


SPANStore: Cost-Effective Geo-Replicated Storage Spanning Multiple Cloud Services

Zhe Wu*, Michael Butkiewicz*, Dorian J. Perkins*,
Ethan Katz-Bassett+, and Harsha V. Madhyastha*

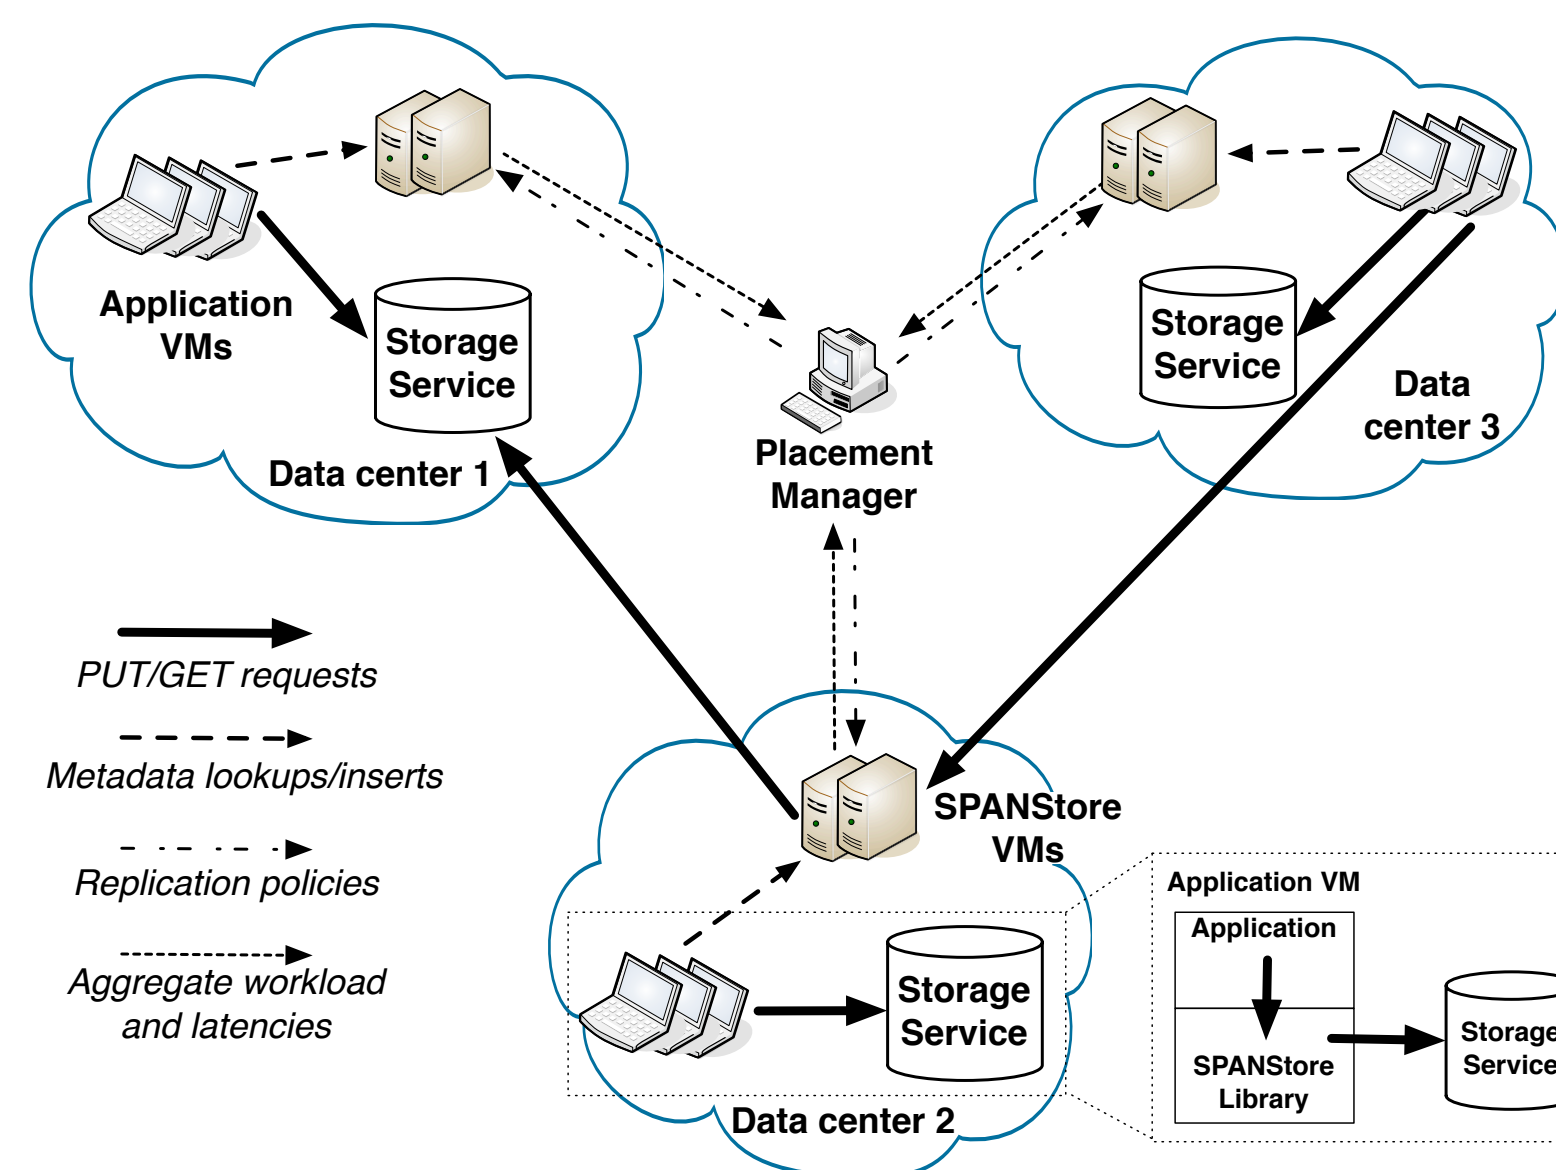
*UC Riverside, +Univ. of Southern California



Problem statement

- Cloud services enable latency-sensitive applications to serve their applications from geographically distributed set of data centers
- However, isolated storage service is offered in every data center, leaving replication to applications
- Simple replication strategies (e.g., replicate to all data centers or maintain single replica) incur high cost

SPANStore: Cost-minimizing unified view of cloud storage services



1. Inter-DC latencies
2. Pricing policies of storage services

SPANStore characterization

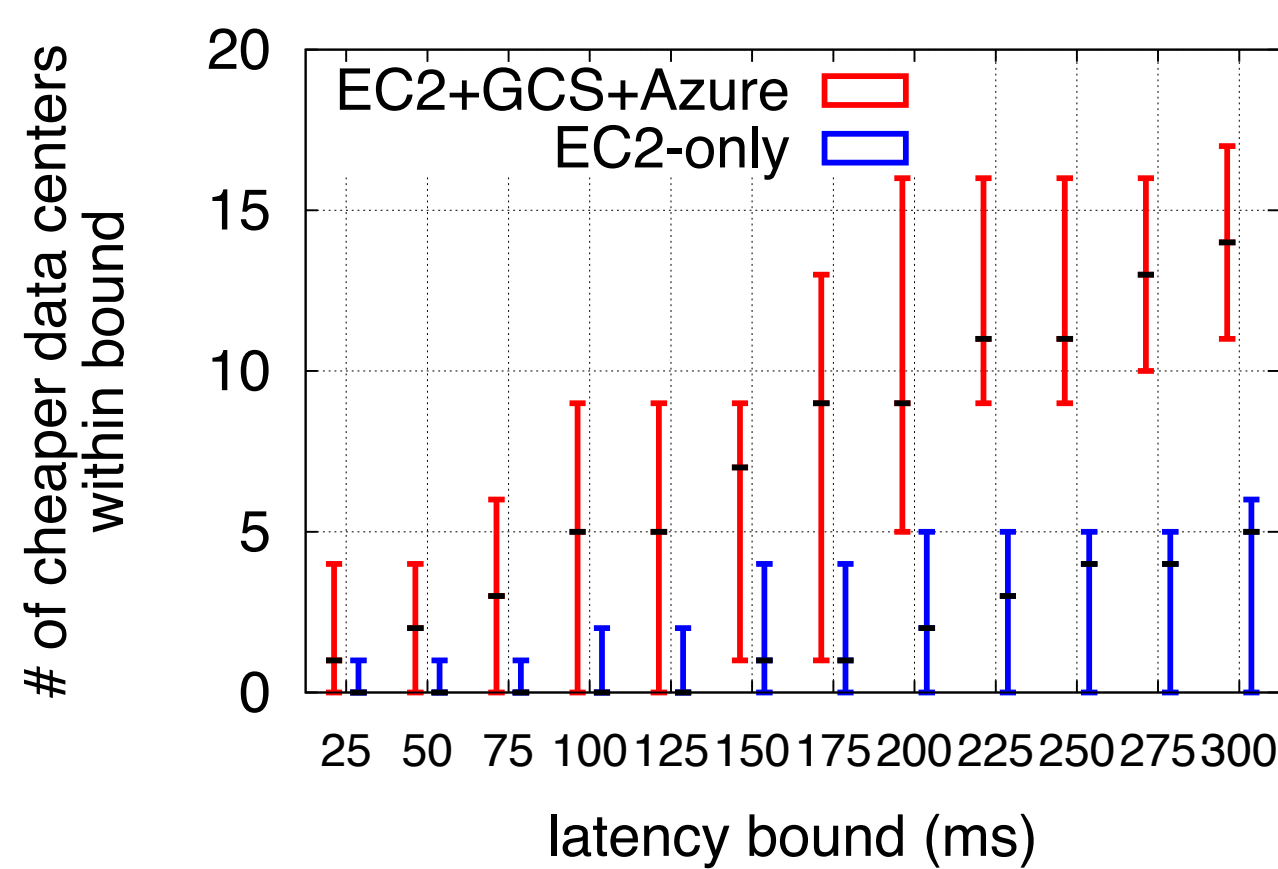
3. Latency SLOs
4. Consistency requirement
5. Fault tolerance requirements
6. Aggregate workload per access set

Application-specific inputs

Placement Manager

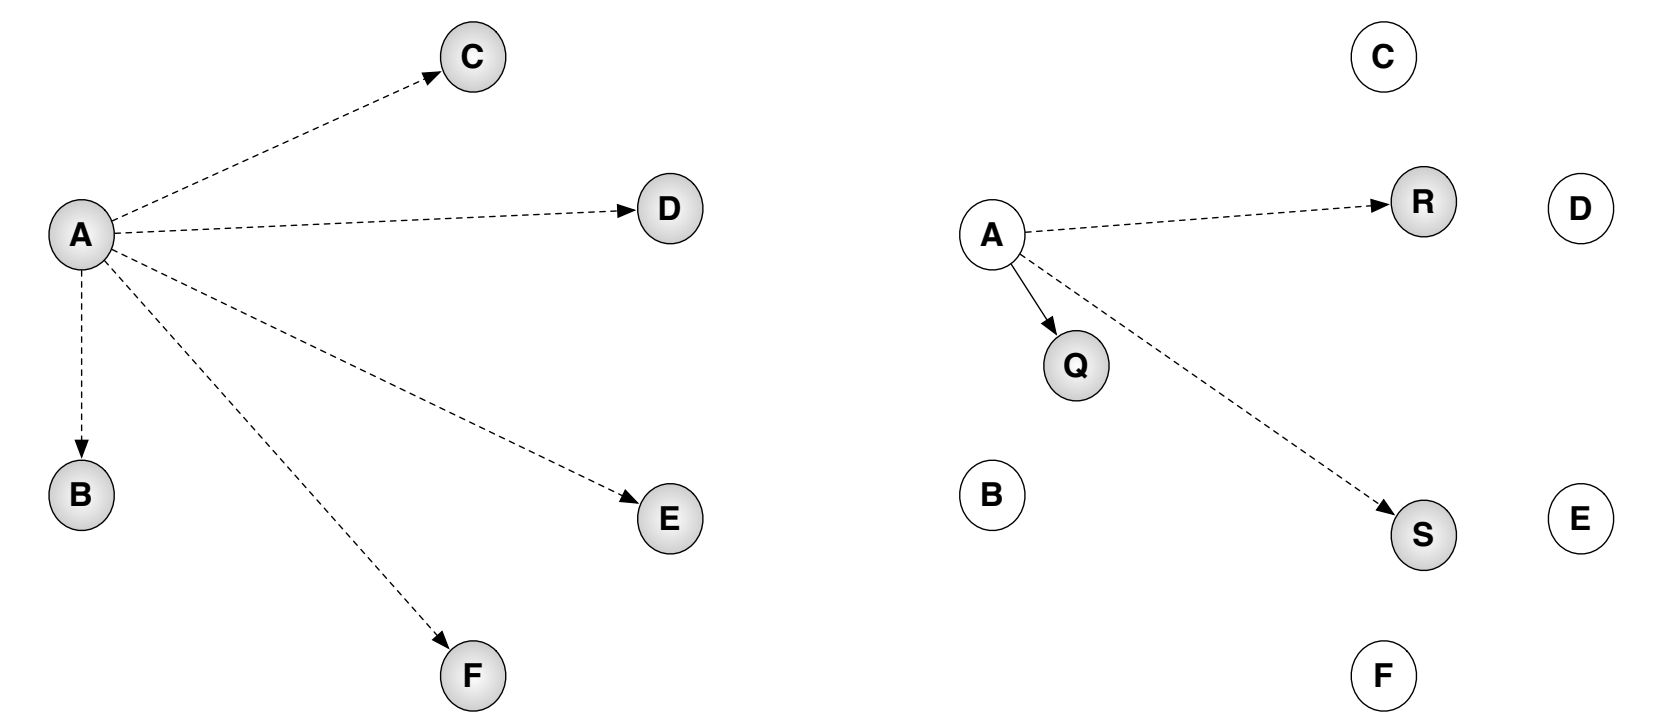
Replication policy for each access set

Use multiple cloud providers to minimize cost



Significant pricing differences between providers.
No provider is the cheapest along all dimensions.

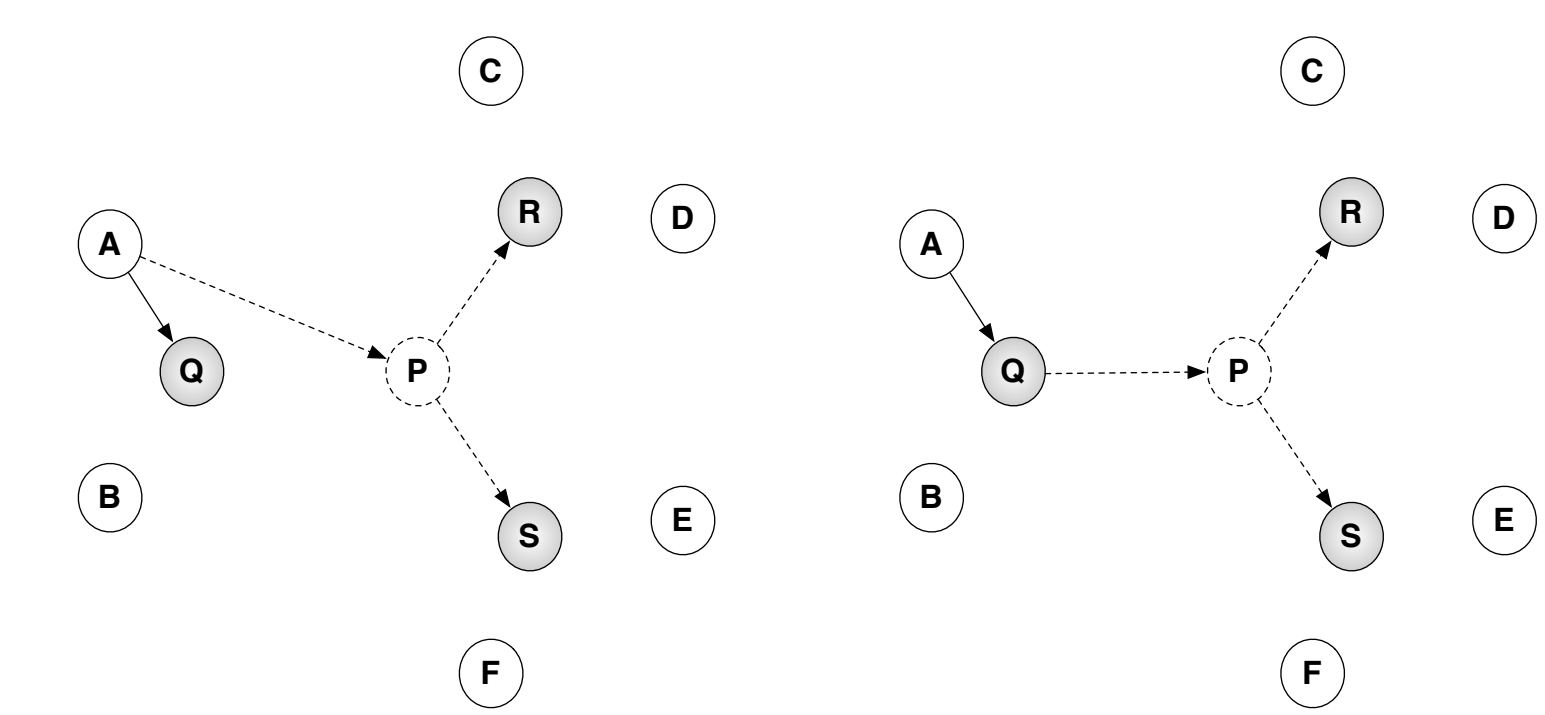
Consistency and latency dependent replication



Replicate everywhere:
+ Low GET latencies
- High networking cost

Store few copies:
+ Low storage cost
- High GET/PUT latencies

Relay data propagation to reduce costs

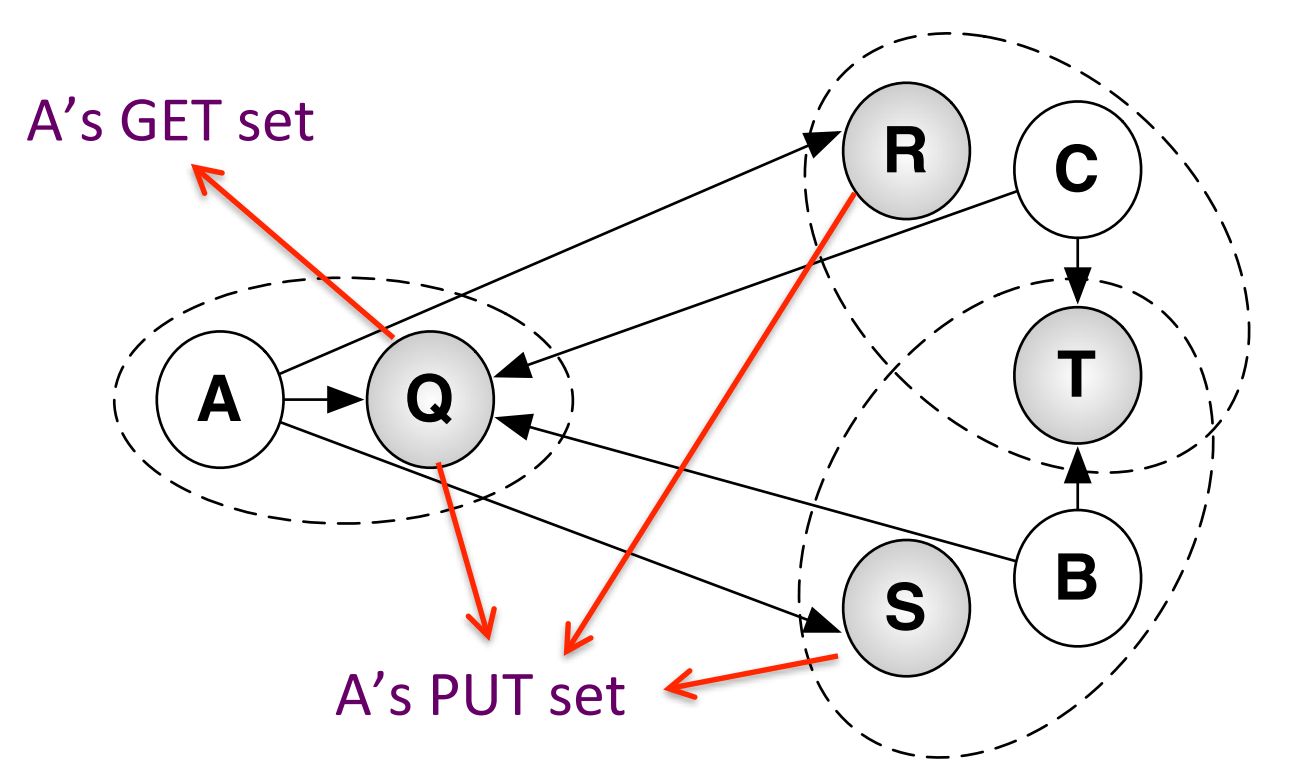


Use multicast trees to exploit bandwidth pricing discrepancies across regions.

Key techniques

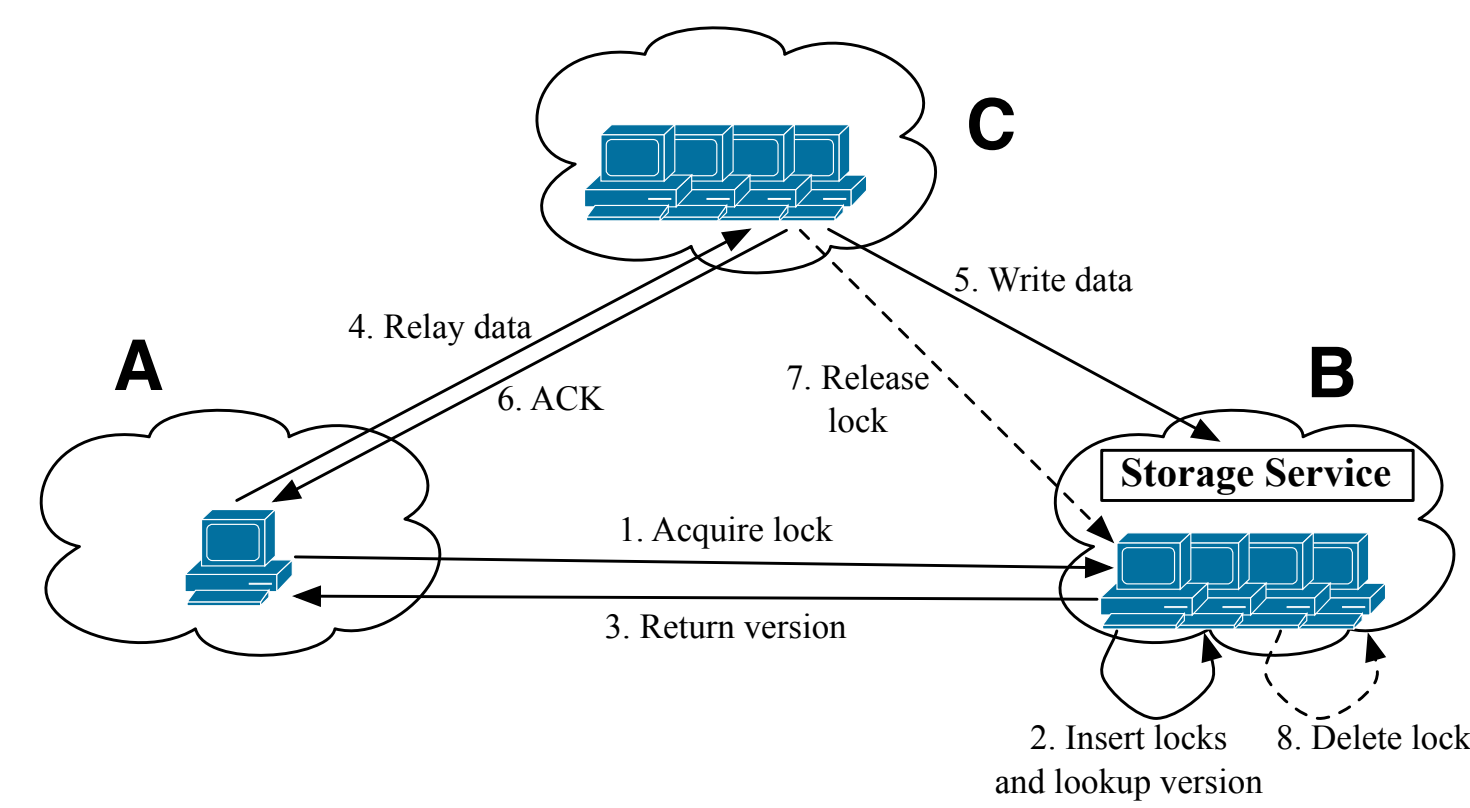
- Integer program run by placement manager to determine replication policy for every access set
- Workload considered per access set since aggregate workloads are more stable
- Asymmetric quorum sets to minimize cost and meet lower latency SLOs
- Novel two-phase commit protocol

Asymmetric quorum sets



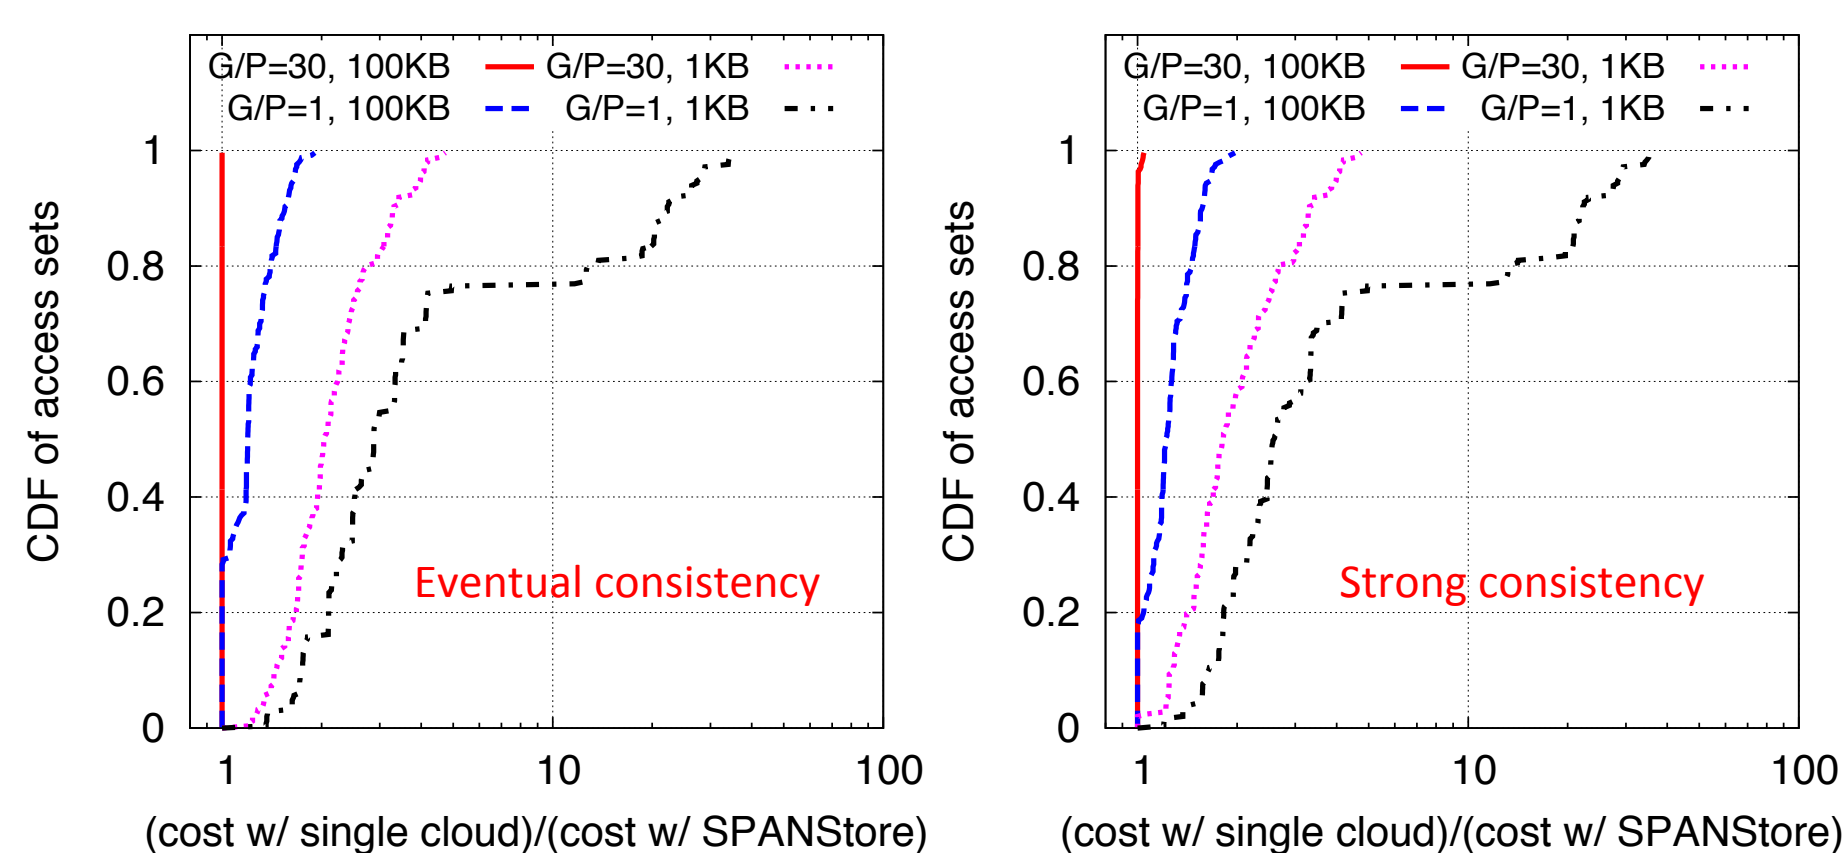
Lower latency SLOs can be met and reduces cost.

Two-phase commit protocol



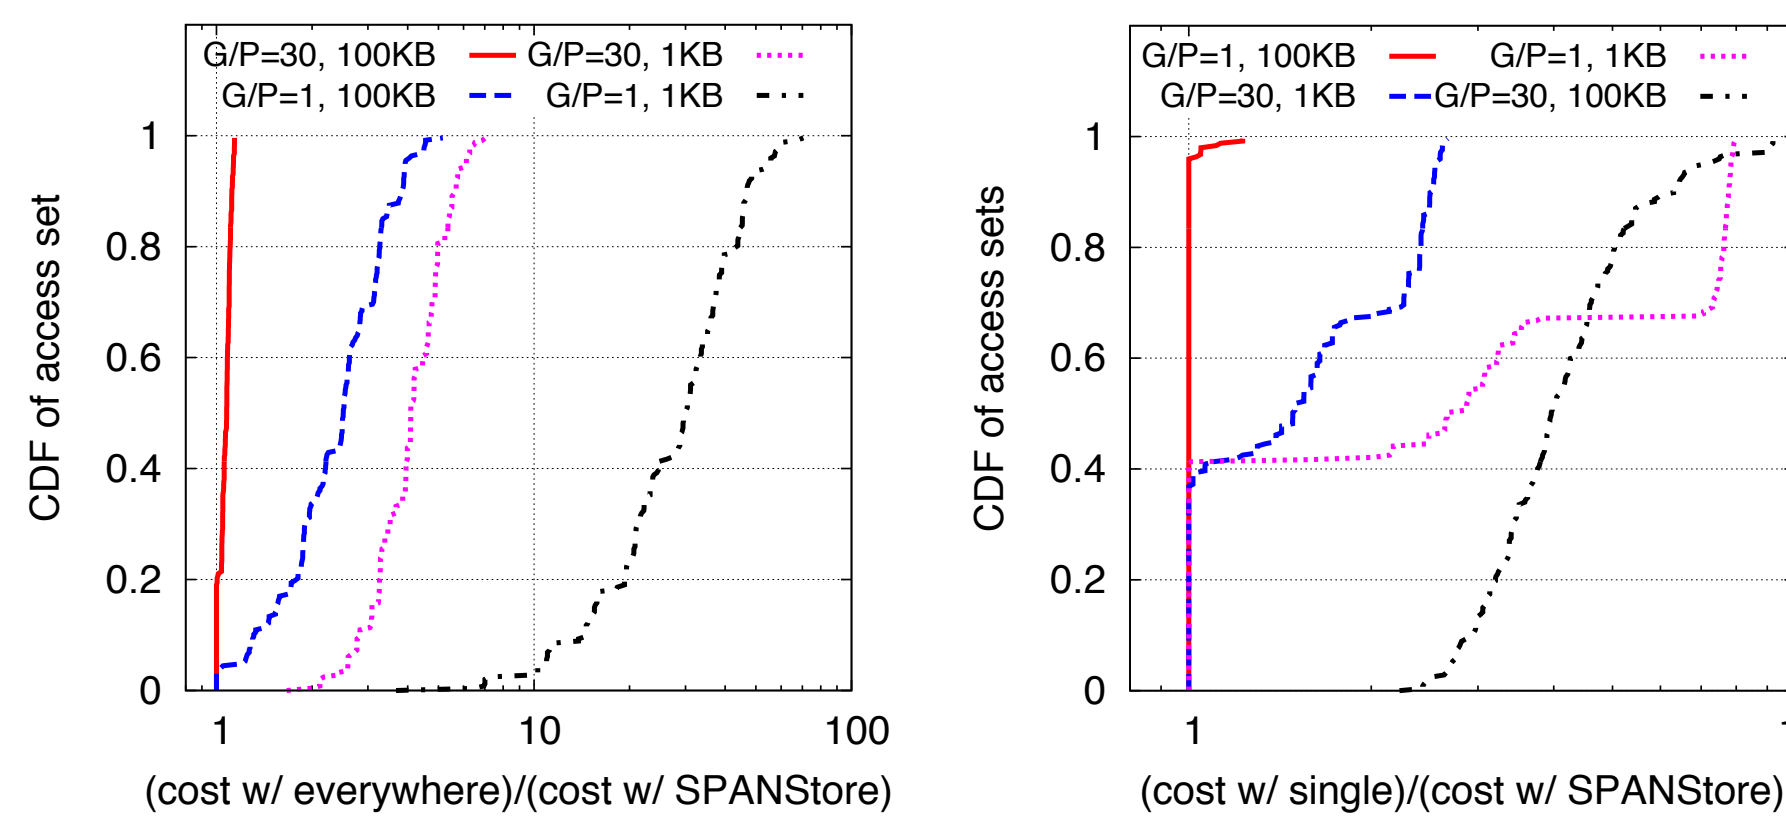
Reduces cost for compute resources by sending data directly from application VMs to storage services. SPANStore's VMs only involved in metadata operations.

SPANStore vs. single cloud



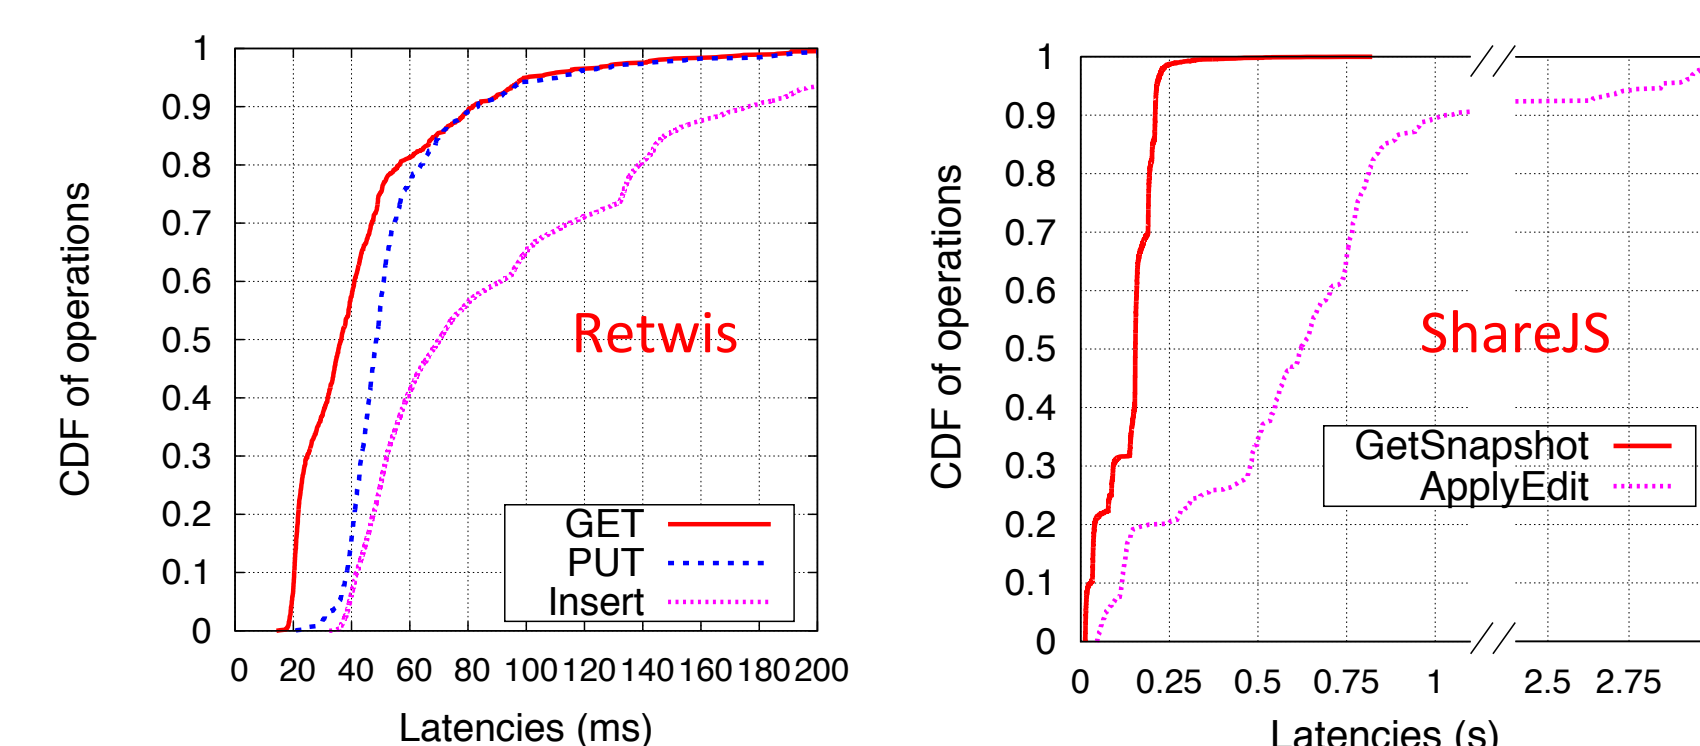
SPANStore can lower cost by over 10x.

SPANStore vs. fixed policies



SPANStore intelligently exploits price discrepancies and several other factors, which fixed policies cannot.

Using SPANStore in applications



SPANStore meets specified latency SLOs for example social networking and collaborative document editing webservices.