DynaMast

Adaptive Dynamic Mastering for Replicated Systems

Michael Abebe mtabebe@uwaterloo.ca
Brad Glasbergen
Khuzaima Daudjee

ICDE (April 2020)

tiny.cc/dynamast

UNIVERSITY OF WATERLOO

Single Database

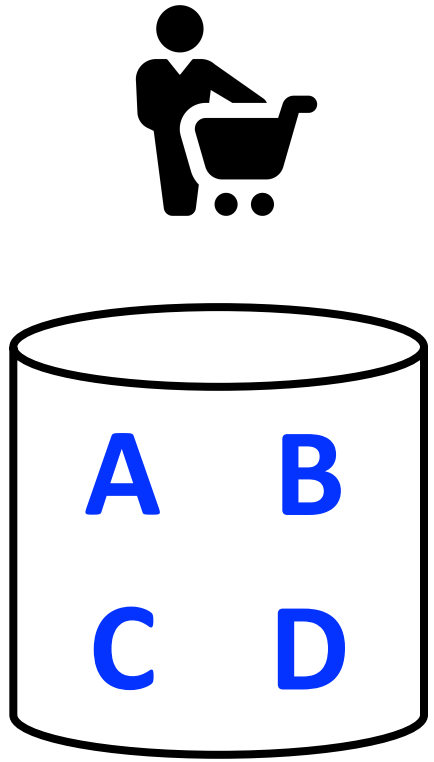

Single Database

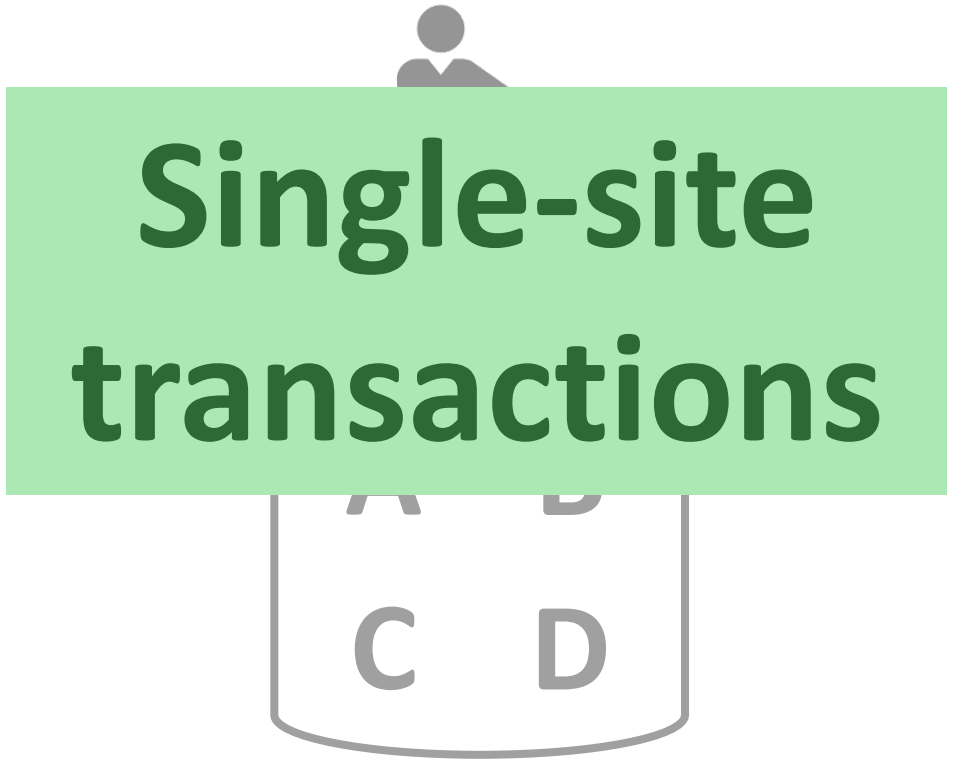

Single Database

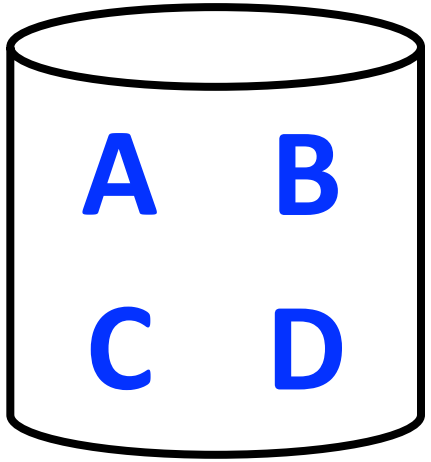

Single Database

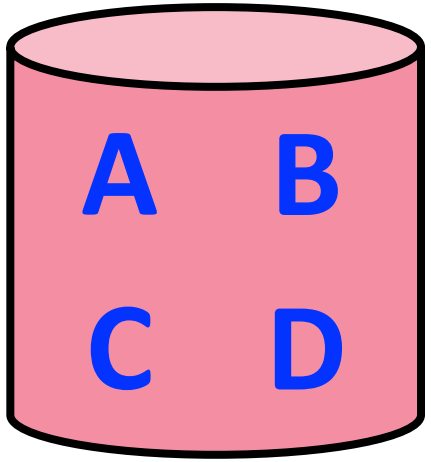

Single Database

Overloads site

C D

How to scale the database?

Single-master

Multi-master

How to scale the database?

Single-master

Multi-master

Single-Master

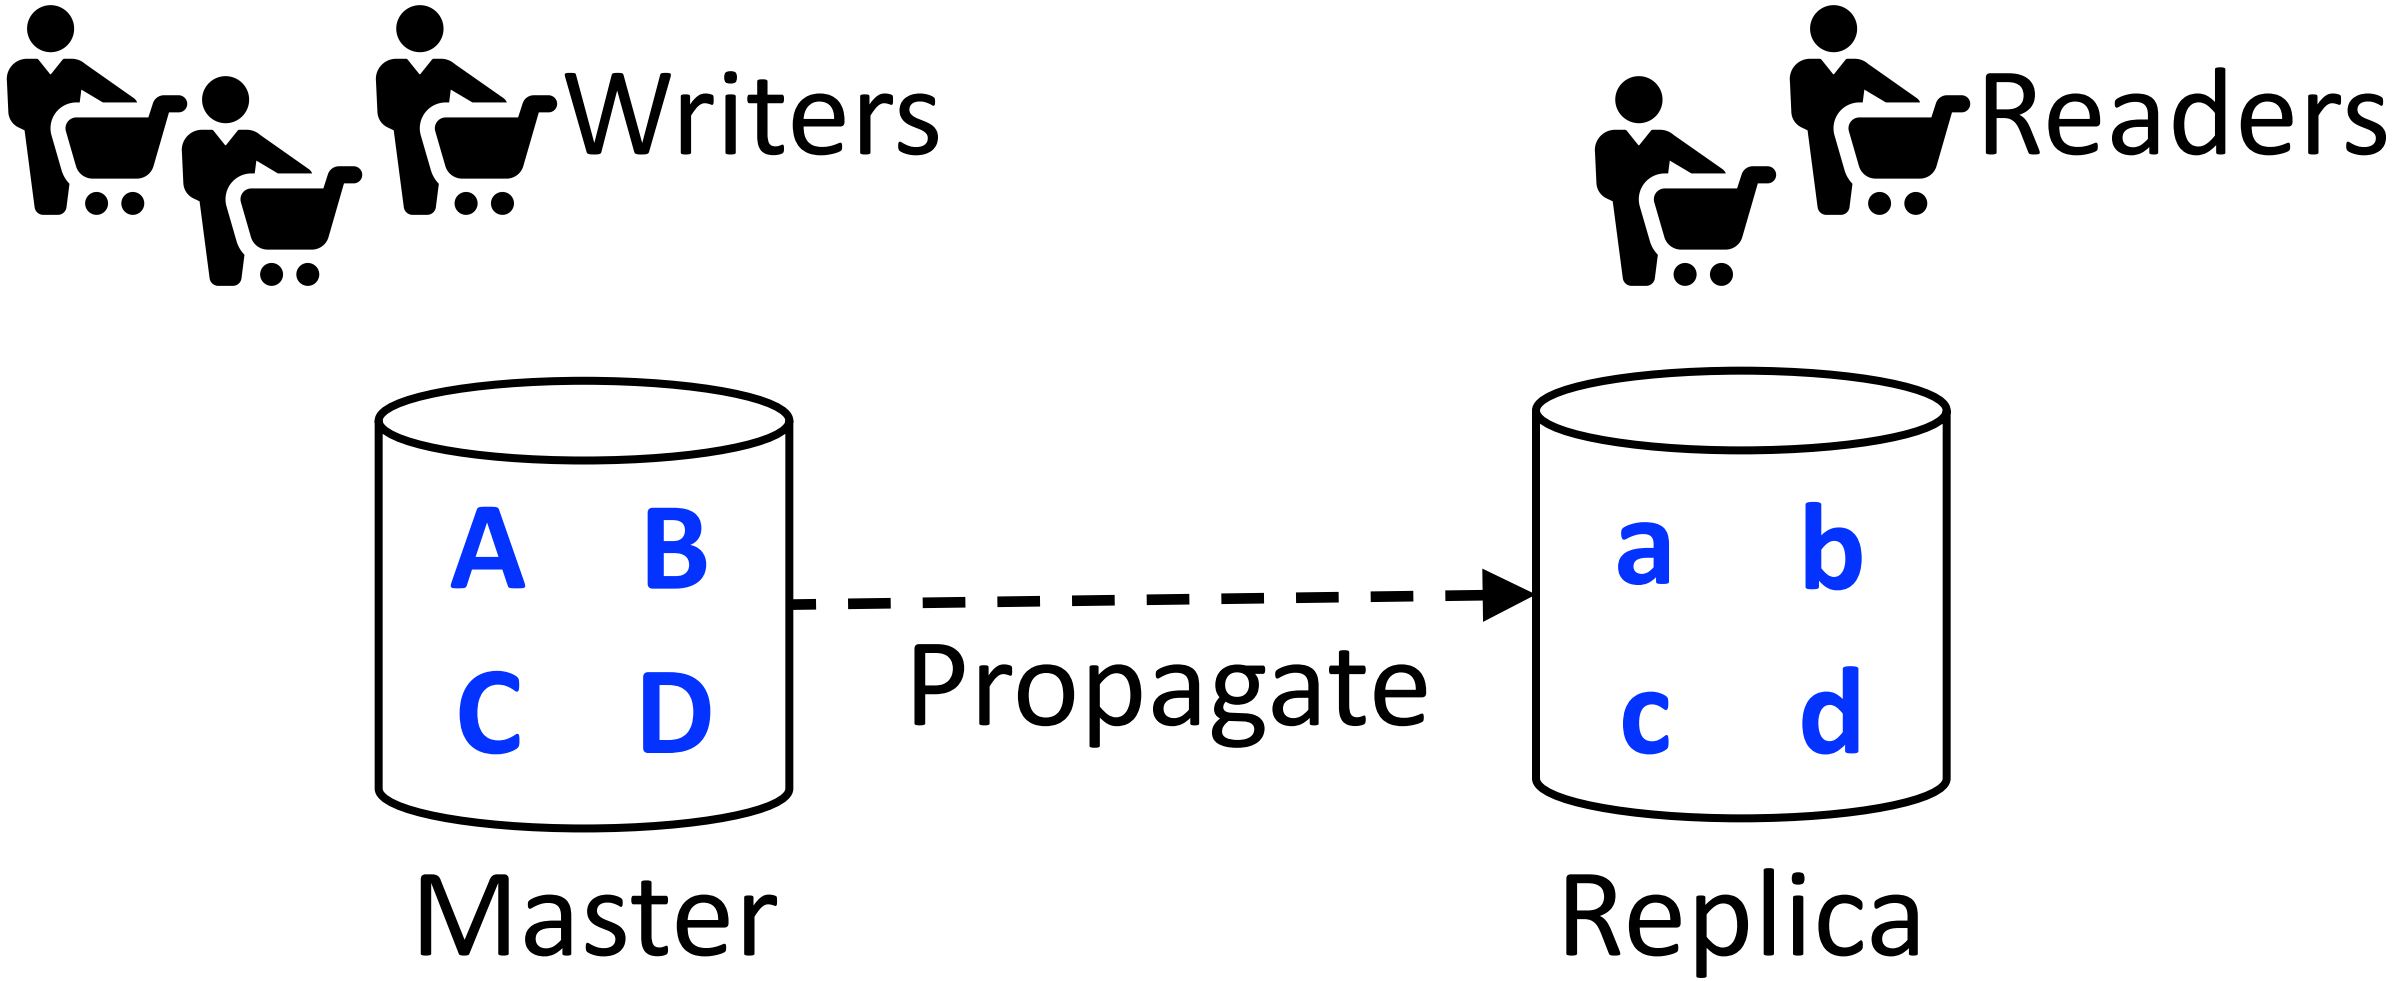

UNIVERSITY OF WATERLOO

Single-Master

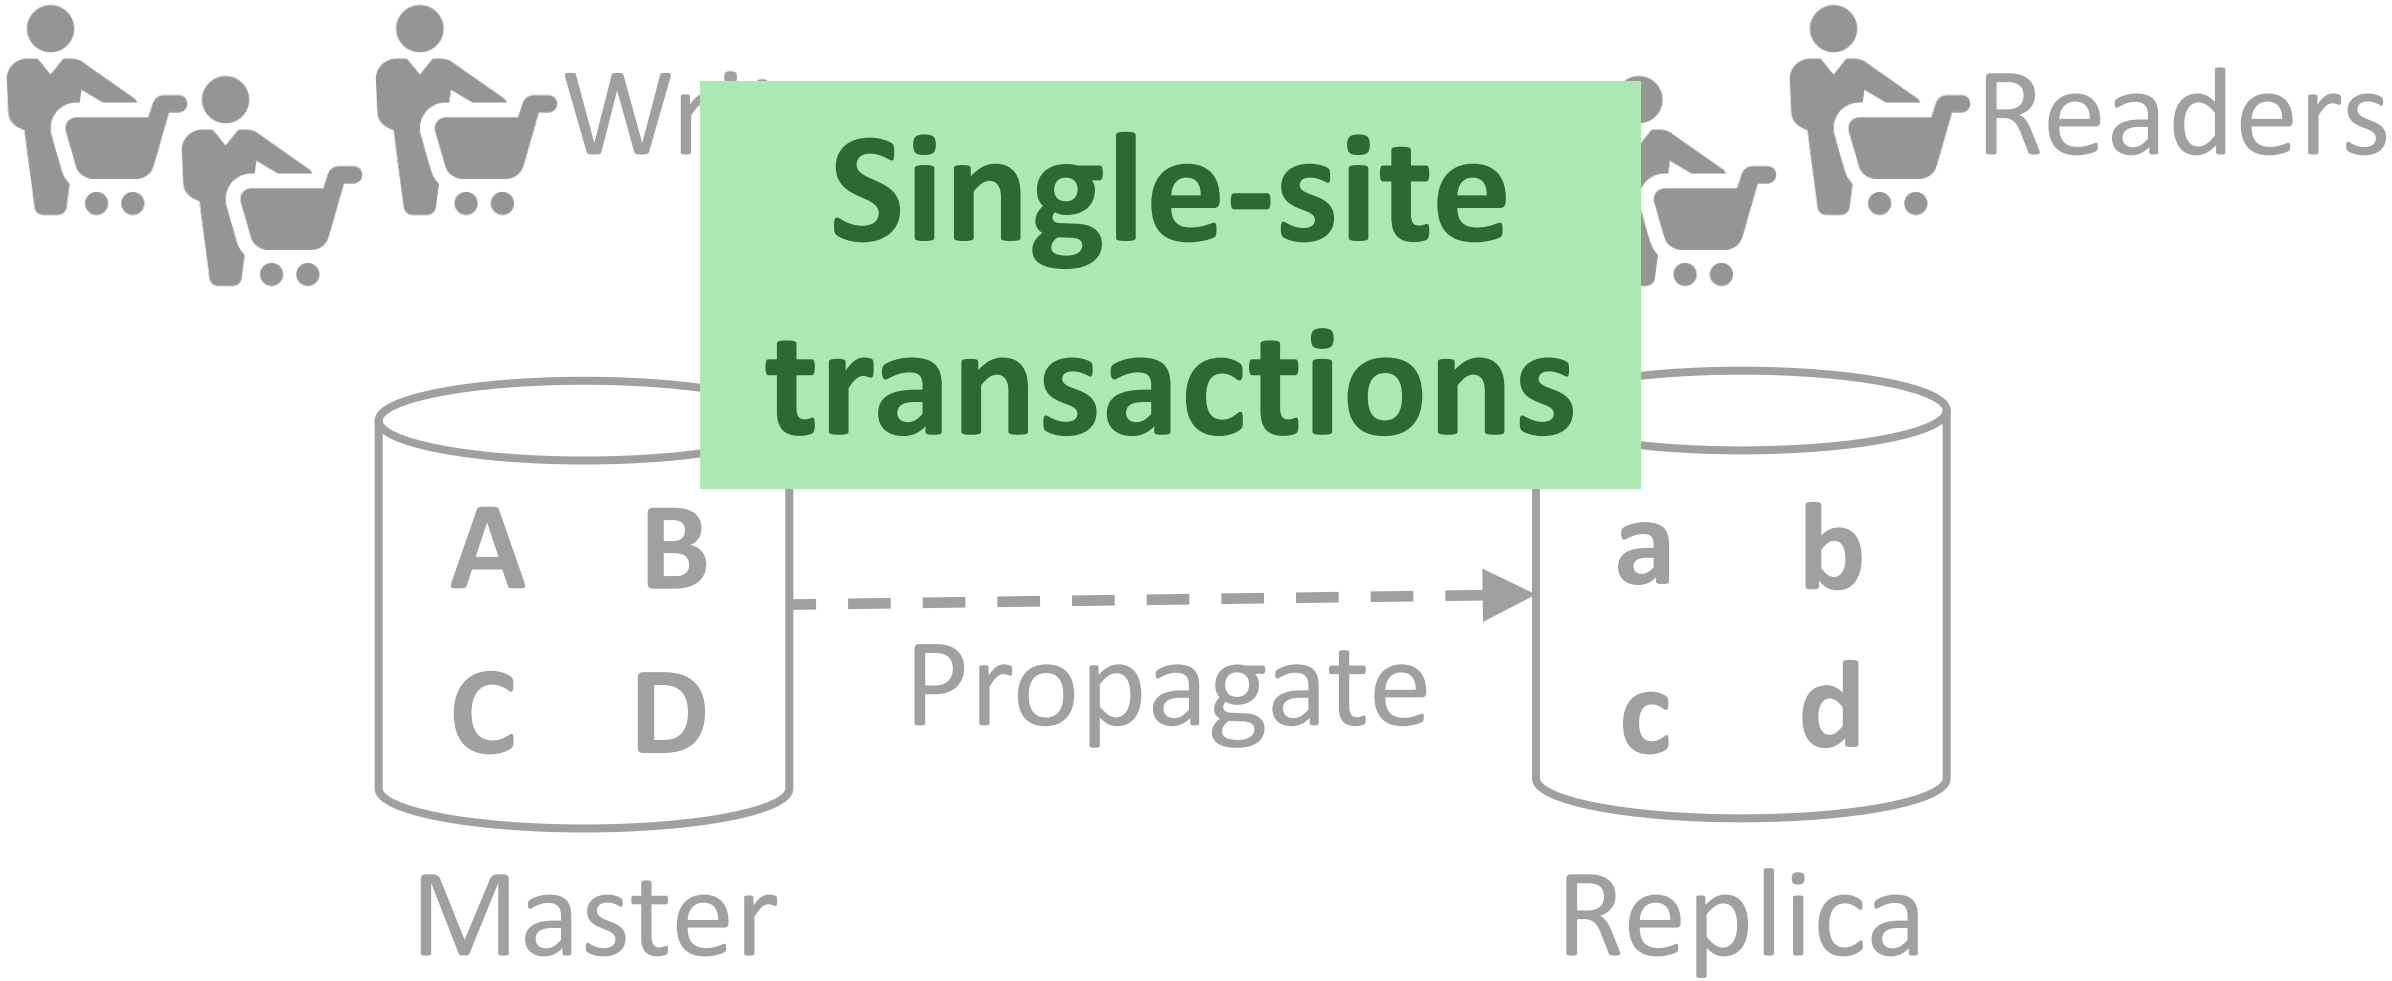

Single-Master

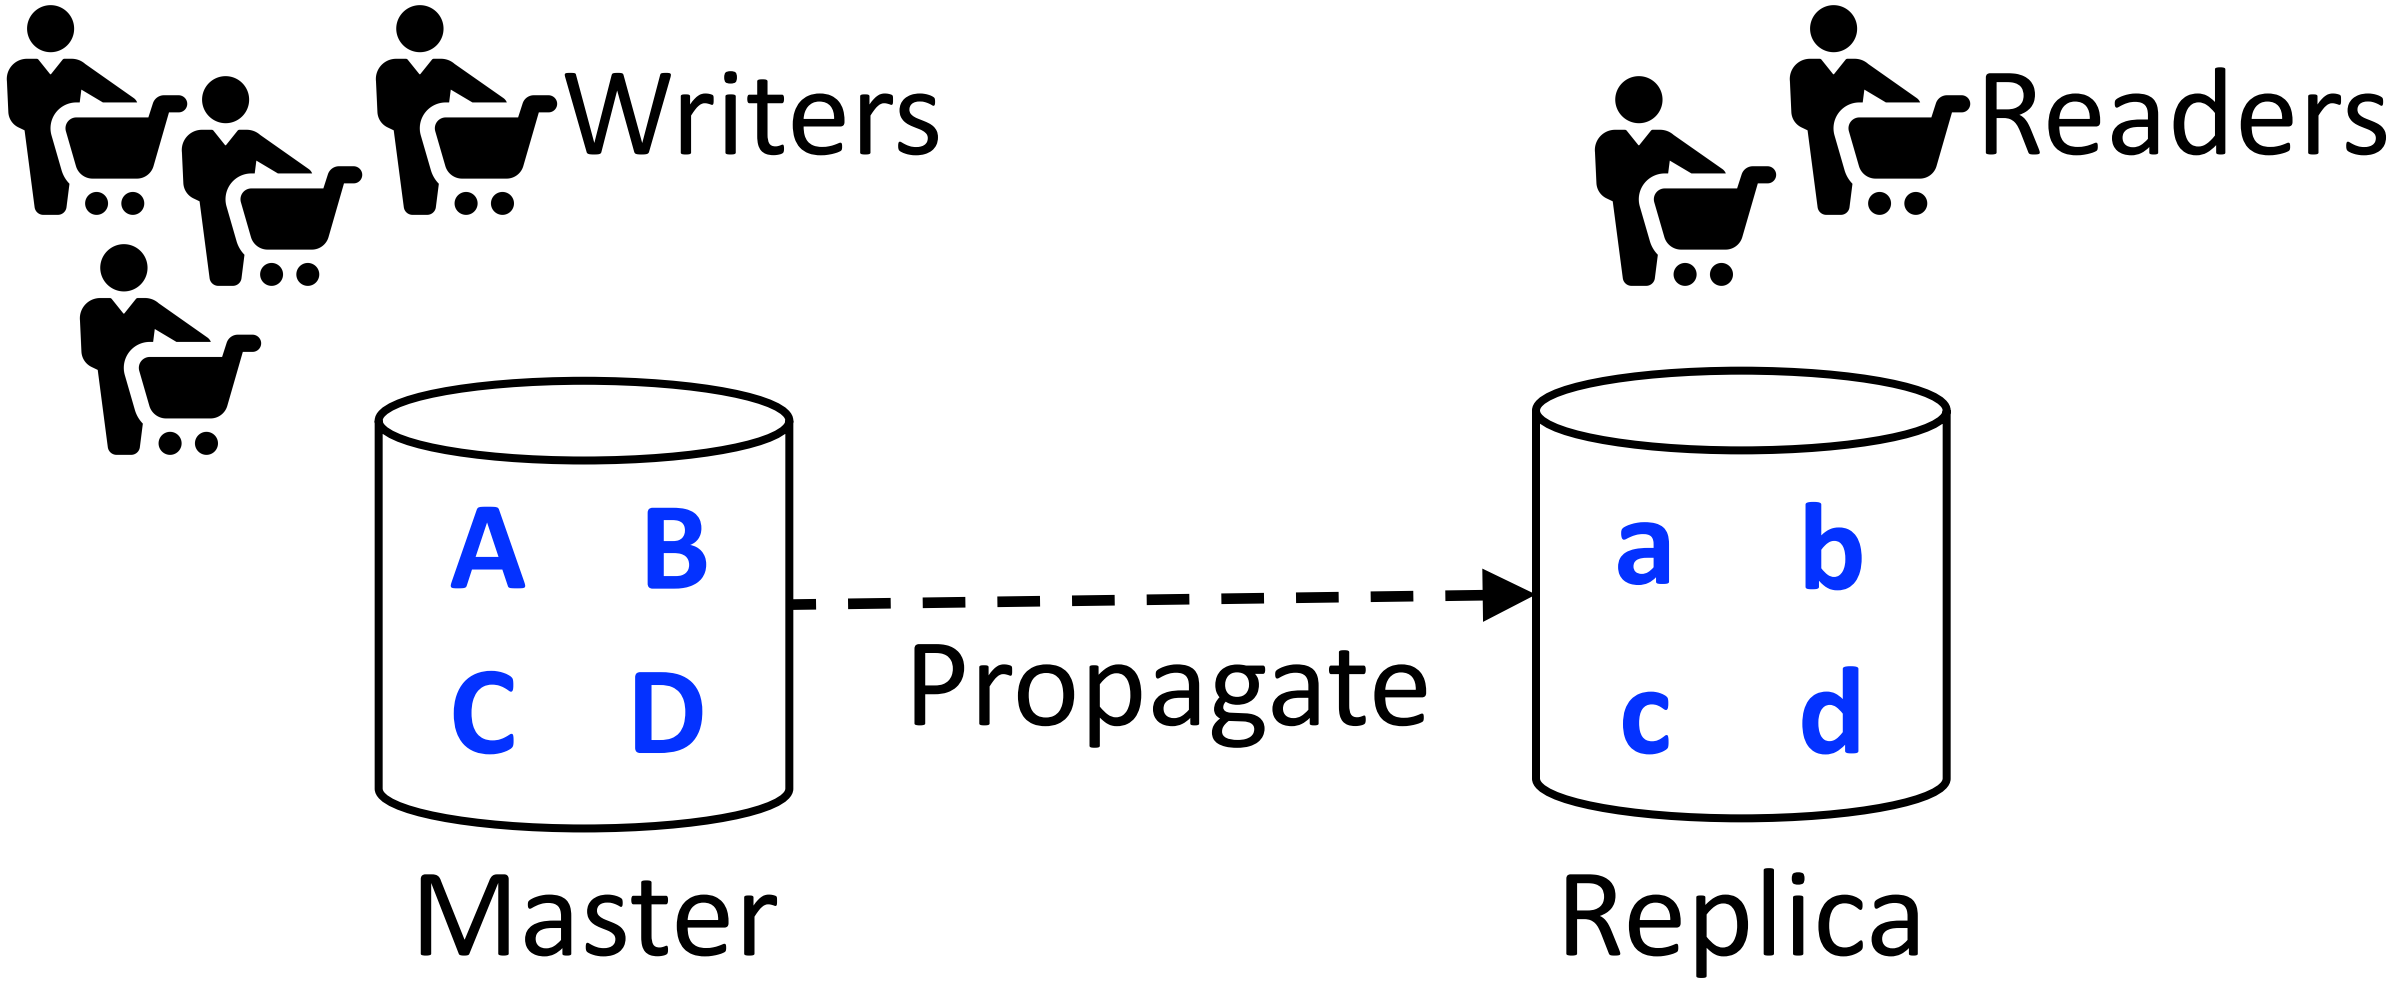

UNIVERSITY OF WATERLOO

Single-Master

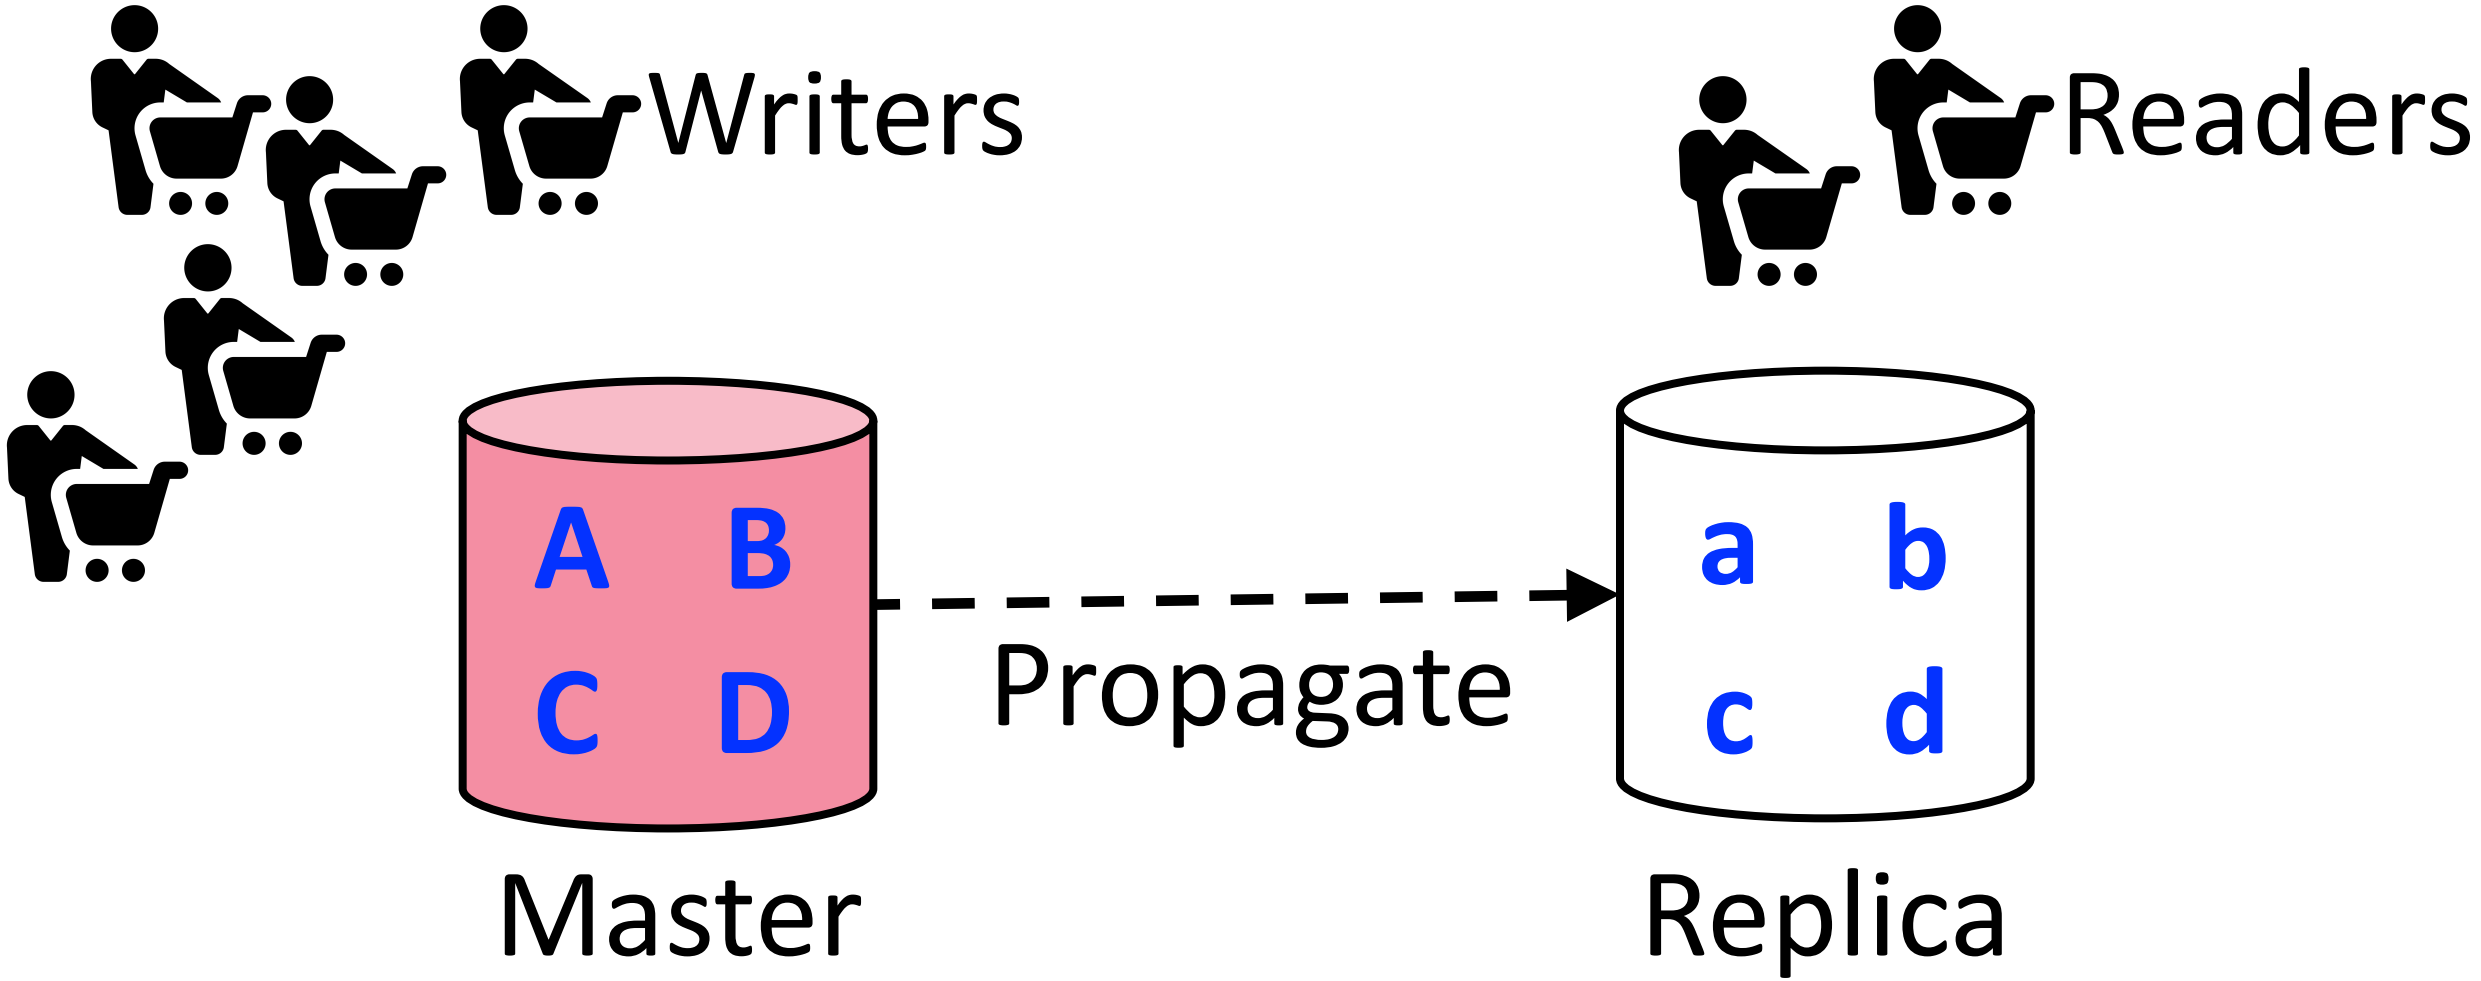

Single-Master

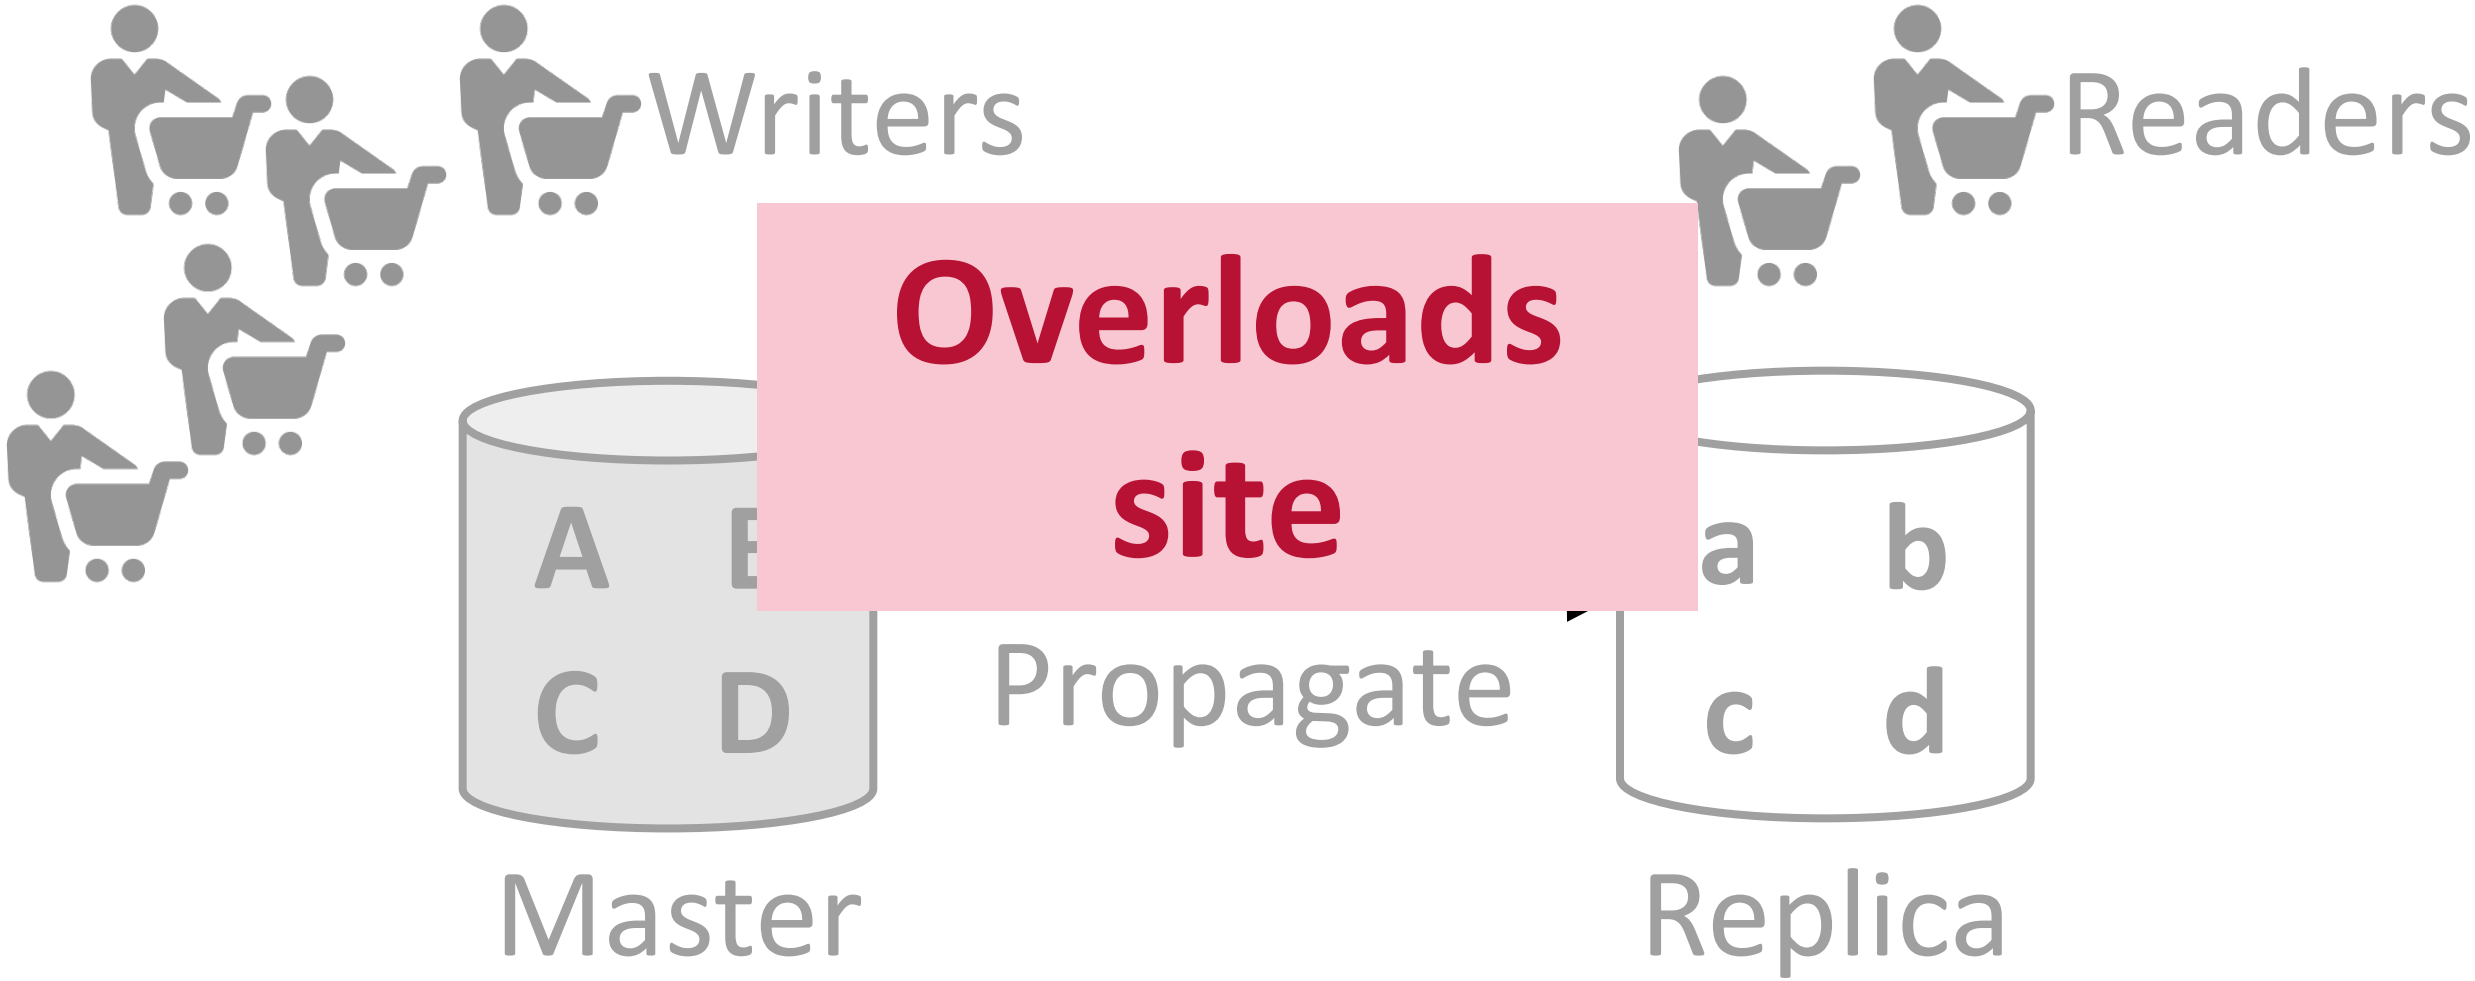

How to scale the database?

Single-master

Multi-master

Multi-Master

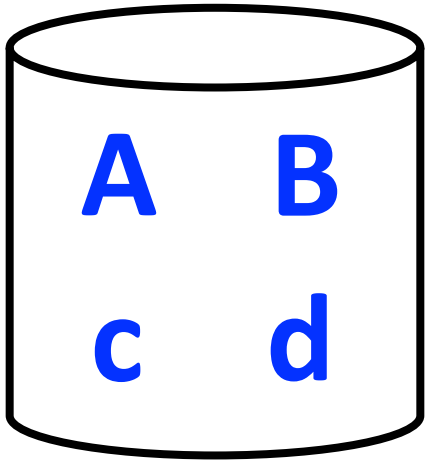

Site 1

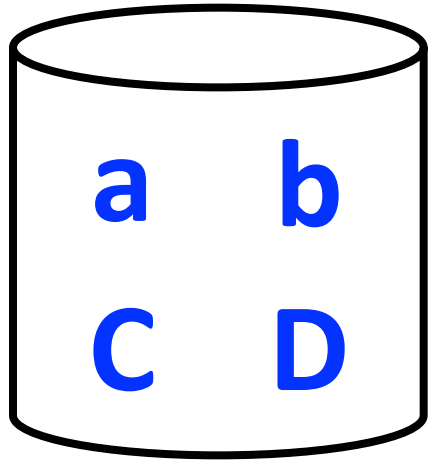

Site 2

UNIVERSITY OF WATERLOO

Multi-Master

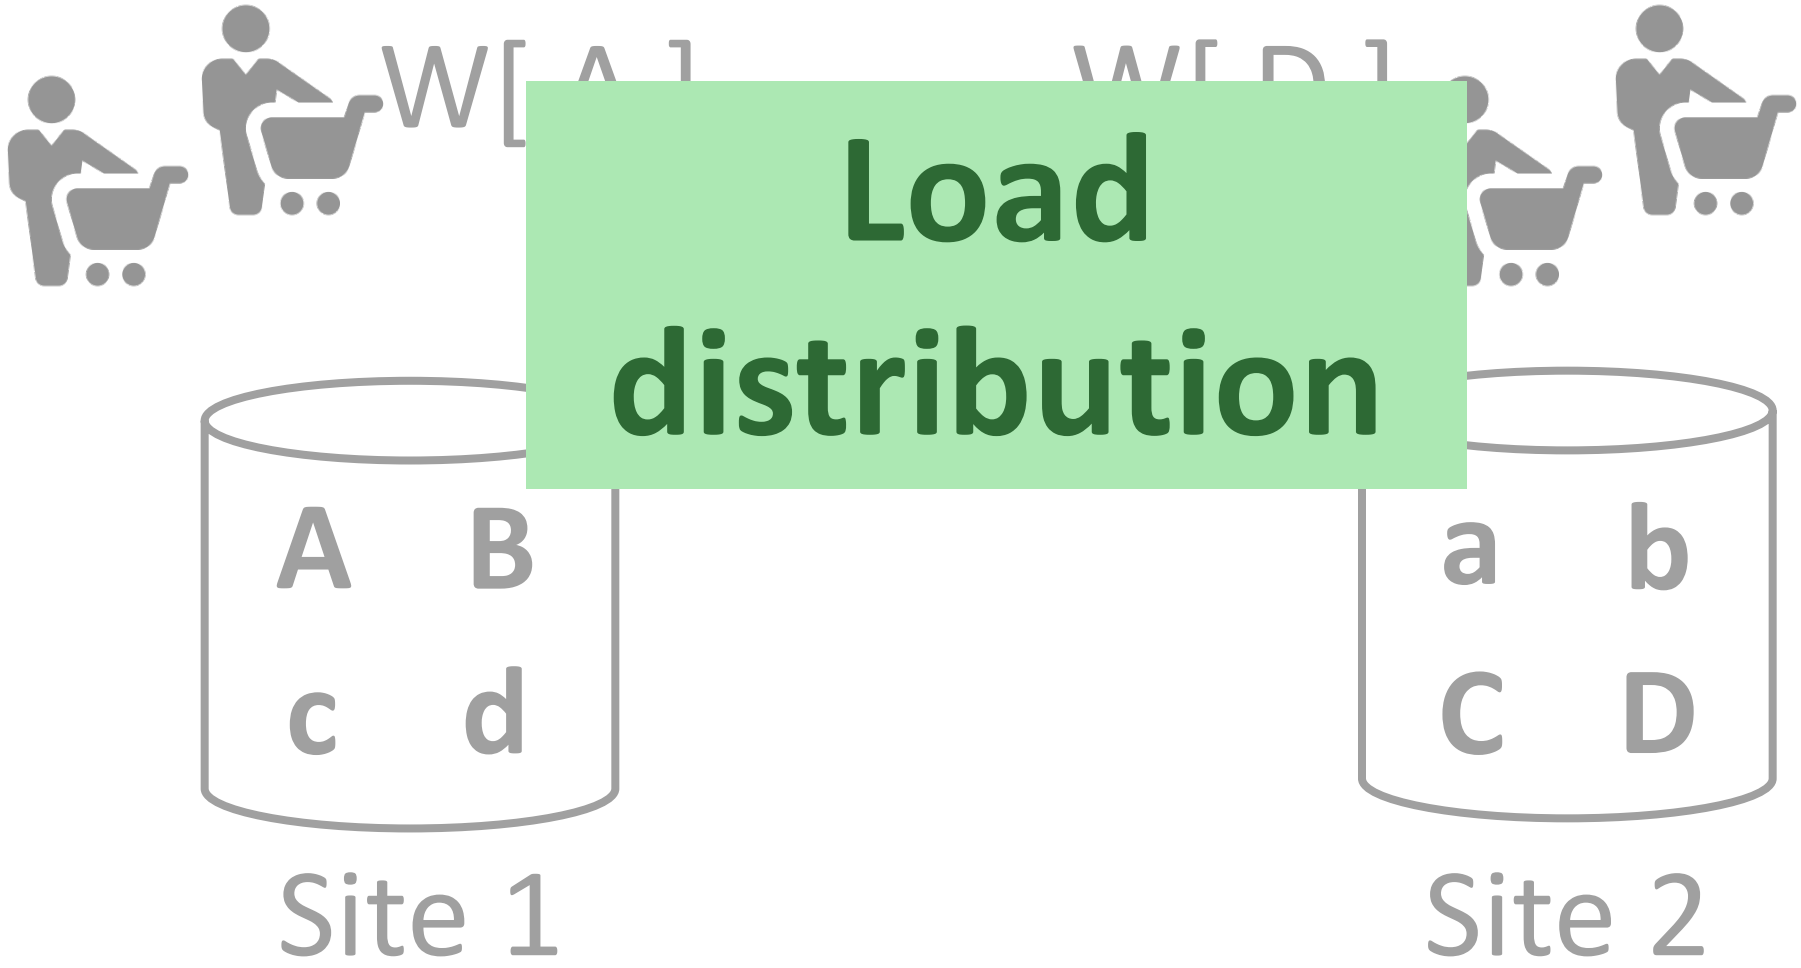

Multi-Master

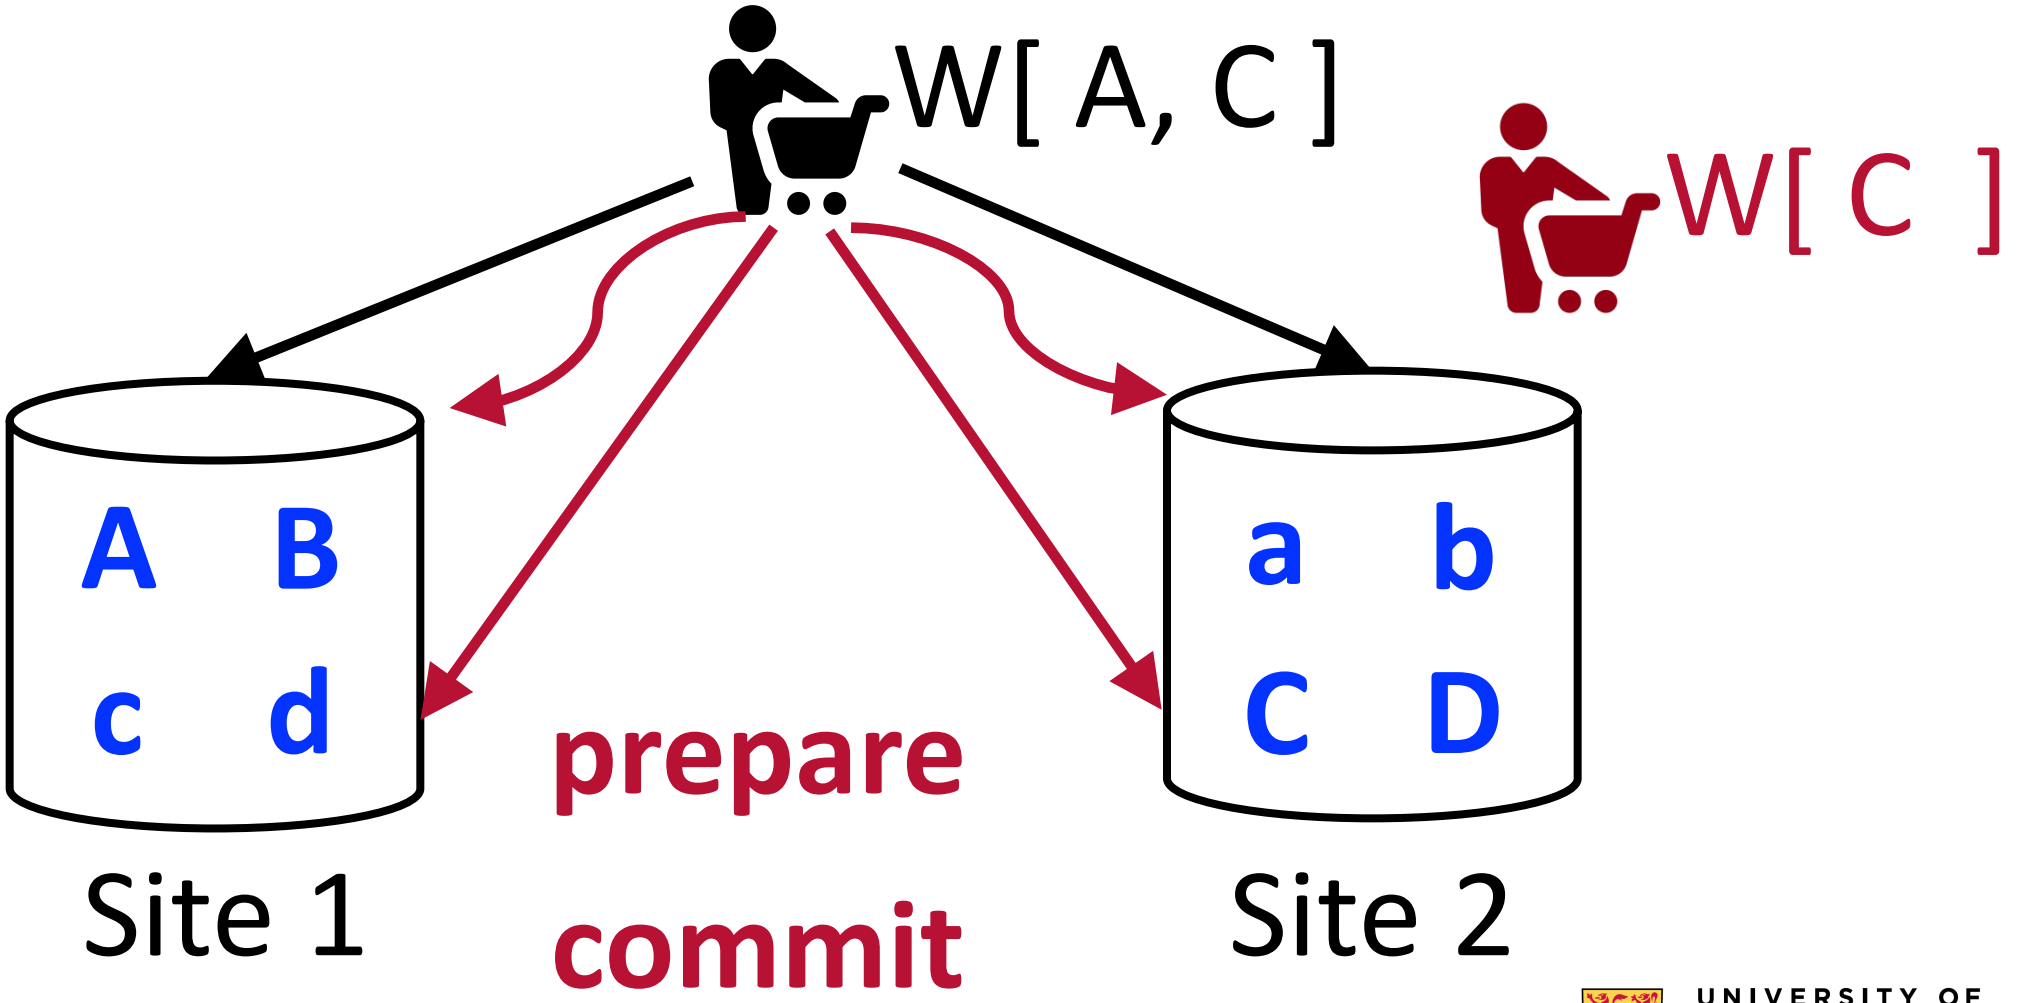

Multi-Master

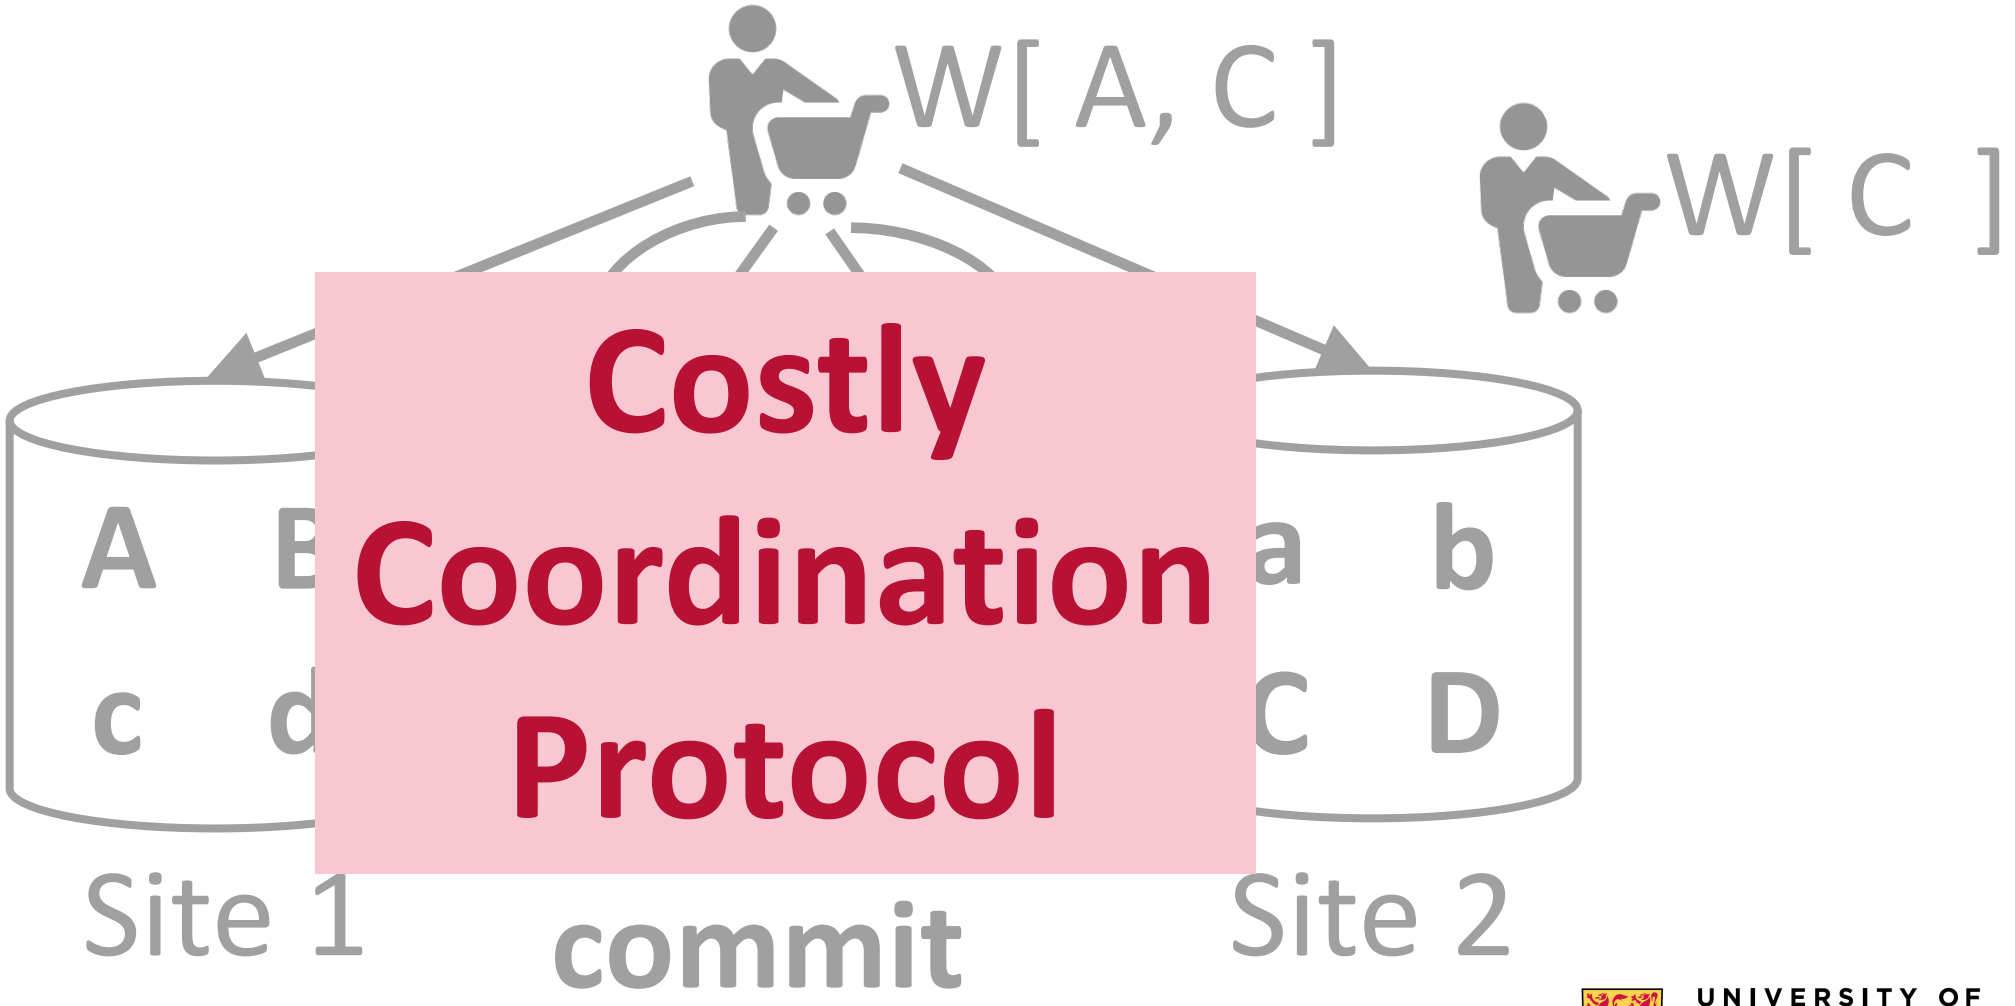

Costly Coordination Protocol

Multi-Master	Single-Master
Load distribution	Single-site transactions
Costly Coordination Protocol	Overloads site

How to provide:

Load distribution

Single-site transactions

Dynamic Mastering

Dynamic Mastering

W[A, C]

UNIVERSITY OF WATERLOO

Dynamic Mastering

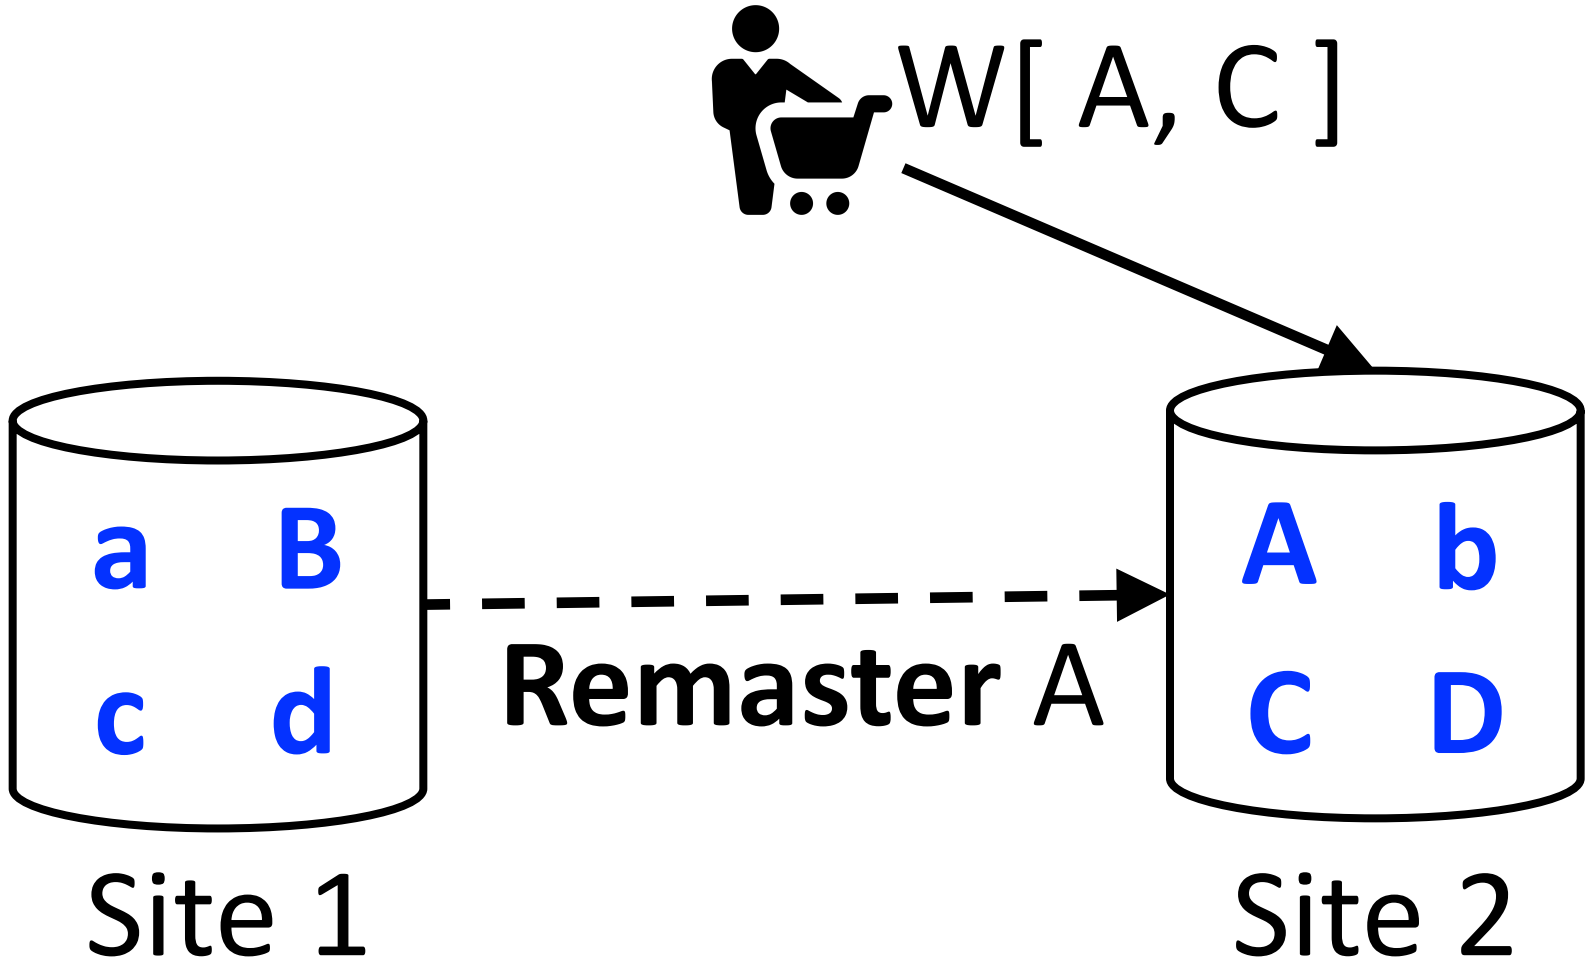

UNIVERSITY OF WATERLOO

Dynamic Mastering

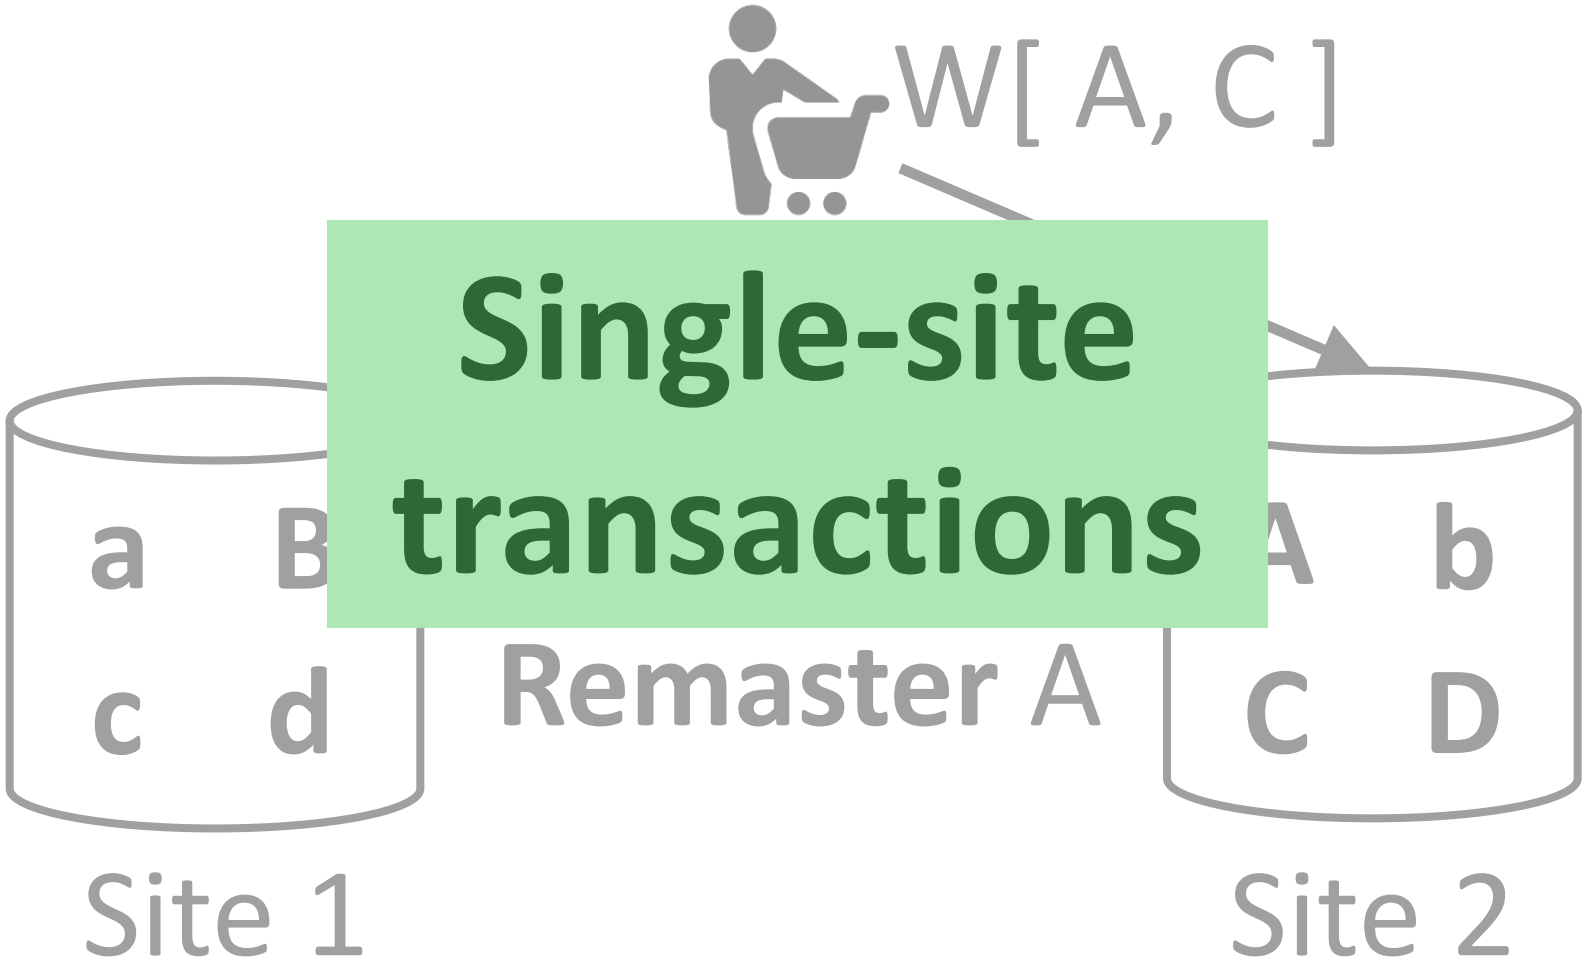

Dynamic Mastering

W[B]

W[A, C]

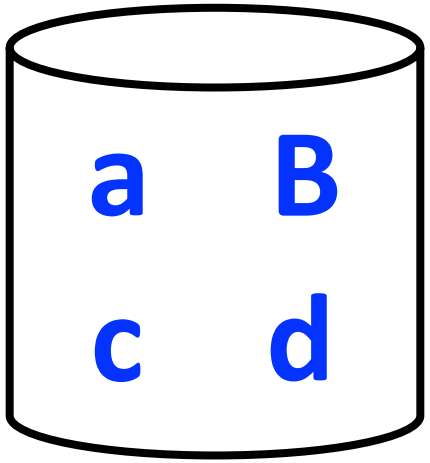

Site 1

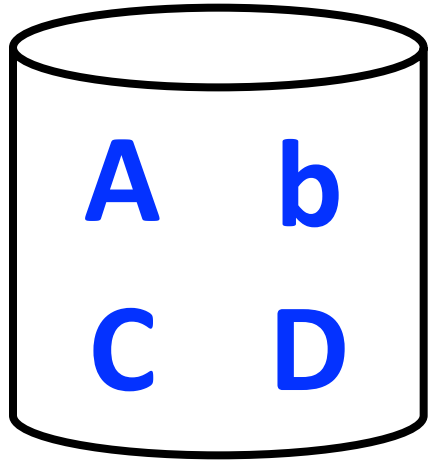

Site 2

UNIVERSITY OF WATERLOO

Dynamic Mastering

W[B]

W[A, C]

Load distribution

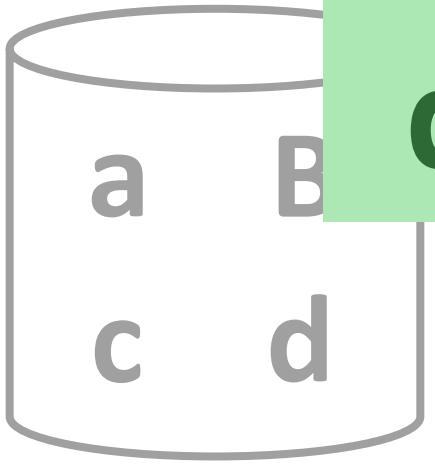

Site 1

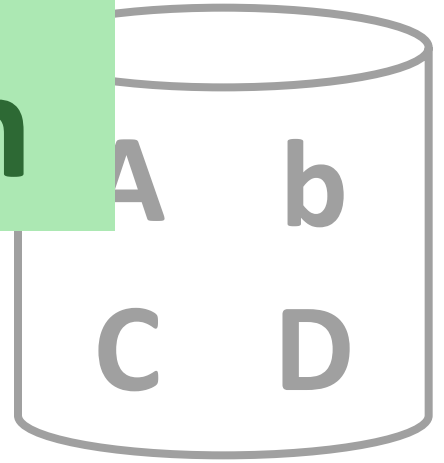

Site 2

Dynamic Mastering

Outside transaction boundaries

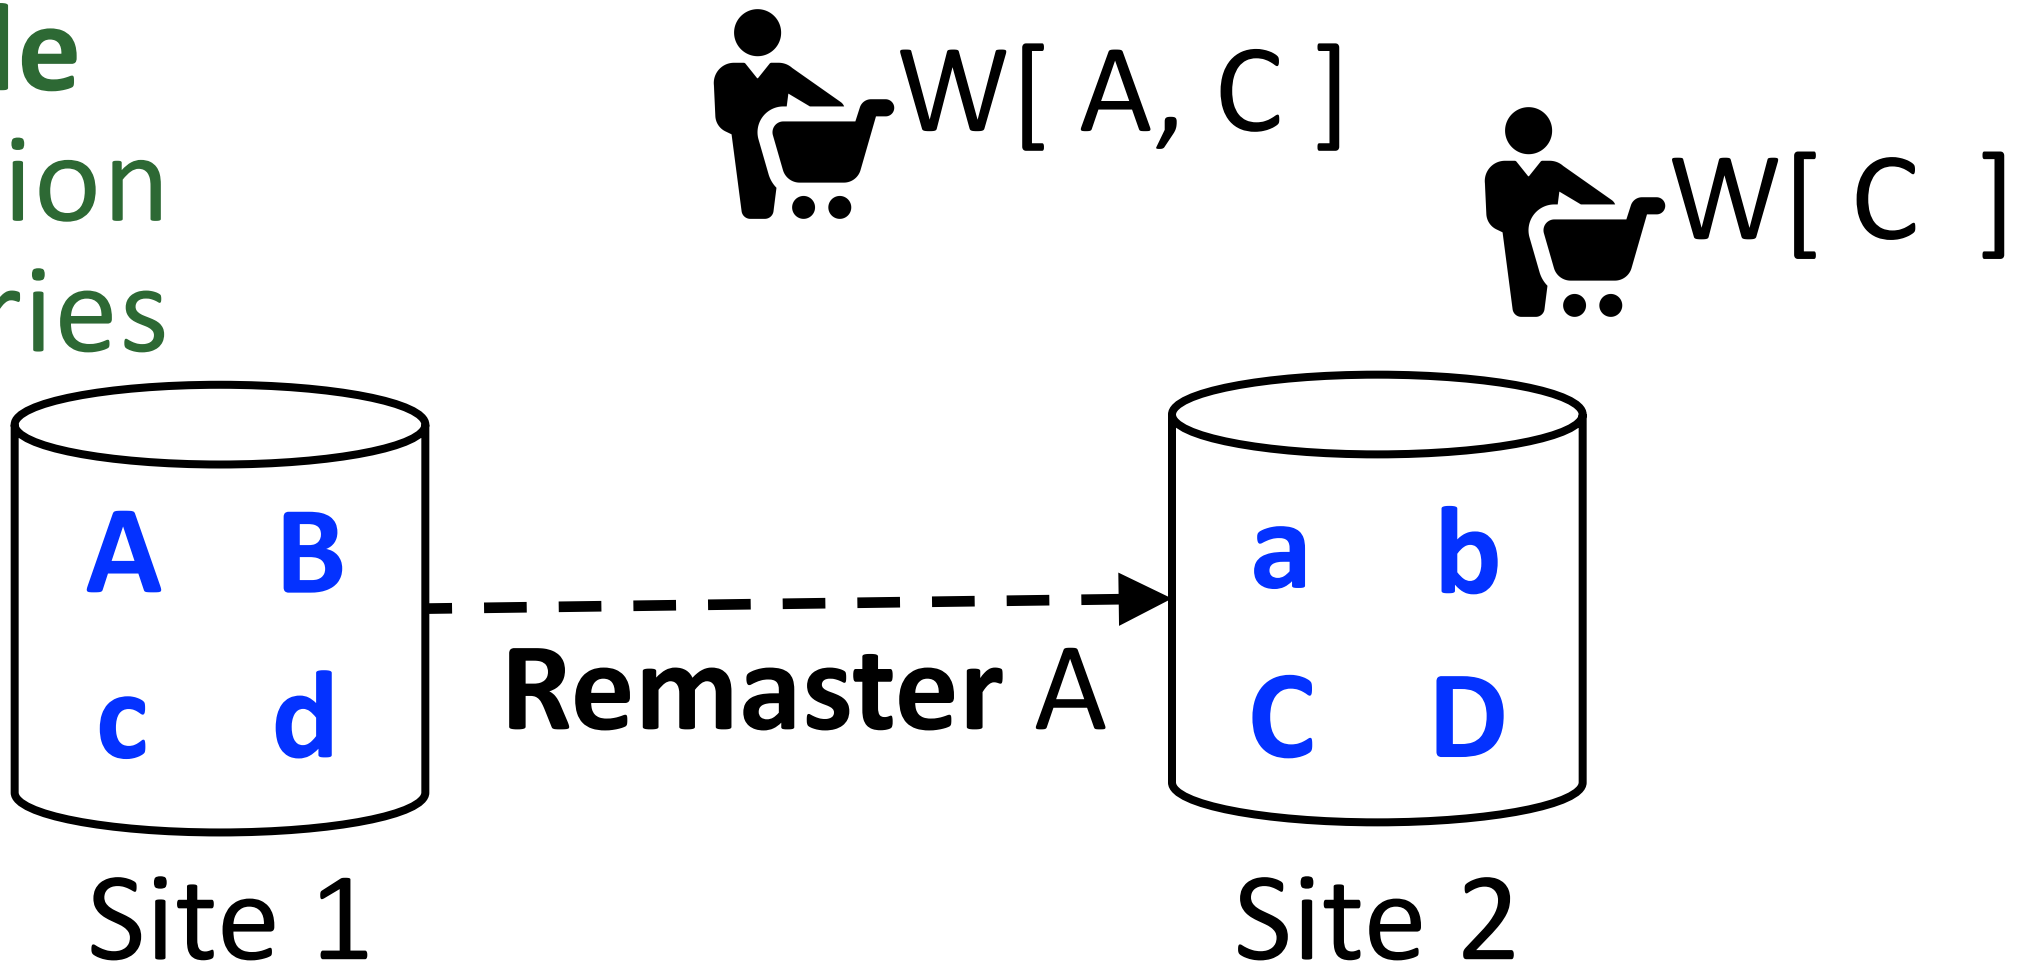

UNIVERSITY OF WATERLOO

DynaMast

Distributed and replicated database system

Employs ***adaptive*** **dynamic mastering**

Provides both **single-site transactions** and **load balance**

Dynamic Mastering Challenges

How to perform remastering efficiently?

How to decide where to master data?

Dynamic Mastering Challenges

How to perform remastering efficiently?

How to decide where to master data?

UNIVERSITY OF WATERLOO

Dynamic Mastering

W[A, C]

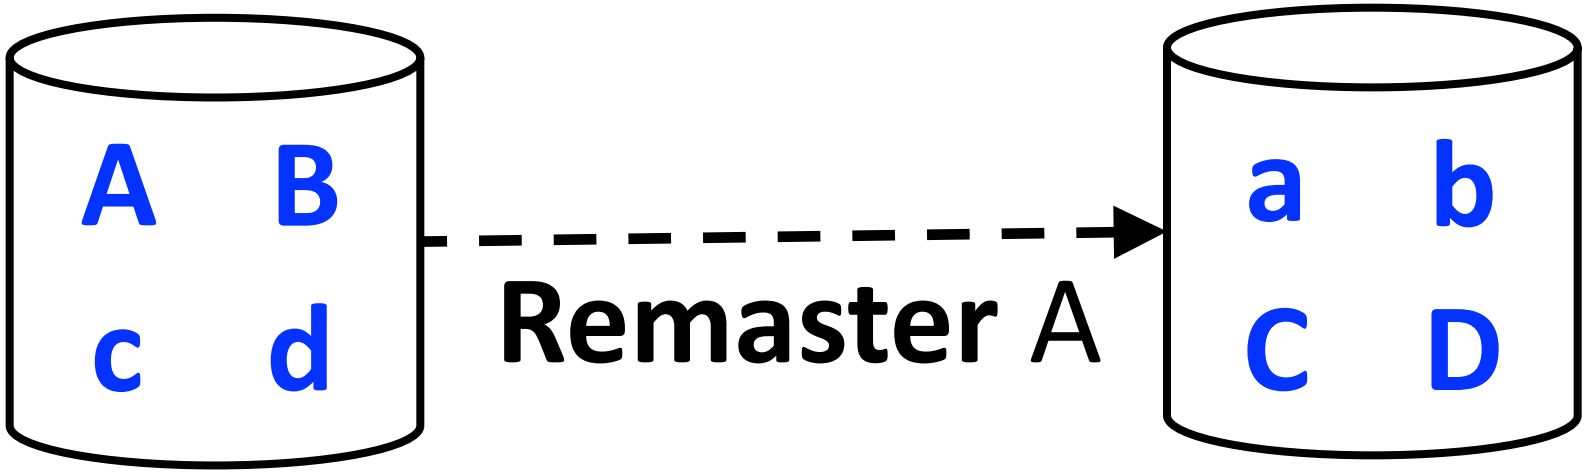

UNIVERSITY OF WATERLOO

Dynamic Mastering

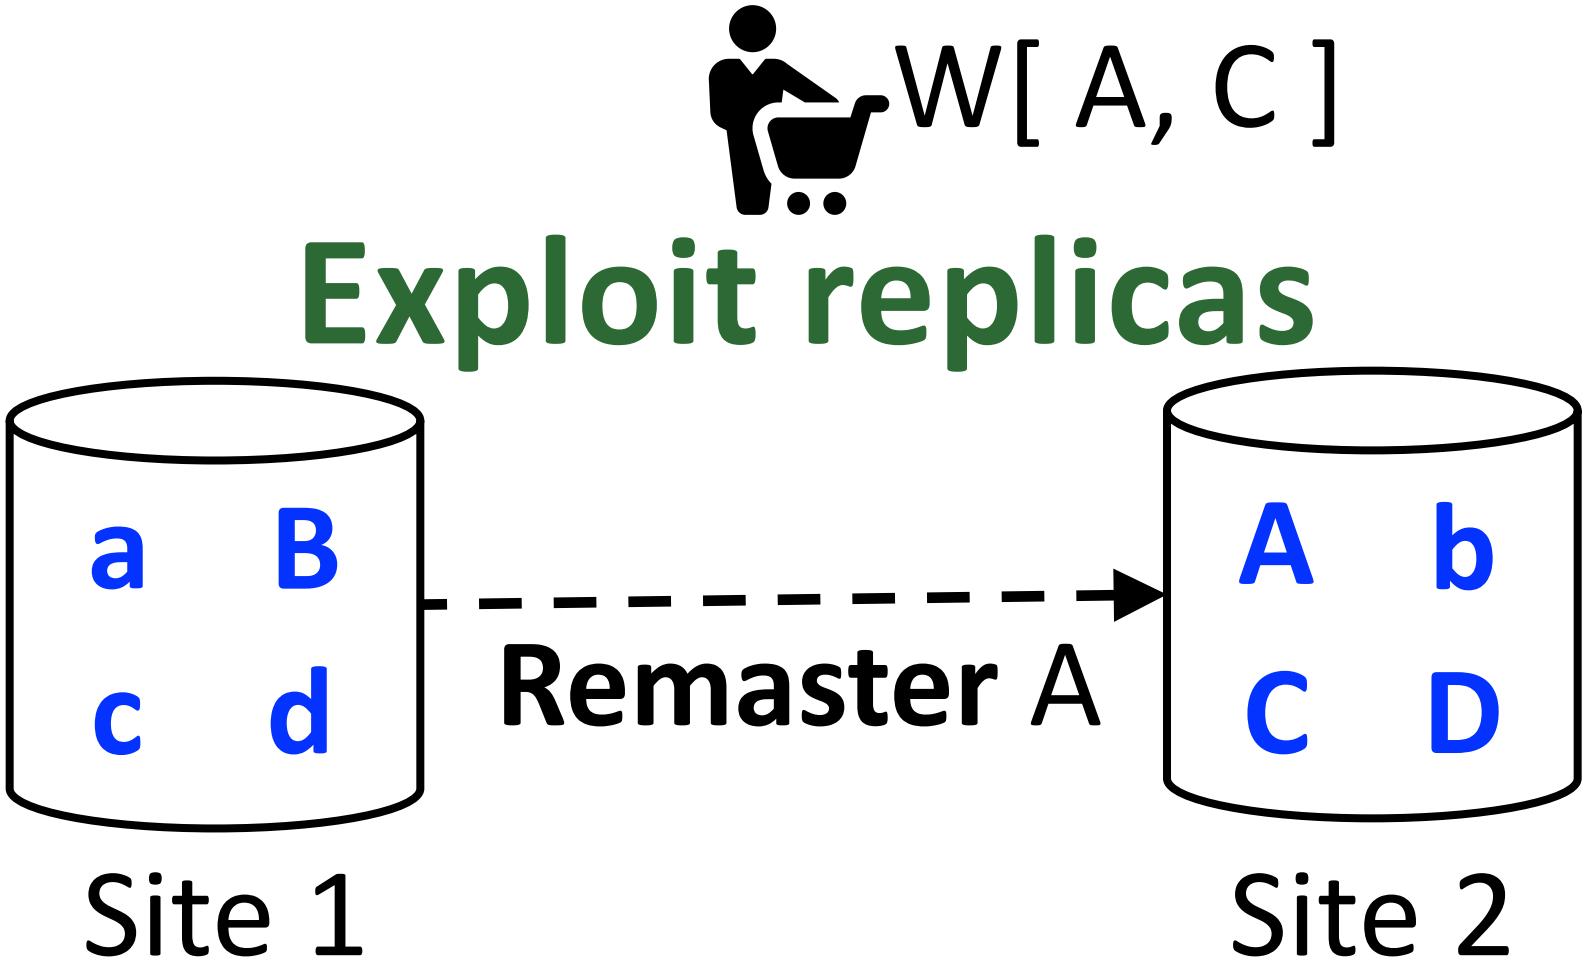

Exploiting Replicas

W[A]

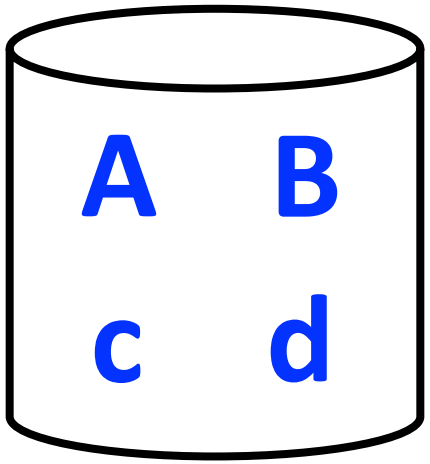

Site 1

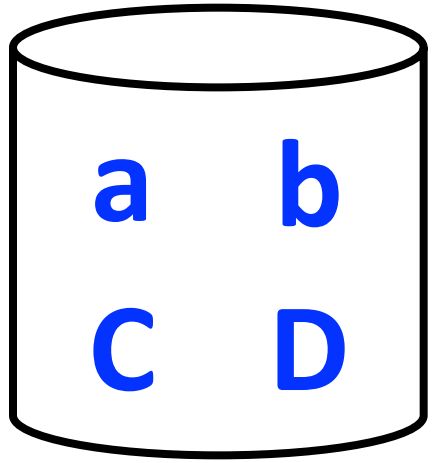

Site 2

Exploiting Replicas

W[A]

W[A, C]

Exploiting Replicas

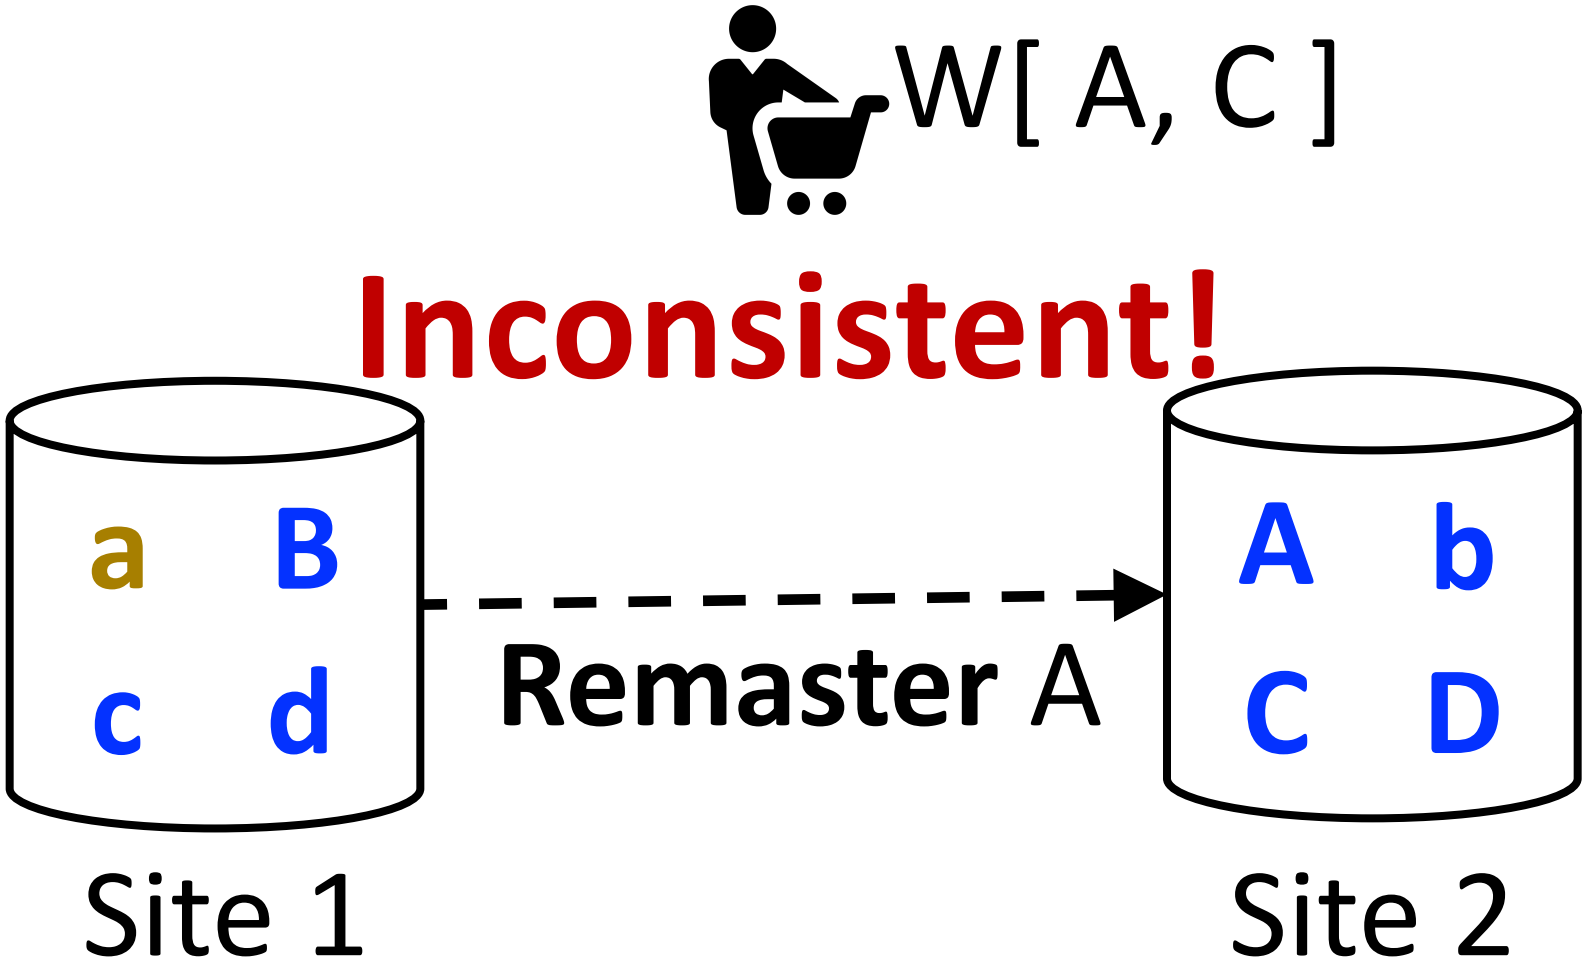

Ensuring Consistency

Old master must **not allow new updates**

release mastership

New master must have **all previous updates**

grant mastership

Ensuring Consistency

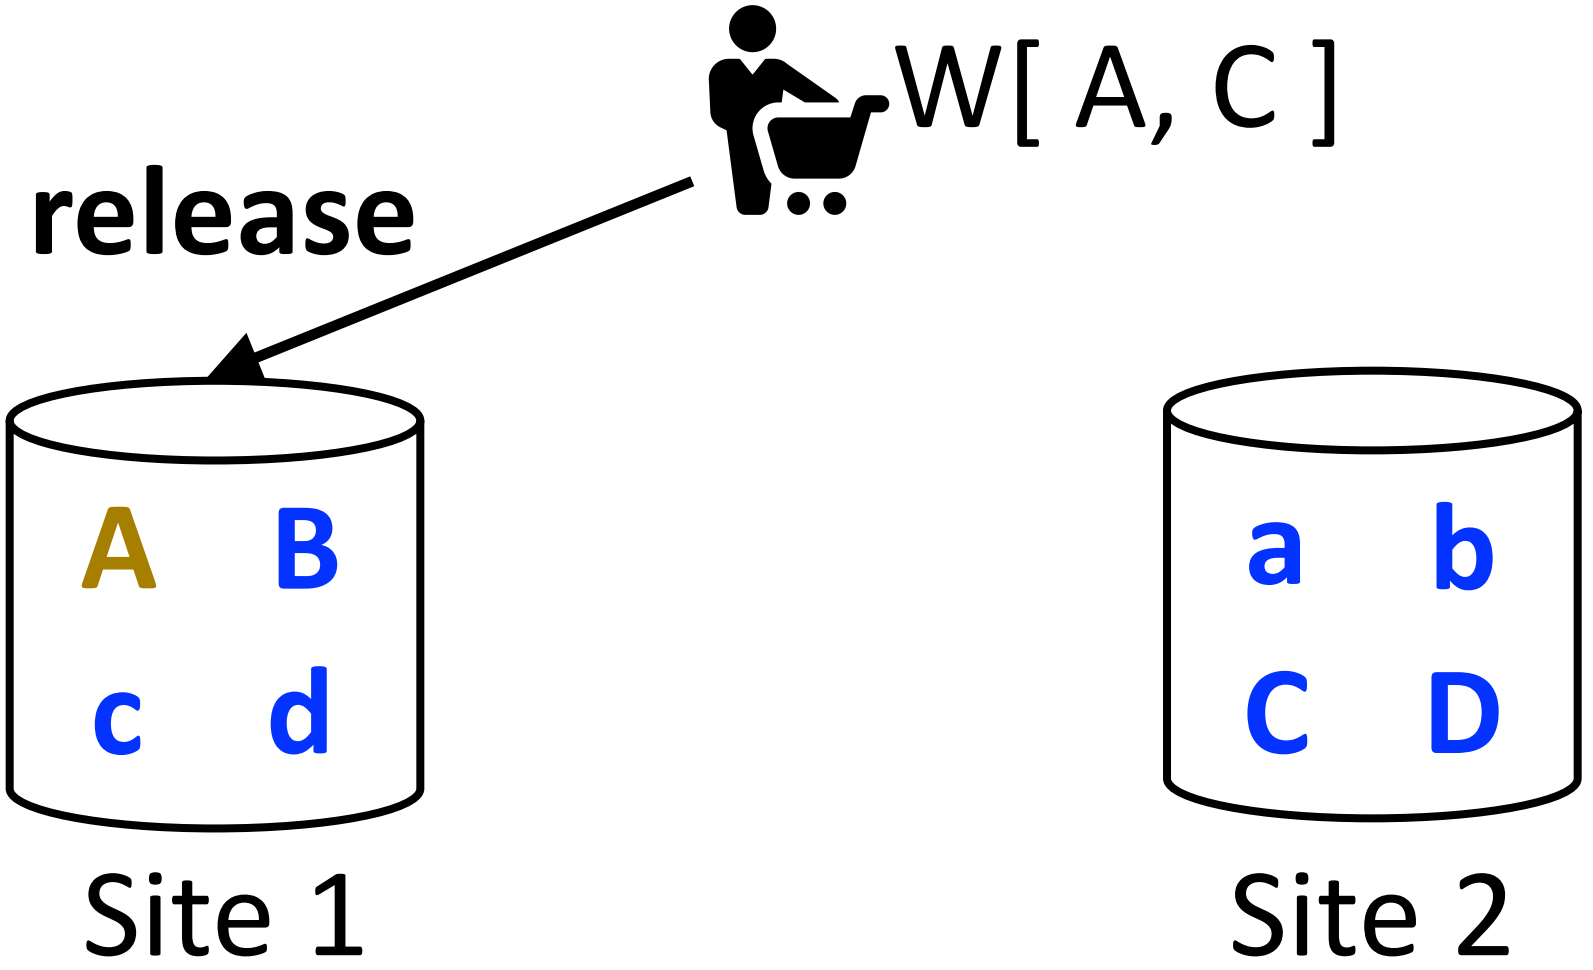

Ensuring Consistency

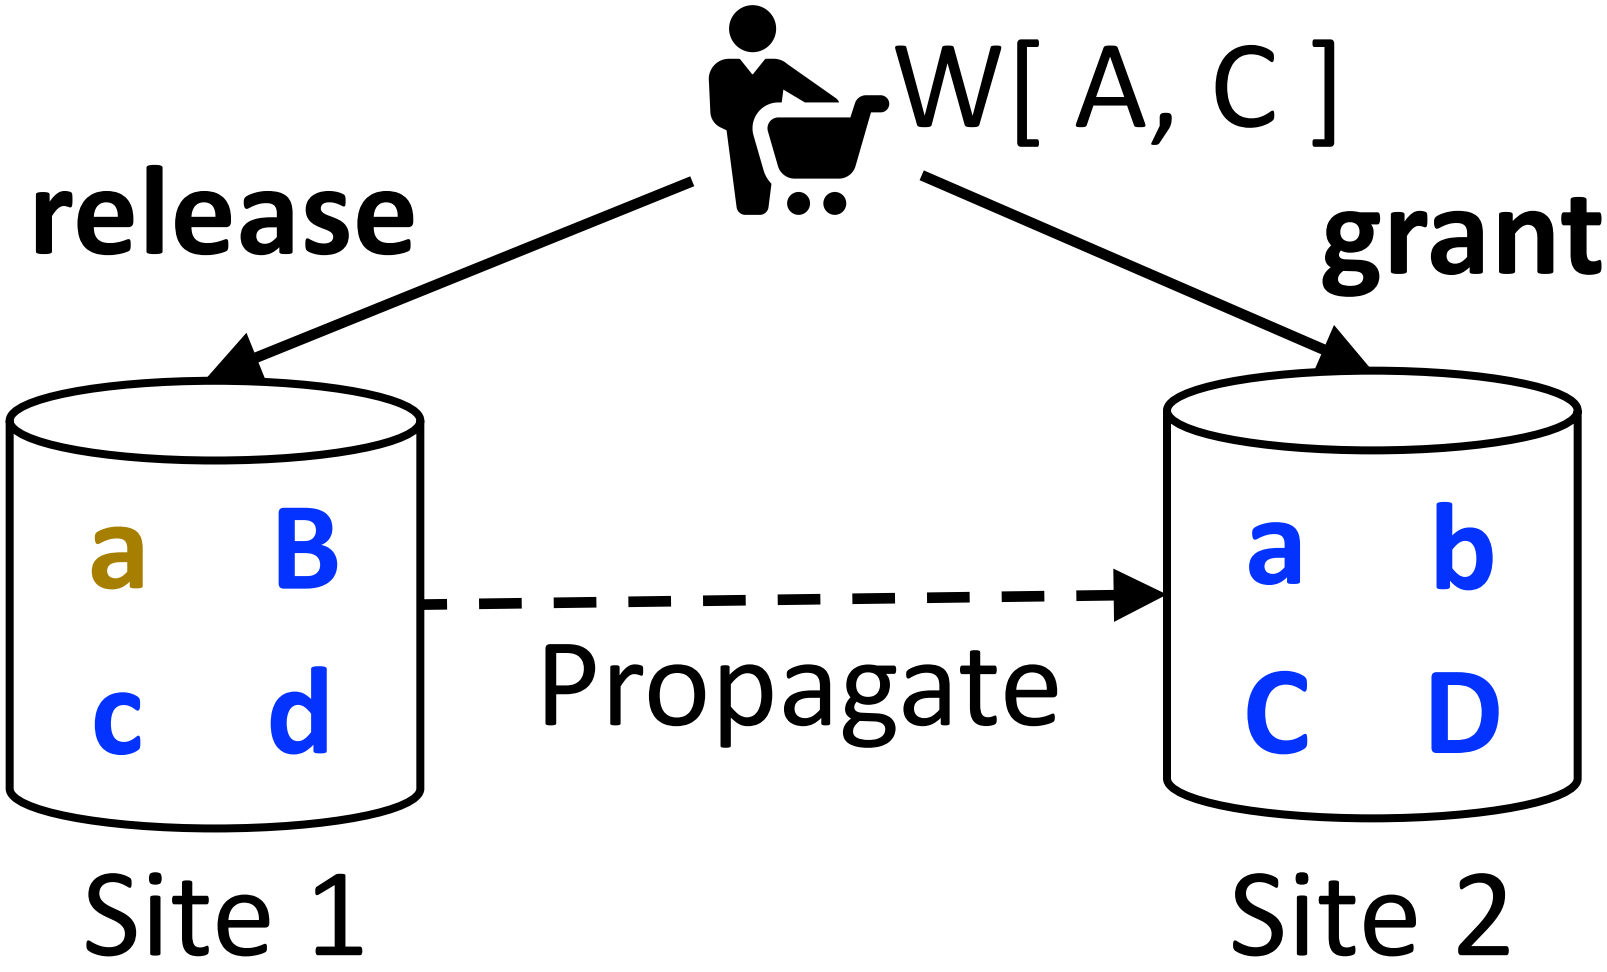

Ensuring Consistency

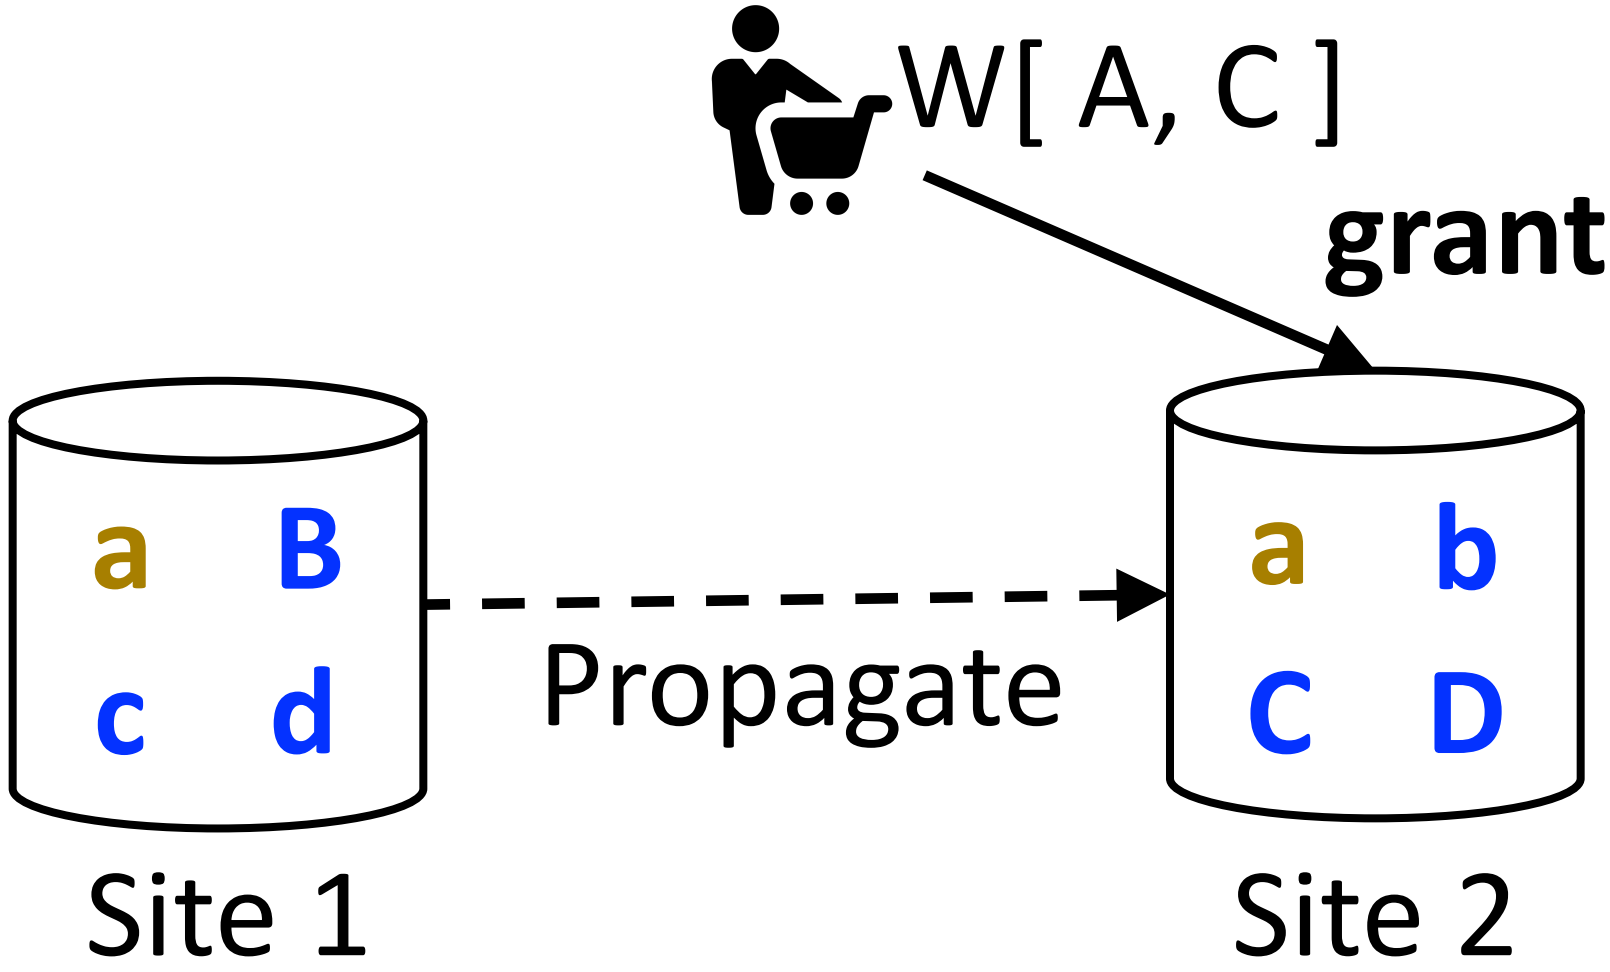

Ensuring Consistency

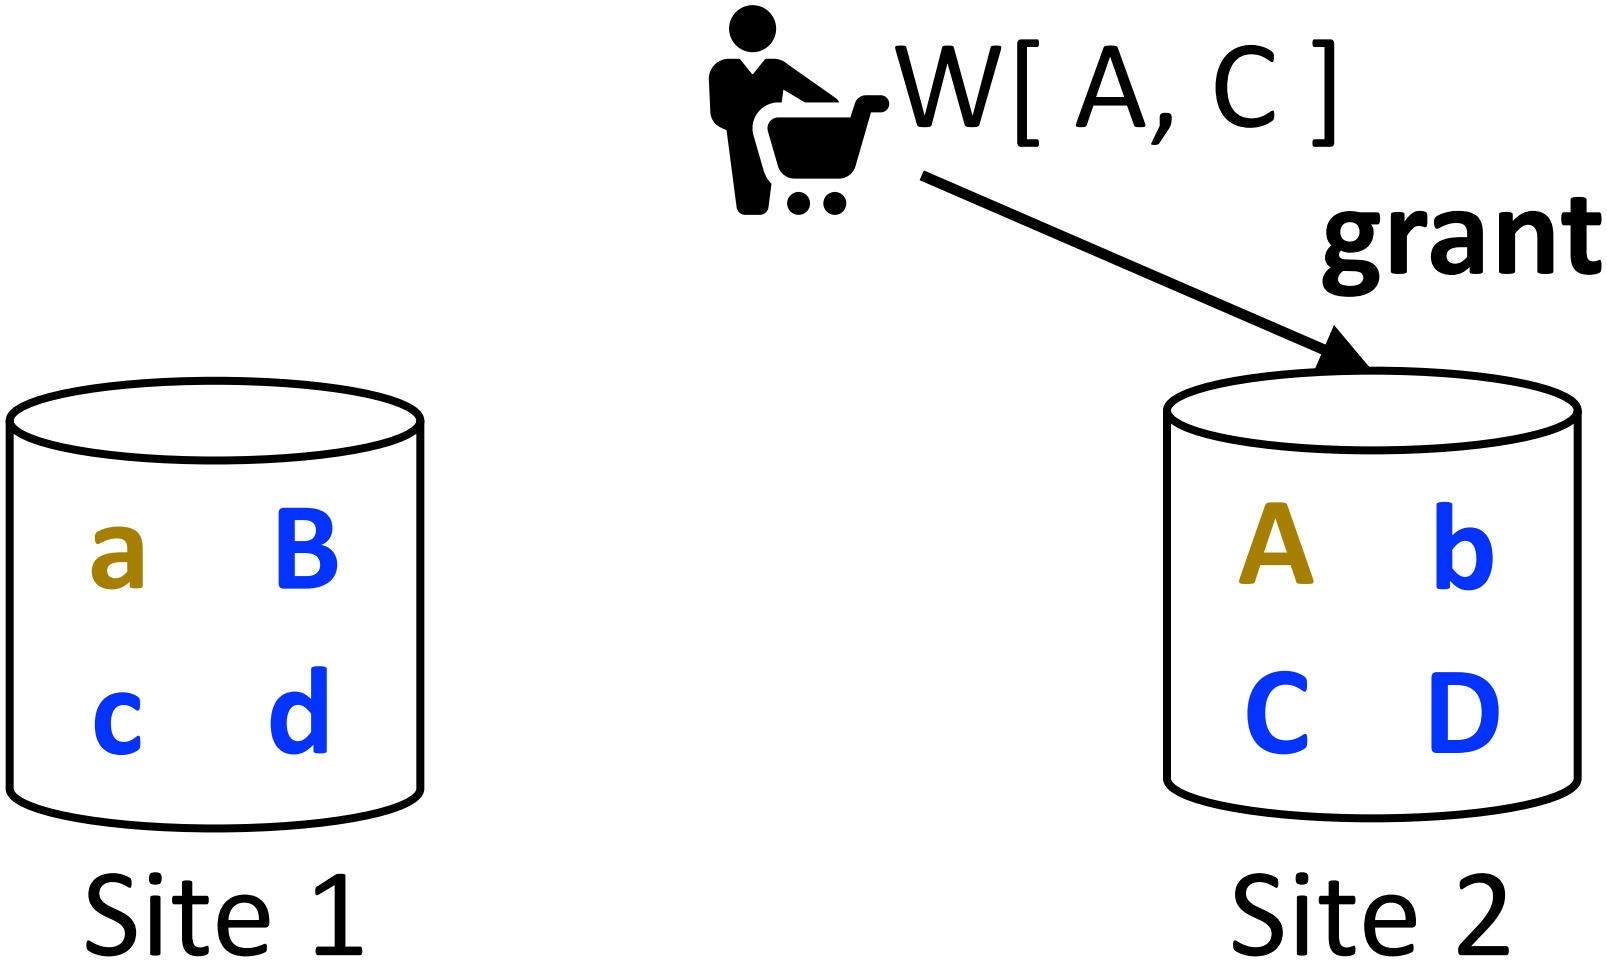

Ensuring Consistency

New master must have **all previous updates**

New master was **a lazy replica of old master**

Little time spent waiting for updates

tiny.cc/dynamast

UNIVERSITY OF WATERLOO

Dynamic Mastering Challenges

How to perform remastering efficiently?

How to decide where to master data?

Where to master?

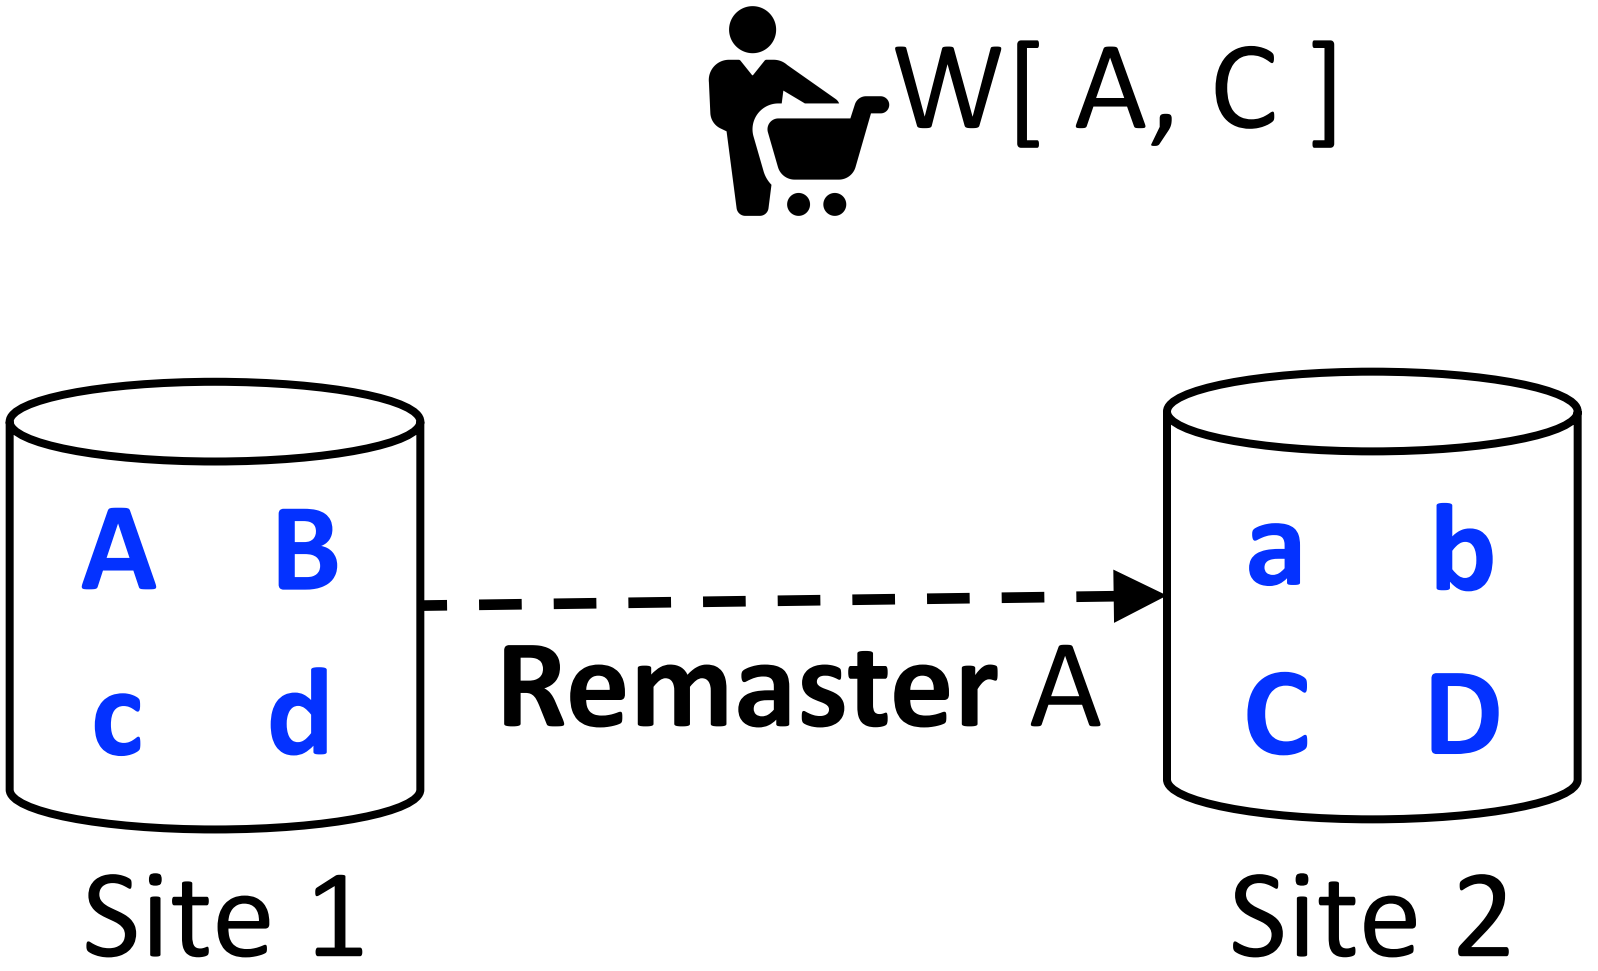

Where to master?

W[A, C]

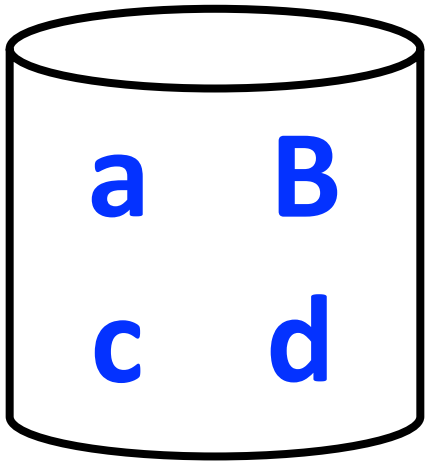

Site 1

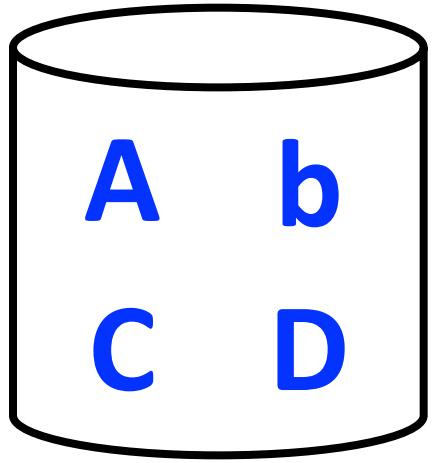

Site 2

Where to master?

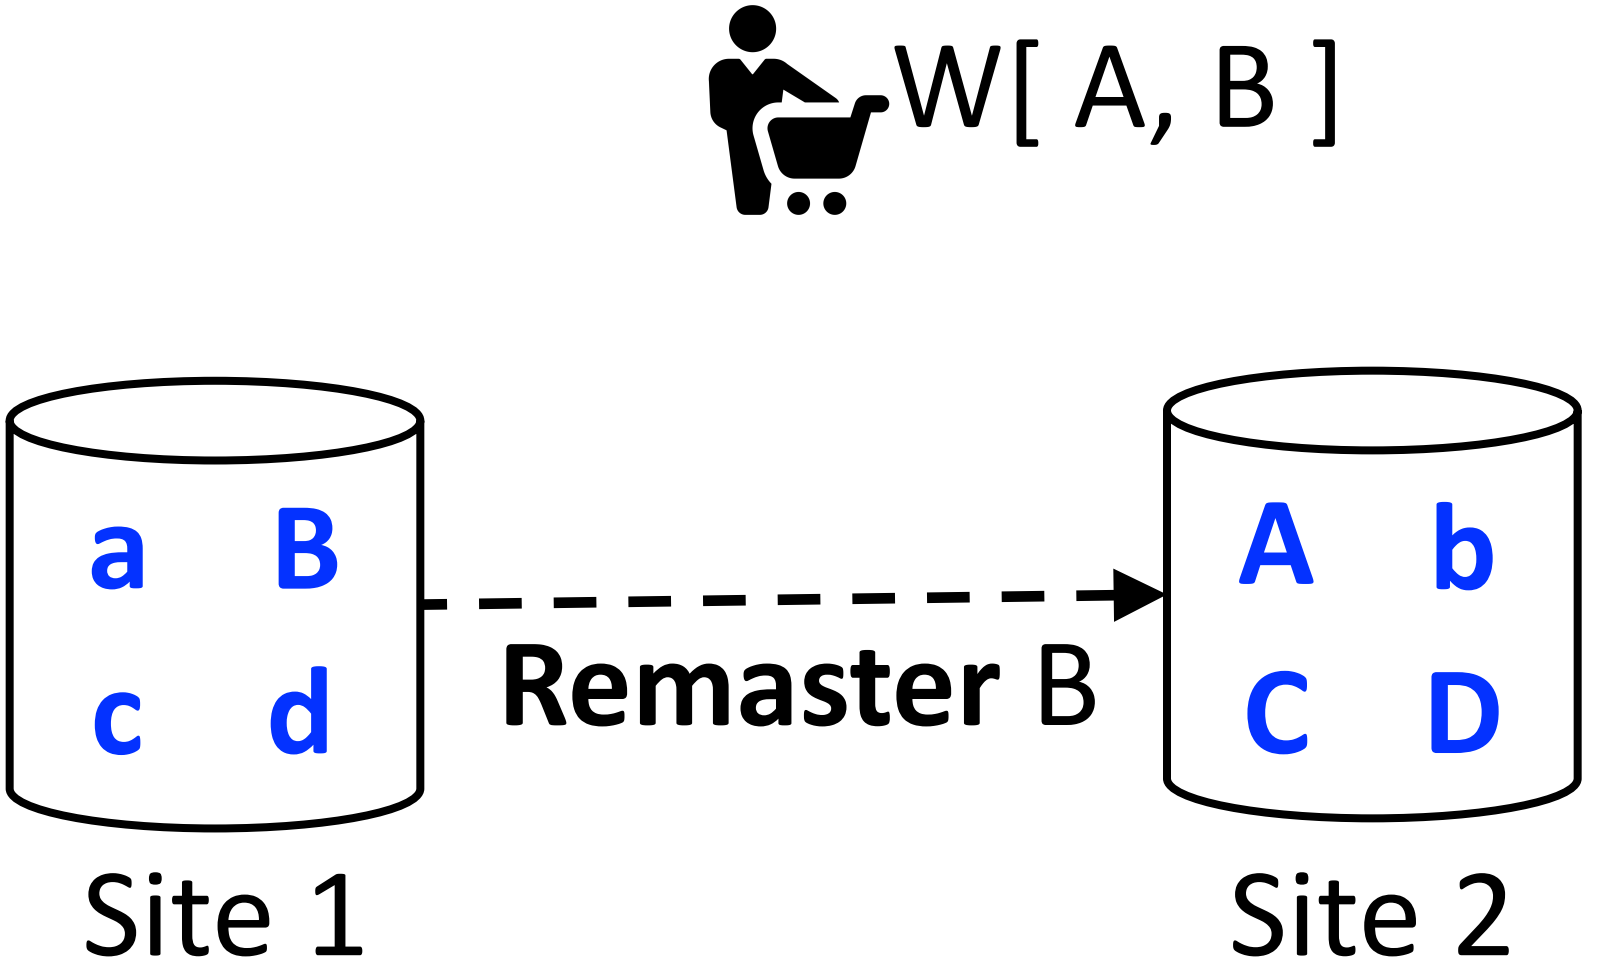

Where to master?

W[A, B]

Single-master!

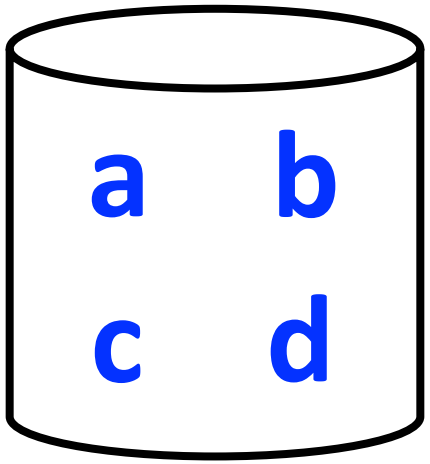

Site 1

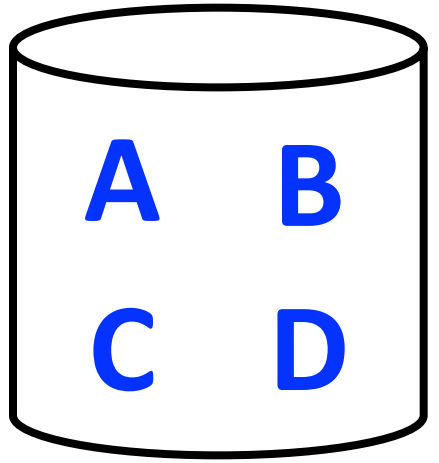

Site 2

UNIVERSITY OF WATERLOO

Where to master?

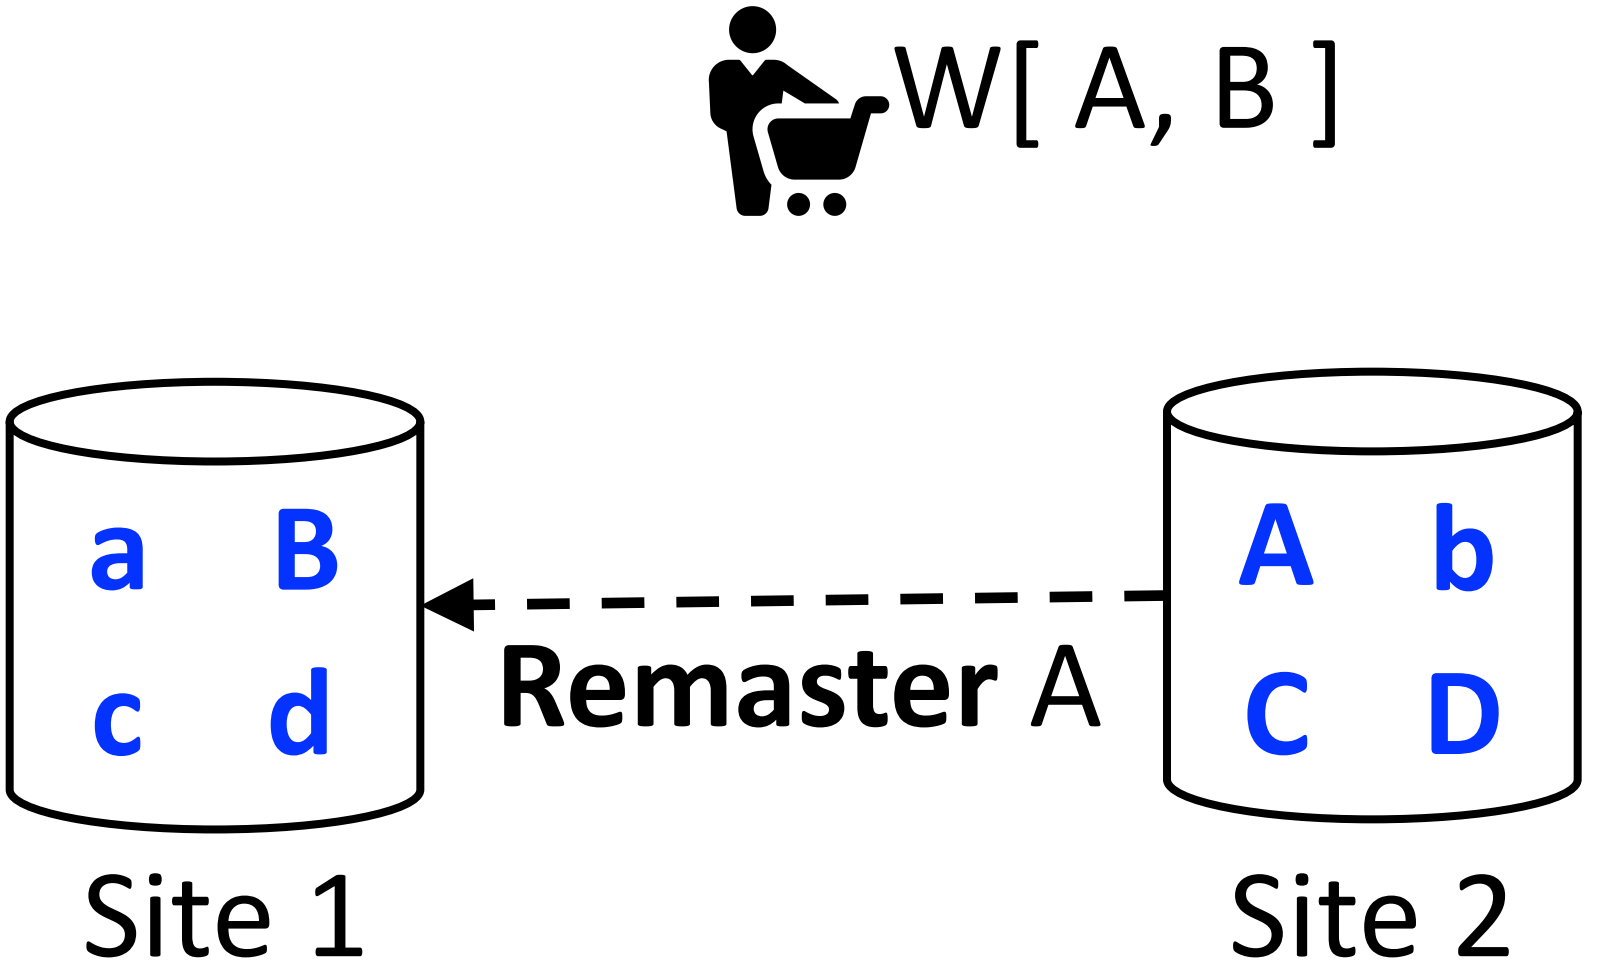

Where to master?

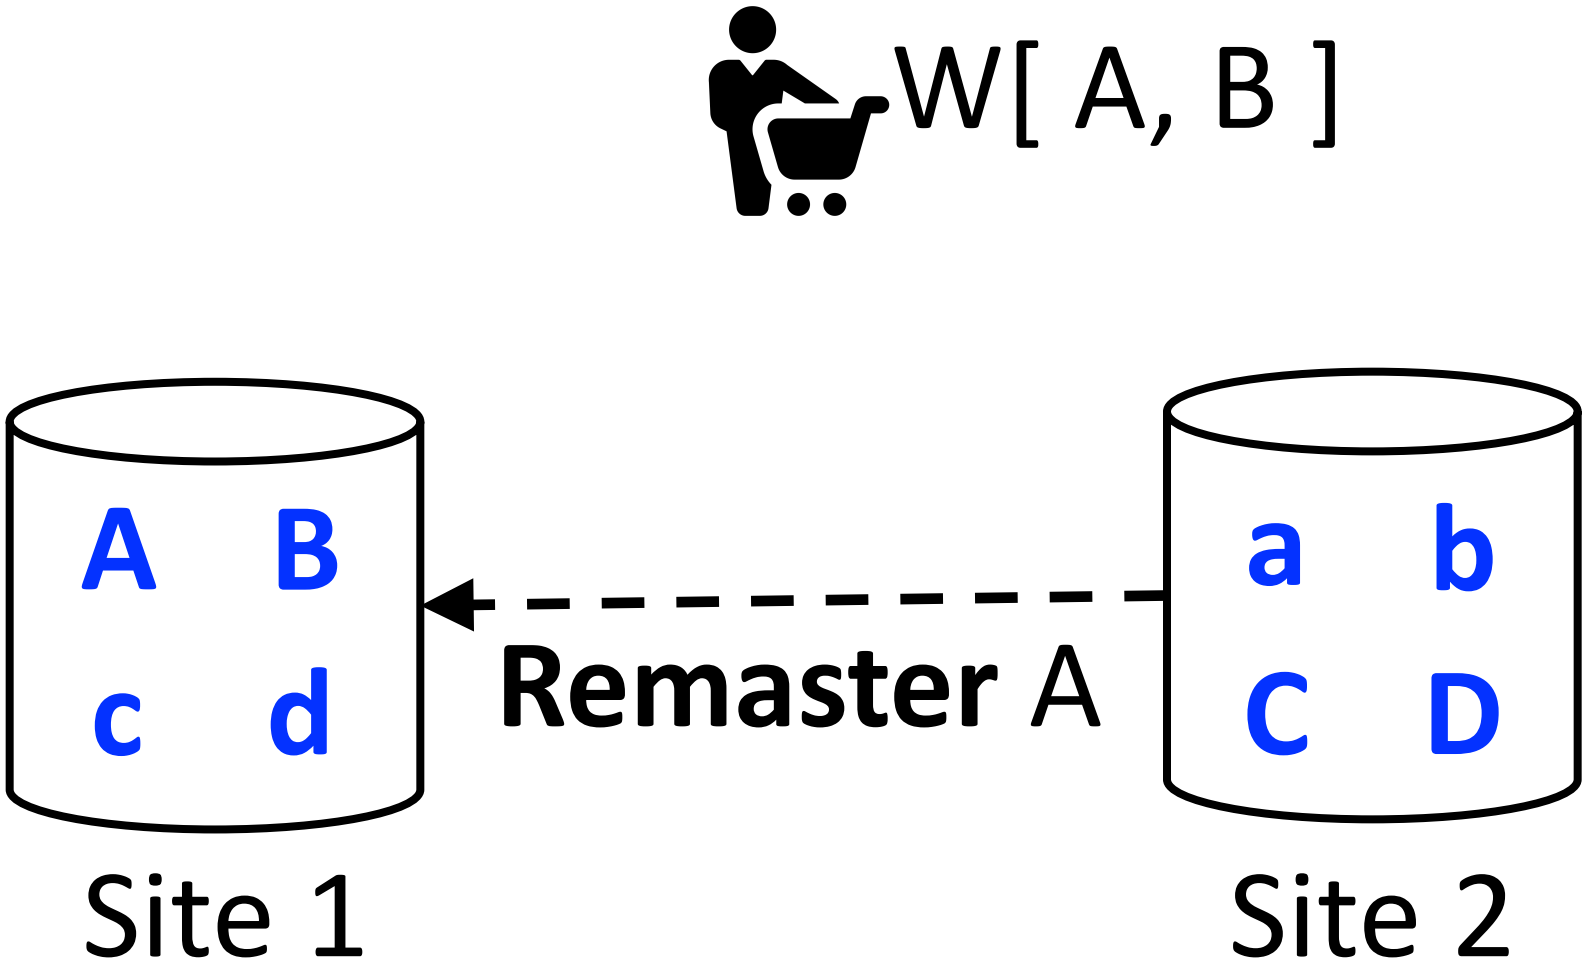

Where to master?

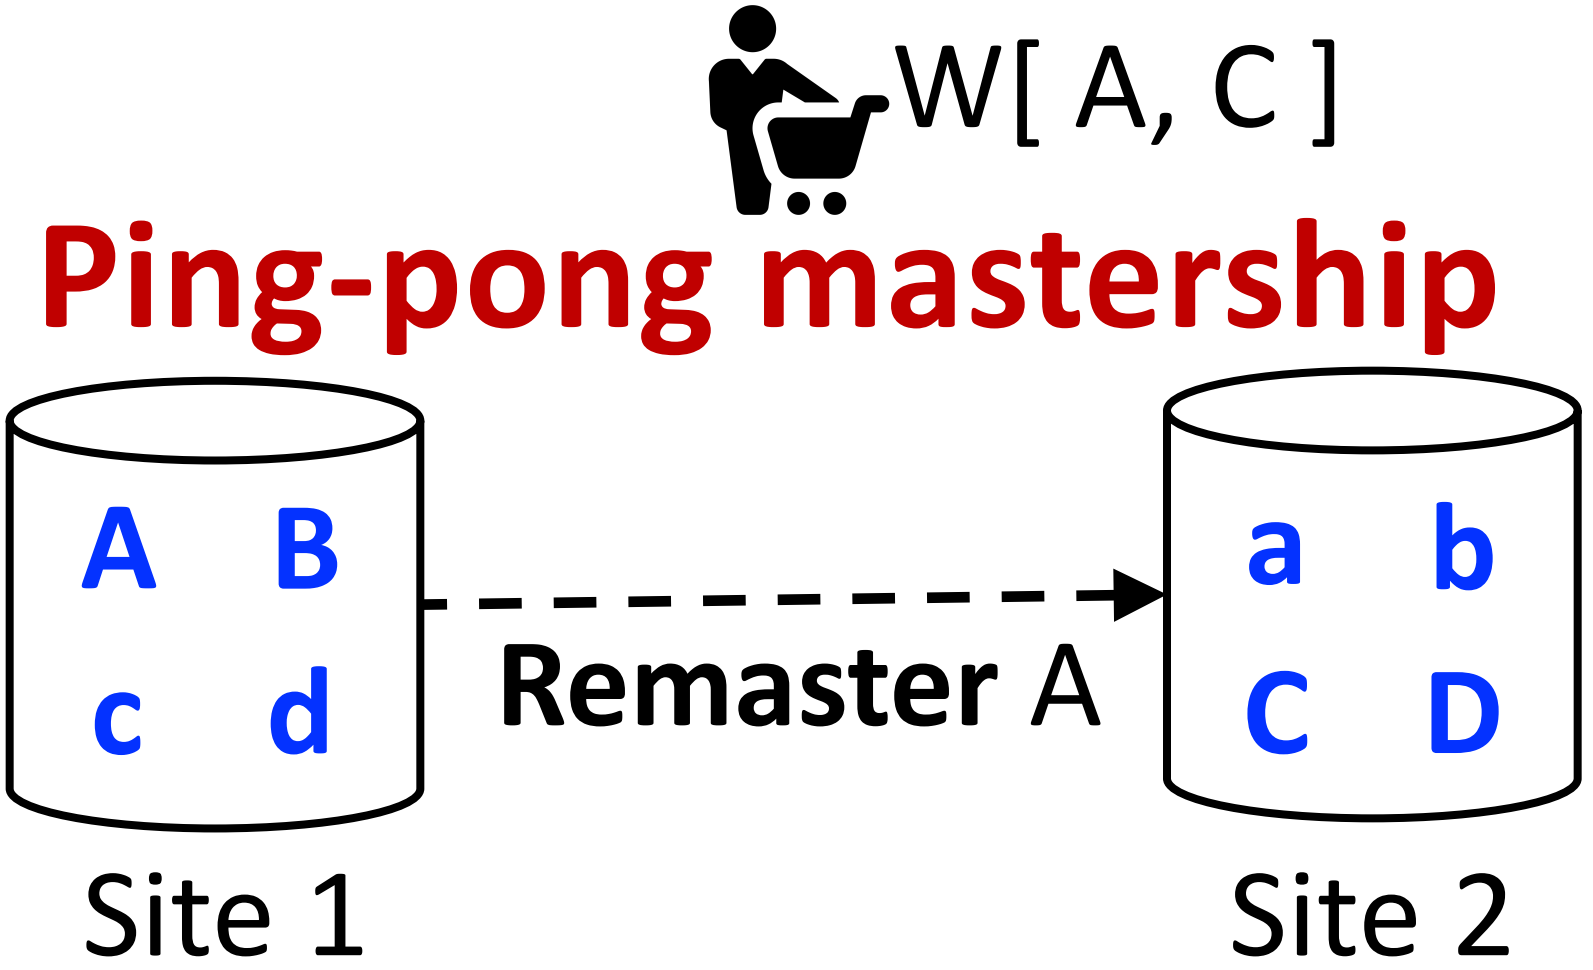

Dynamic Mastering Strategies

Site Selector

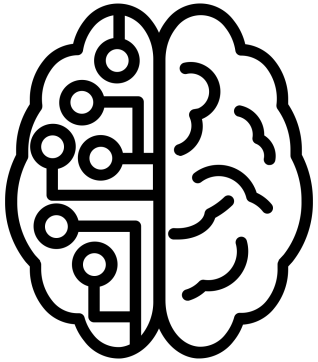

Makes **adaptive** decisions

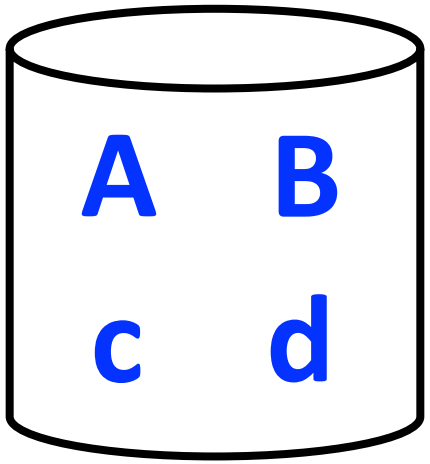

Site 1

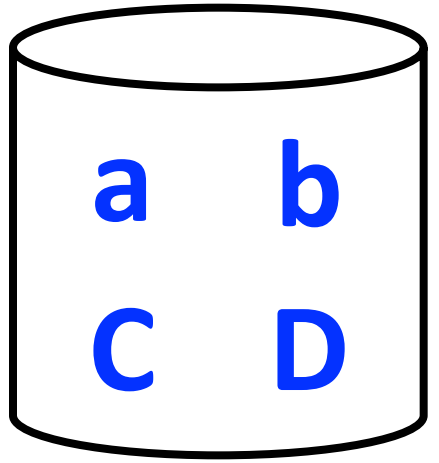

Site 2

UNIVERSITY OF WATERLOO

Dynamic Mastering Strategies

Track: data access patterns, site load

Quantify **benefit** of remastering to each site

- Load distribution
- Future remastering
- Update propagation

Remaster to site that **maximizes benefit**

tiny.cc/dynamast

UNIVERSITY OF WATERLOO

How well does it work?

Workloads

YCSB Scans & Multi-Key Read-Modify-Write

Uniform and Skew Access Patterns

TPC-C Complex updates and reads

UNIVERSITY OF WATERLOO

How well does it work?

Comparisons

Single-Master
Multi-Master
} *Replicated*

Partition-Store
LEAP
} *Single Copy*

UNIVERSITY OF WATERLOO

YCSB with Skew - Throughput

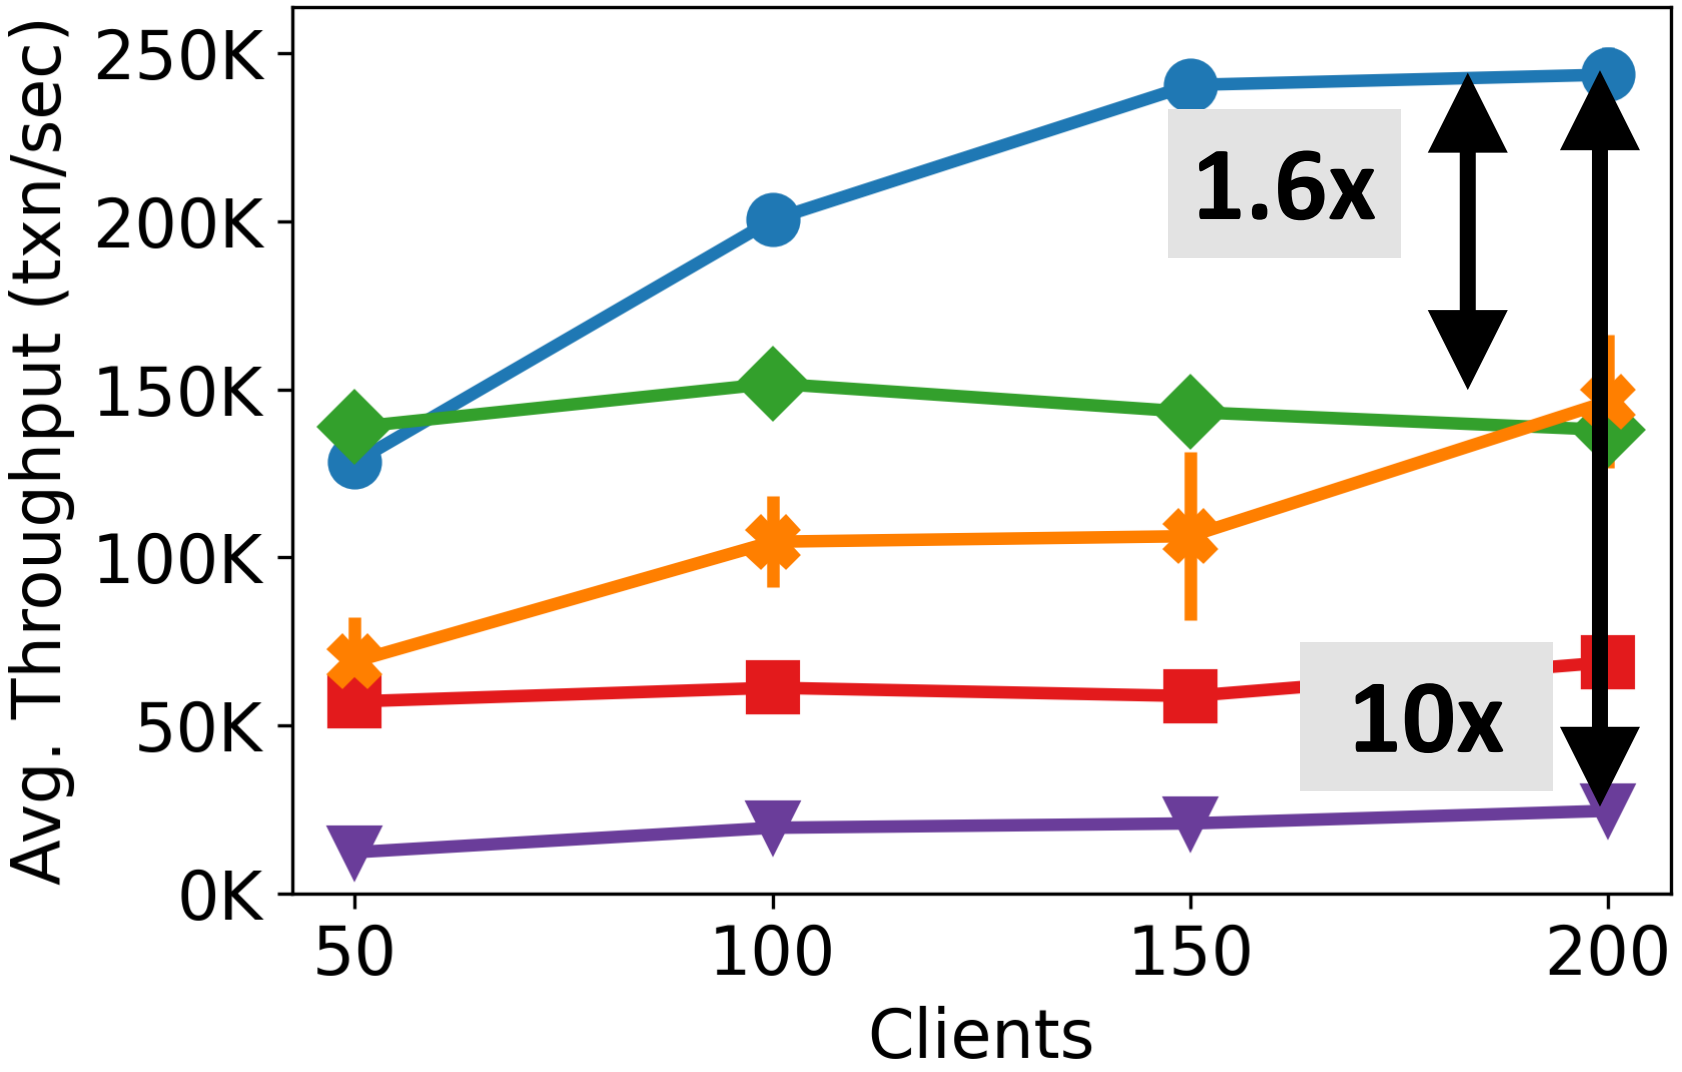

DynaMast

LEAP

Single-Master

Partition-Store

Multi-Master

UNIVERSITY OF WATERLOO

YCSB with Skew - Routing

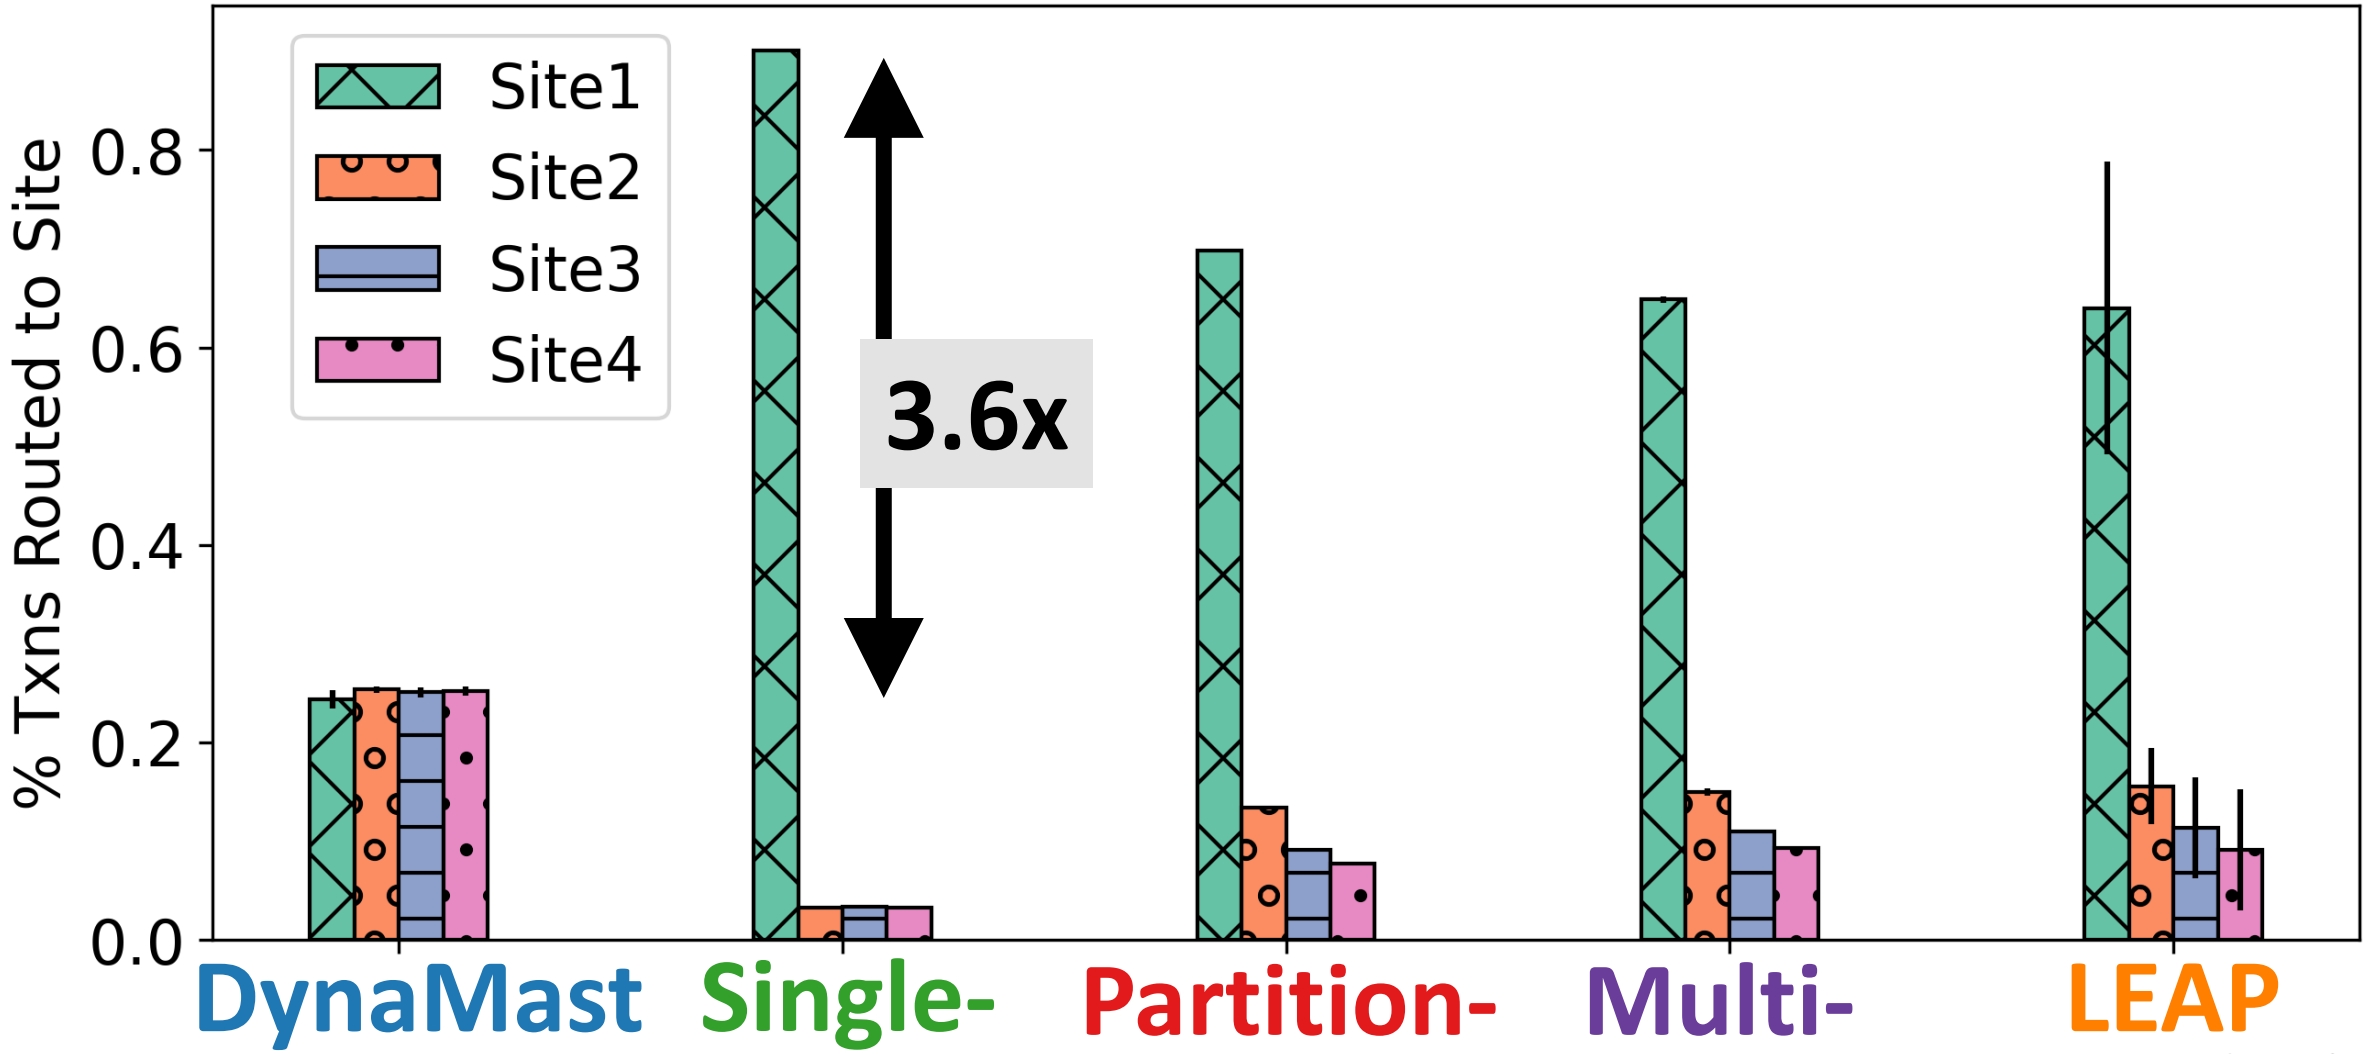

UNIVERSITY OF WATERLOO

TPC-C – New-Order Latency

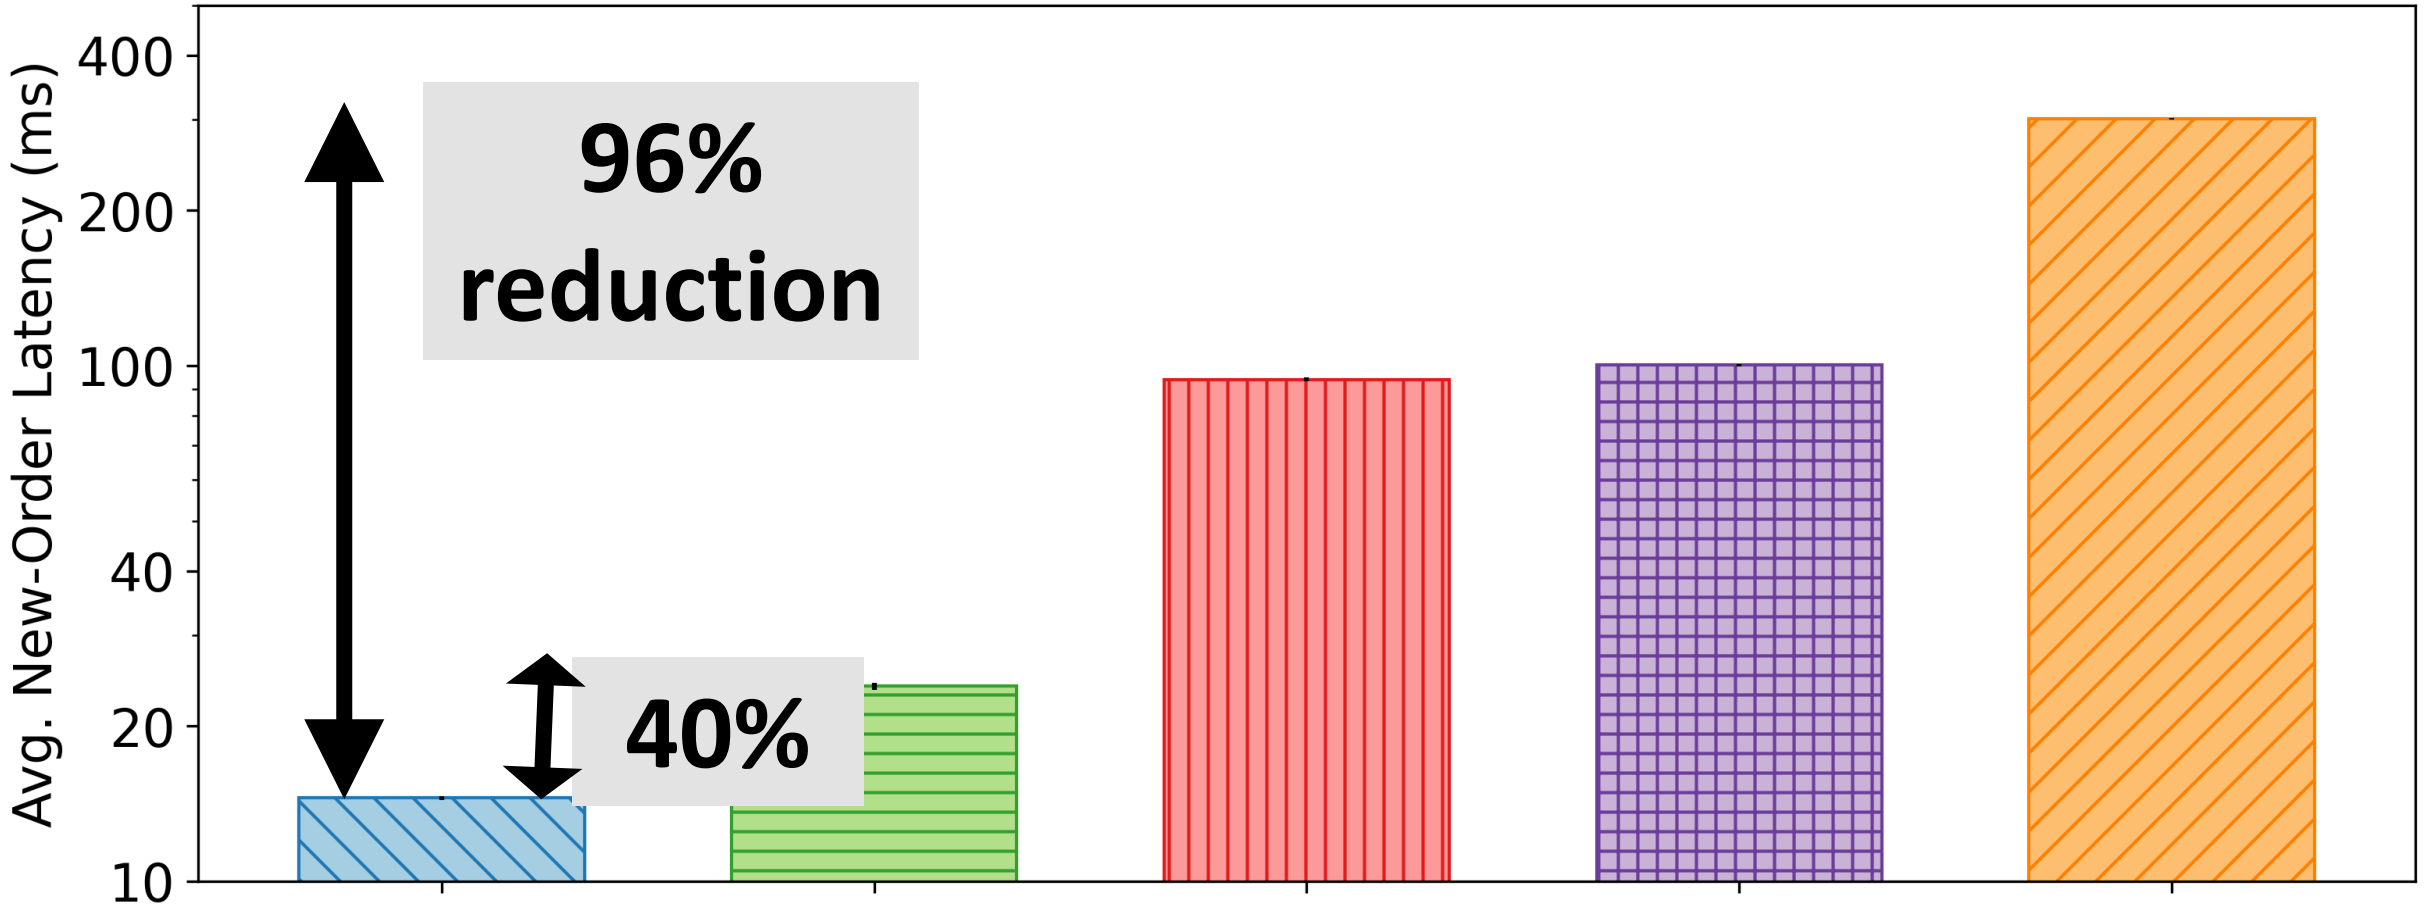

UNIVERSITY OF WATERLOO

DynaMast Takeaways

Dynamic mastering guarantees **single-site transactions**

tiny.cc/dynamast

Use replicas to remaster **efficiently outside transaction boundaries**

Adaptive site selector strategies **balance load** and **minimizes future remastering**

UNIVERSITY OF WATERLOO

Extra slides

DynaMast Learns Over Time

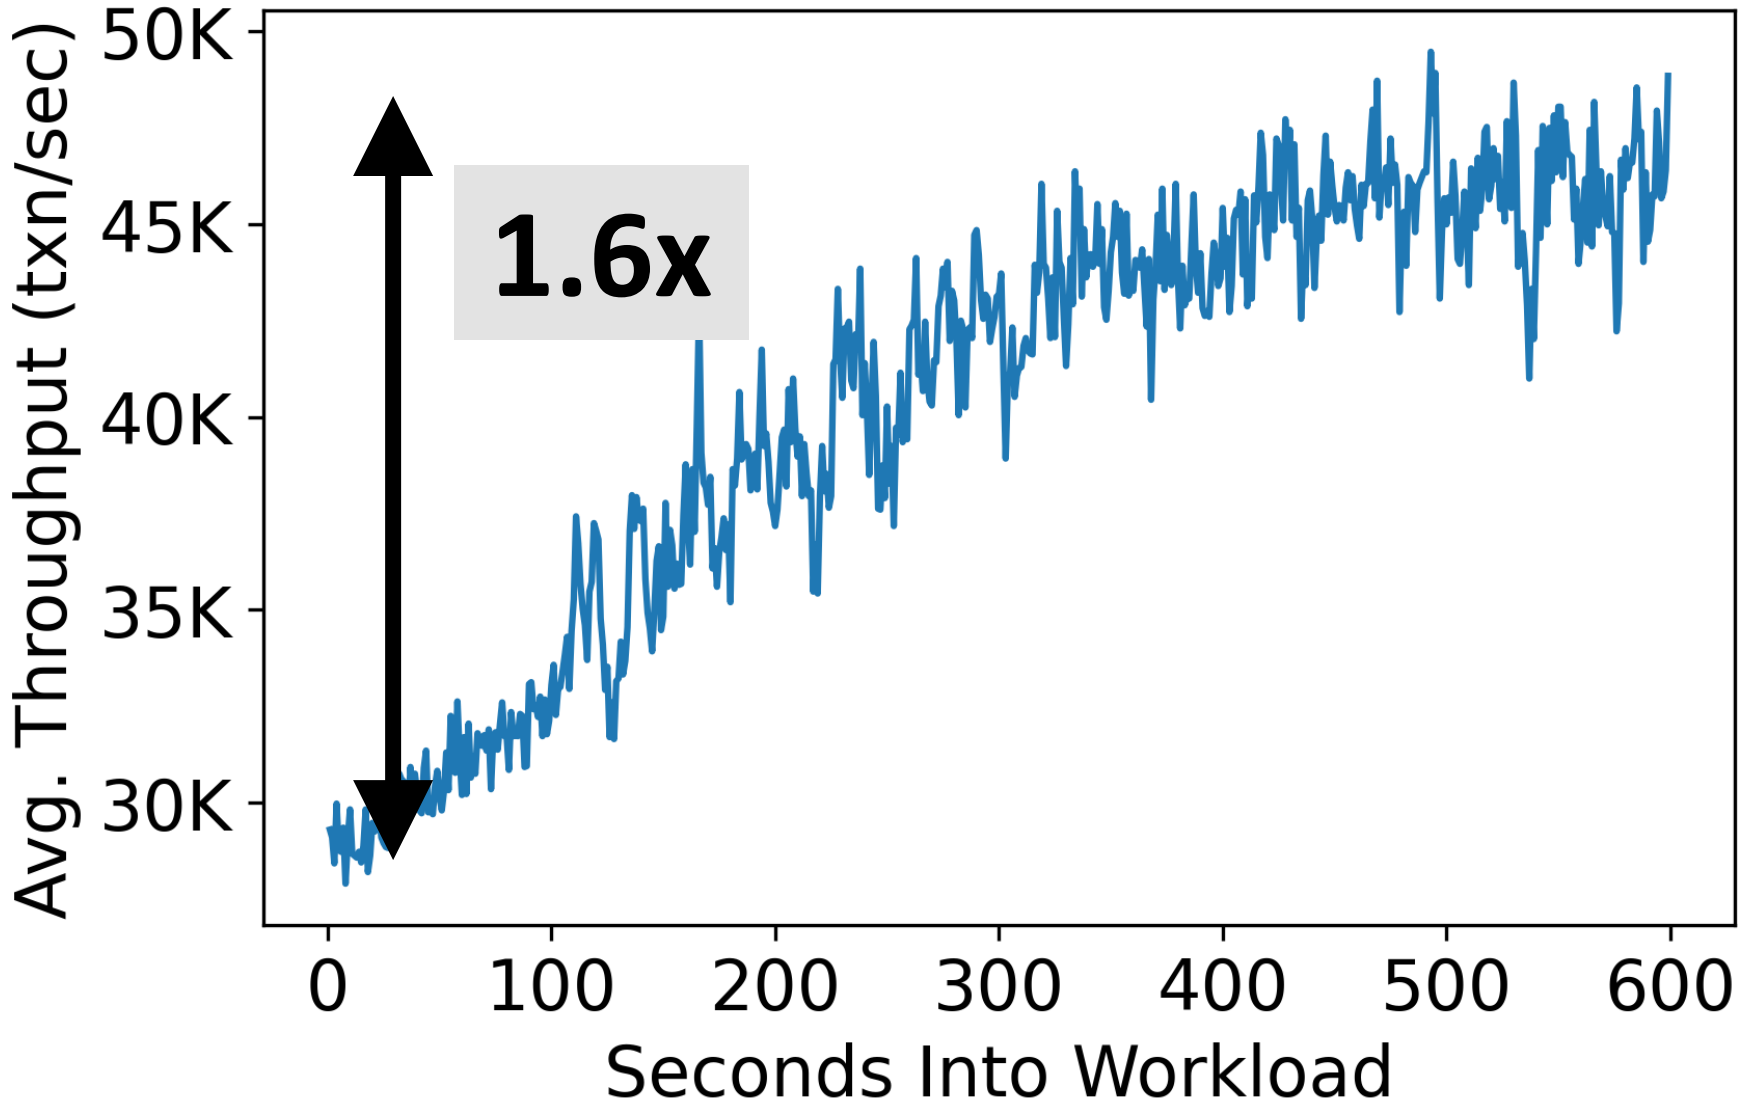

UNIVERSITY OF WATERLOO

TPC-C – New-Order Latency

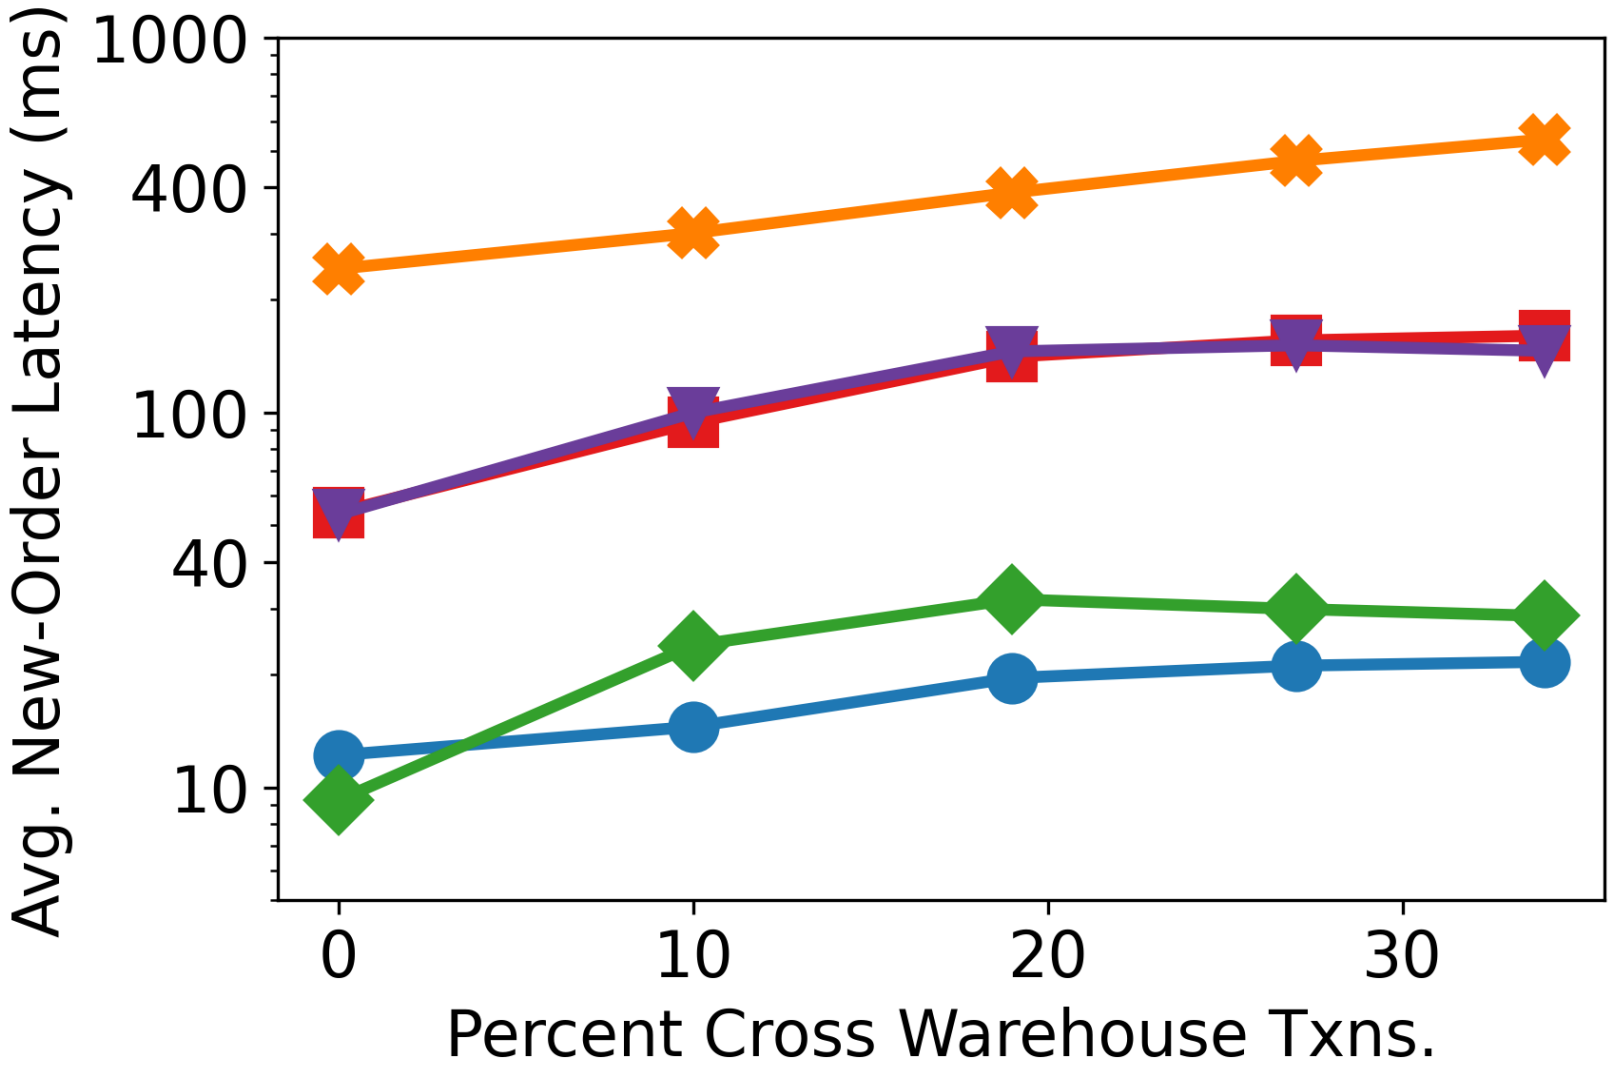

LEAP

Partition-Store

Multi-Master

Single-Master

DynaMast

UNIVERSITY OF WATERLOO

Comparisons – Single-Master

Writers

Readers

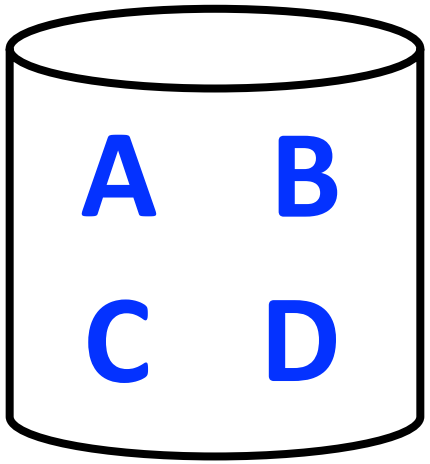

Site 1

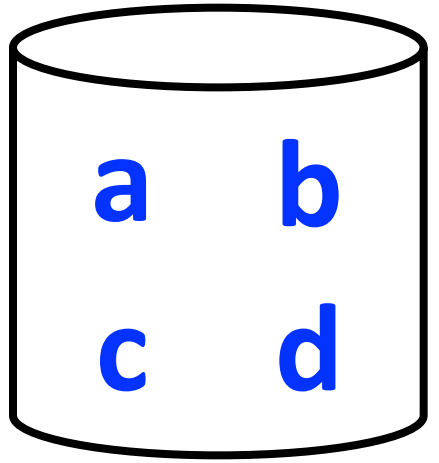

Site 2

Comparisons – Multi-Master

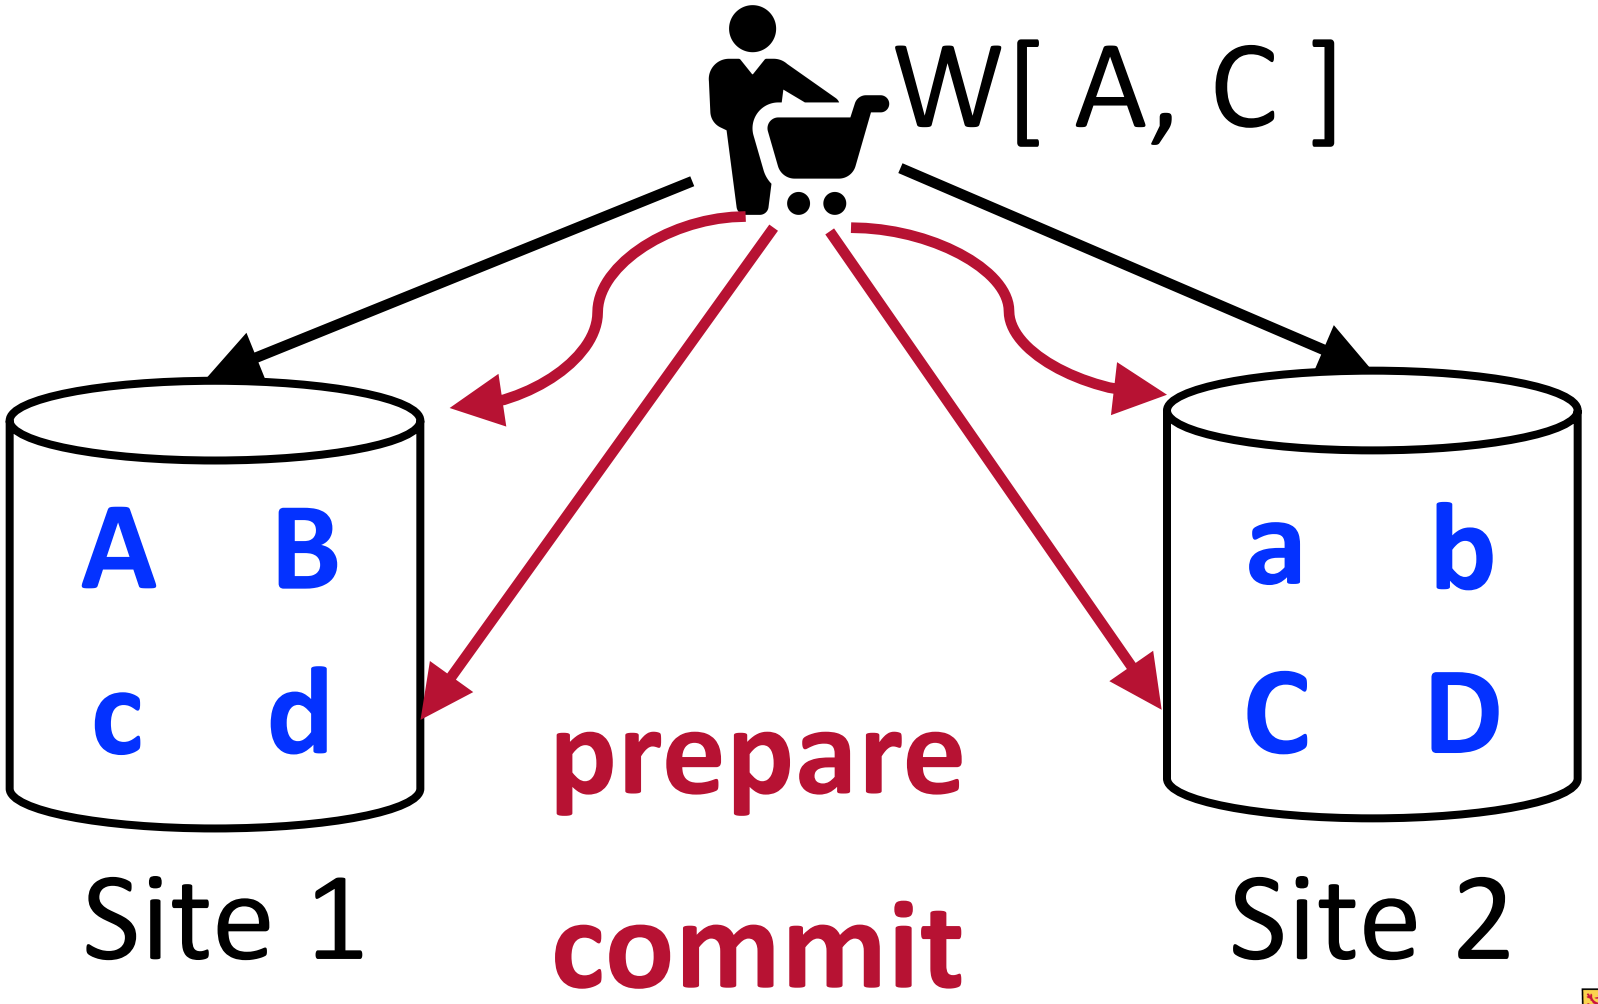

Comparisons – Partition-Store

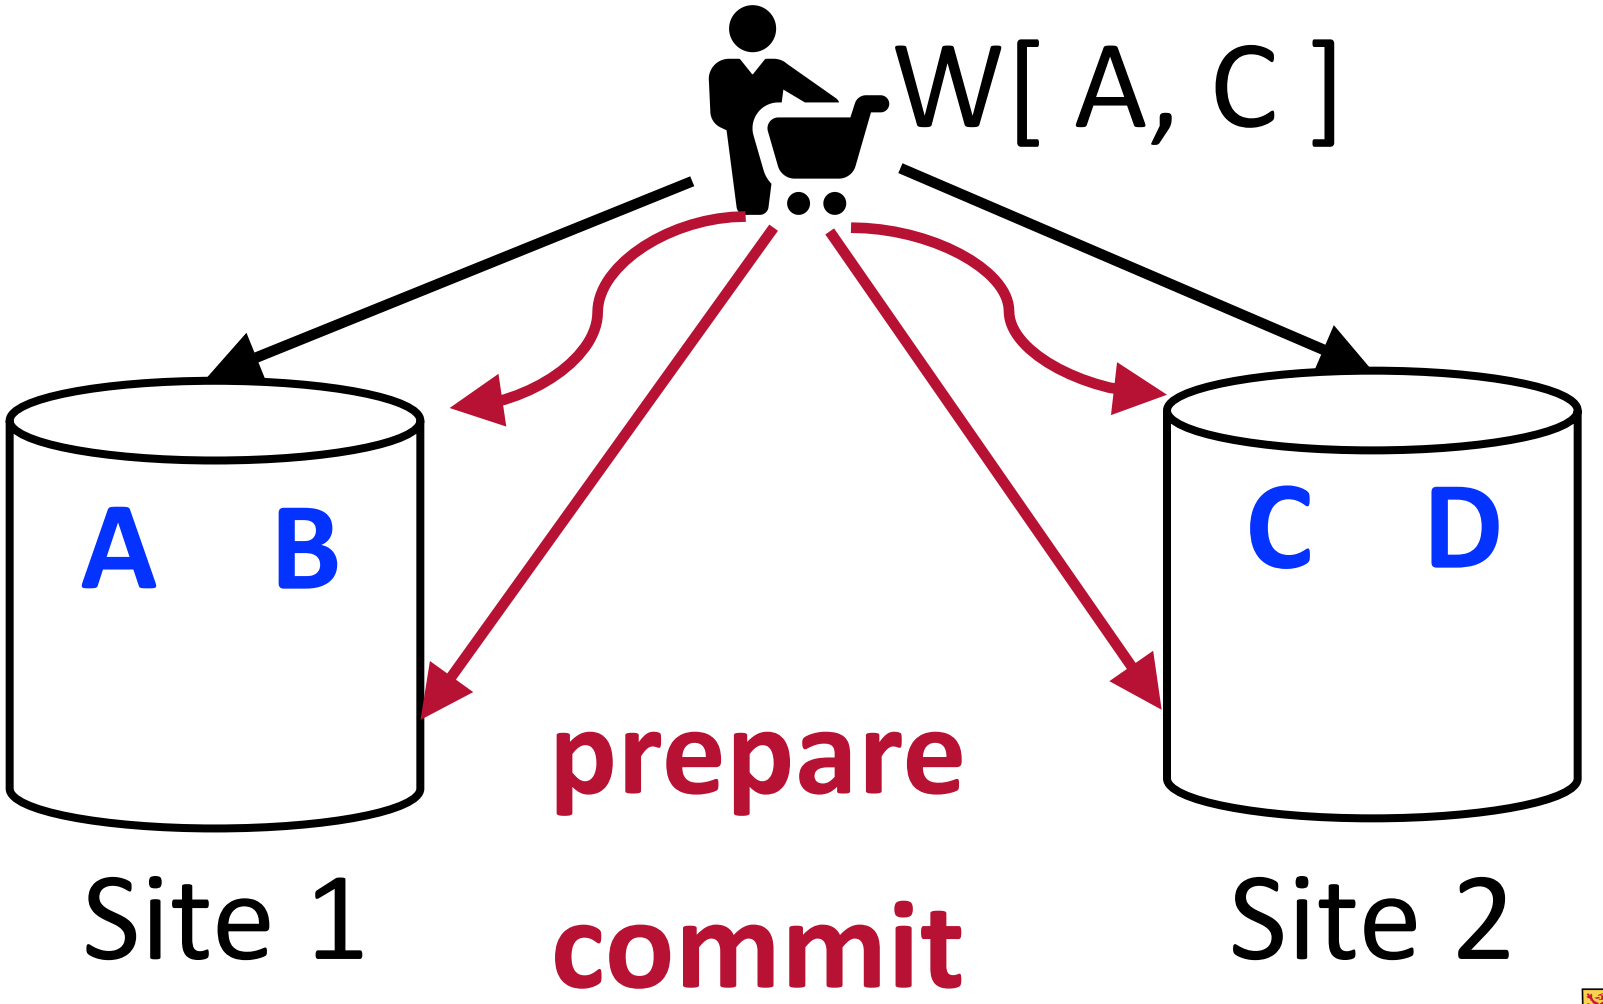

Comparisons – LEAP

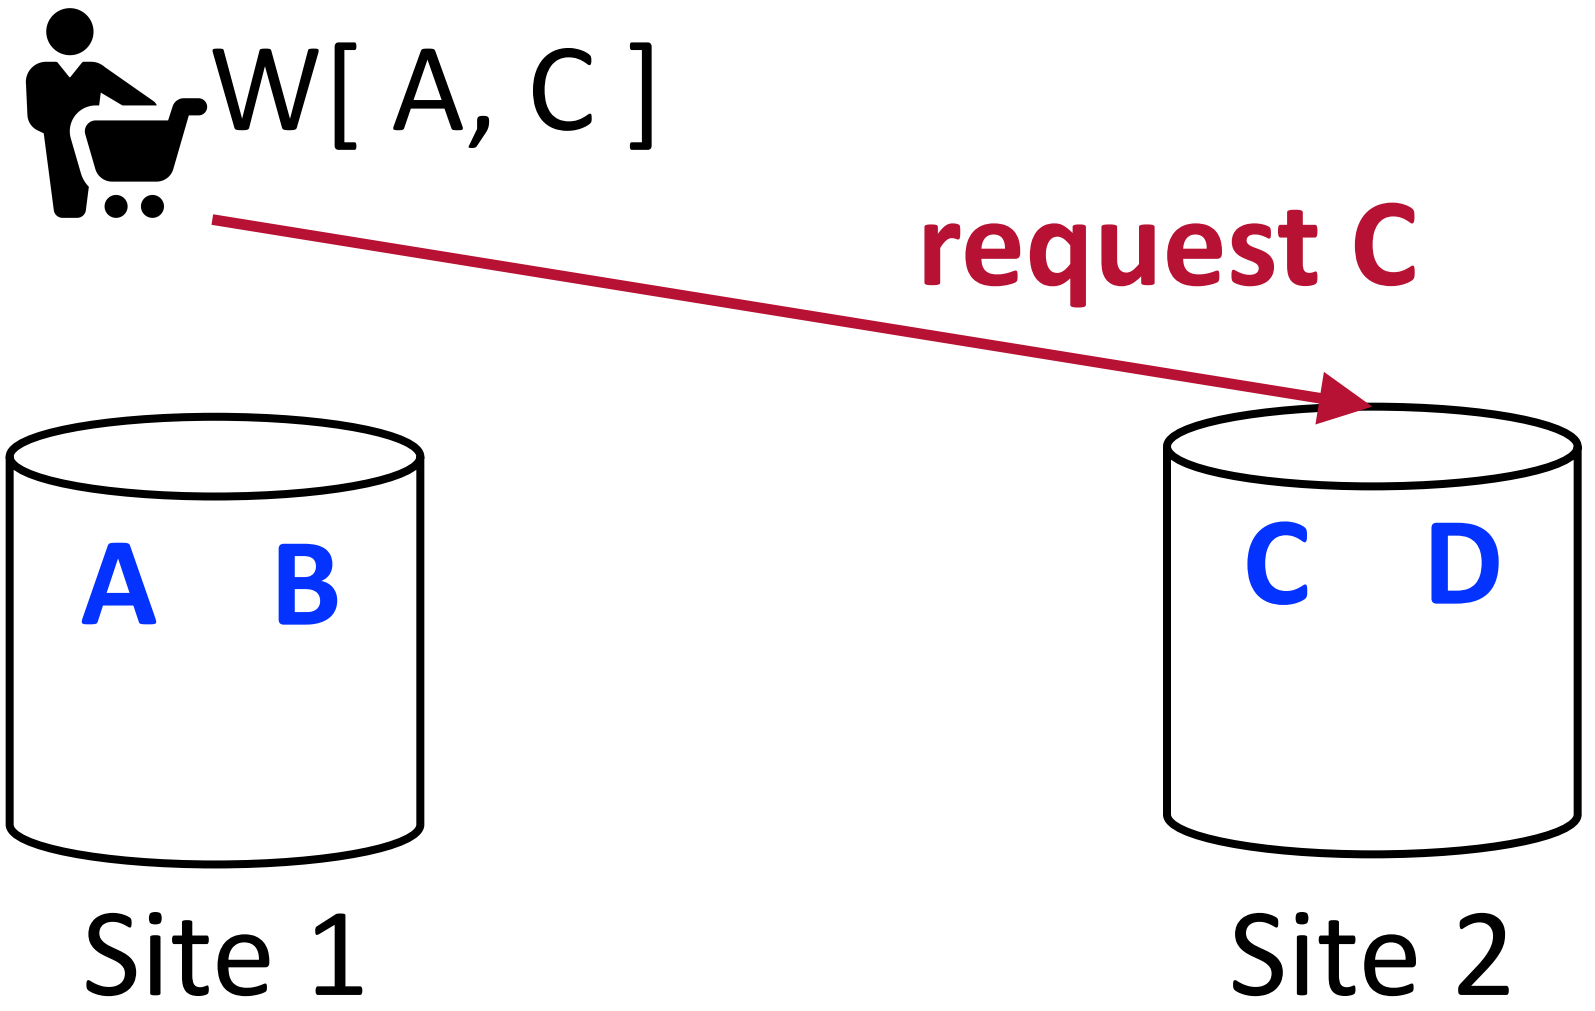

Comparisons – LEAP

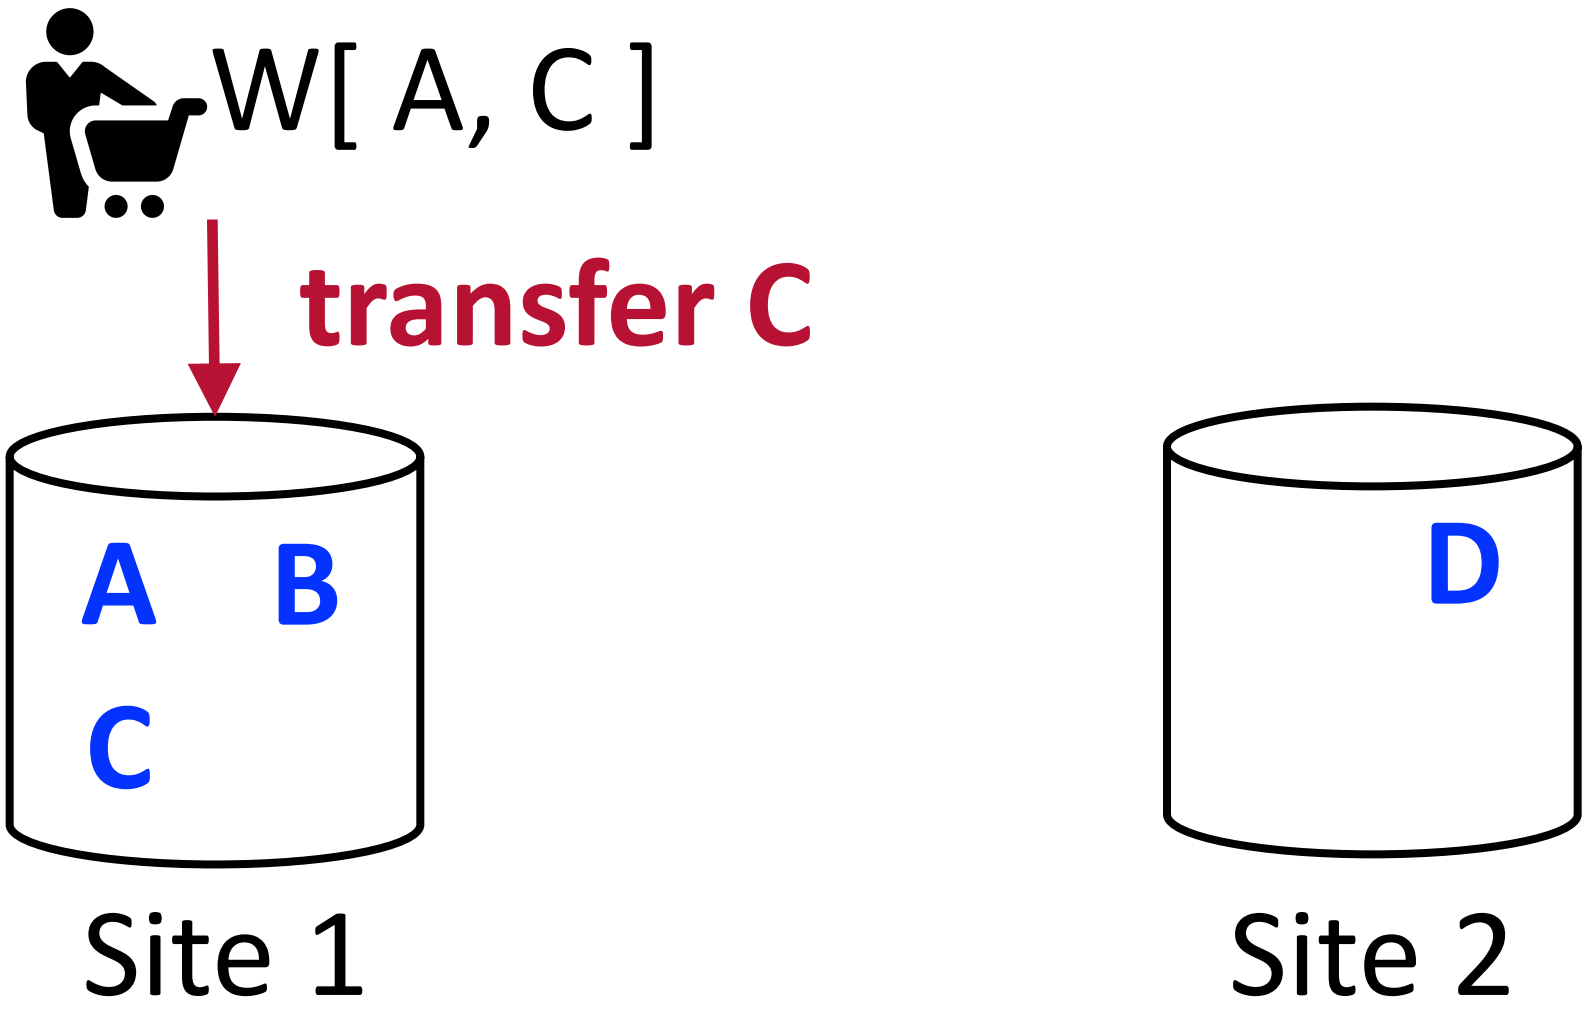

Comparisons – LEAP