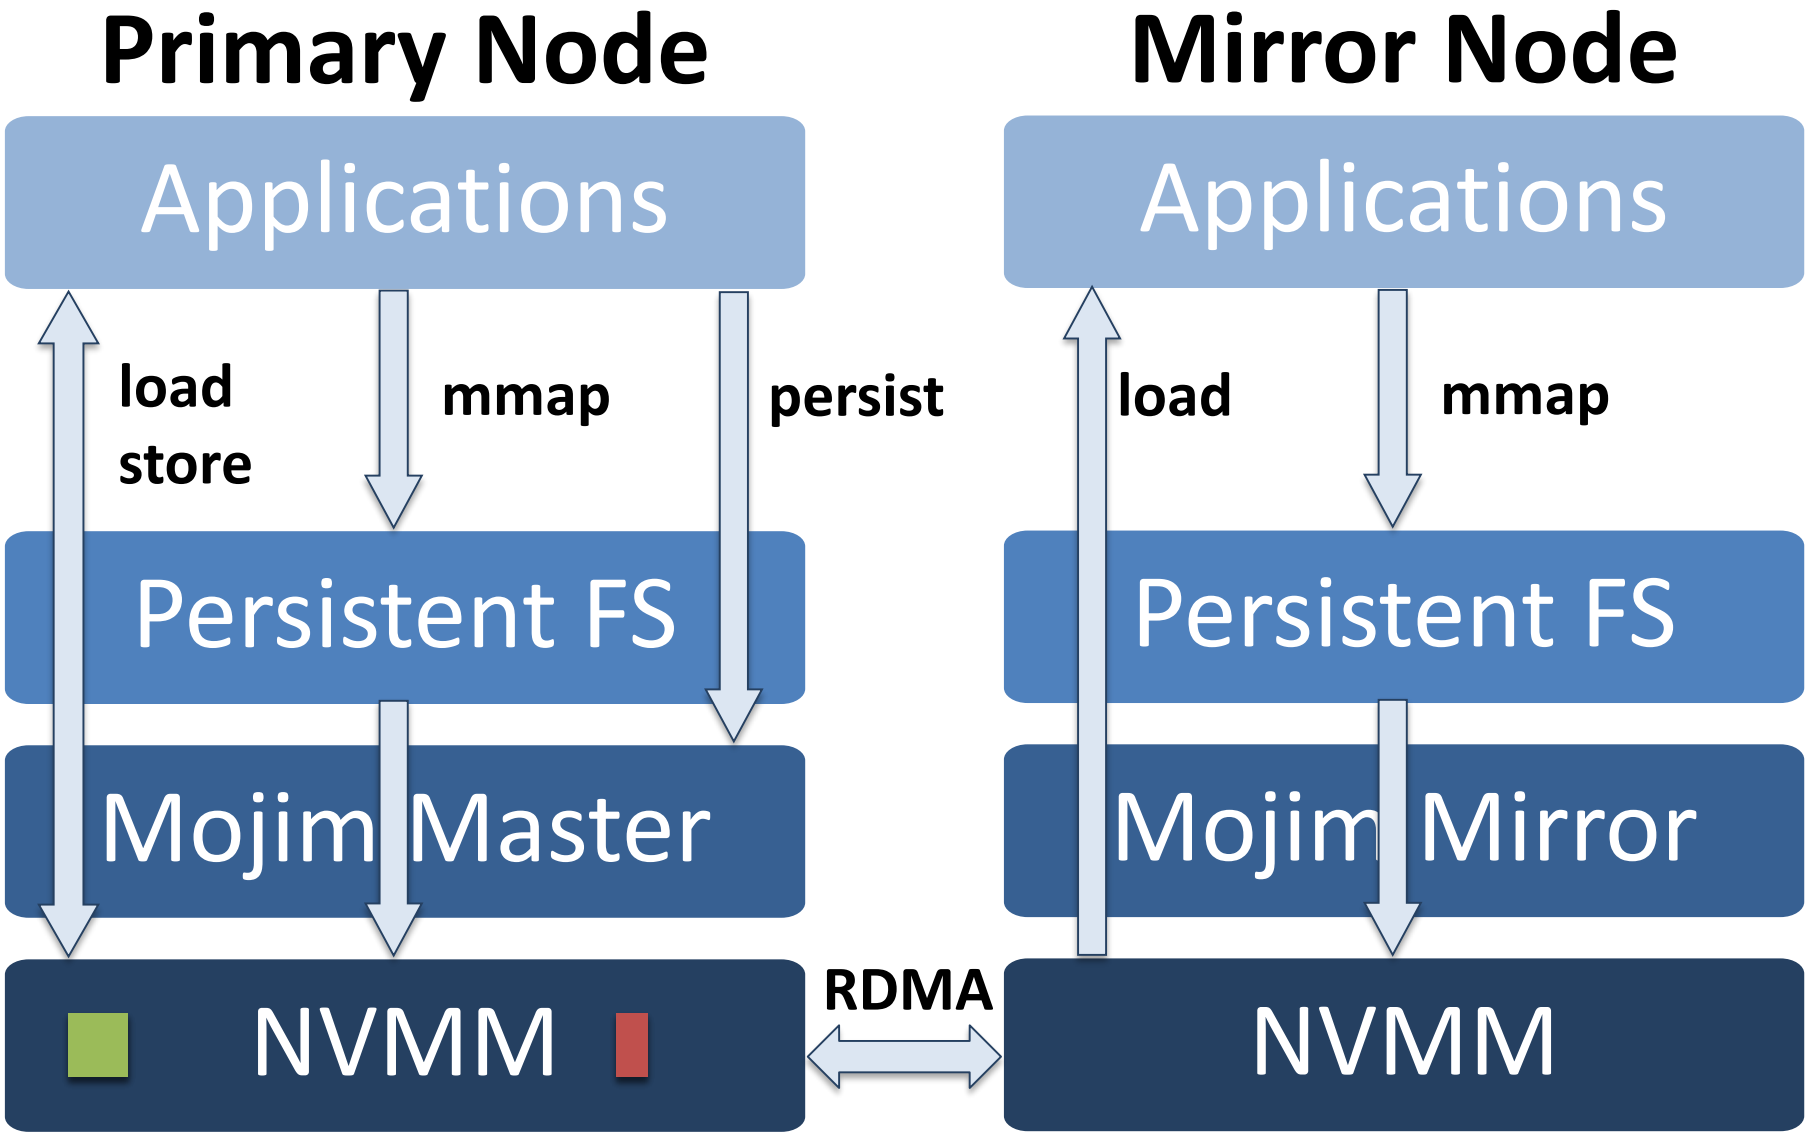

Mojim Architecture

opt1

Primary Node

Mirror Node

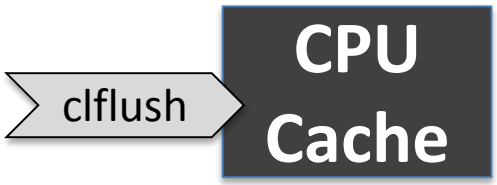

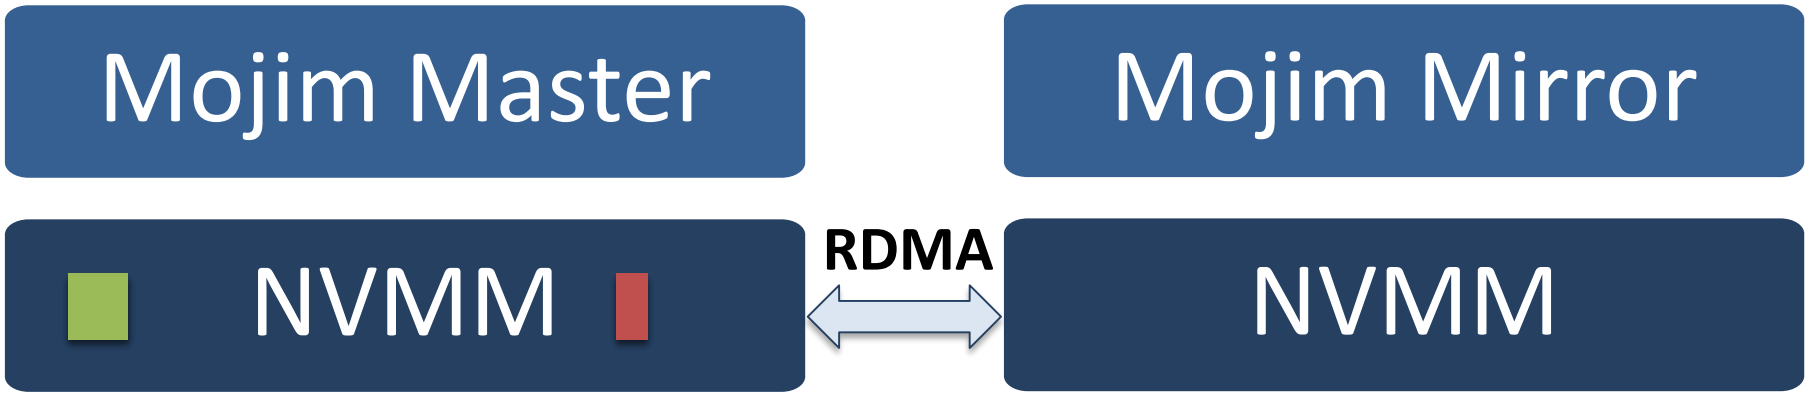

Mojim Architecture

opt1 opt2

Primary Node

Mirror Node

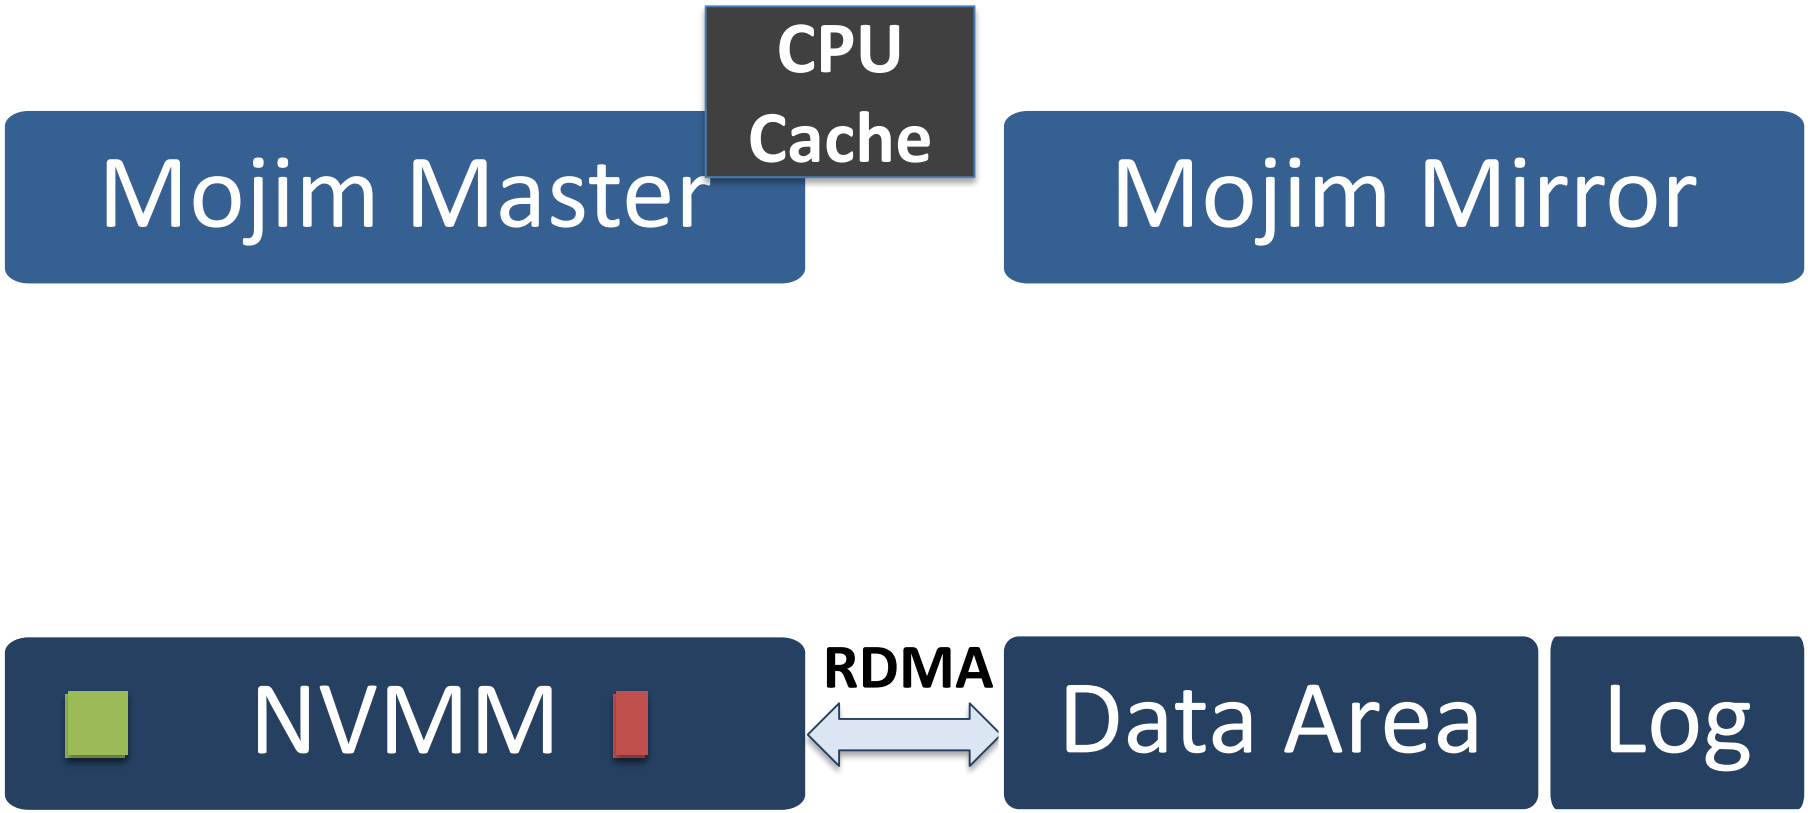

Mojim Architecture

opt1 opt2 opt3

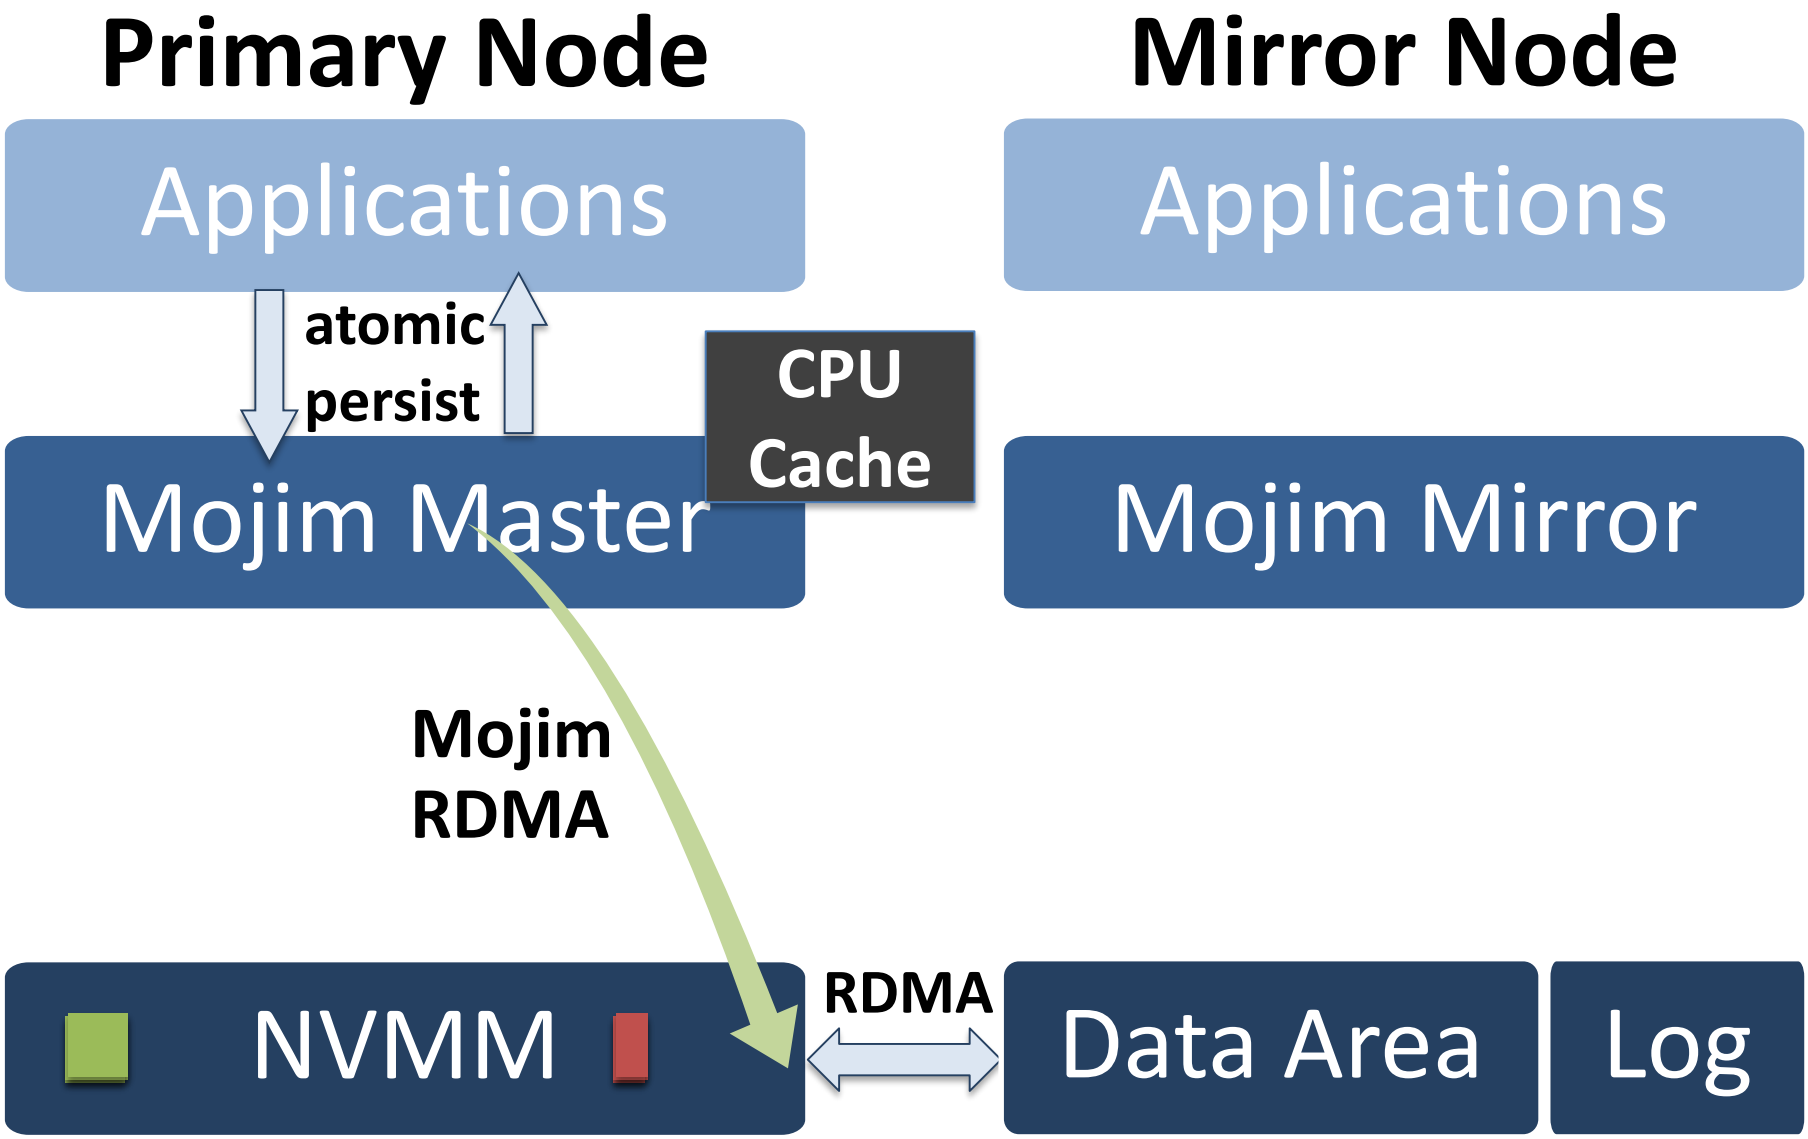

Mojim Architecture

opt1 opt2 opt3 opt4 opts ...

Primary Node

Mirror Node

Applications

Applications

Primary Tier

Mojim Master

Mojim Mirror

IB/Ethernet

Backup

Mojim Backup

Secondary Tier

NVMM

Flash/Disk

Flexible Modes

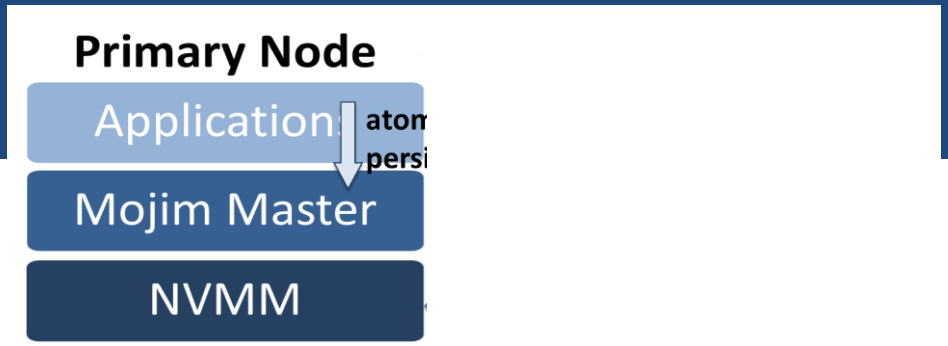

	Scheme	Performance	Reliability	Availability	Consistency	$ Cost
	Un-replicated	Good	0	Worst	N/A	Low
Mojim	Async	Good	1	Good	Weak	Fair
Mojim	Sync	Good	1	Good	Strong	Fair
Mojim	Sync-disk	Good	1	OK	Strong	Low
Mojim	Sync-two-tier	Good	N-1	Best	Strong+Weak	High
Mojim	Sync-twotier-ETH	Bad	N-1	Good	Strong+Weak	Fair
Existing	Write-all	Bad	N-1	Best	Strong	High
Existing	Chain-rep	OK	N-1	Best	Strong	High
Existing	Broadcast-rep	OK	N-1	Best	Strong	High

Outline

- Introduction
- Mojim Design and Architecture
- Mojim Implementation
- Evaluation Results
- Conclusion

Implementation

- Mojim as a generic layer in Linux kernel
- Networking
  - Optimized implementation of IBV-verbs in kernel
  - Zero copy, reliable
  - Multiple connections, multiple receiving threads polling
- Replication and recovery
  - Redo logs on mirror and backup nodes
  - Atomic operation support
  - Fast recovery (ensured by thresholds)

Primary Tier Replication

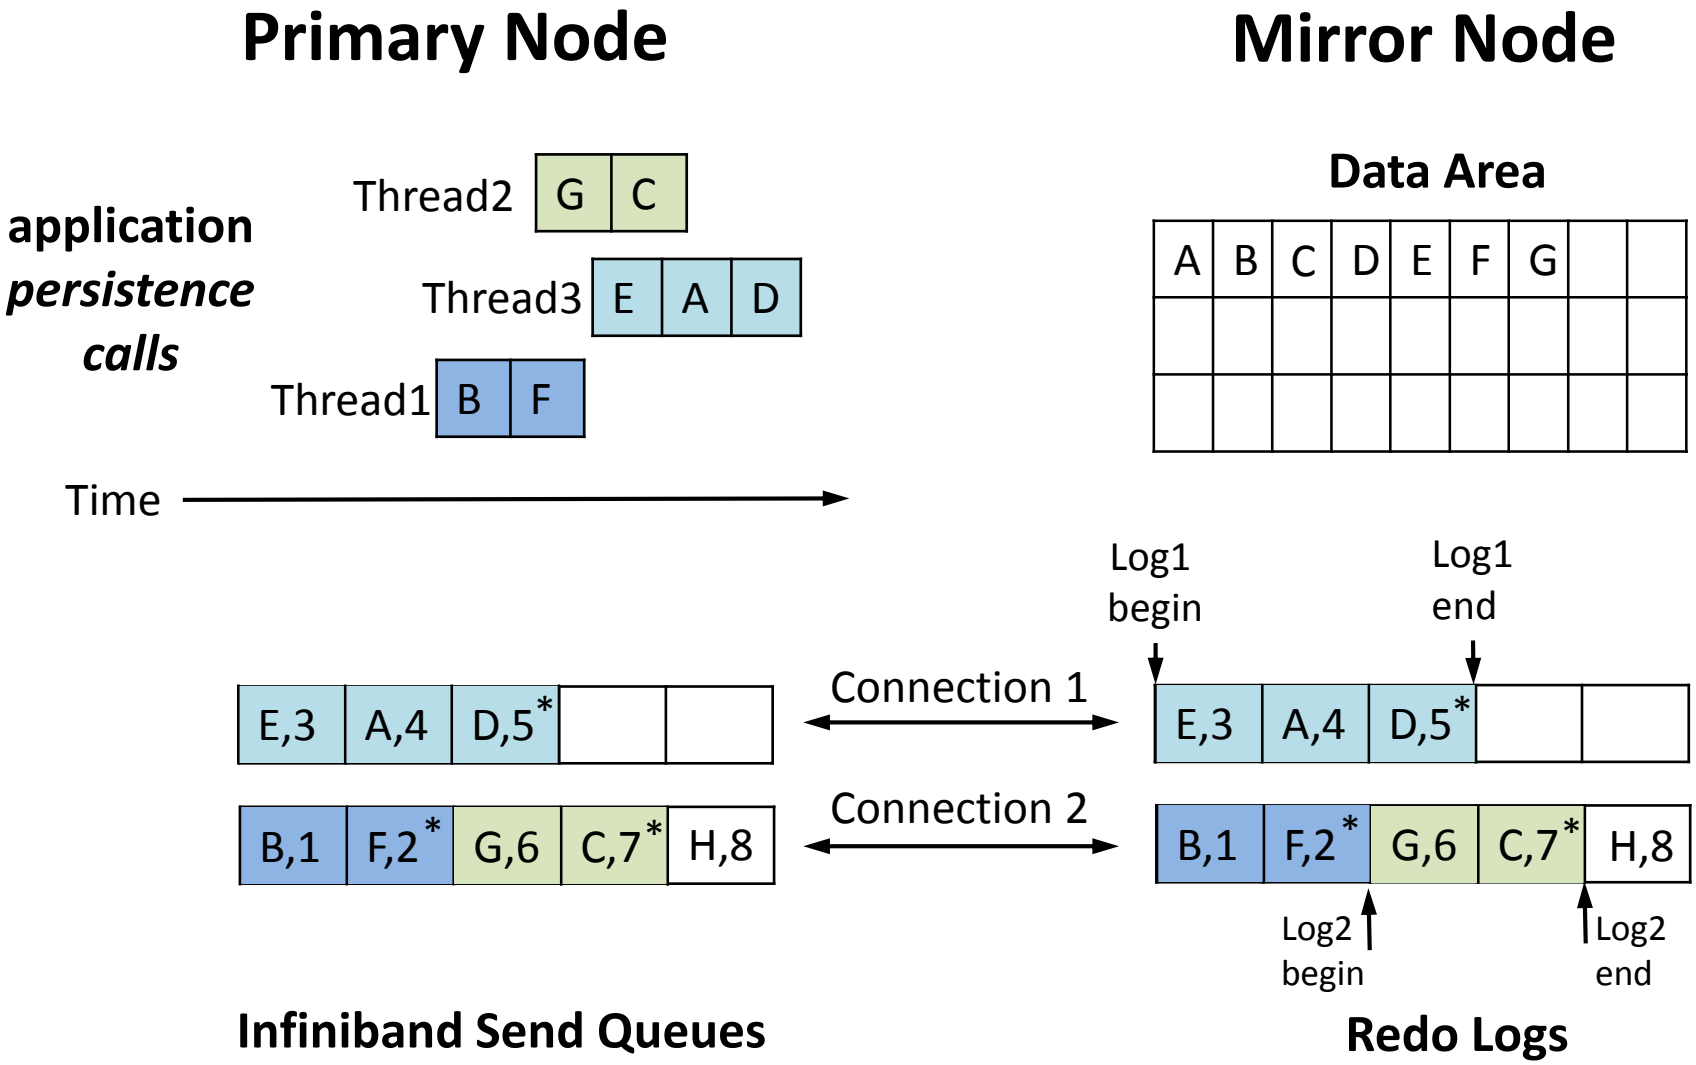

Mojim Applications

- Persistent Memory File System
- Google Hash Table
- MongoDB
- No or small change to applications

Outline

- Introduction
- Mojim Design and Architecture
- Mojim Implementation
- Evaluation Results
- Conclusion

Data Persistence Latency

- Testbed: DRAM as NVM, 40Gbps Infiniband, 1Gbps Ethernet
- Workload: Persist random 4KB regions in a 4GB *mmap*'d file

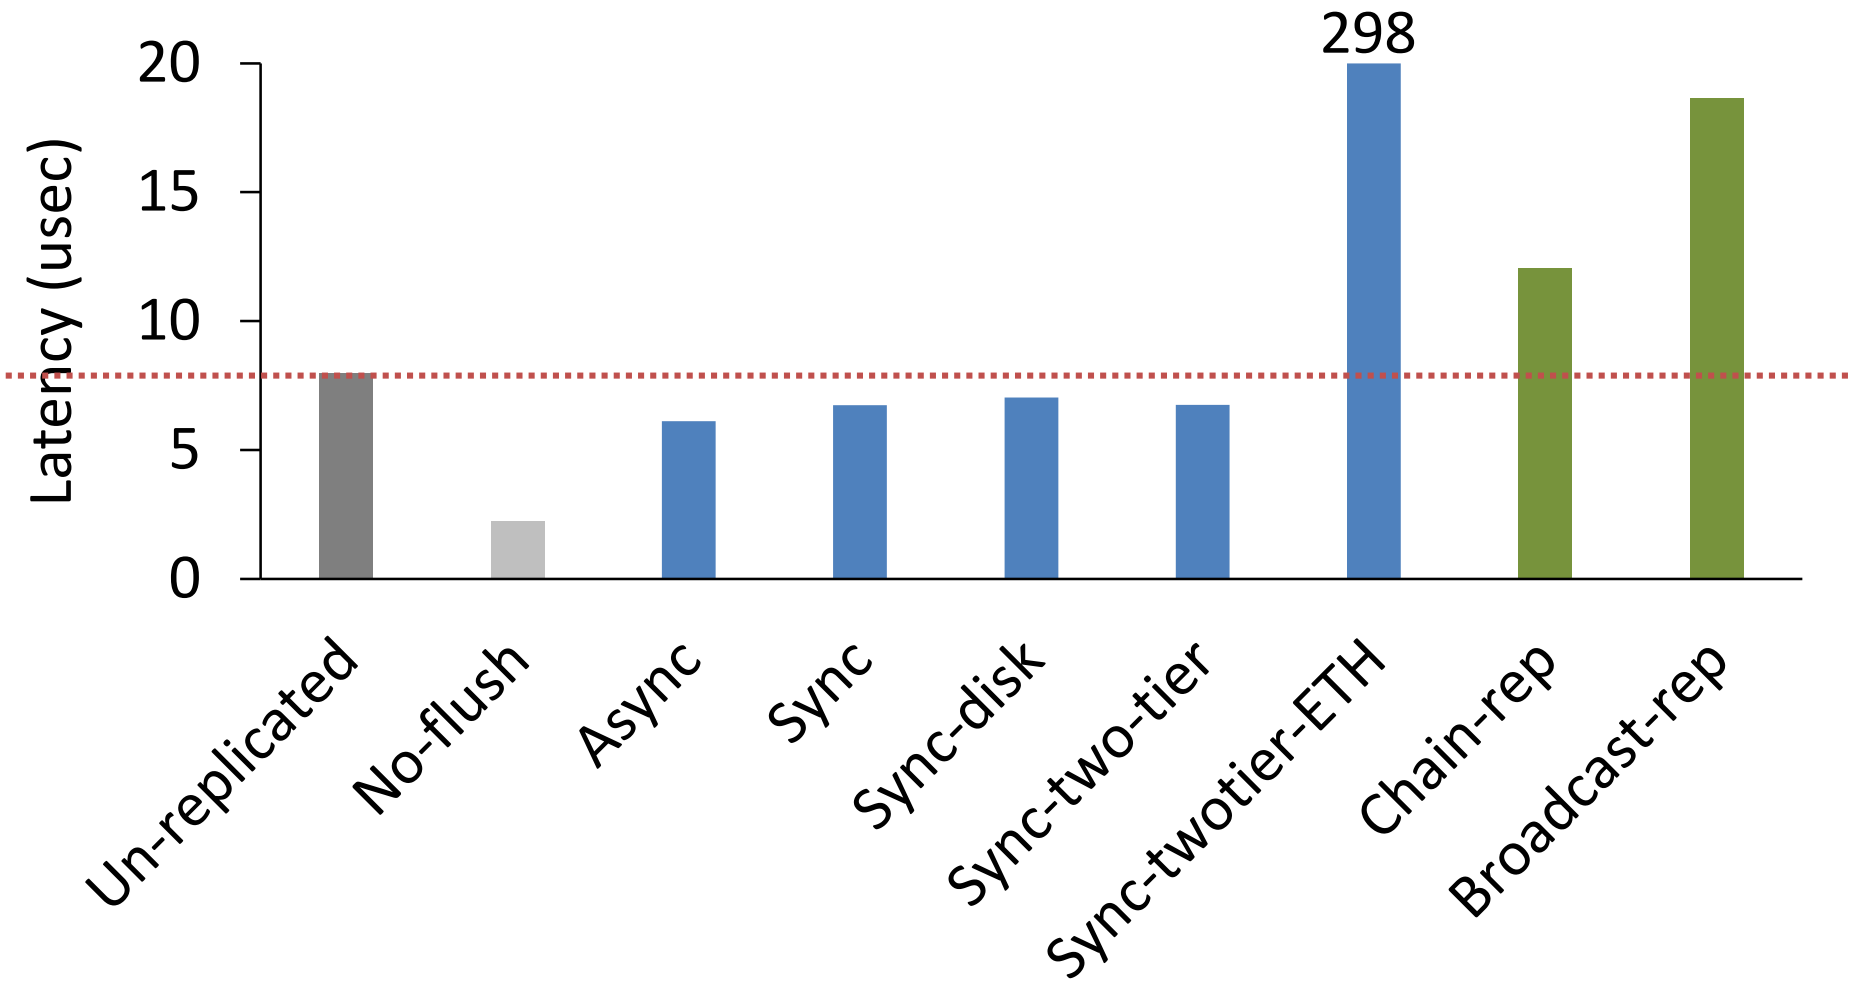

Data Persistence Throughput

- Workload: Persist random 12KB regions in a 4GB *mmap*'d file

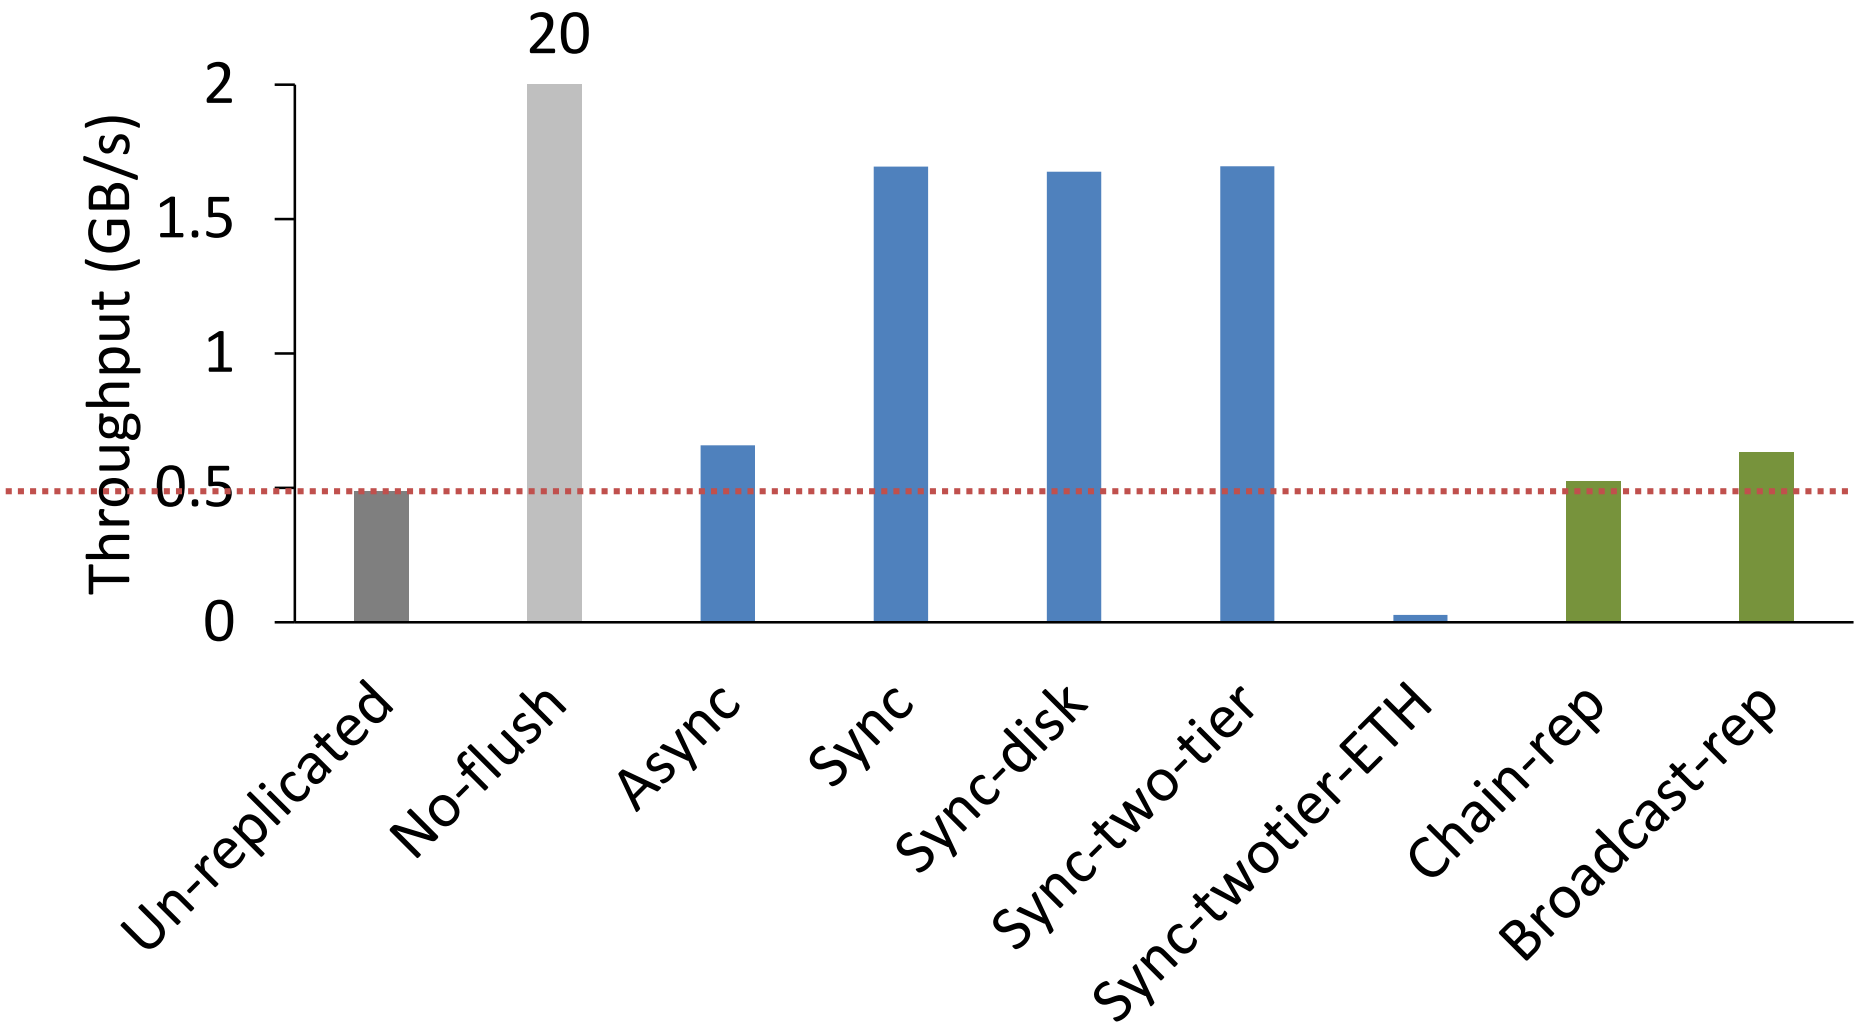

Effect of Persisted Data Size

Mojim latency **29% – 73%** of un-replicated

Mojim throughput **0.5 – 3.5x** un-replicated

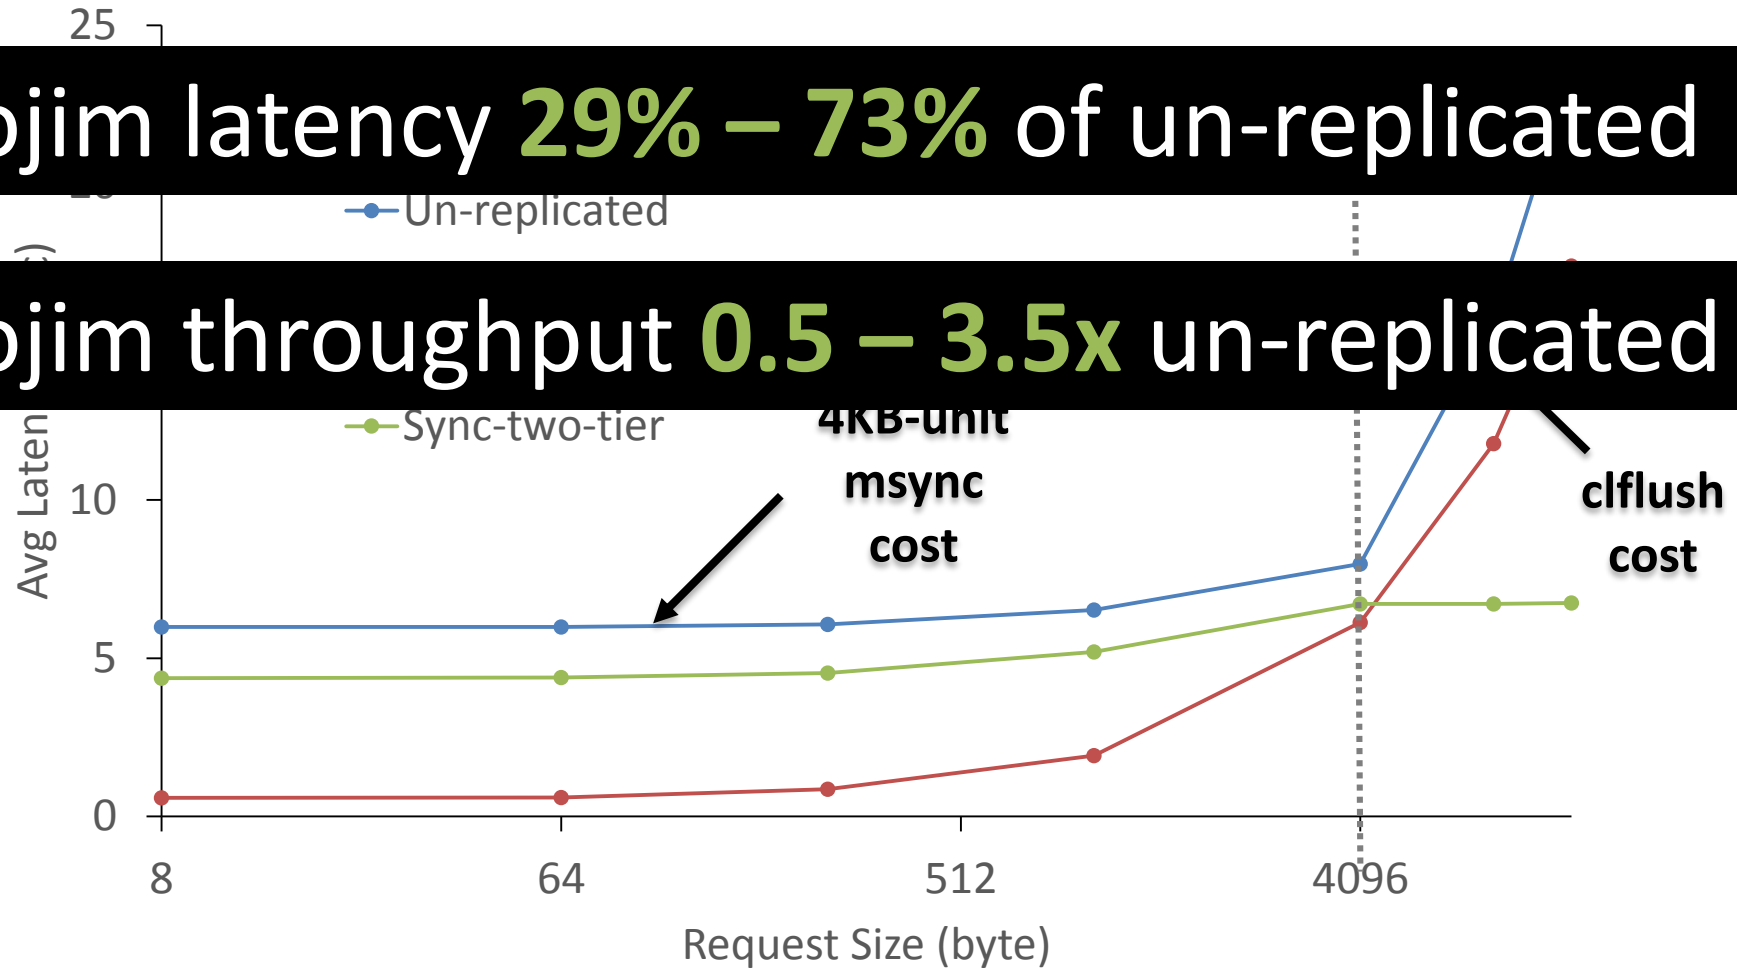

MongoDB Key-Value Pair Load

Mojim much faster than existing replication

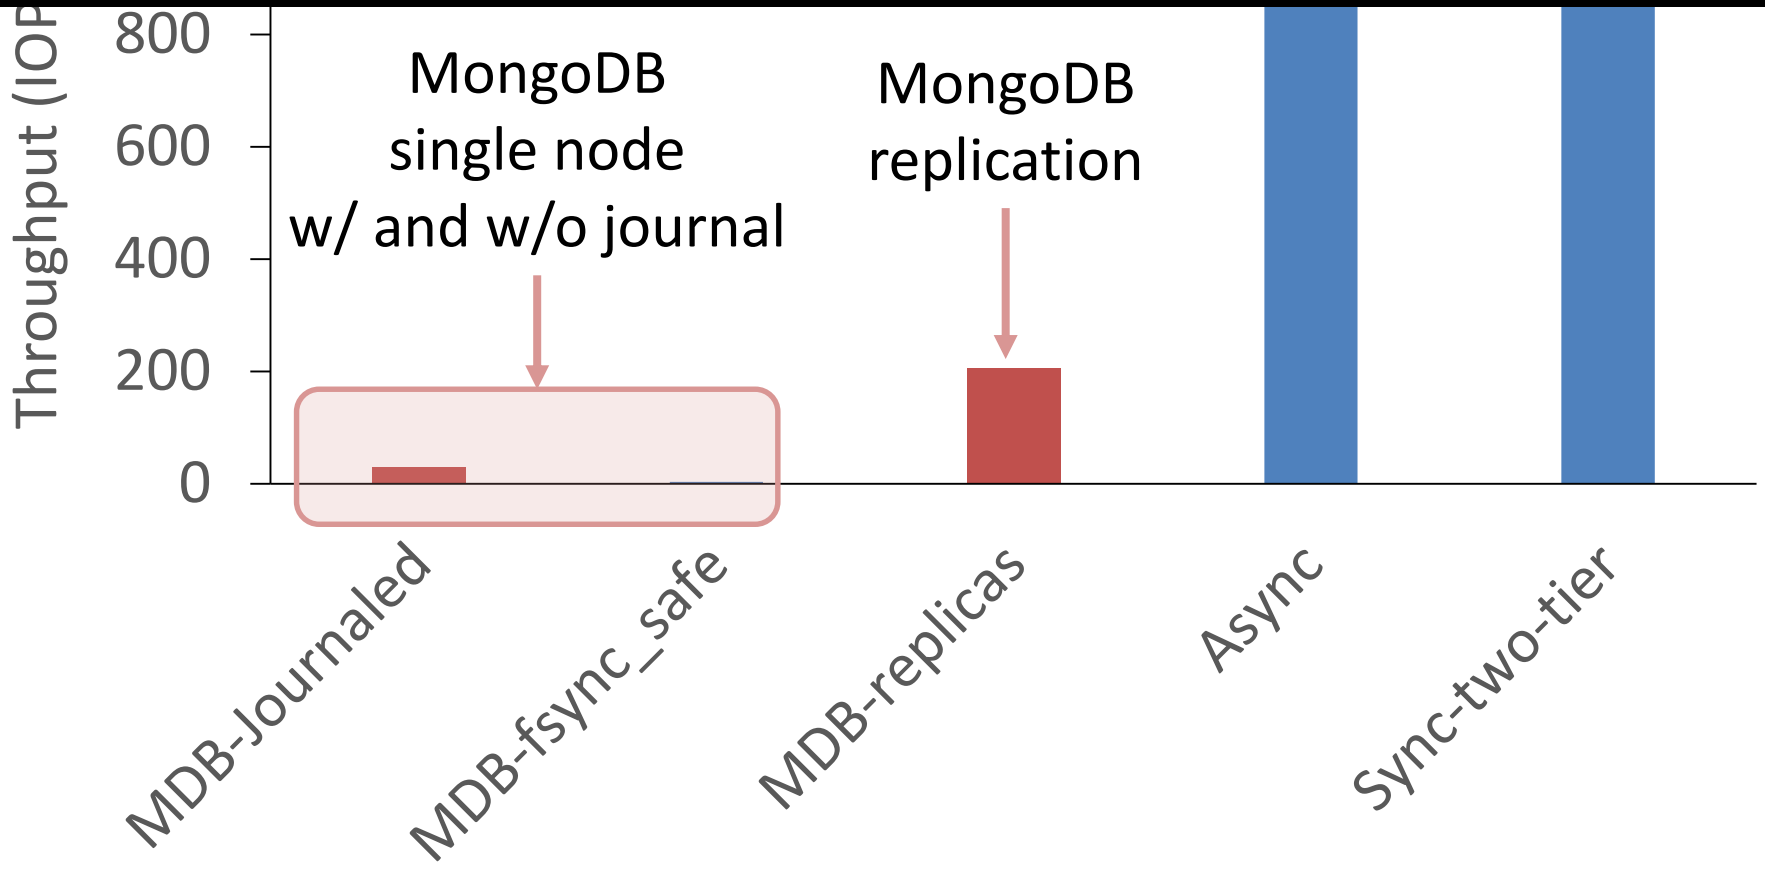

Mojim Summary

- Provide reliability and high availability to NVMM
- RDMA-based replication optimized for NVMM
- Two-tier architecture
- Flexible modes offering different guarantees
- Performance even better than un-replicated

Conclusion

- Gap between storage and memory getting smaller
- Time to rethink traditional software/networking
- More problems to be solved
  - Virtualization (e.g., replication/migration/snapshot)
  - Distributed systems (built on top of Mojim?)
  - Abstraction, programming language
  - Mobile devices

Thank you !
Questions ?

Yiying Zhang yiyingzhang@cs.ucsd.edu
Steven Swanson swanson@cs.ucsd.edu