CRDTs in Production

Dmitry Martyanov, Software Engineer @ PayPal

QCon, 2018

Geo-Distributed Datastore

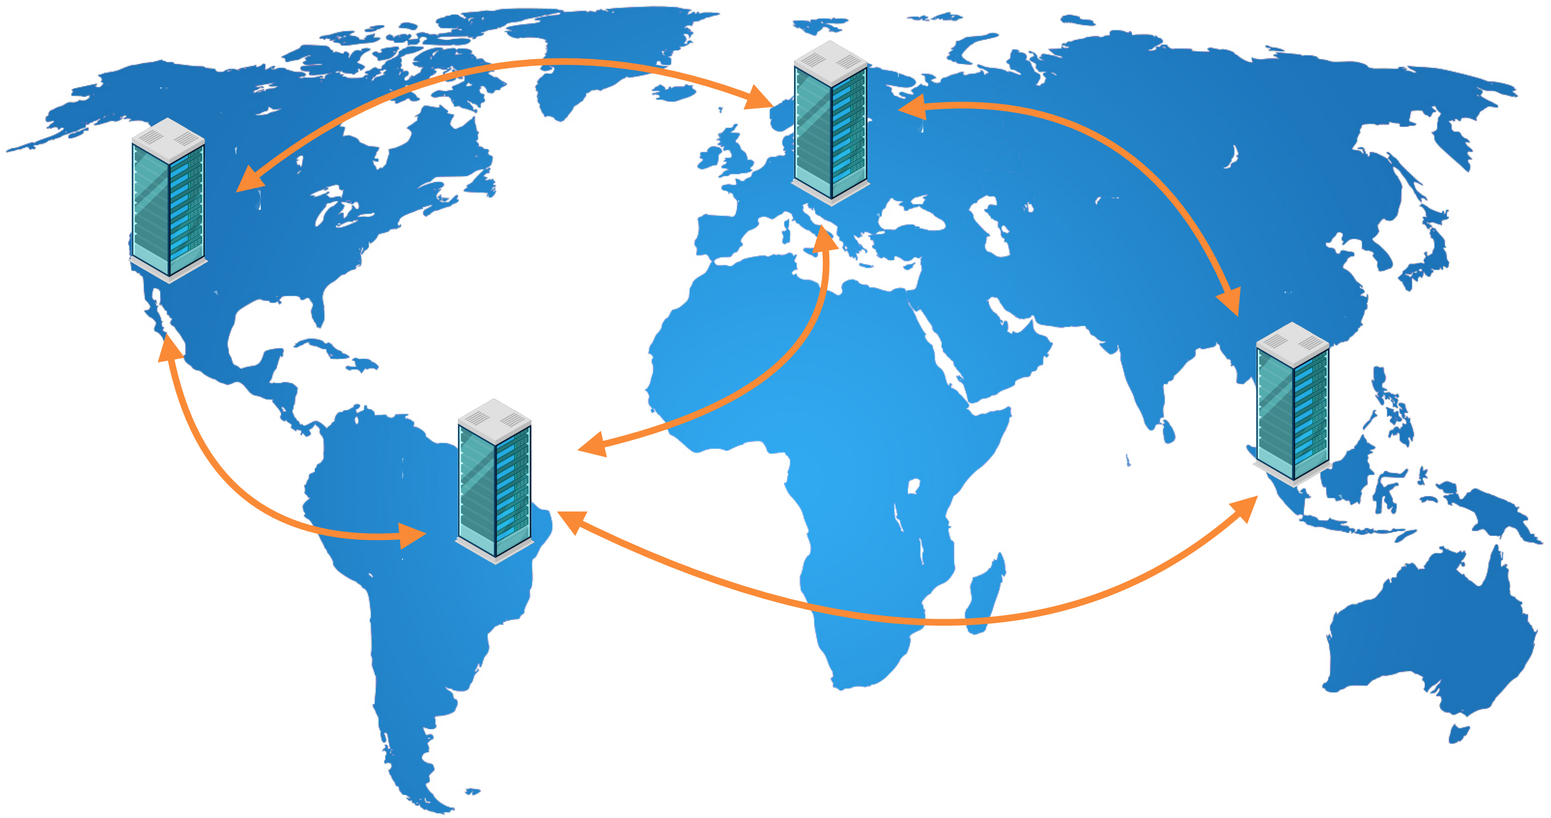

Context

- More than 200 countries
- Regulatory requirements
- State Machine of Compliance Status
- Modified by multiple Actors

Shared Mutable State

Shared Mutable State

Mutex

Shared Mutable State

Mutex

Transactions

Geo-Distributed Datastore

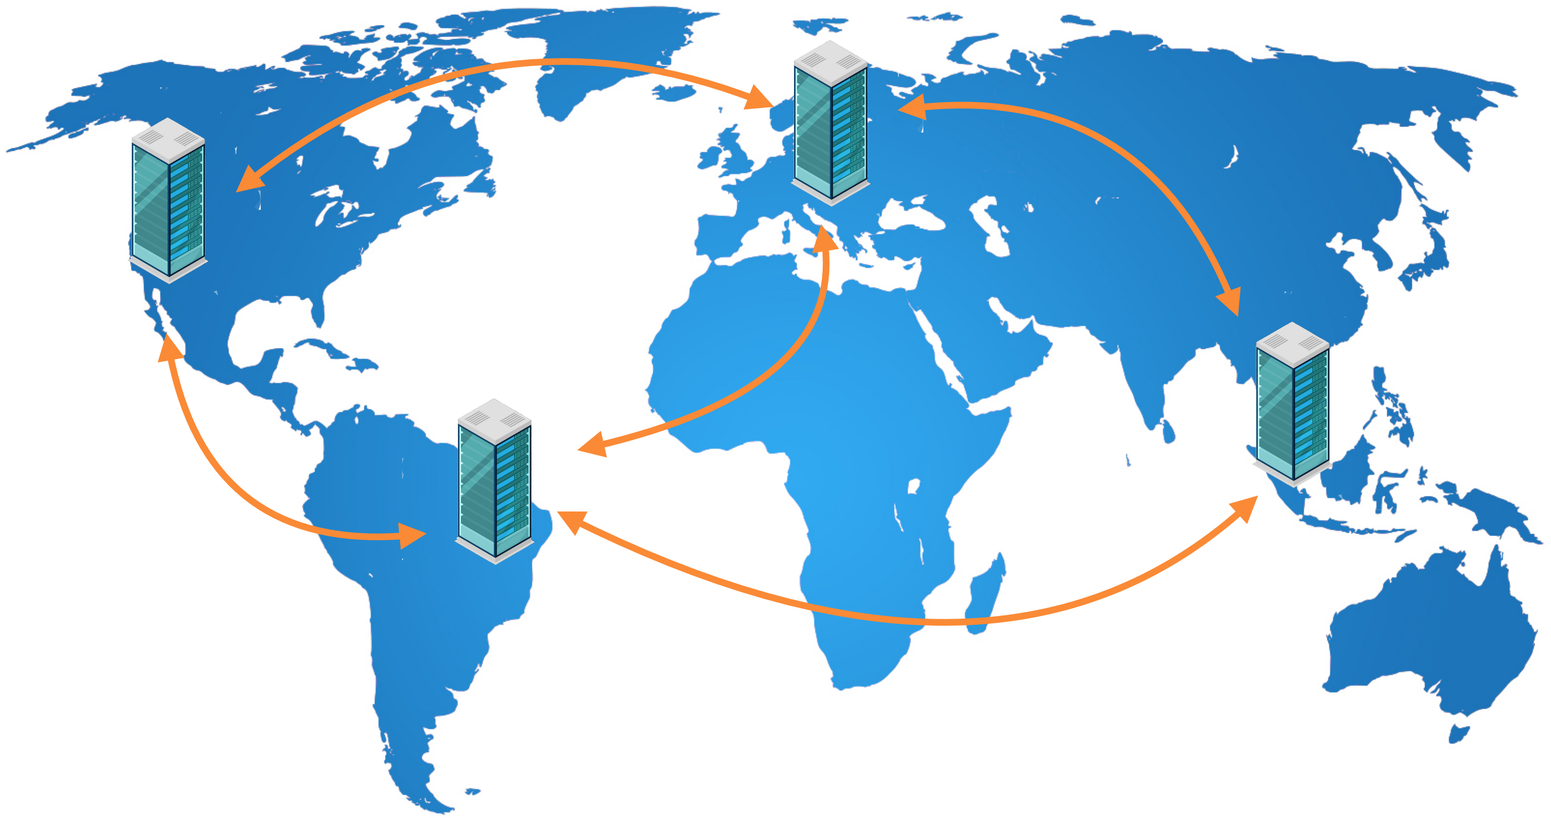

Eventual Consistency

Distributed System

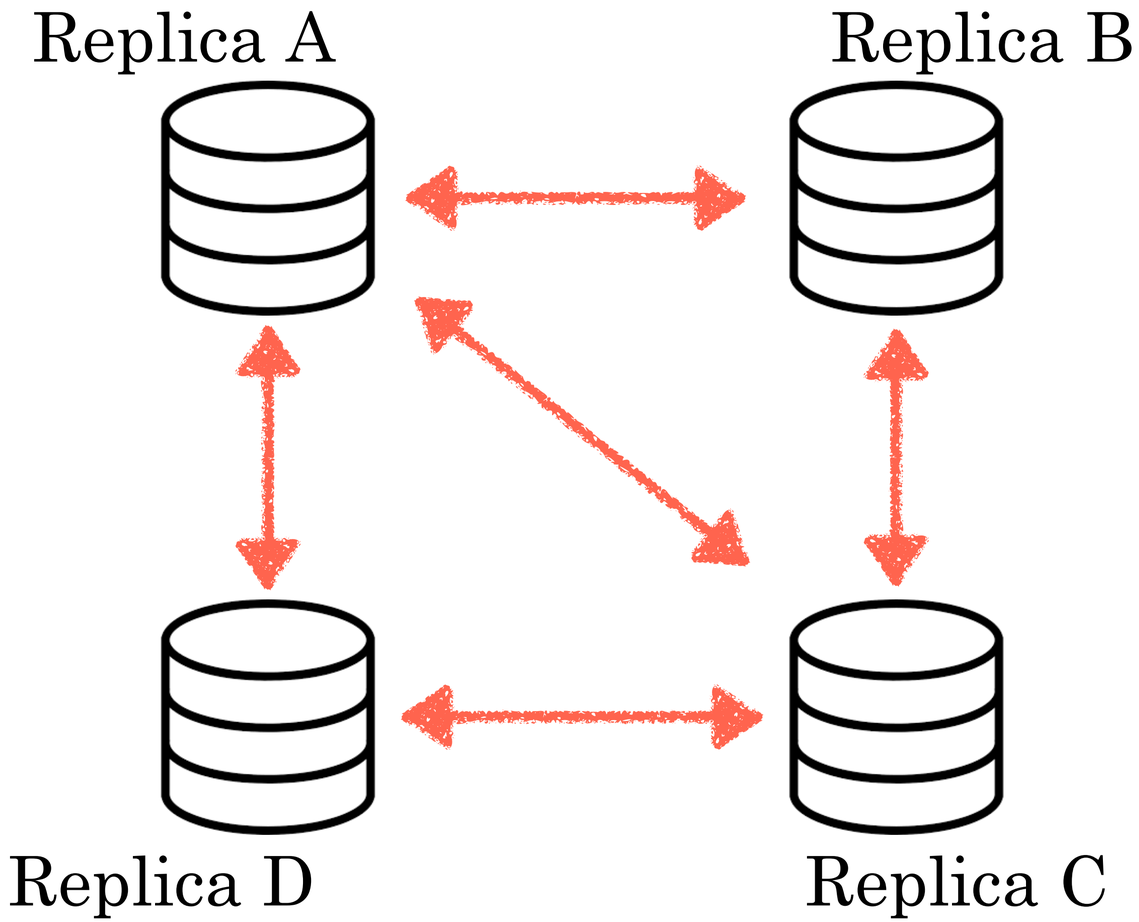

Distributed System

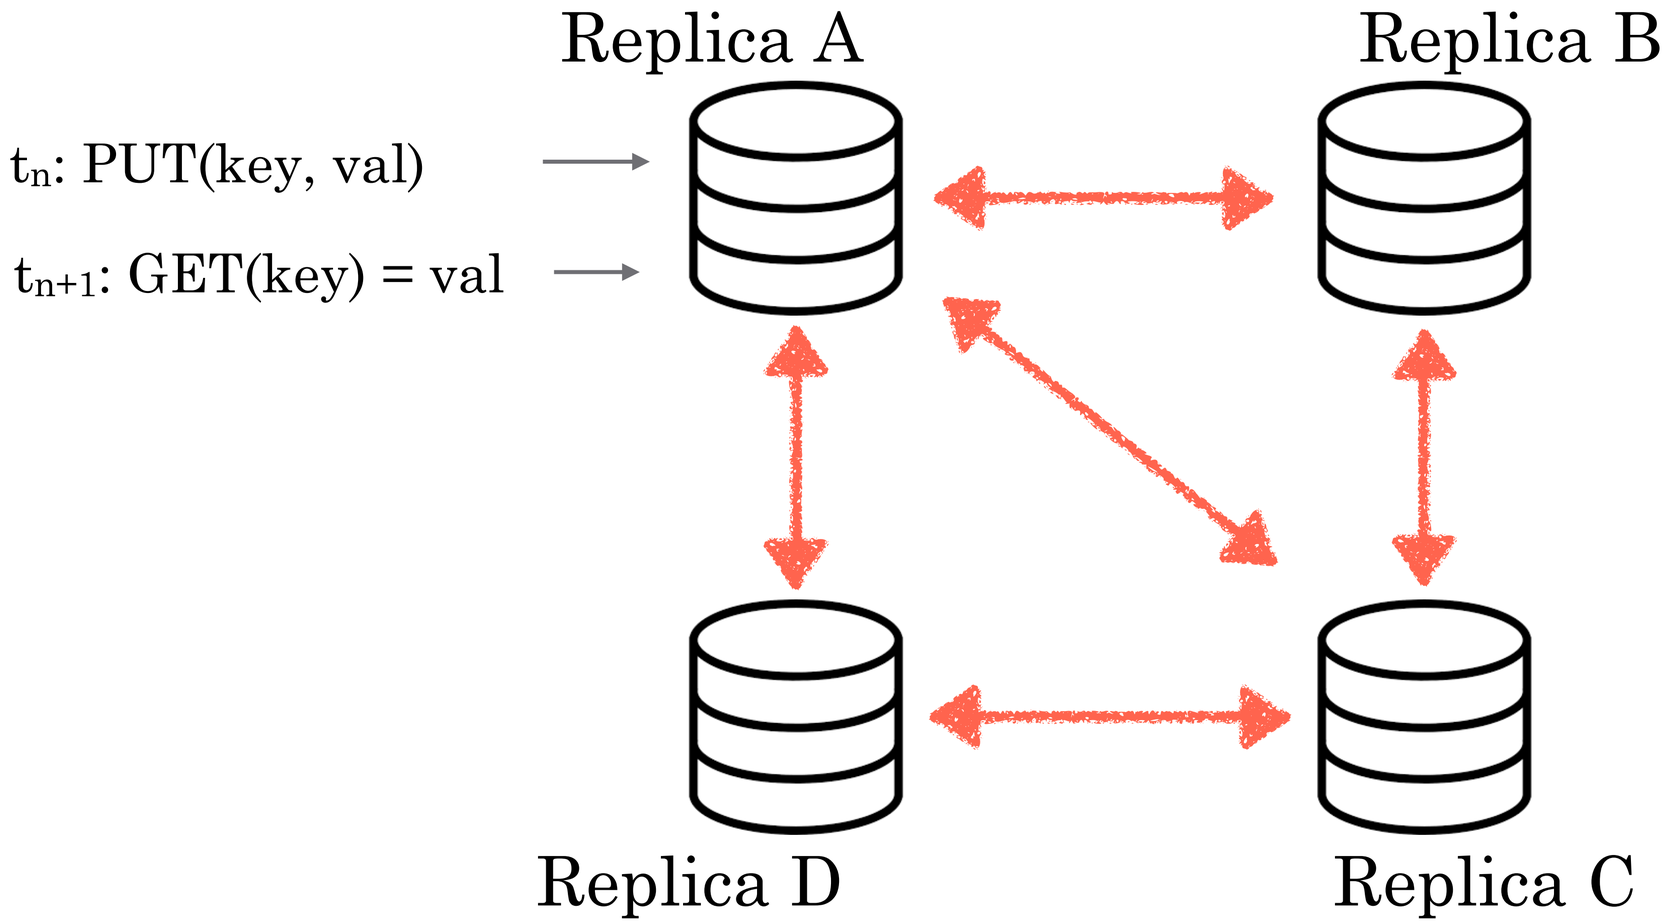

Distributed System

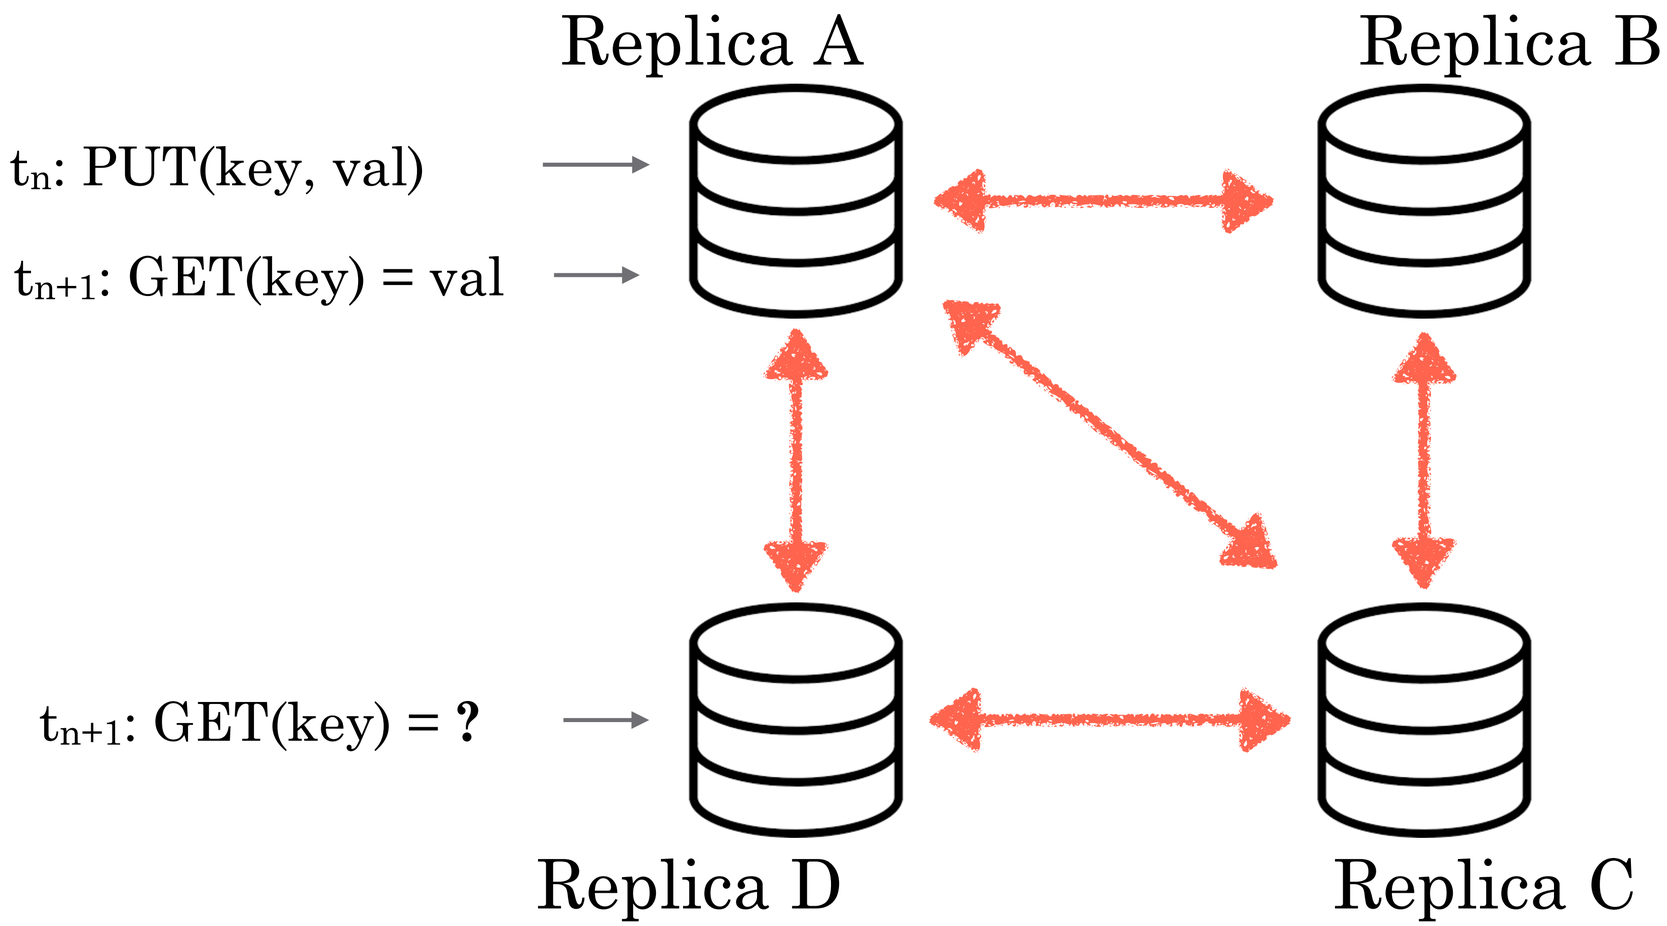

Affinity Based Approaches

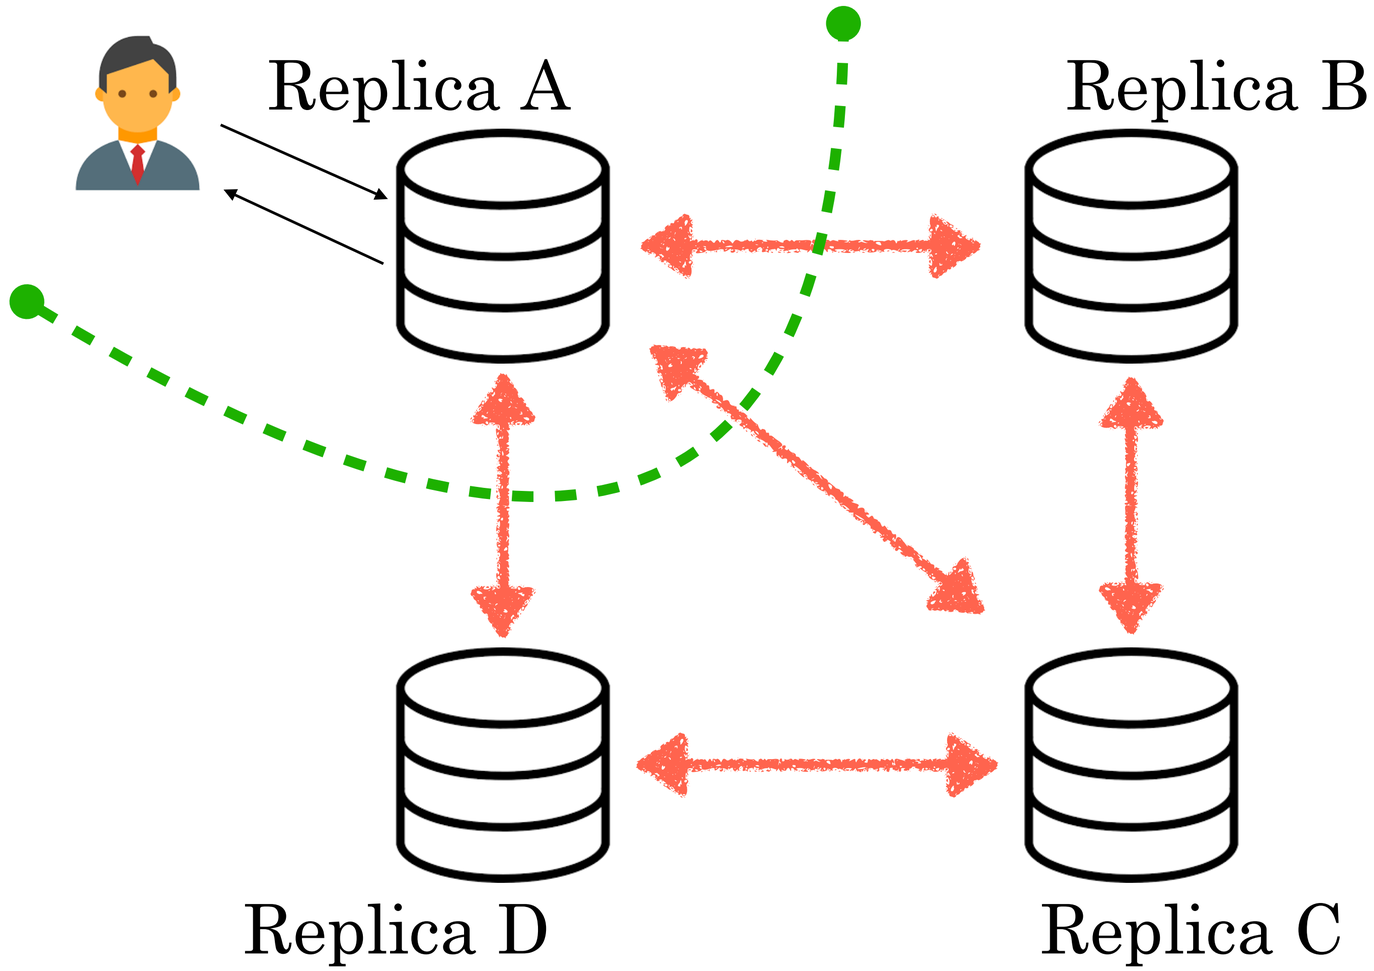

Affinity Based Approaches

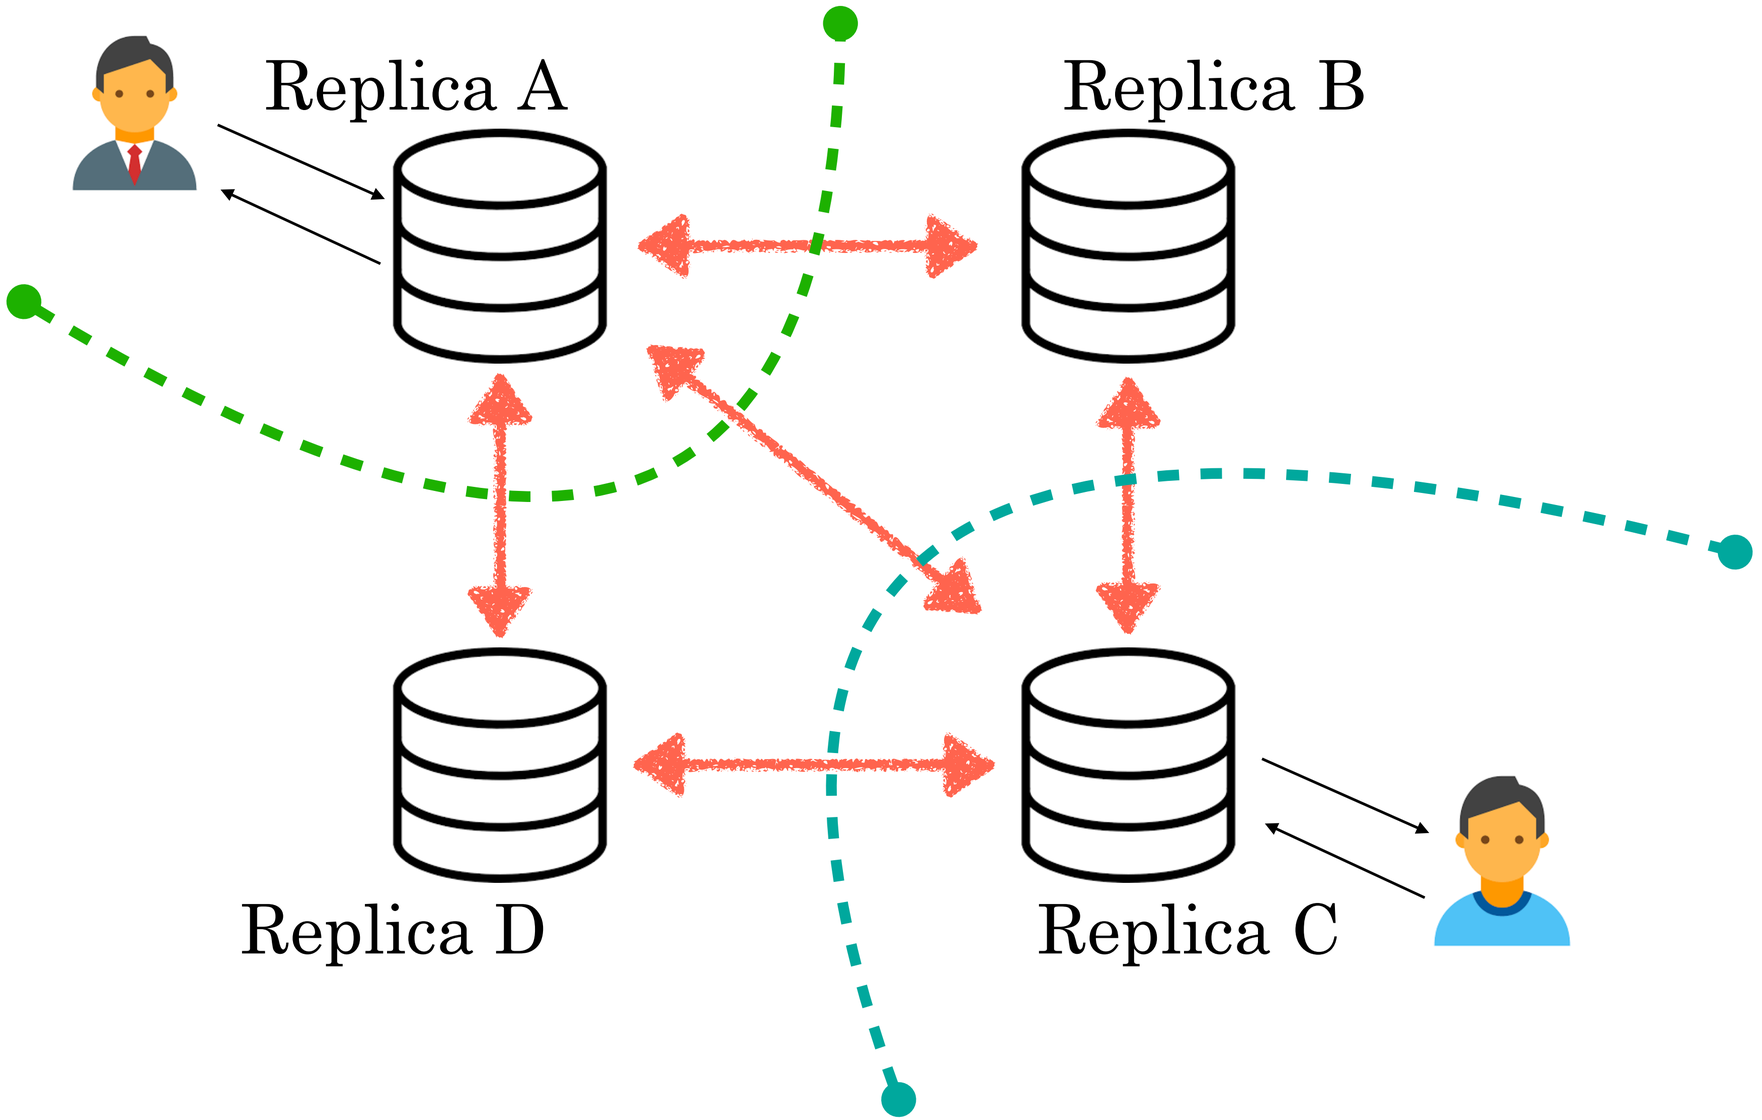

Coordinator Based Approaches

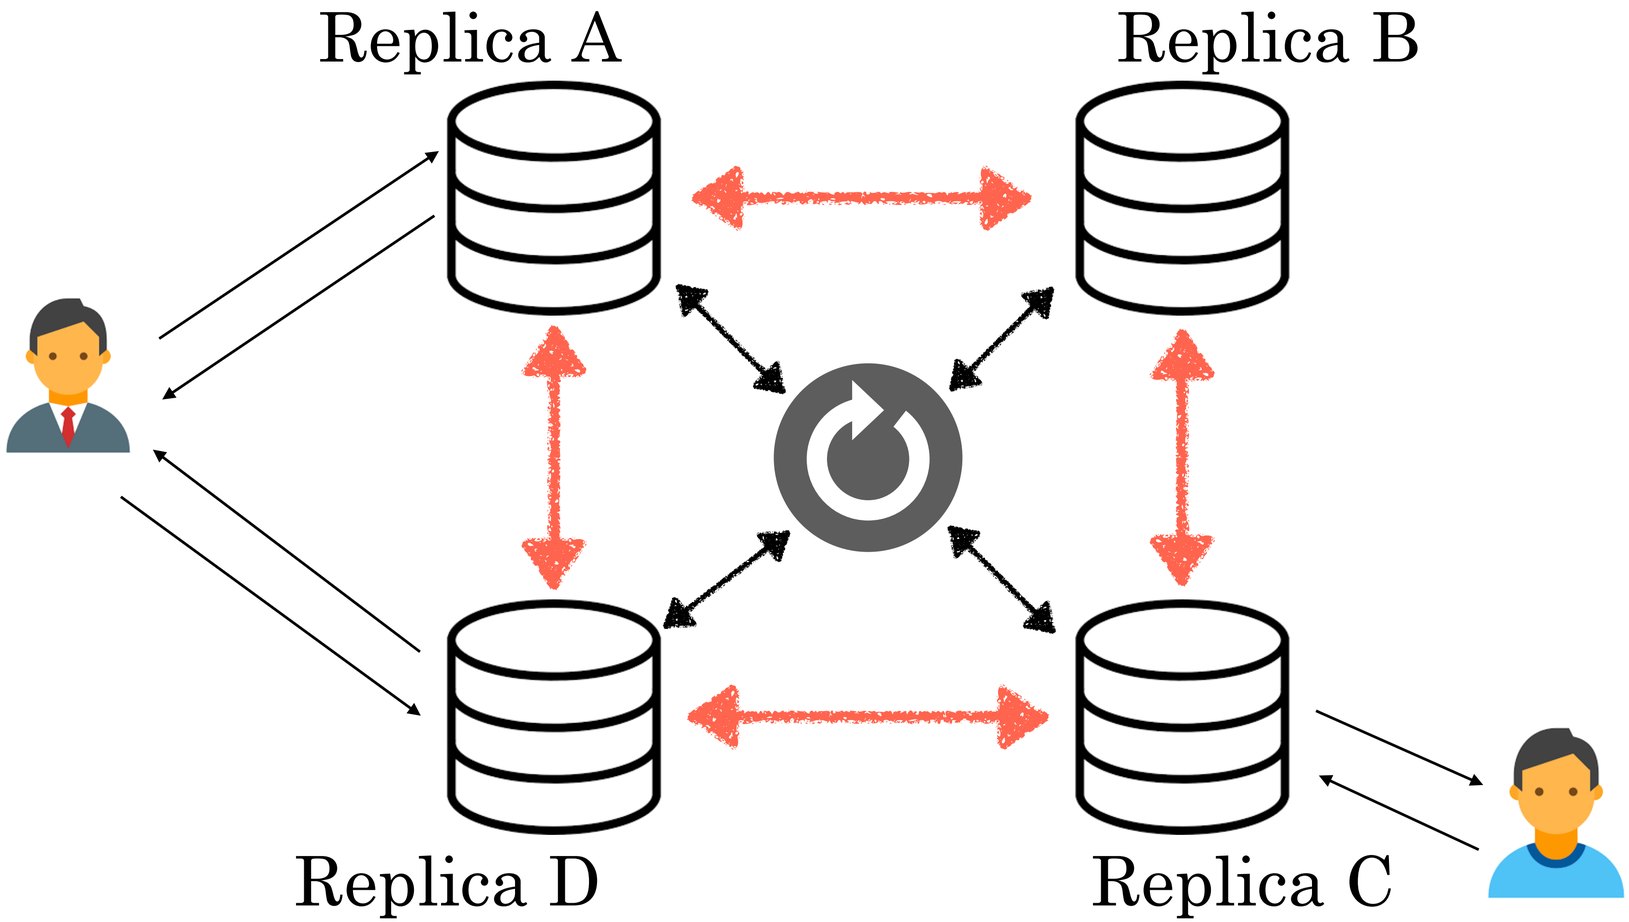

Consensus Based Approaches

Paxos, Raft, etc.

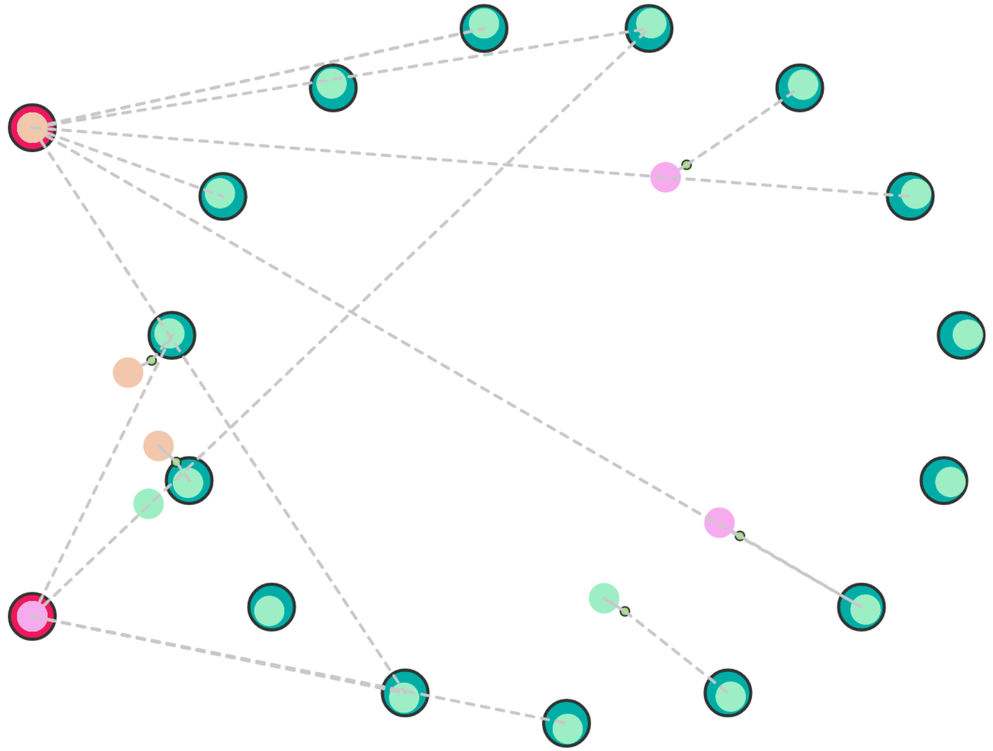

Service Stack

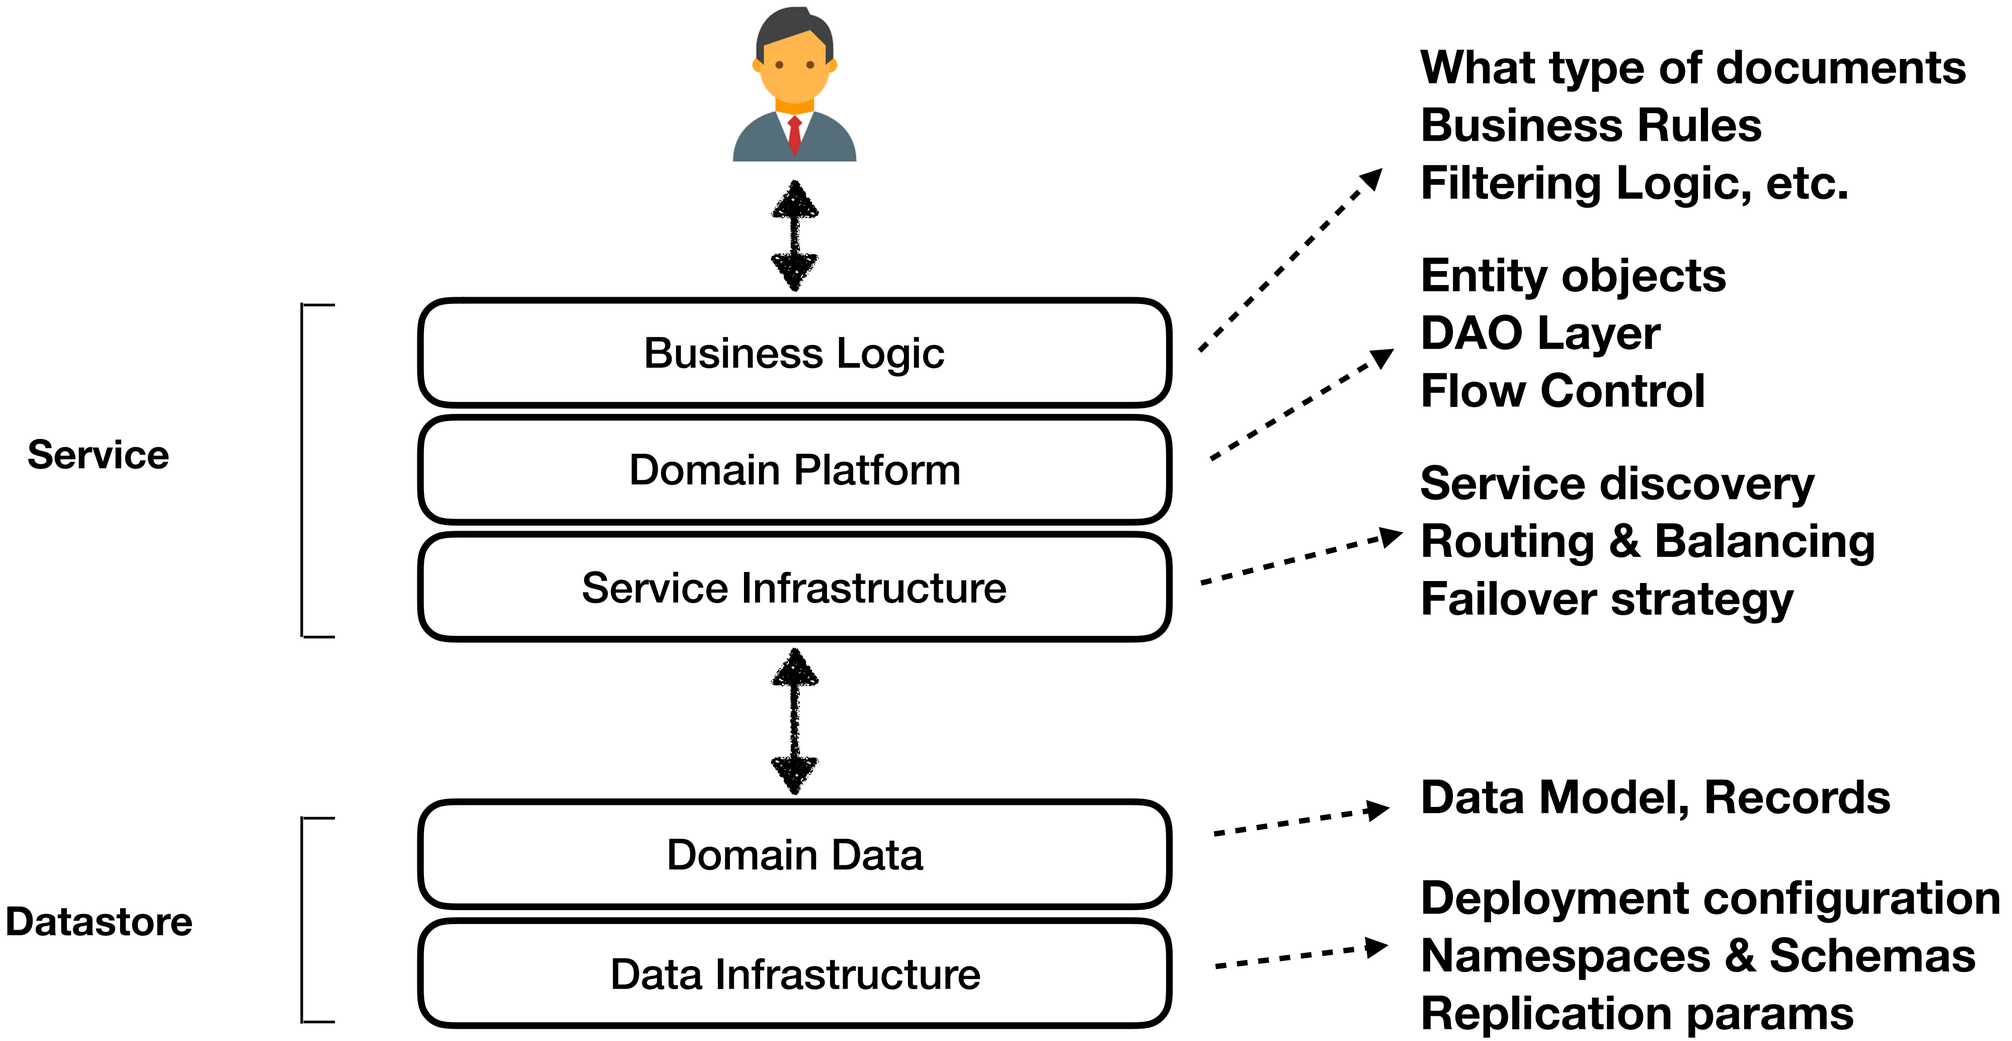

Service Stack

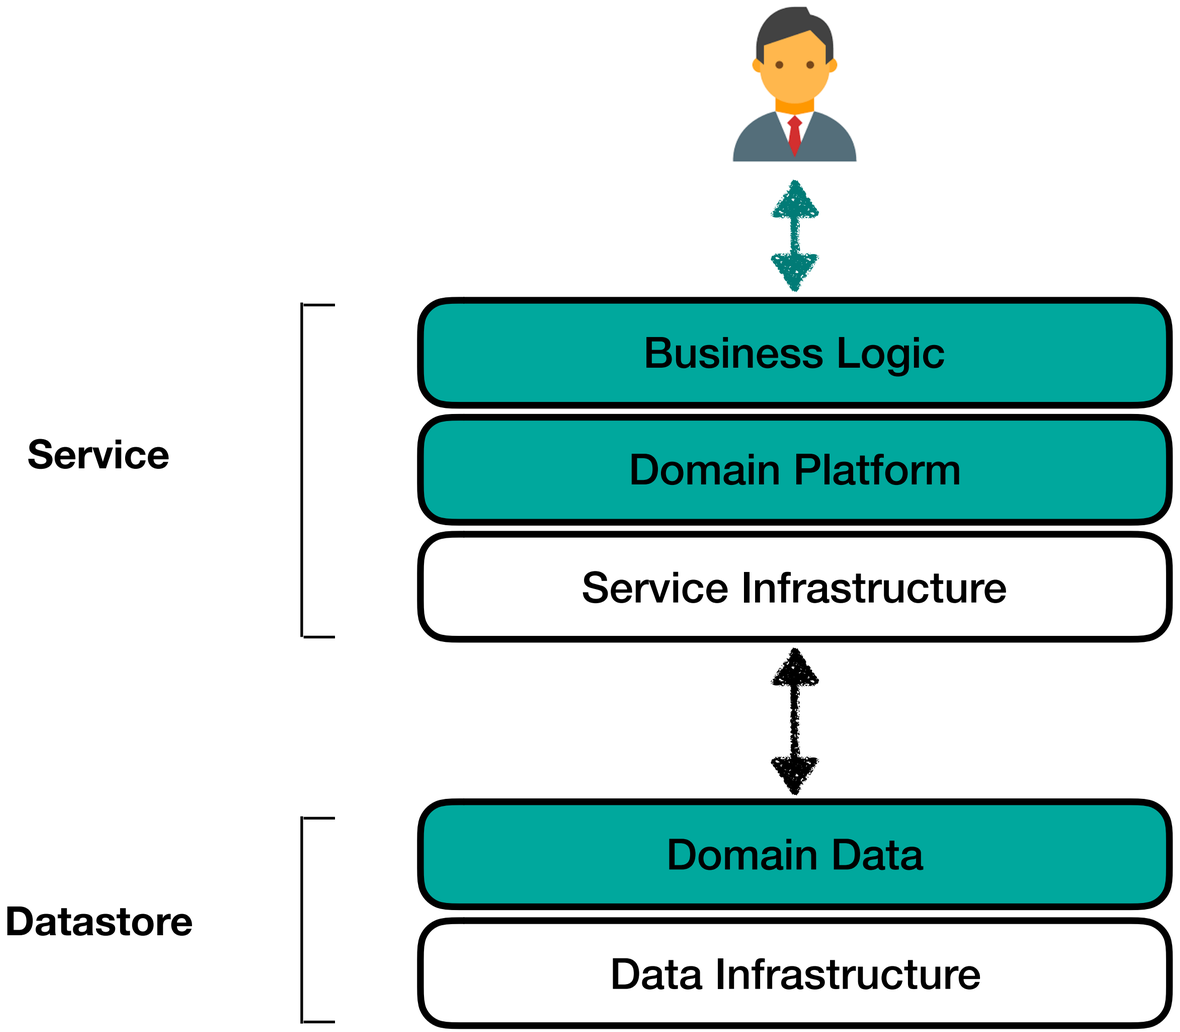

Service Stack

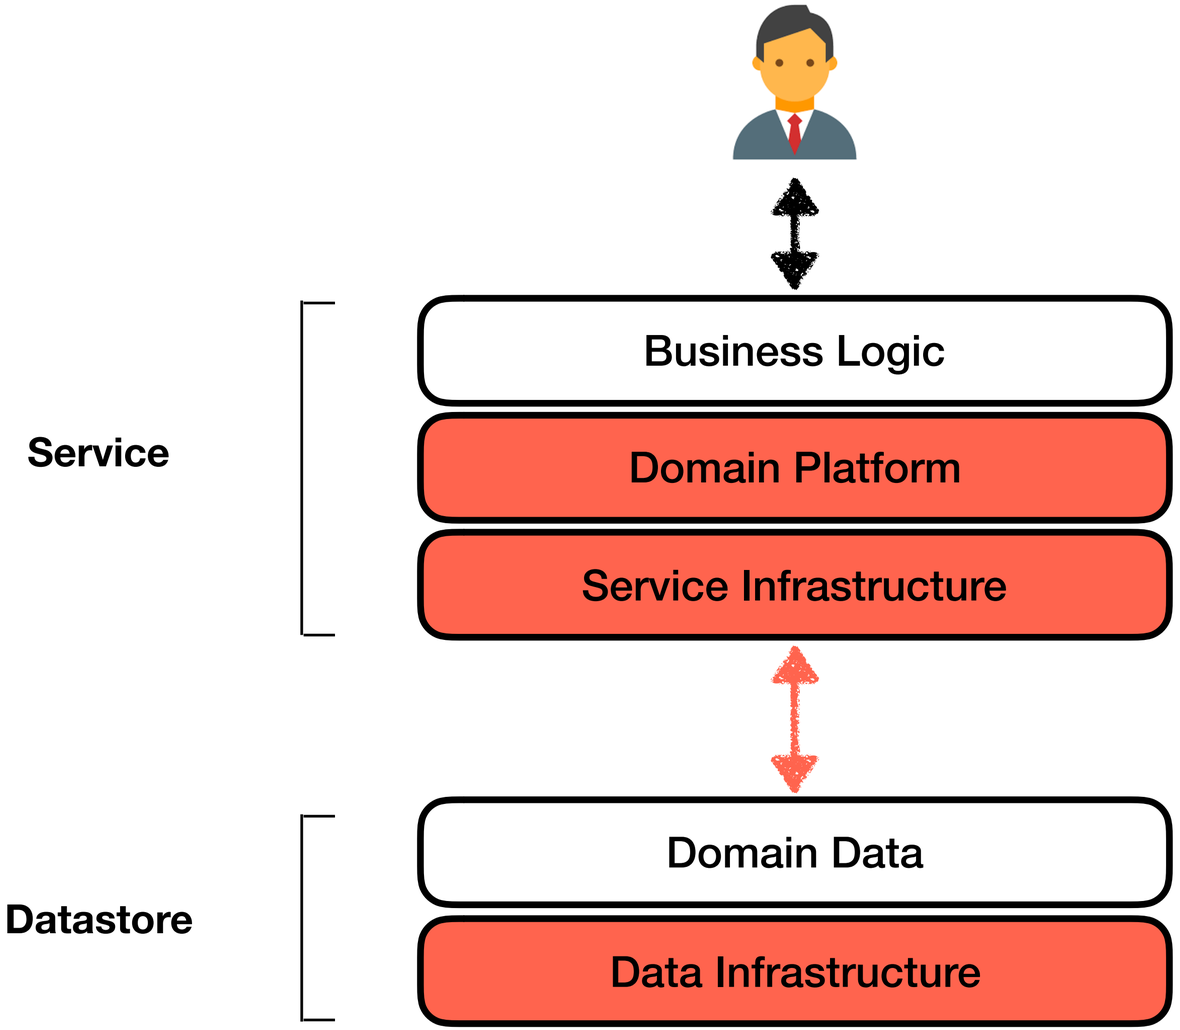

Conflict-free Replicated Data Types

CRDTs

commutative

Requirements:

- + Commutativity
- + Associativity
- + Exactly once delivery
- - Idempotence

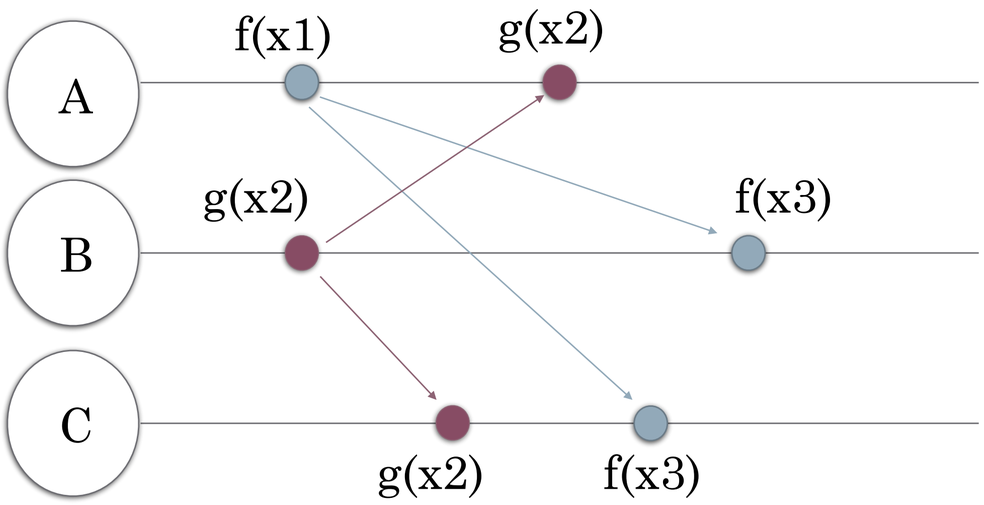

convergent

Requirements:

- + Commutativity
- + Associativity
- + Idempotence
- - Exactly once delivery

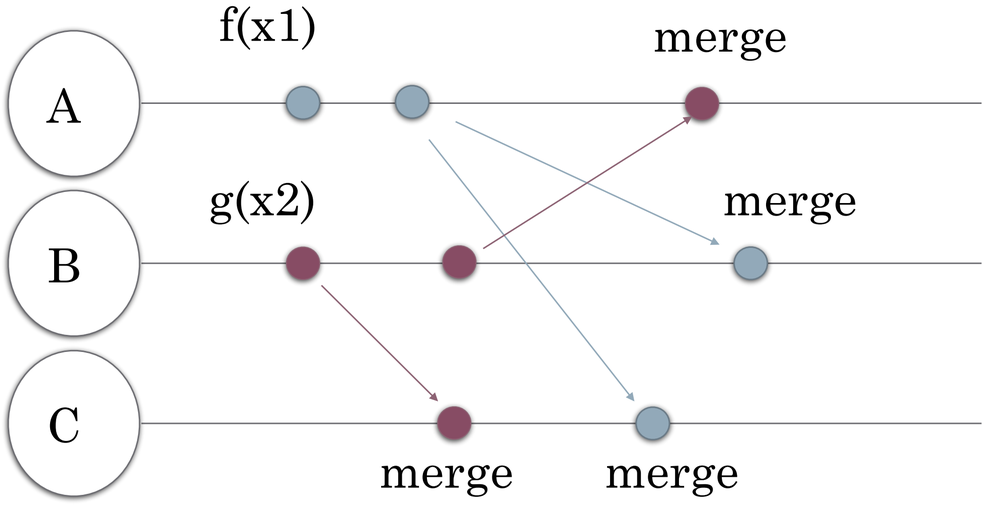

Convergent CRDTs

- M(a, b) = M(b, a)
- M(M(a, b), c) = M(a, M(b, c))
- M(a, b) = M(M(a, b), b) = M(M(M(a, b), b), b)

Impacted Components for CRDTs

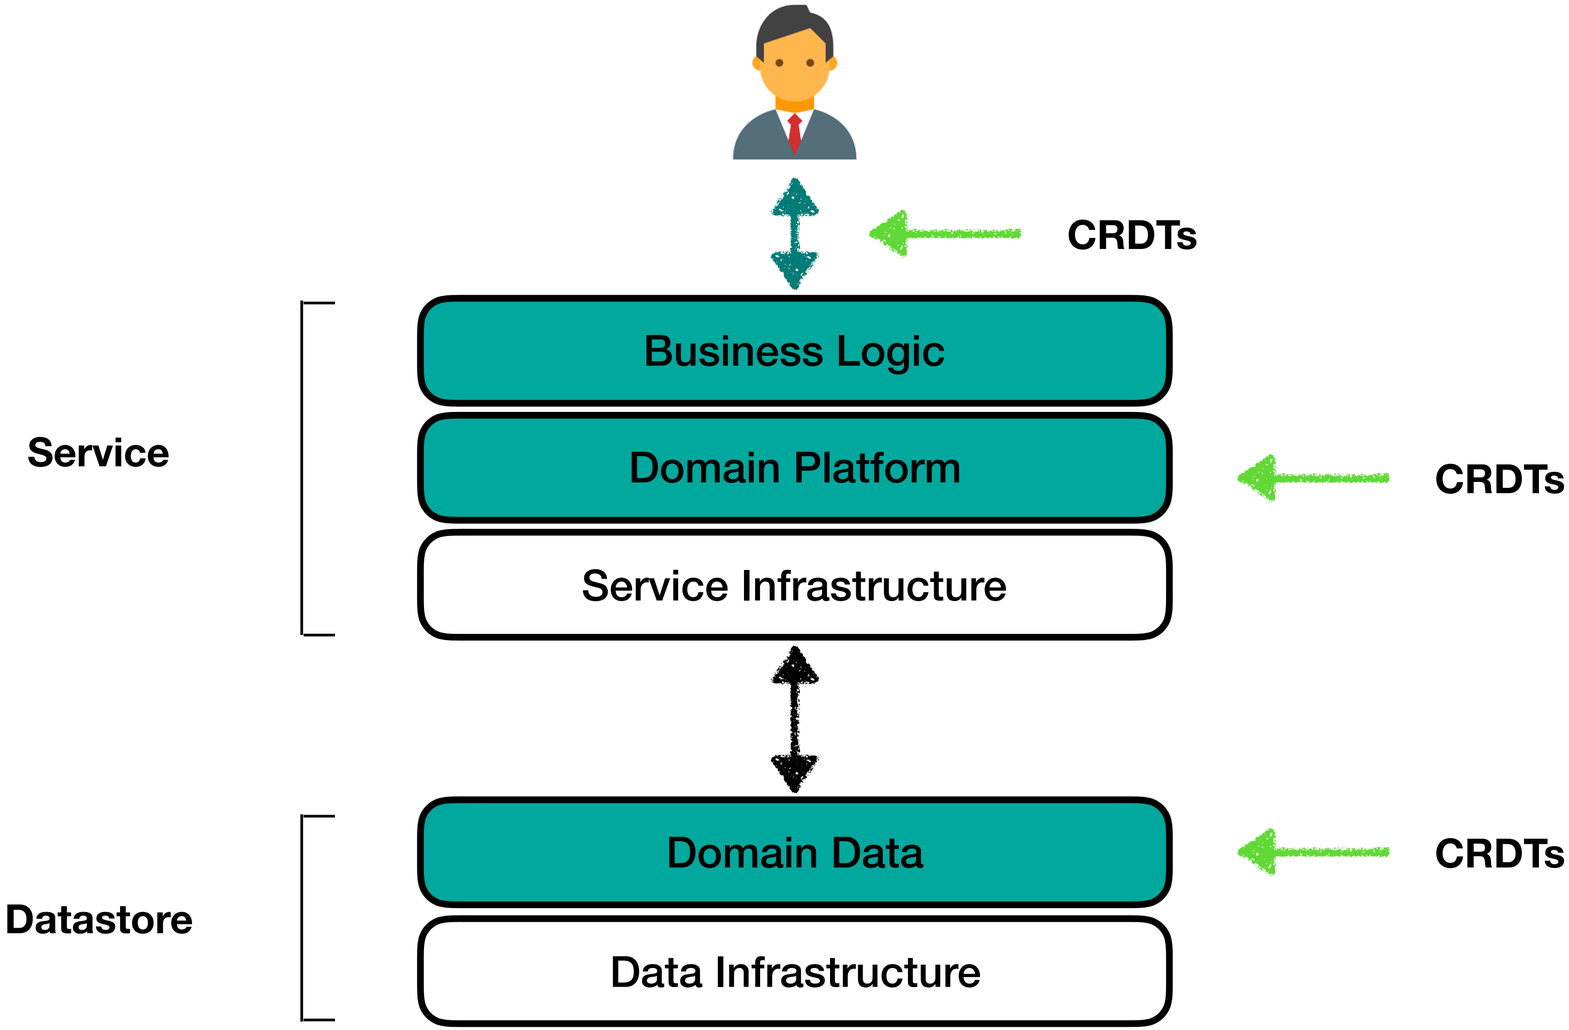

Online Flight Check-in System

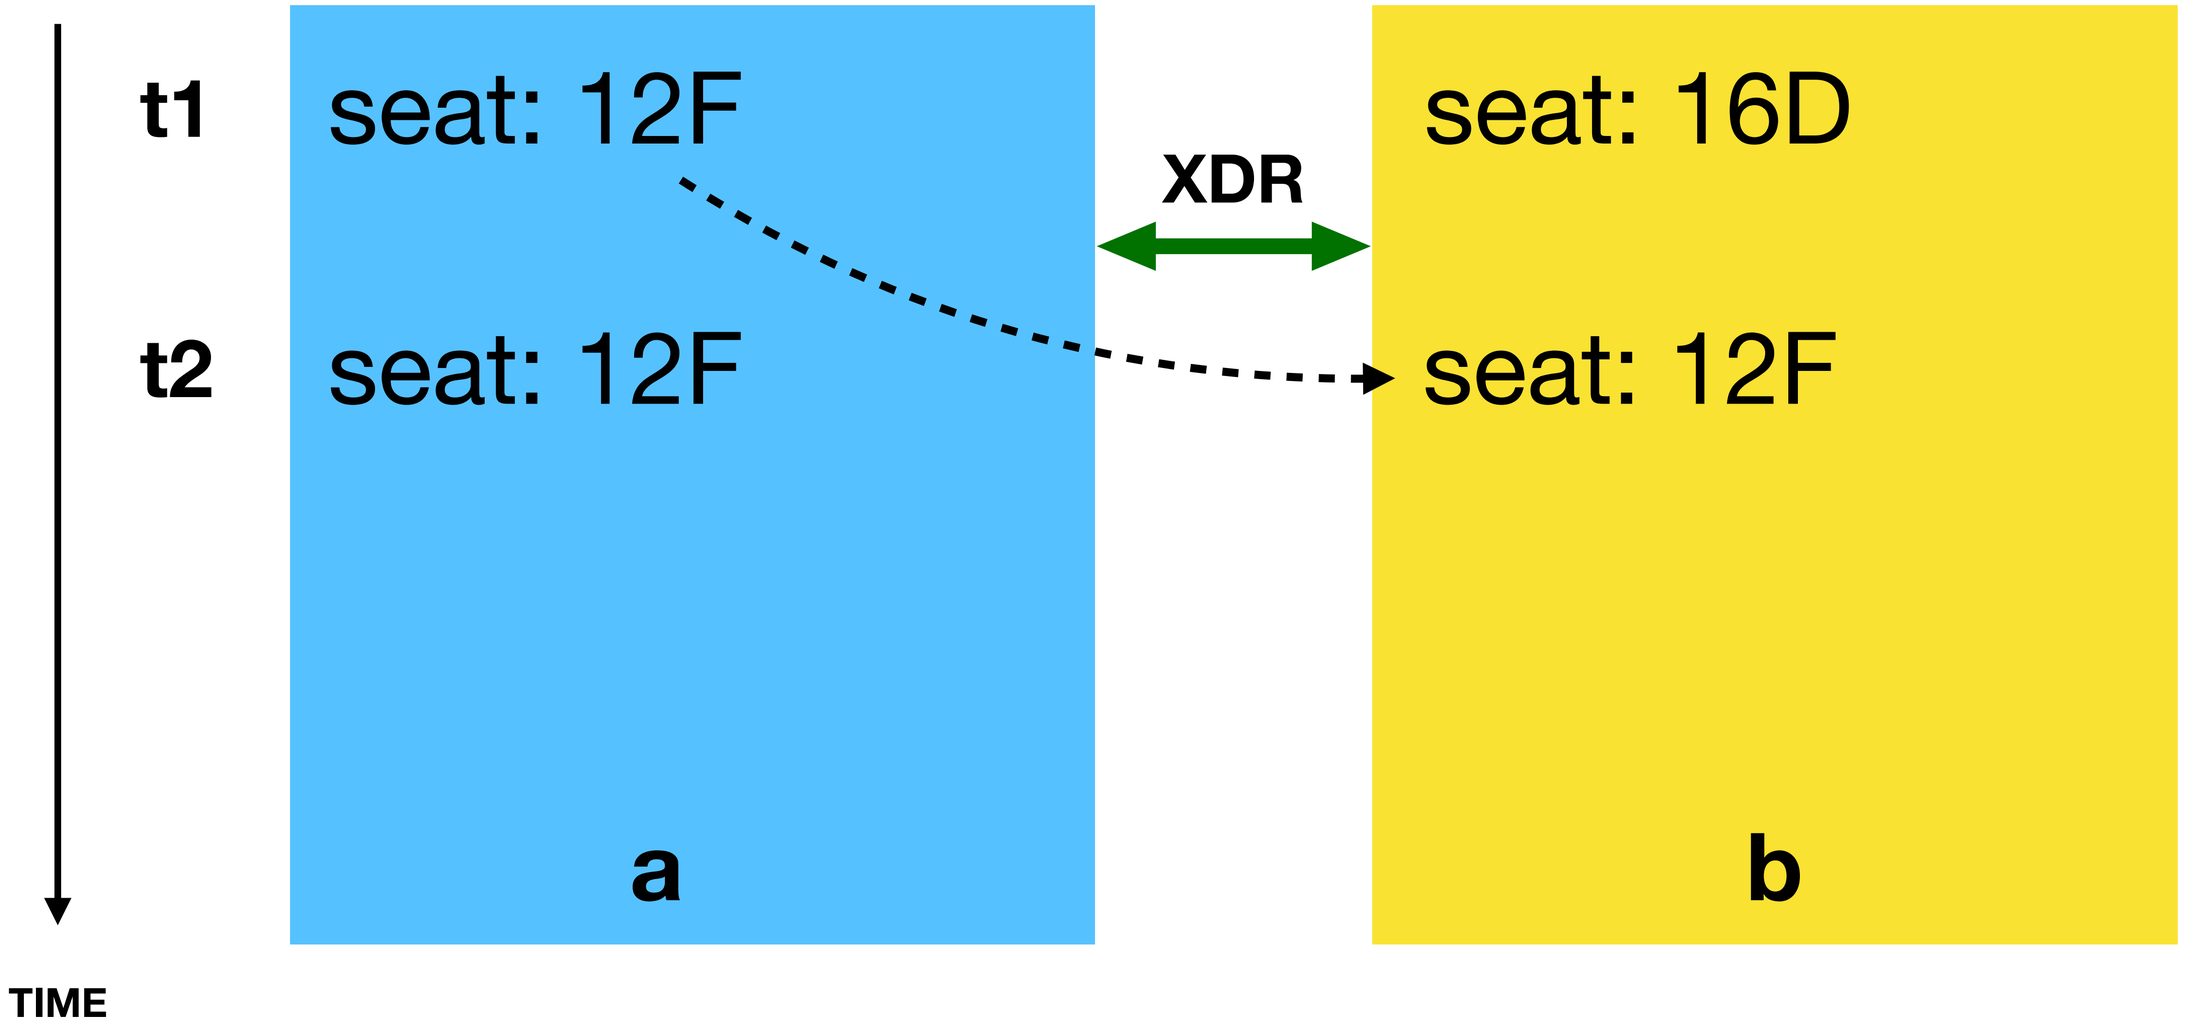

Online Flight Check-in System

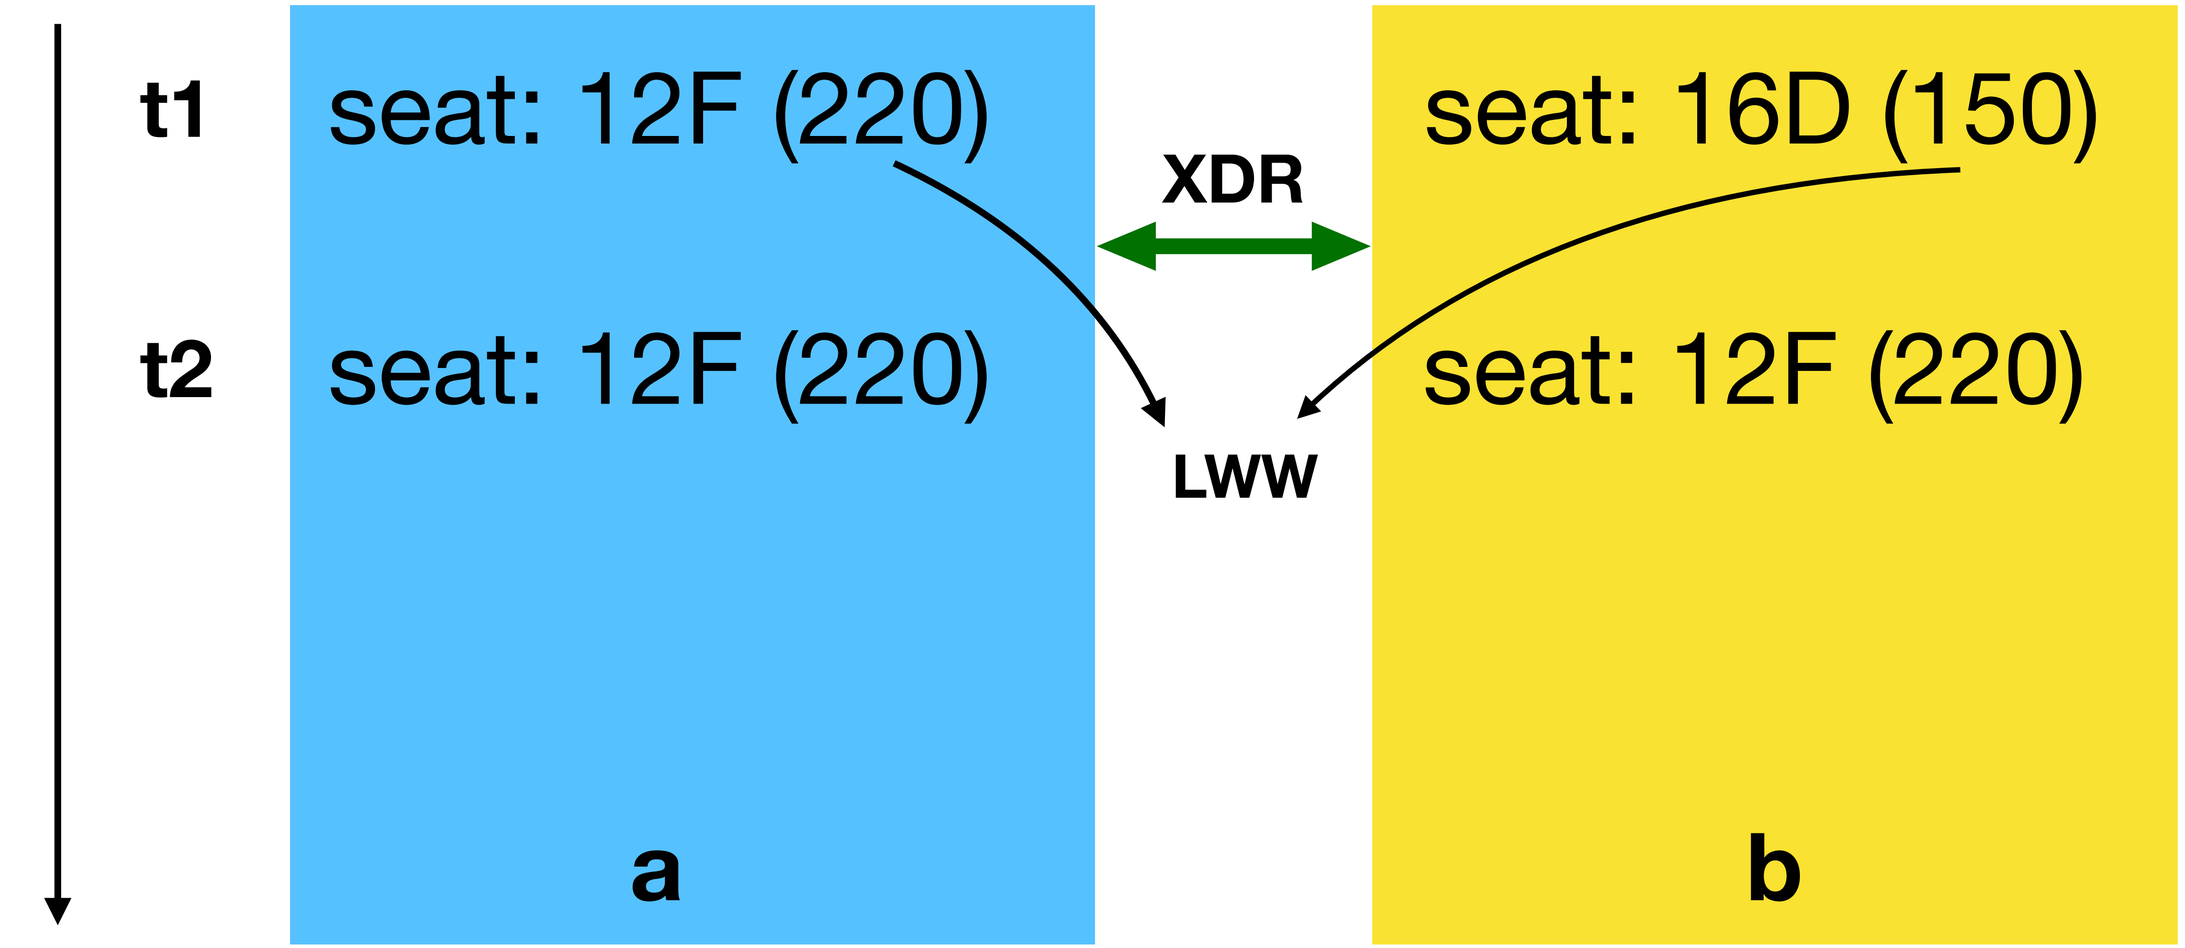

M(a, b) for LWW = MAX(a, b)

Online Flight Check-in System

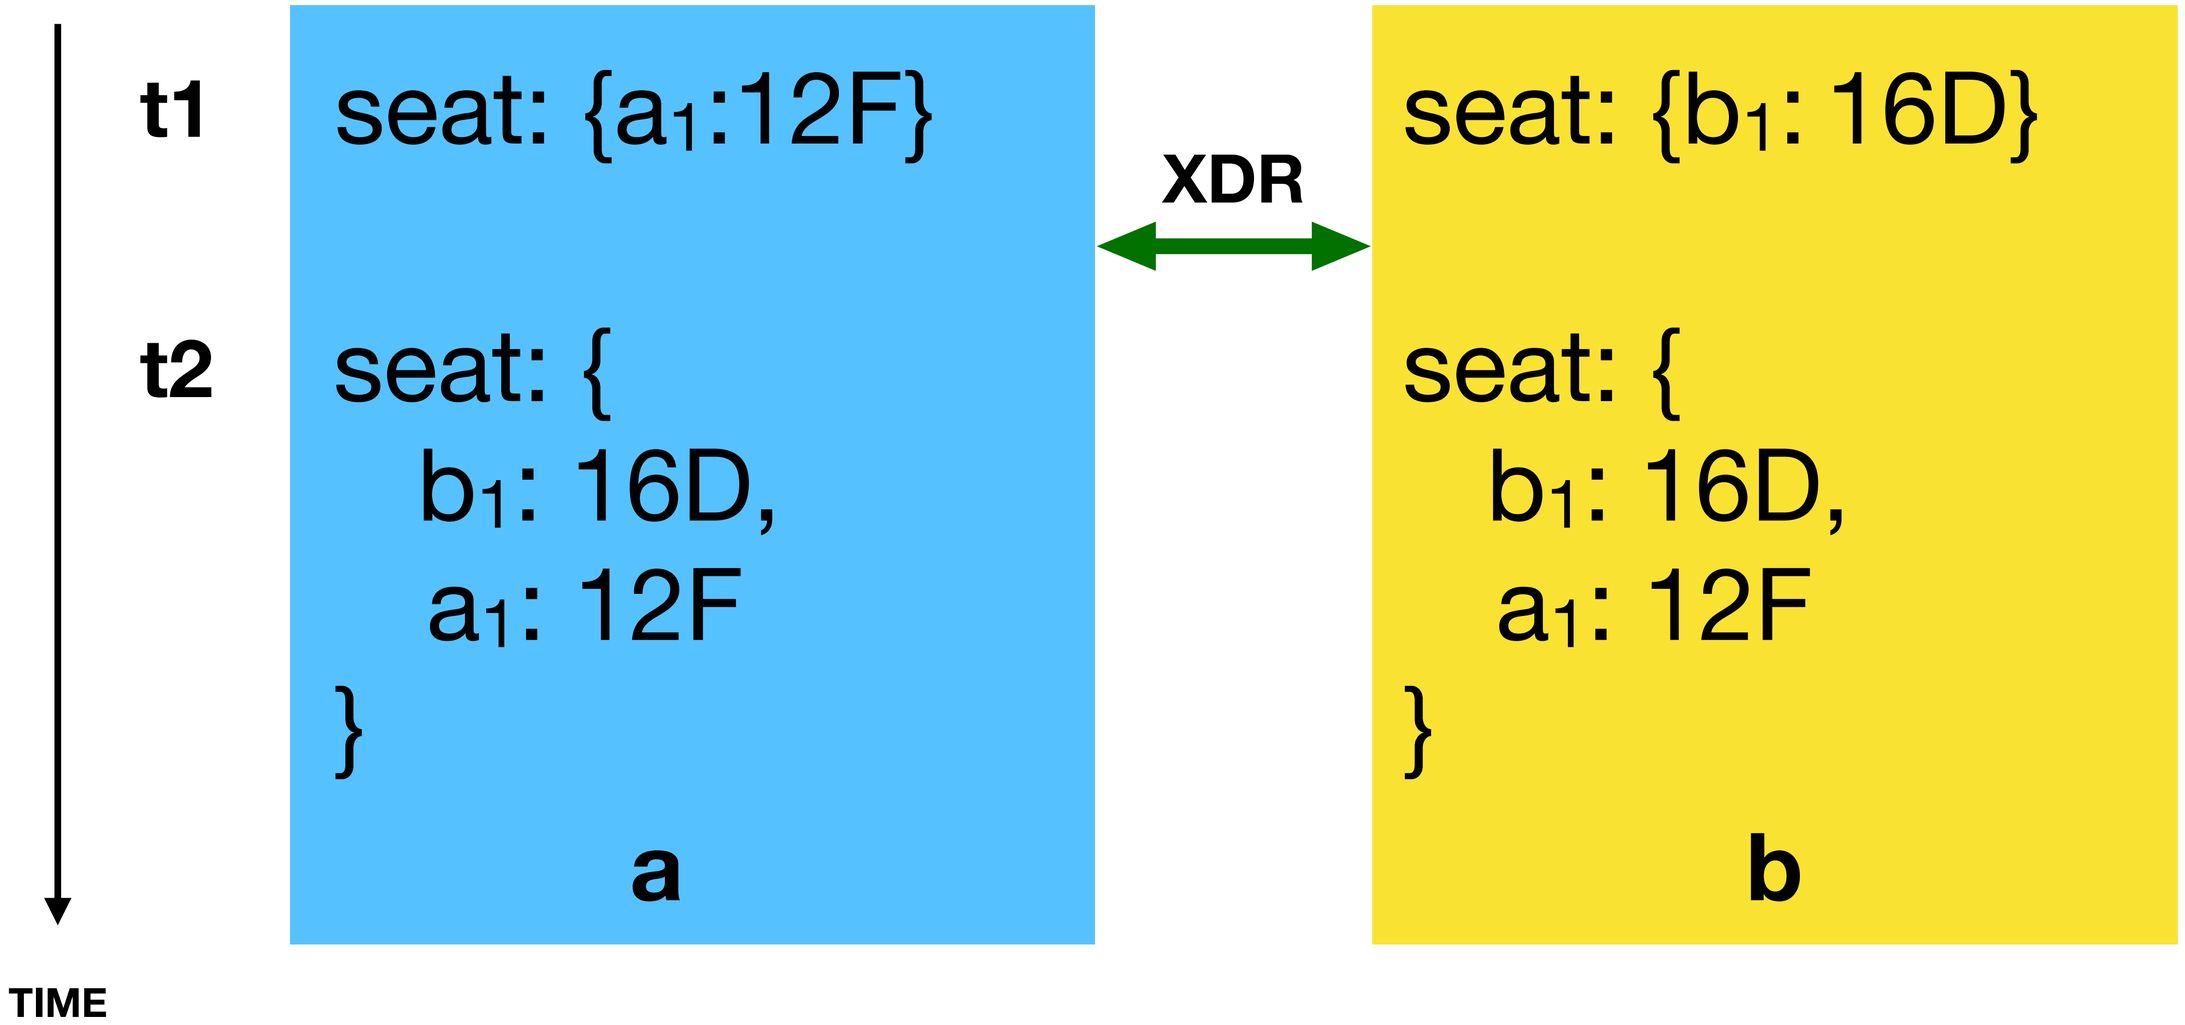

Online Flight Check-in System

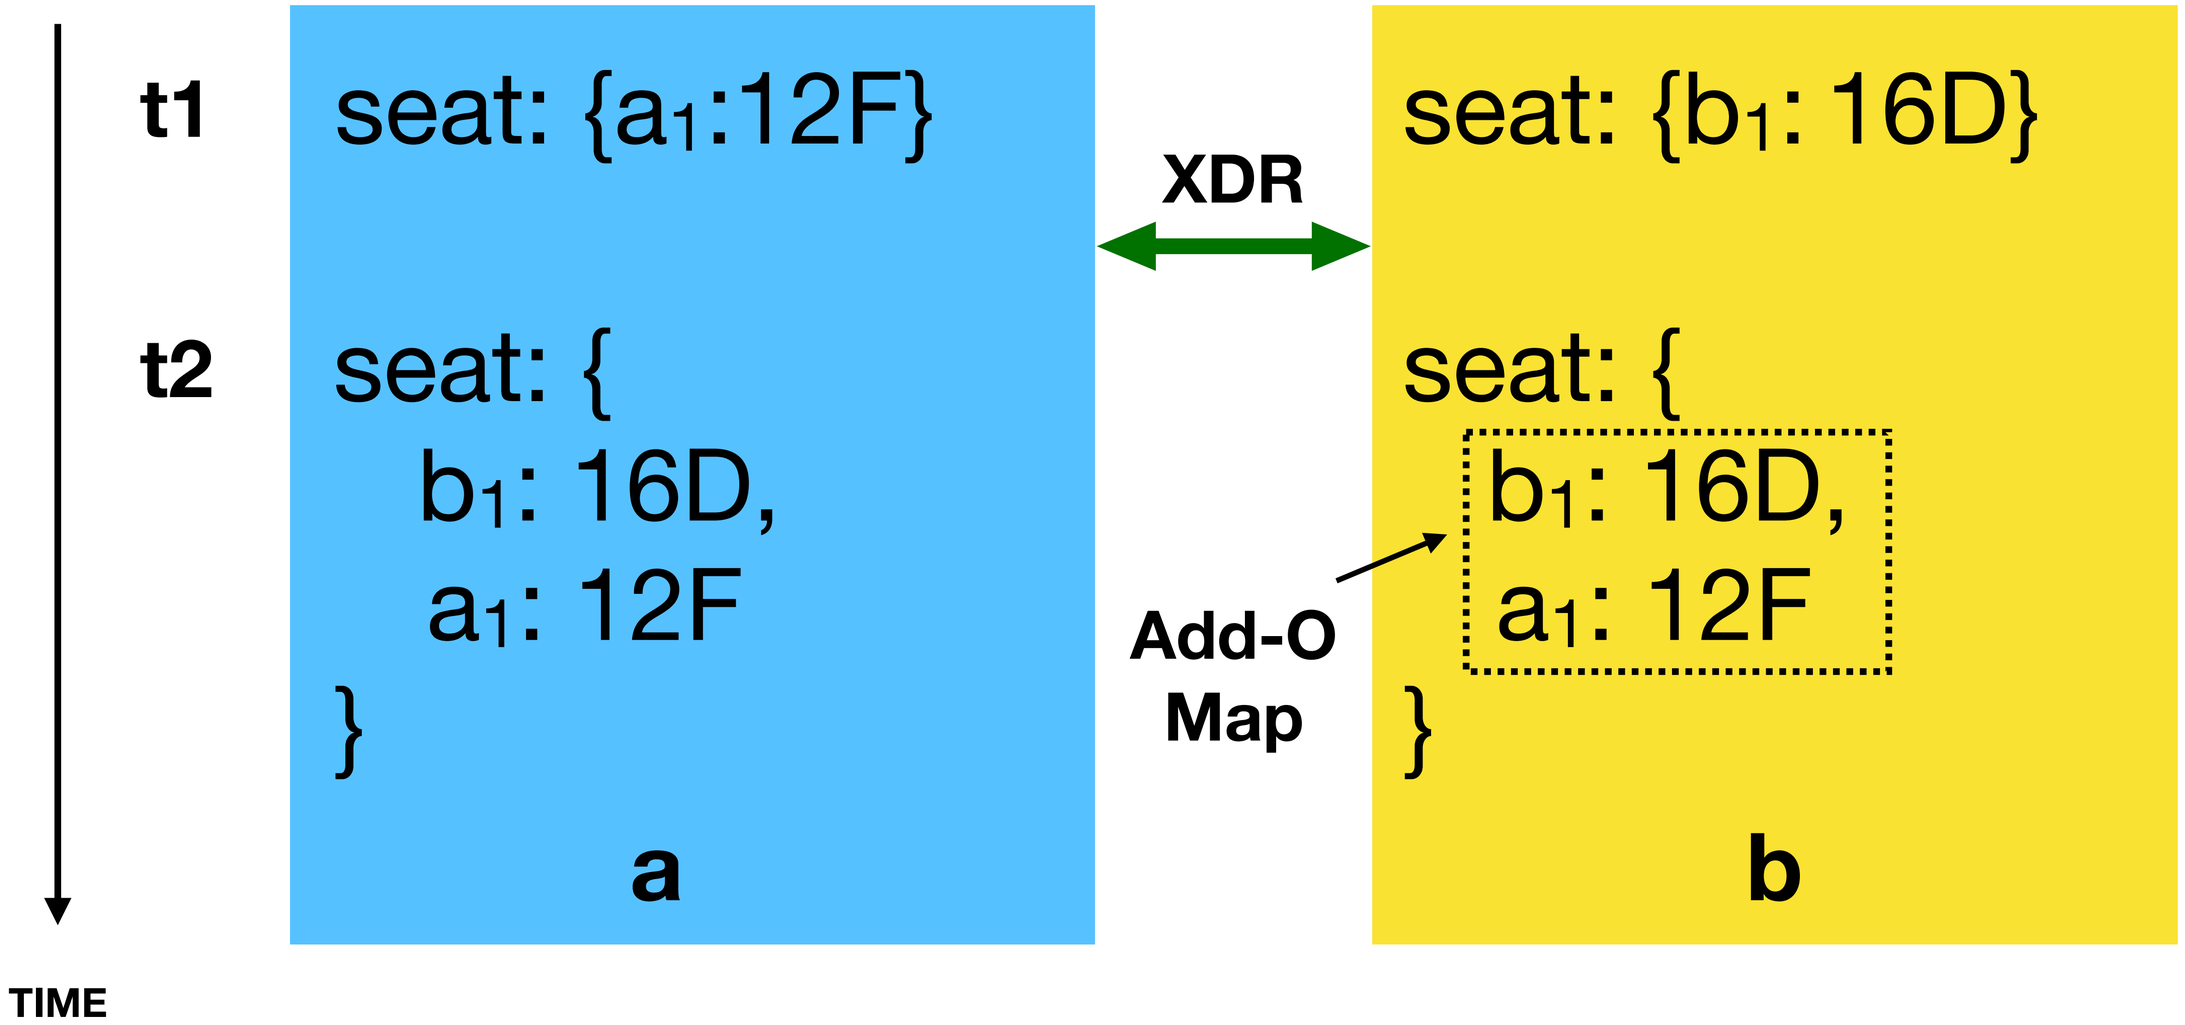

Causality

seat: {

b_1: 16D,

a_1: 12F

}

a1;b1

Causality Vector (cv)

Causality

seat: {

b_1: 16D,

a_1: 12F

}

a1;b1

Causality Vector (cv)

12F is causal to 10A - we can drop 12F

10A is causal to 5C - we can drop 10A

Causality

seat: {

b_1: 16D,

a_1: 12F

}

a1;b1

Causality Vector (cv)

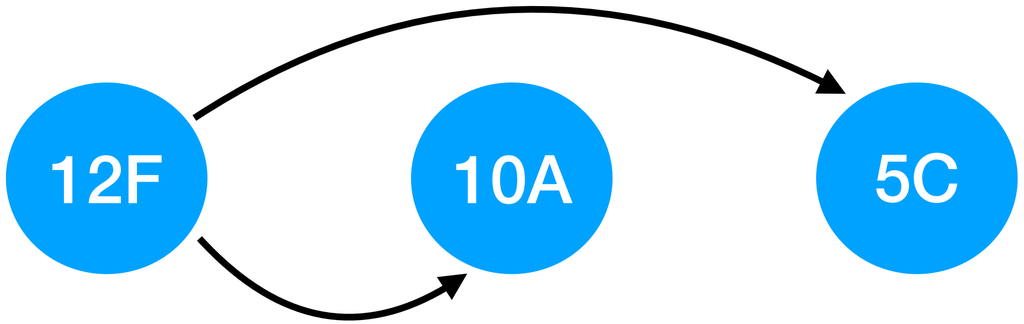

12F is causal to 10A - we can drop 12F

10A is **NOT** causal to 5C - we can **NOT** drop 10A

Causality

seat: {

b_1: 16D,

a_1: 12F

}

a1;b1

Causality Vector (cv)

Client Operations:

GET(key): value => GET(key): (value, **cv**)

PUT(key, value) => PUT(key, value, **cv**)

Causality

seat: {

b_1: (16D, **cv**),

a_1: (12F, **cv**)

}

a1;b1

Causality Vector (cv)

Client Operations:

GET(key): value => GET(key): (value, **cv**)

PUT(key, value) => PUT(key, value, **cv**)

Aerospike Datastore

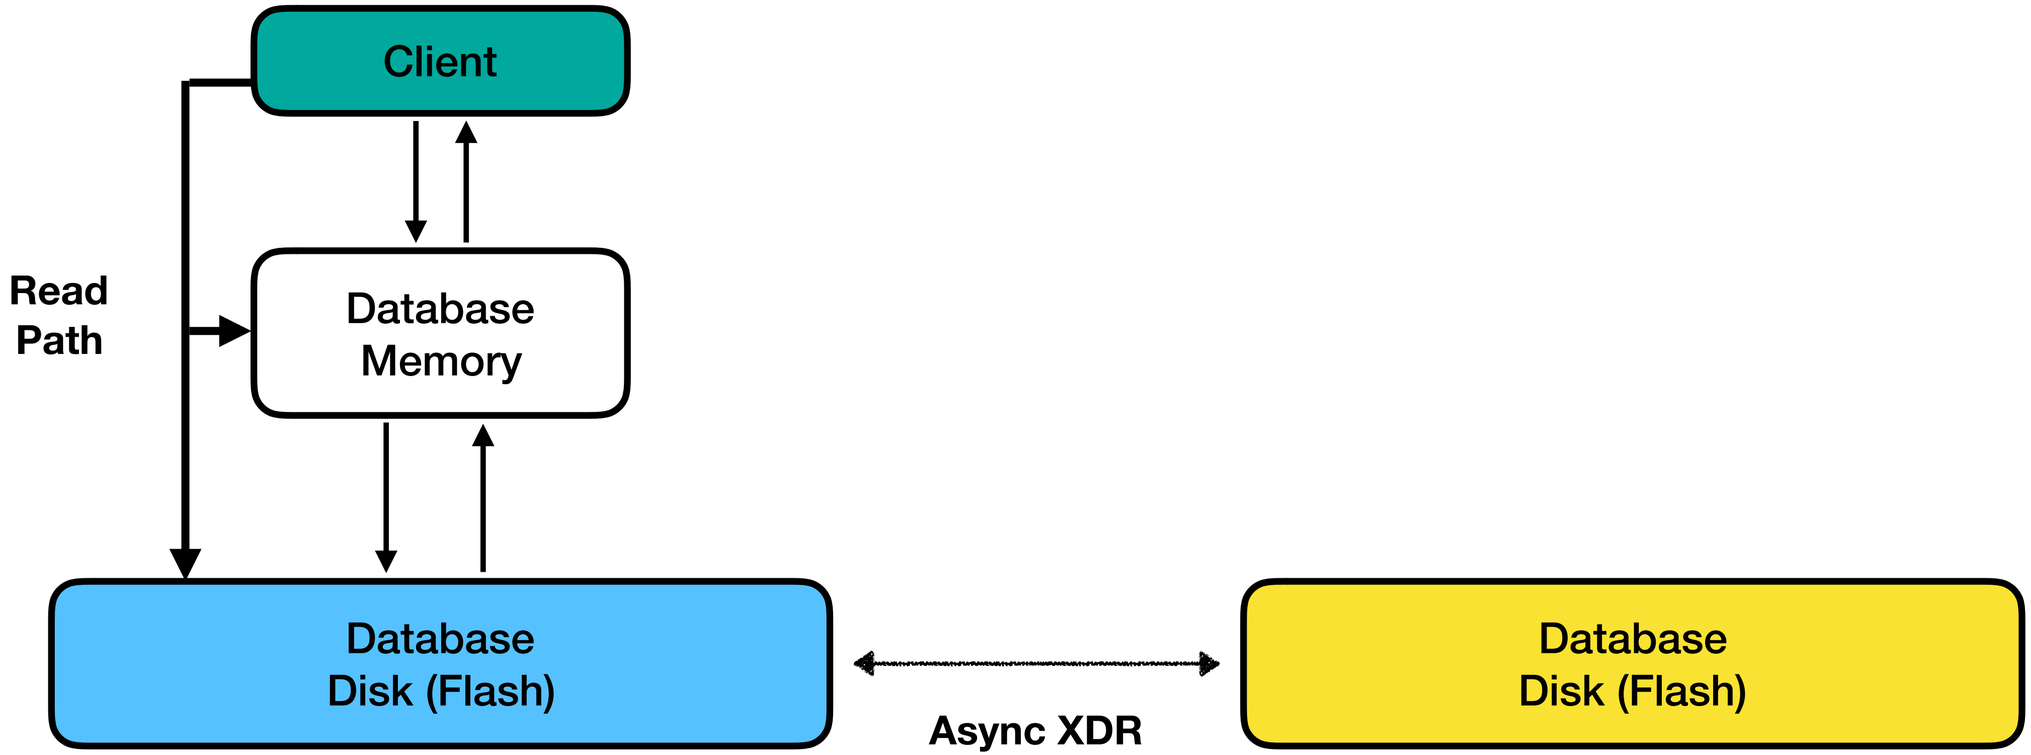

Aerospike Datastore

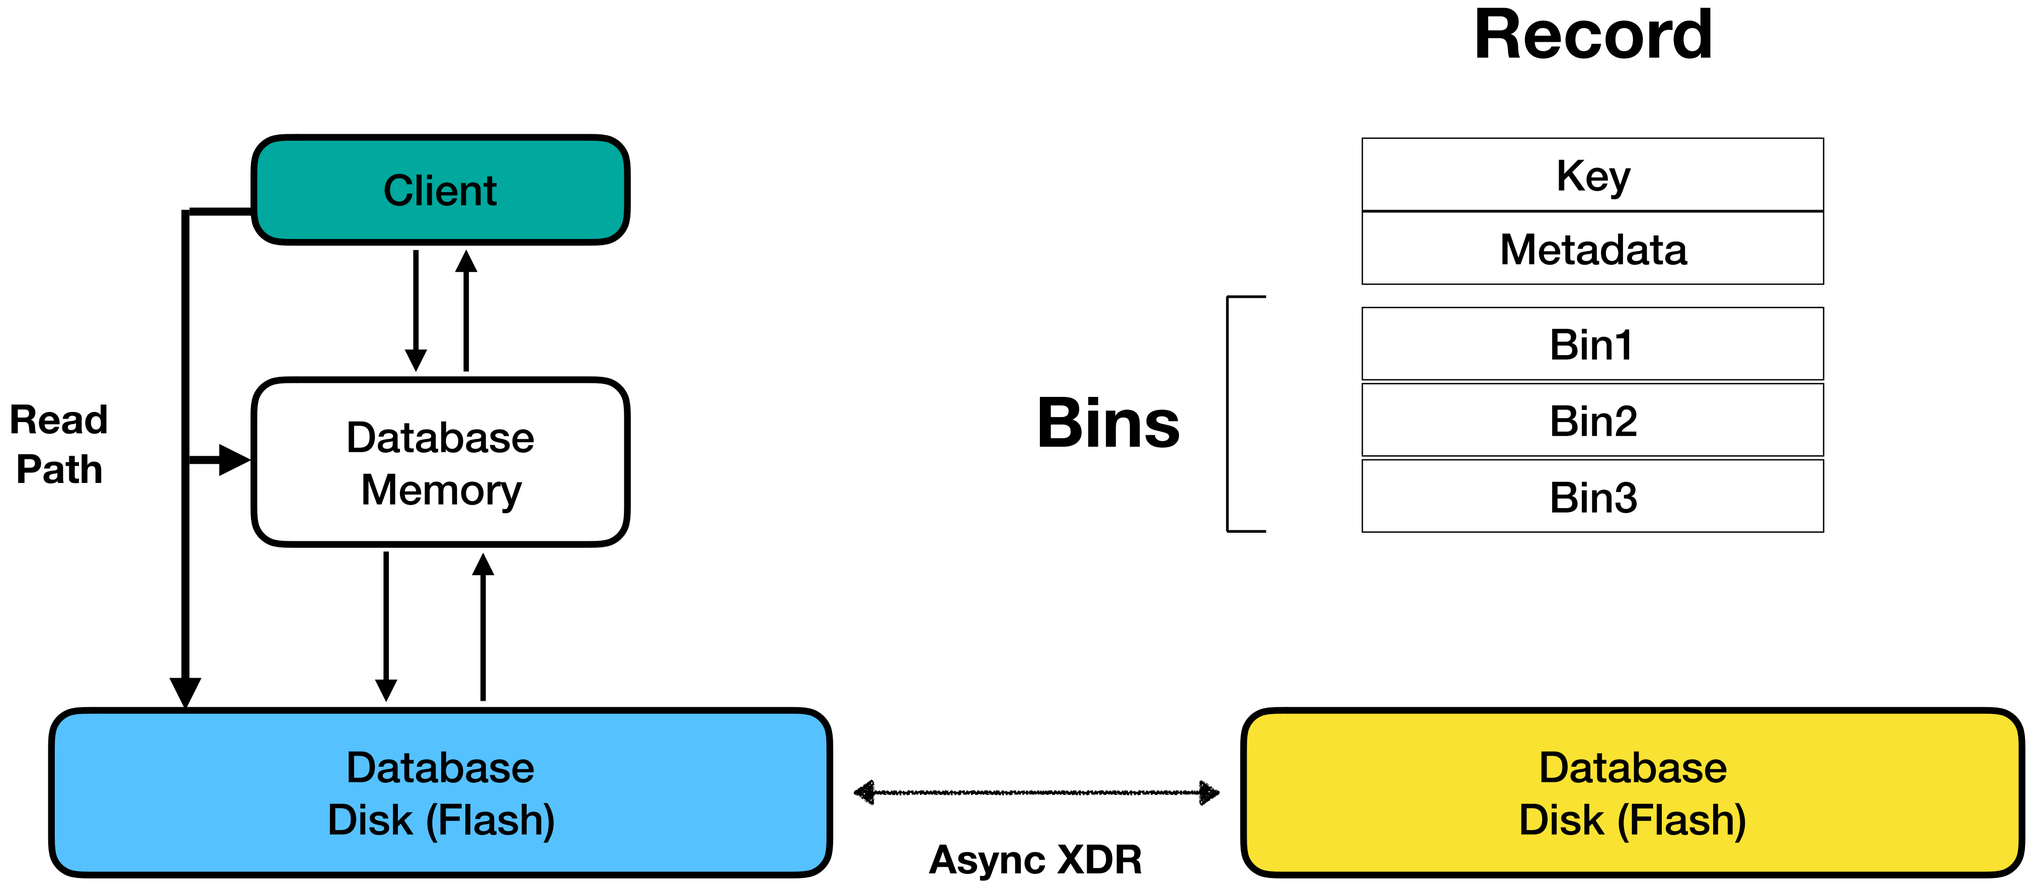

AEROSPIKE

1 User-Defined Functions

2

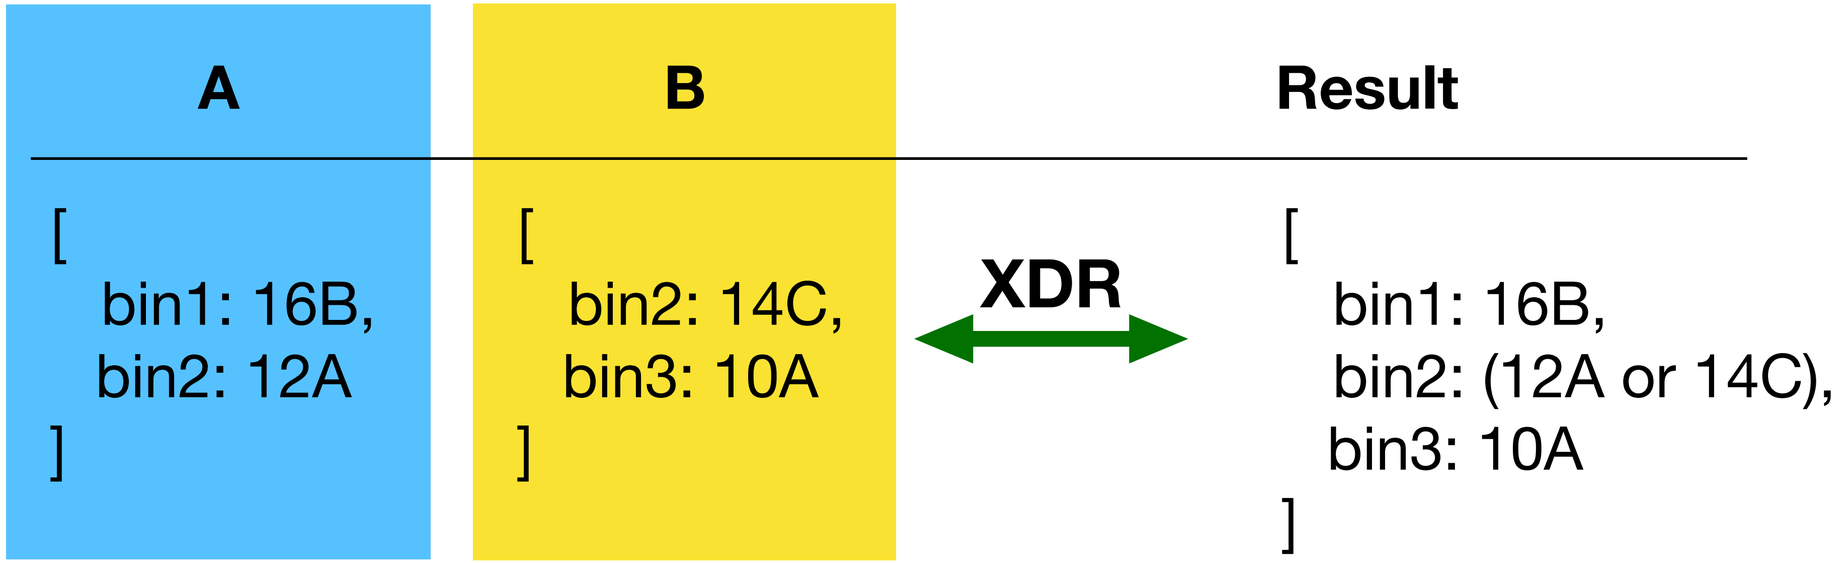

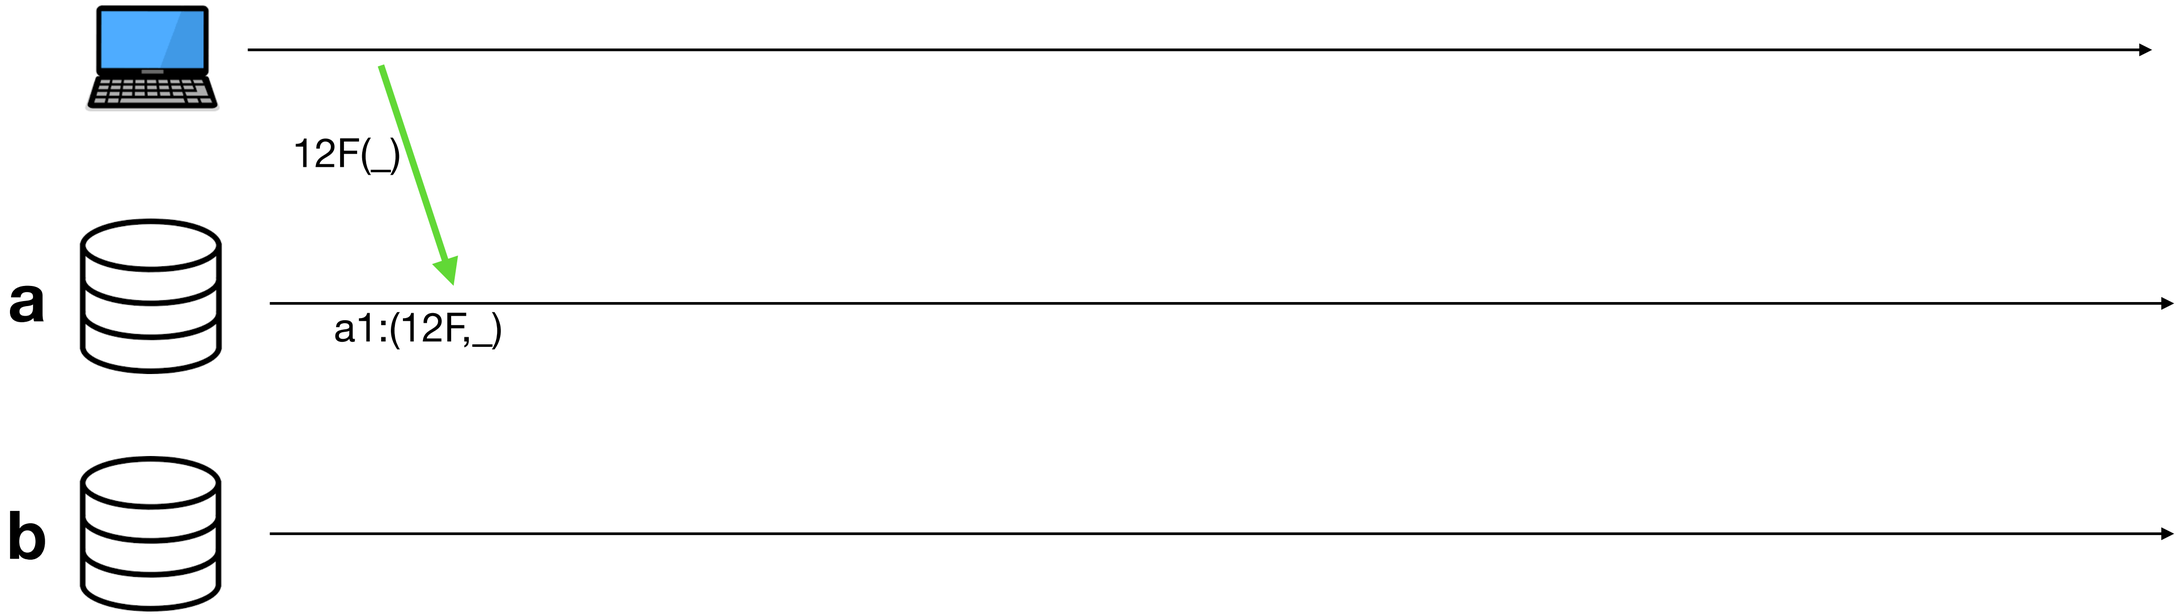

Bins	1
a1	(12F, _)

a

Bins	1

b

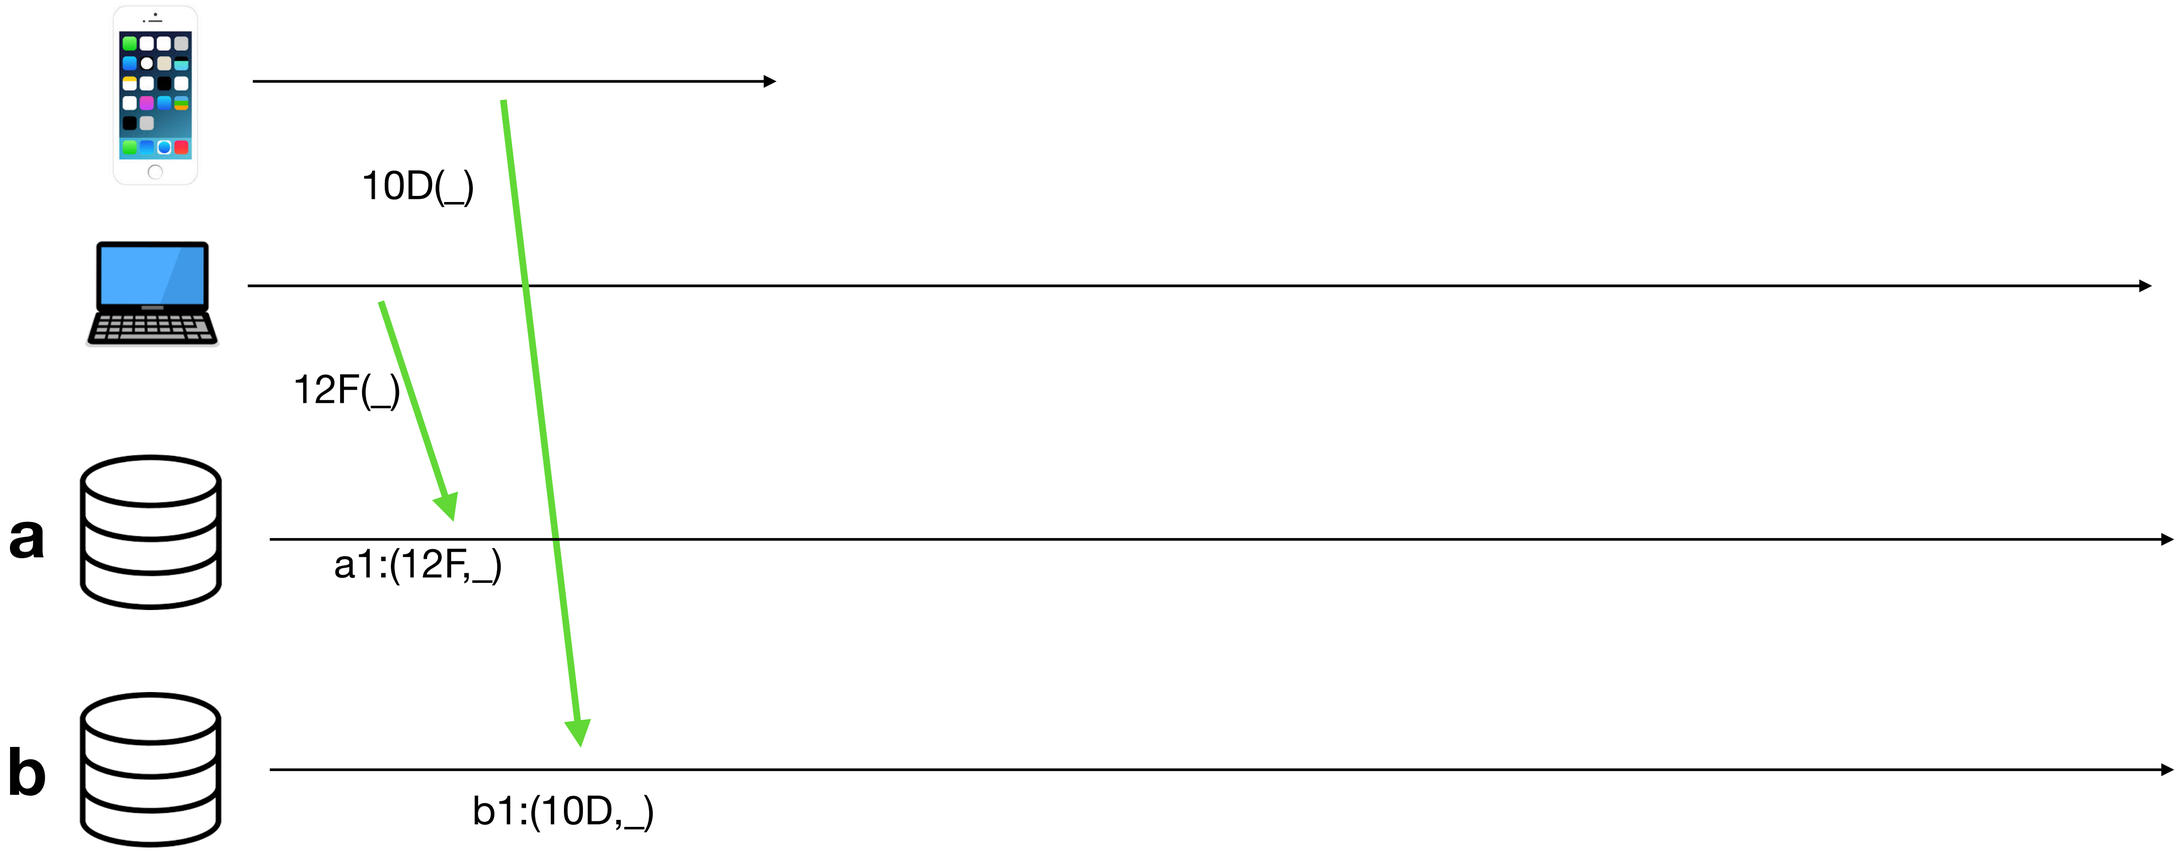

Bins	1	2
a1	(12F, _)	(12F, _)

a

Bins	1	2
b1		(10D, _)

b

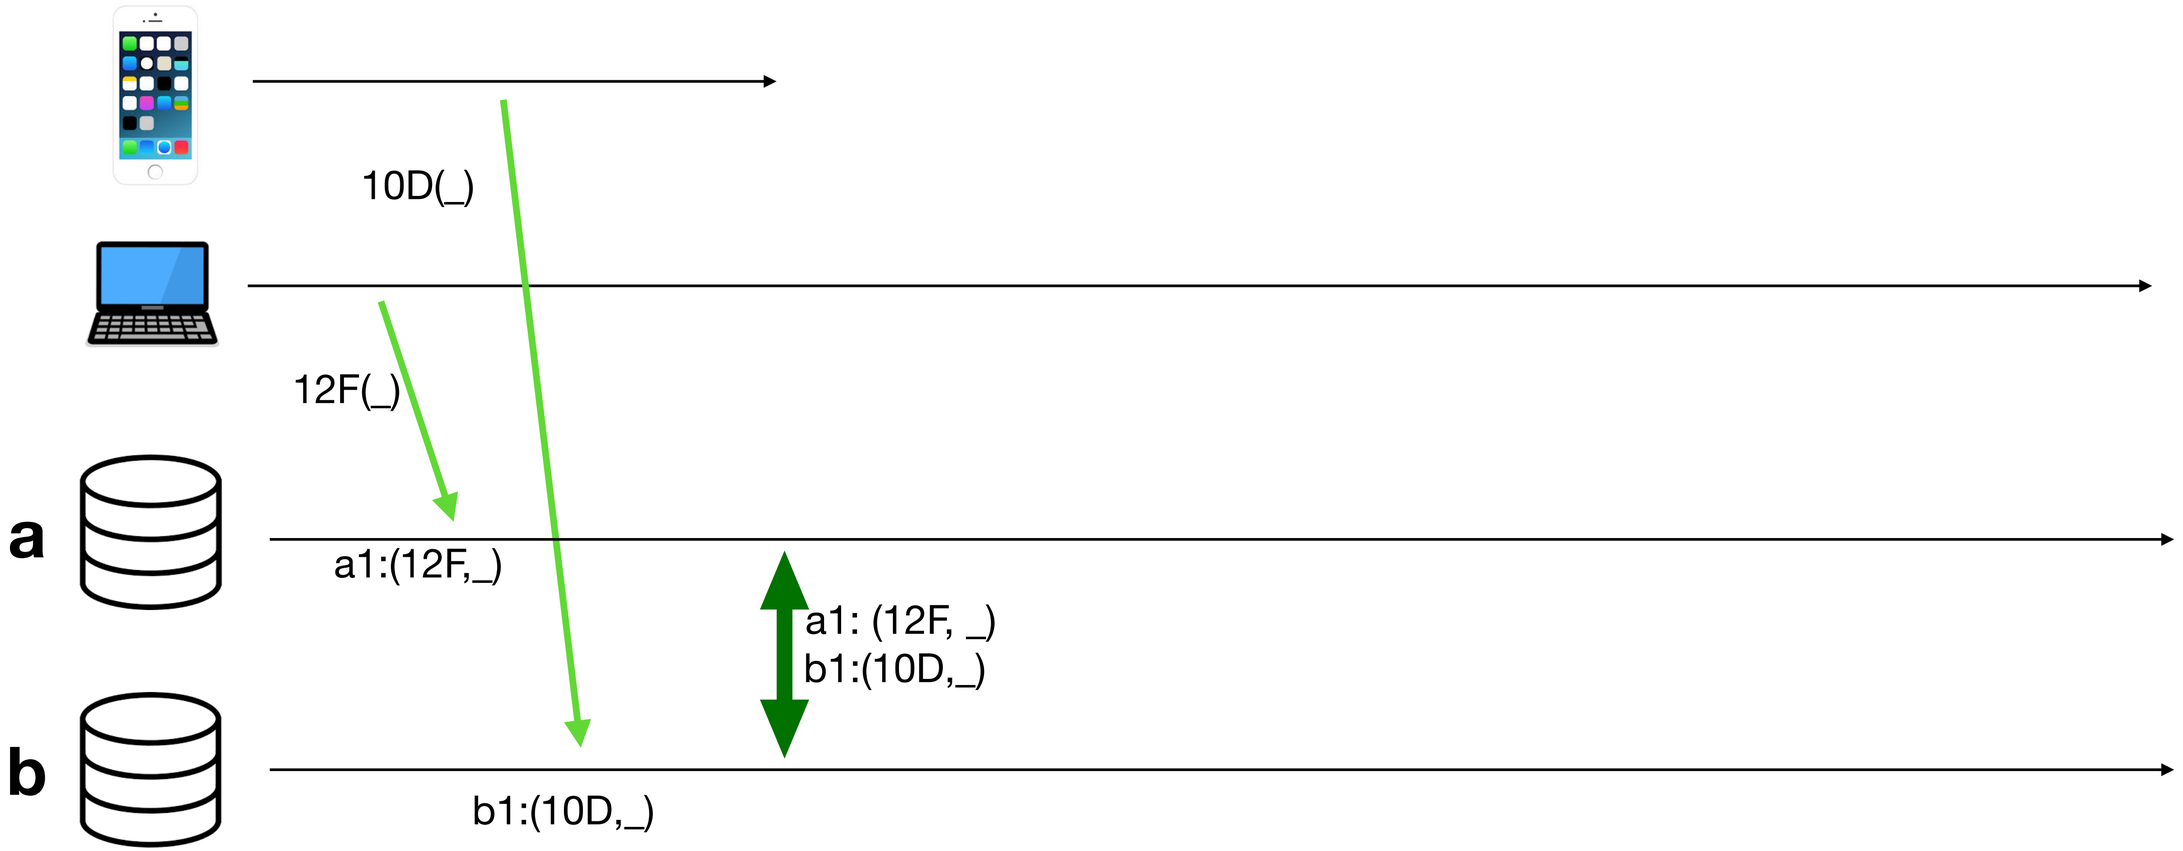

Bins	1	2	3
a1	(12F, _)	(12F, _)	(12F, _)
b1			(10D, _)

a

Bins	1	2	3
a1			(12F, _)
b1		(10D, _)	(10D, _)

b

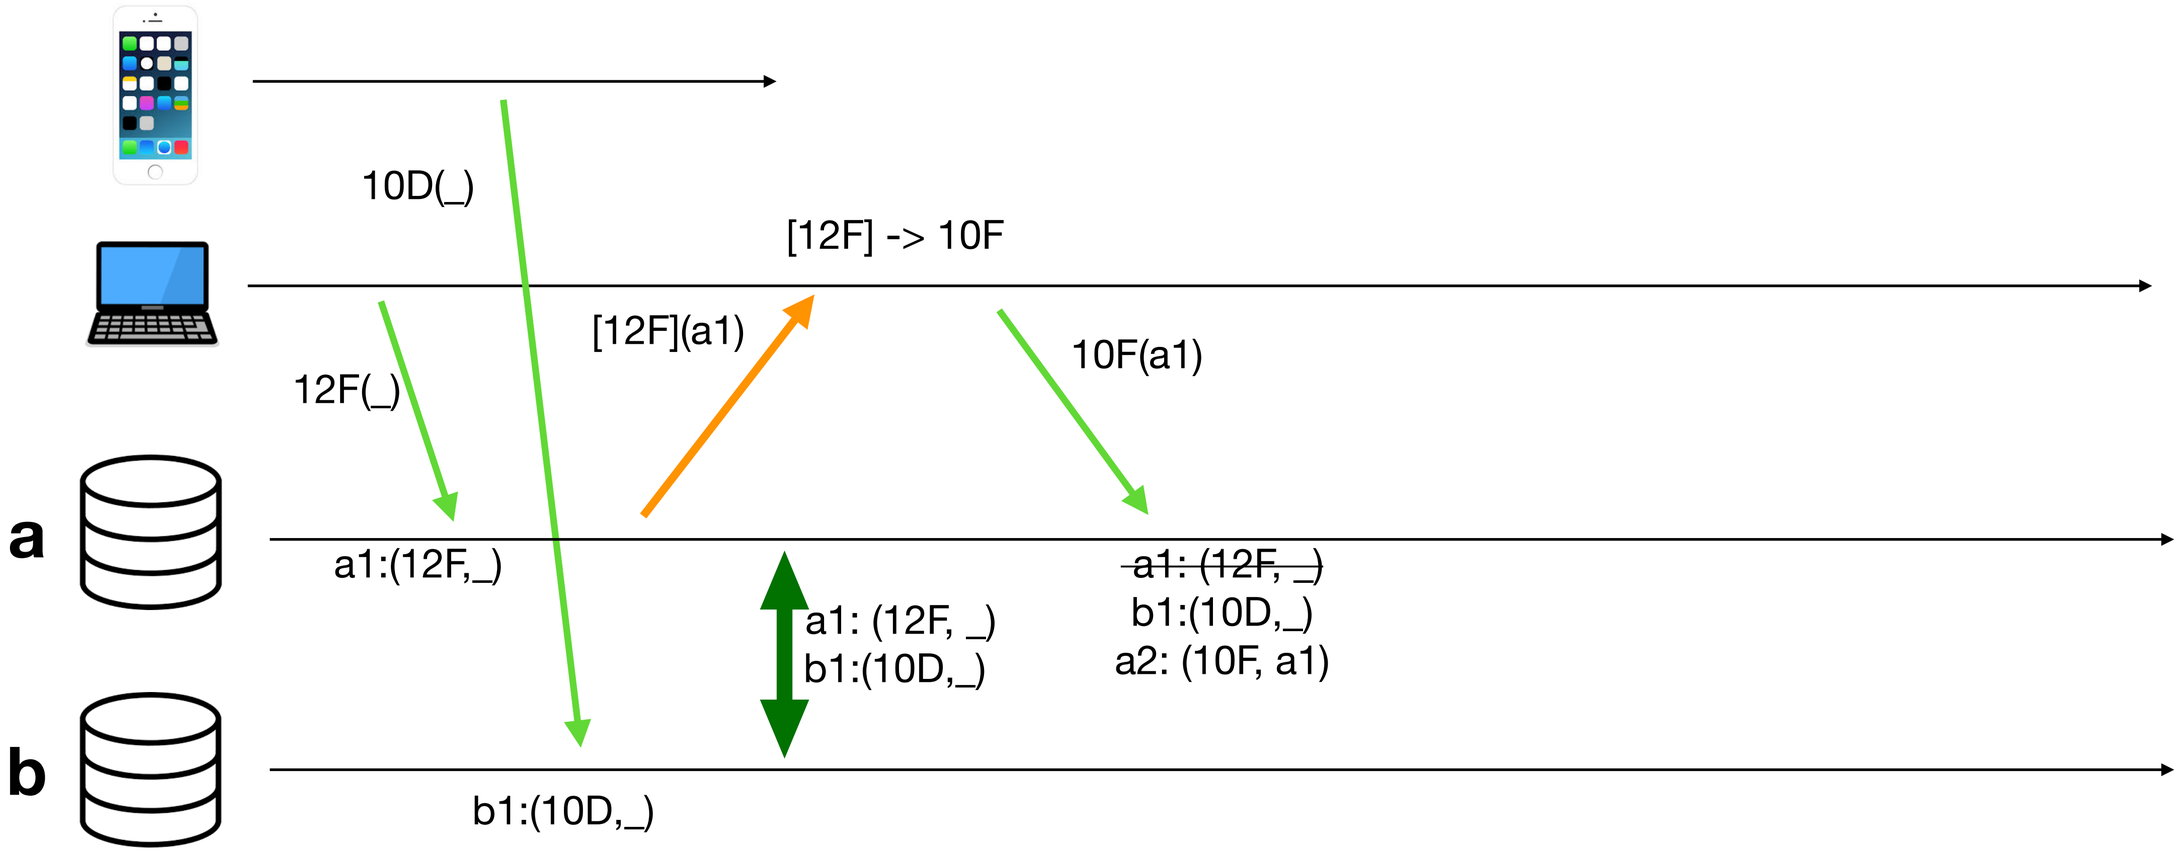

Bins	1	2	3	4
a1	(12F, _)	(12F, _)	(12F, _)	(12F, _)
b1			(10D, _)	(10D, _)
a2				(10F, a1)

a

Bins	1	2	3	4
a1			(12F, _)	(12F, _)
b1		(10D, _)	(10D, _)	(10D, _)

b

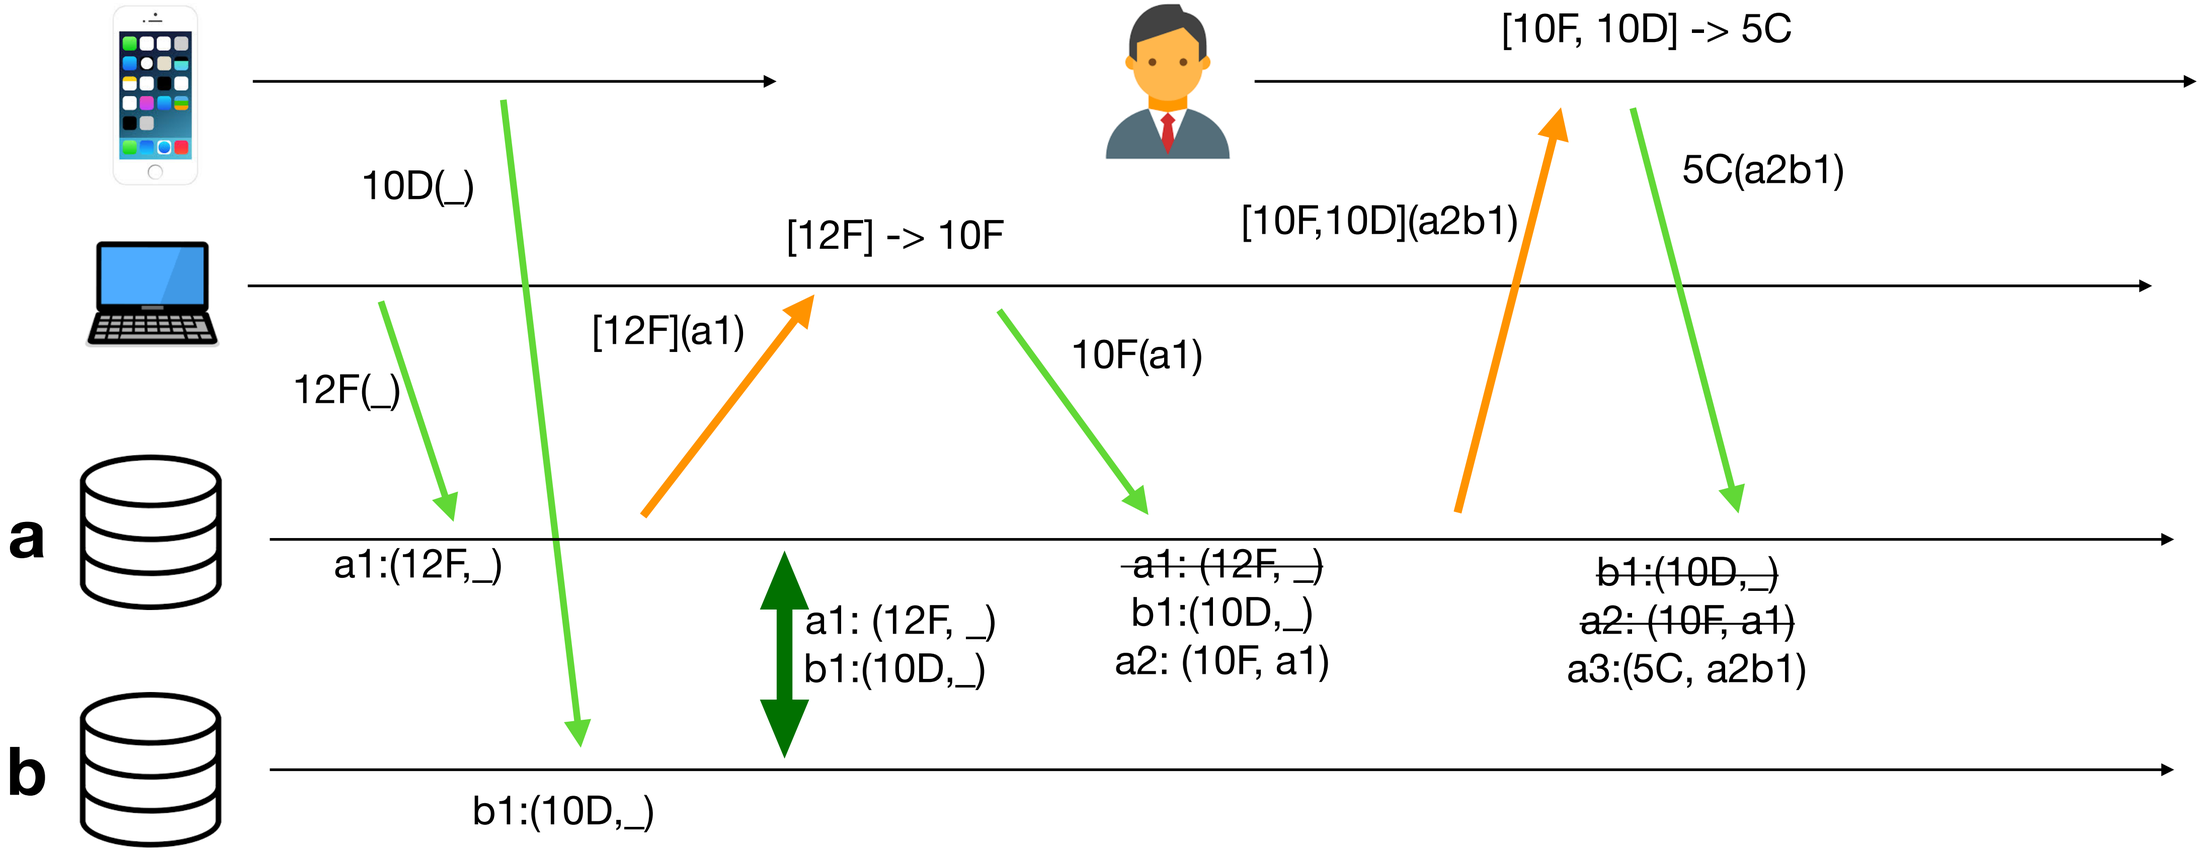

Bins	1	2	3	4	5
a1	(12F, _)	(12F, _)	(12F, _)	(12F, _)	(12F, _)
b1			(10D, _)	(10D, _)	(10D, _)
a2				(10F, a1)	(10F, a1)
a3					(5C, a2b1)

a

Bins	1	2	3	4	5
a1			(12F, _)	(12F, _)	(12F, _)
b1		(10D, _)	(10D, _)	(10D, _)	(10D, _)

b

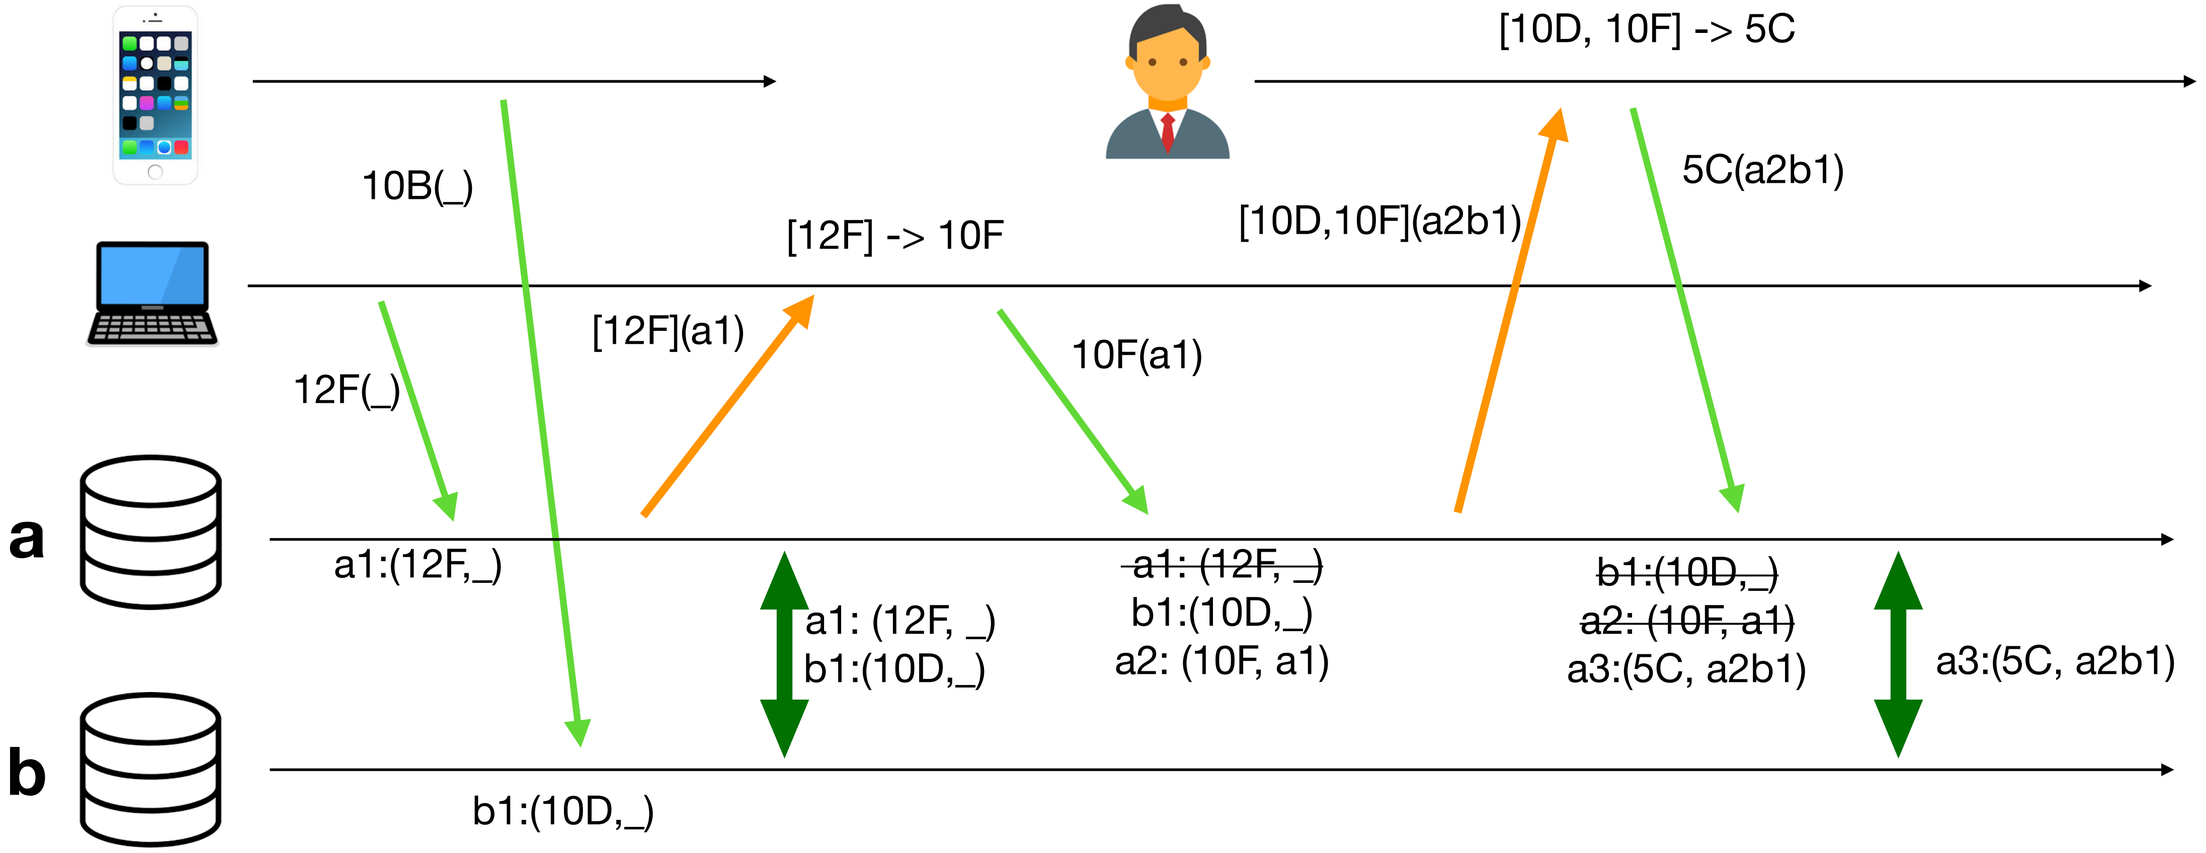

Bins	1	2	3	4	5	6
a1	(12F, _)	(12F, _)	(12F, _)	(12F, _)	(12F, _)	(12F, _)
b1			(10D, _)	(10D, _)	(10D, _)	(10D, _)
a2				(10F, a1)	(10F, a1)	(10F, a1)
a3					(5C, a2b1)	(5C, a2b1)

a

Bins	1	2	3	4	5	6
a1			(12F, _)	(12F, _)	(12F, _)	(12F, _)
b1		(10D, _)	(10D, _)	(10D, _)	(10D, _)	(10D, _)
a3						(5C, a2b1)

b

Learnings

- CRDTs allowed us to achieve convergent **predictable** state of our data

Learnings

- CRDTs allowed us to achieve convergent **predictable** state of our data
- Education about right trade-off between **Consistency** and **Correctness**

Learnings

- CRDTs allowed us to achieve convergent **predictable** state of our data
- Education about right trade-off between **Consistency** and **Correctness**
- Do not **underestimate** concurrent data access

Caveat #1: CV Propagation

1

2

Caveat #1: CV Propagation

1

2

Caveat #2: Siblings Explosion

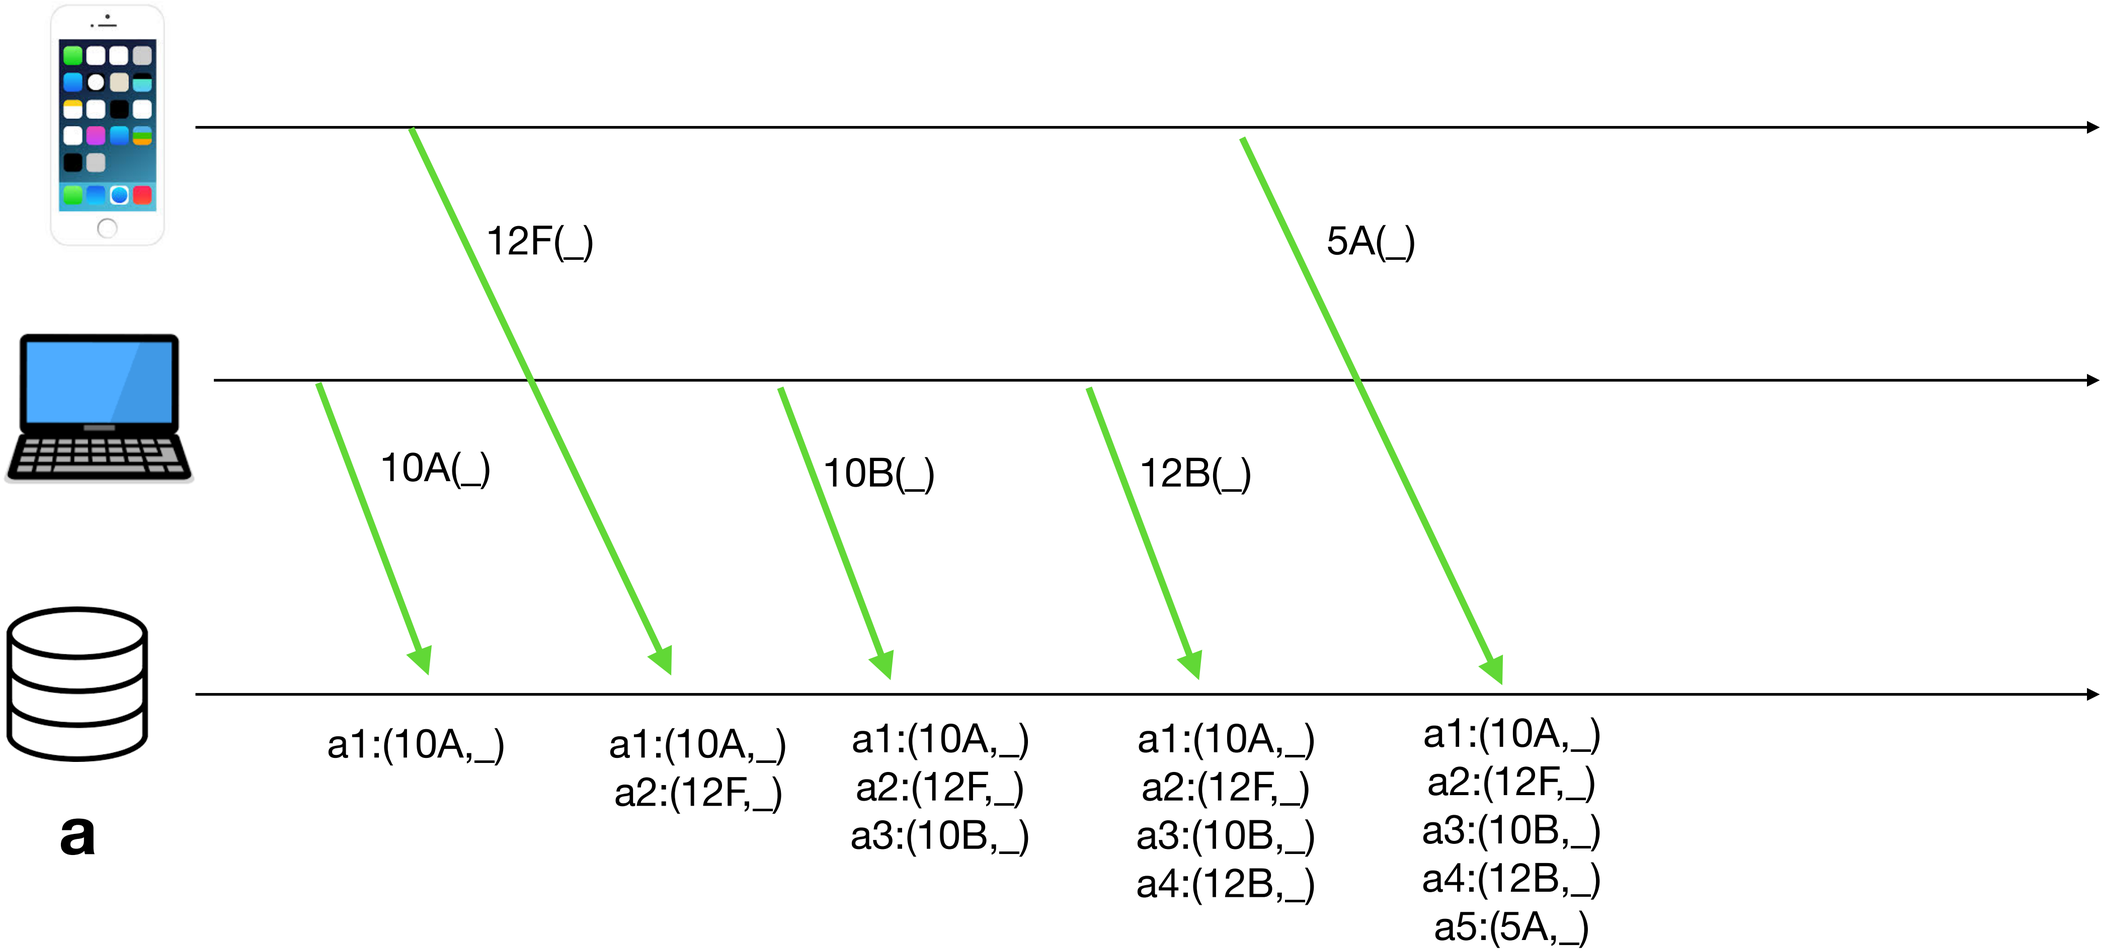

Caveat #3: Wait, Siblings ?

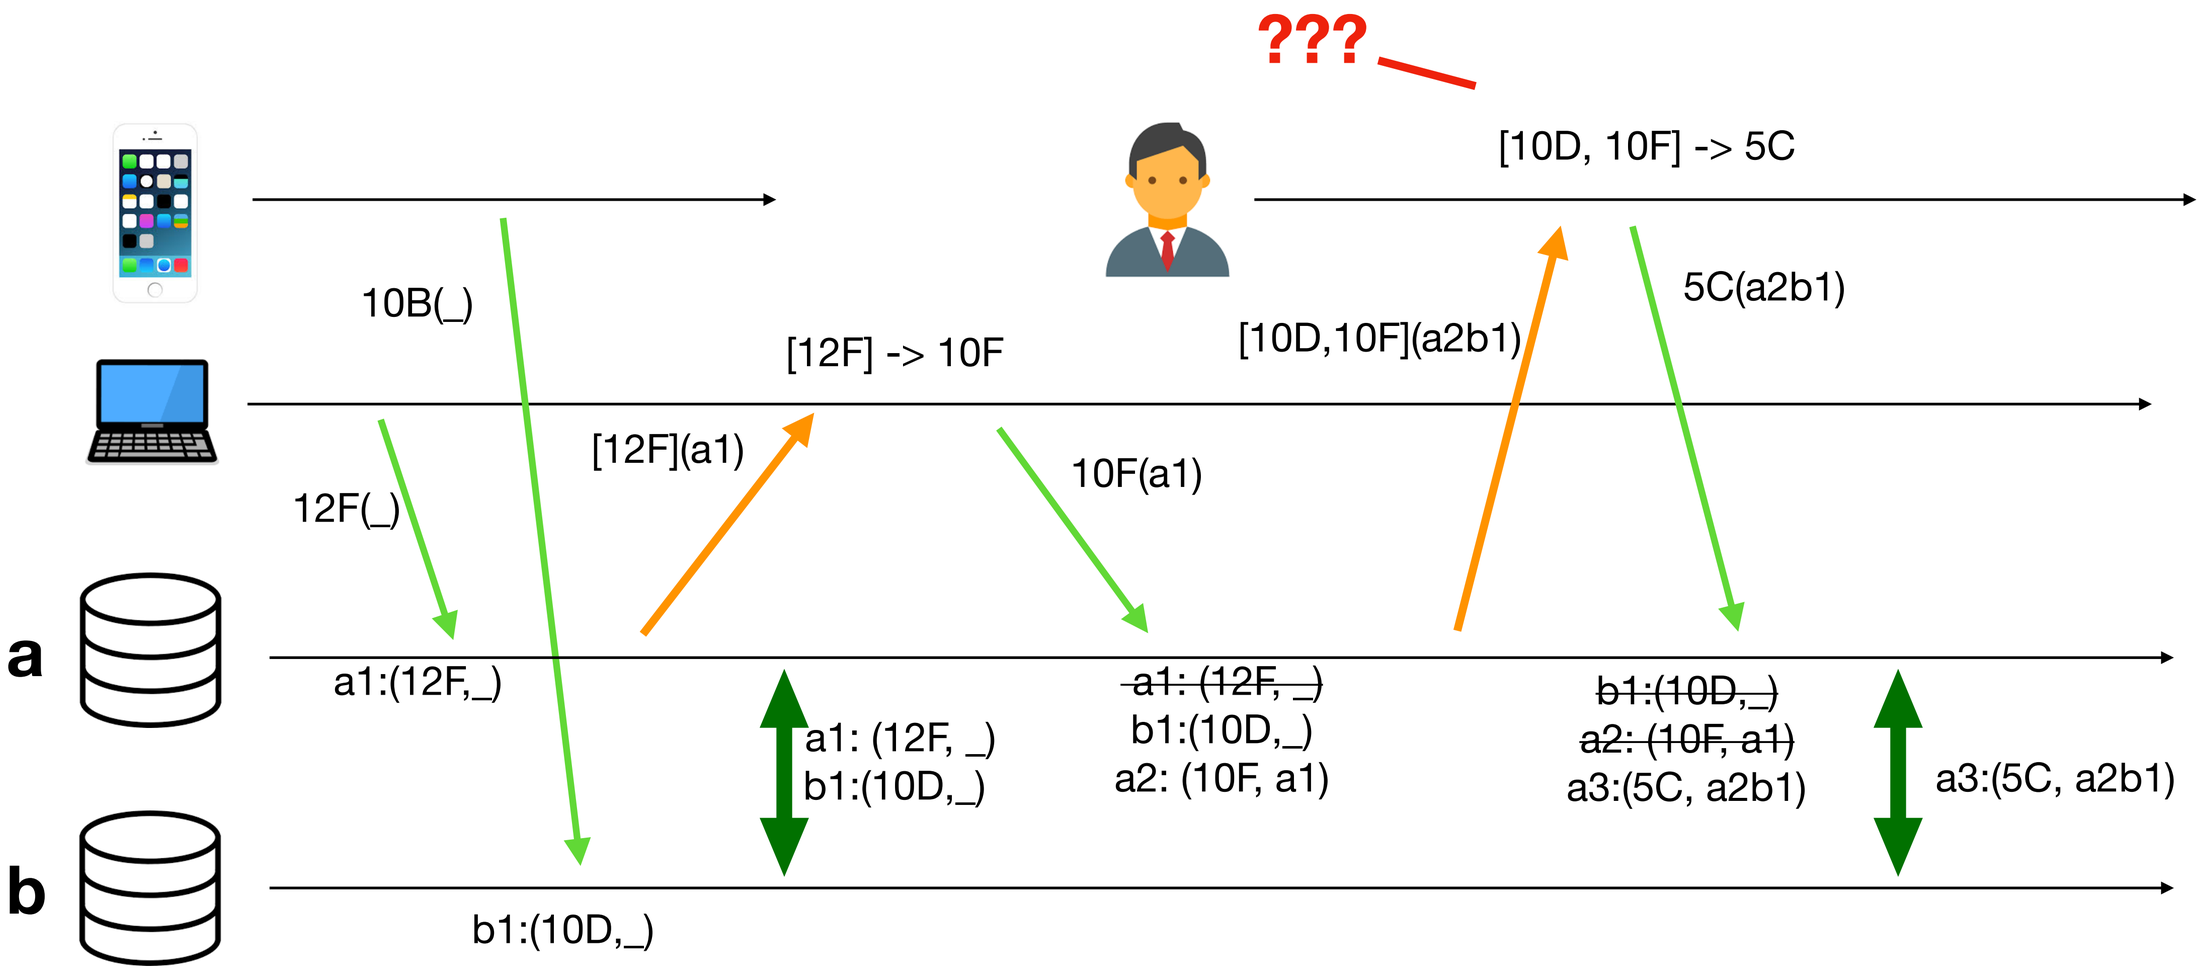

Thanks!