BEYOND THE CLUSTER:
WAN DATA REPLICATION WITH GRIDGAIN

YAKOV ZHDANOV

WHO?

Yakov Zhdanov:
- GridGain's Product Development VP
- With GridGain since 2010
- Apache Ignite committer and PMC
- Passion for performance & scalability
- Finding ways to make product better
- St. Petersburg, Russia

PLAN

1)Why replicate?

PLAN

1)Why replicate?
2)How do DBs solve this?

PLAN

1)Why replicate?
2)How do DBs solve this?
3)Replication: Monolith vs Distributed
4)

PLAN

1)Why replicate?
2)How do DBs solve this?
3)Replication: Monolith vs Distributed
4)GridGain DR overview – roles, features, process

PLAN

1)Why replicate?
2)How do DBs solve this?
3)Replication: Monolith vs Distributed
4)GridGain DR overview – roles, features, process
5)Future plans – Sync/Async TX replication

WHY REPLICATE YOUR DATA?

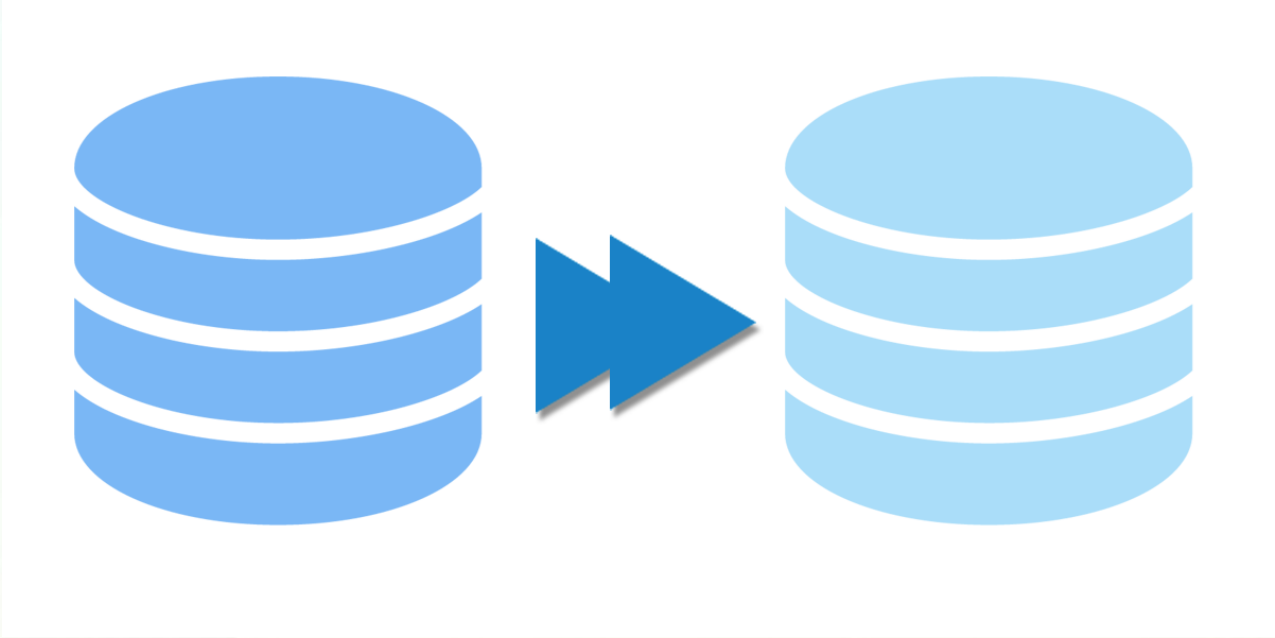

WHY REPLICATE?

- Data security

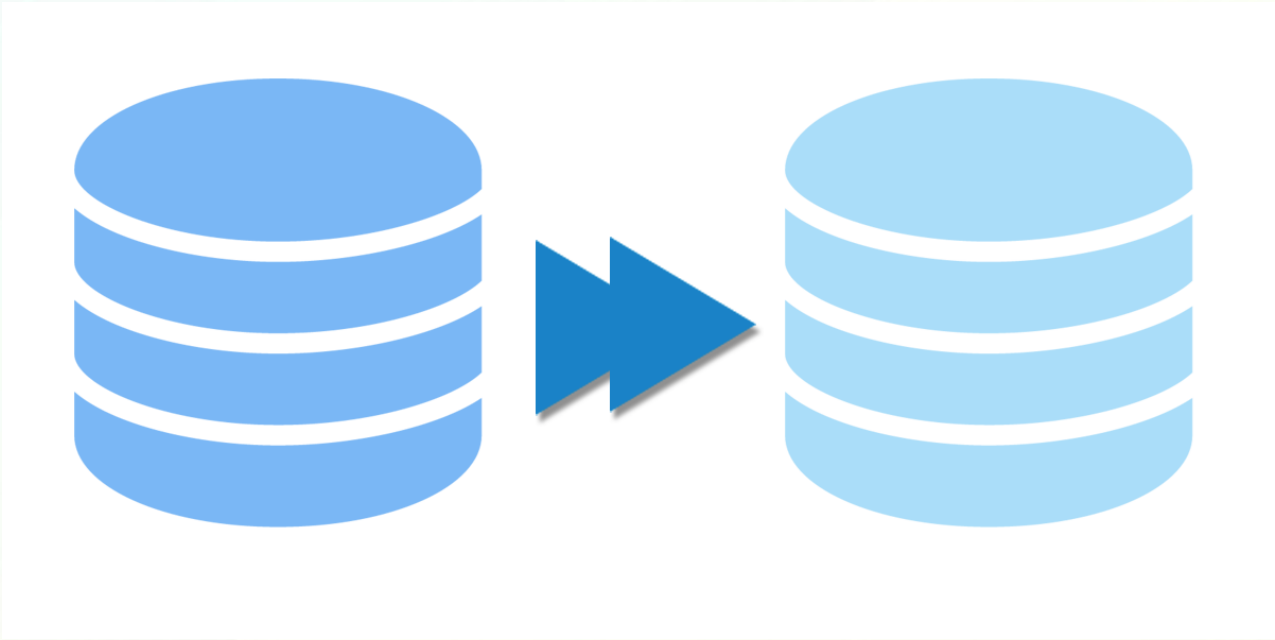

WHY REPLICATE?

- Data security
- Failover

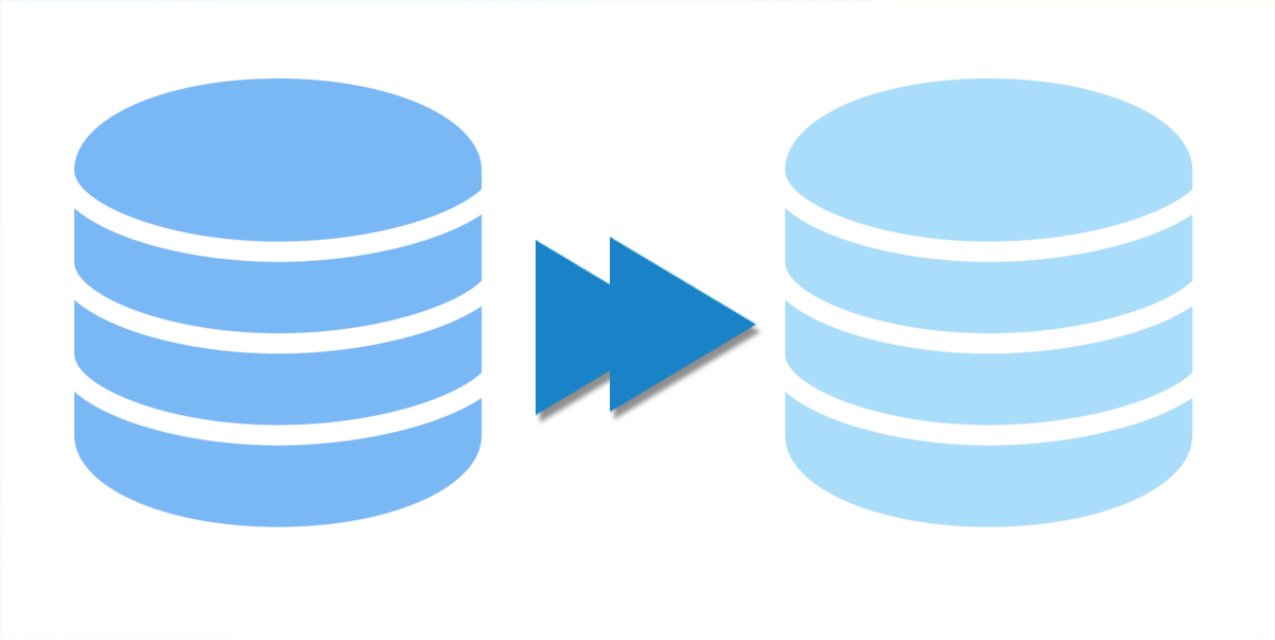

WHY REPLICATE?

- Data security
- Failover
- Data warehousing

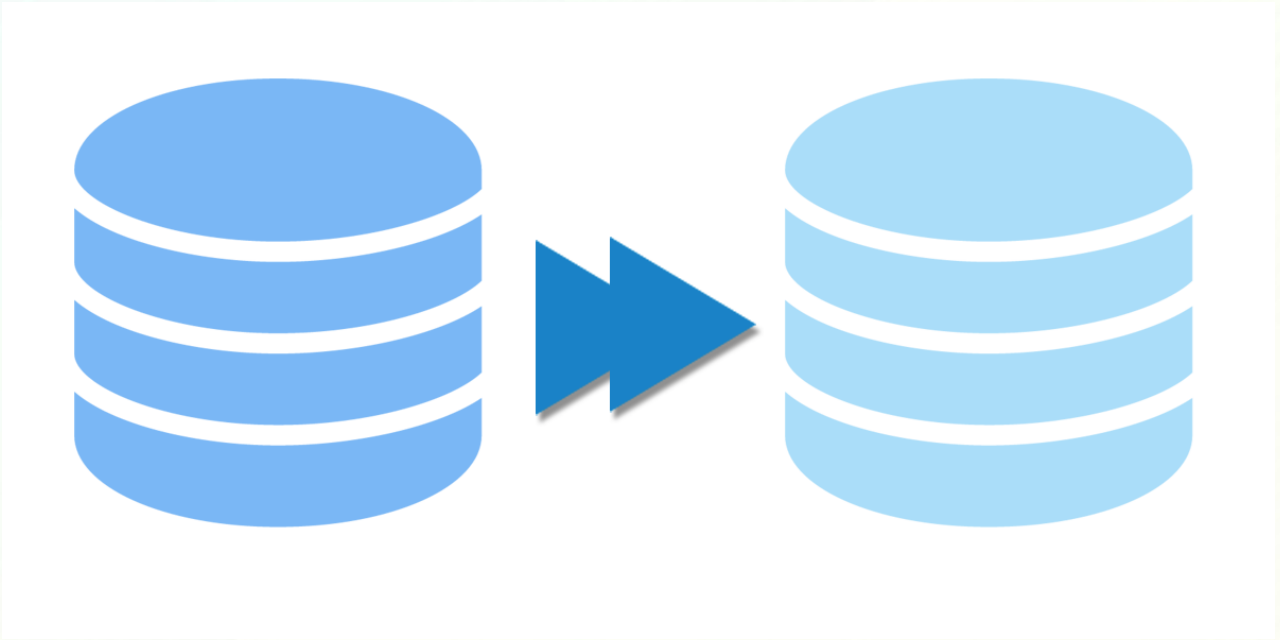

WHY REPLICATE?

- Data security
- Failover
- Data warehousing
- Load balancing

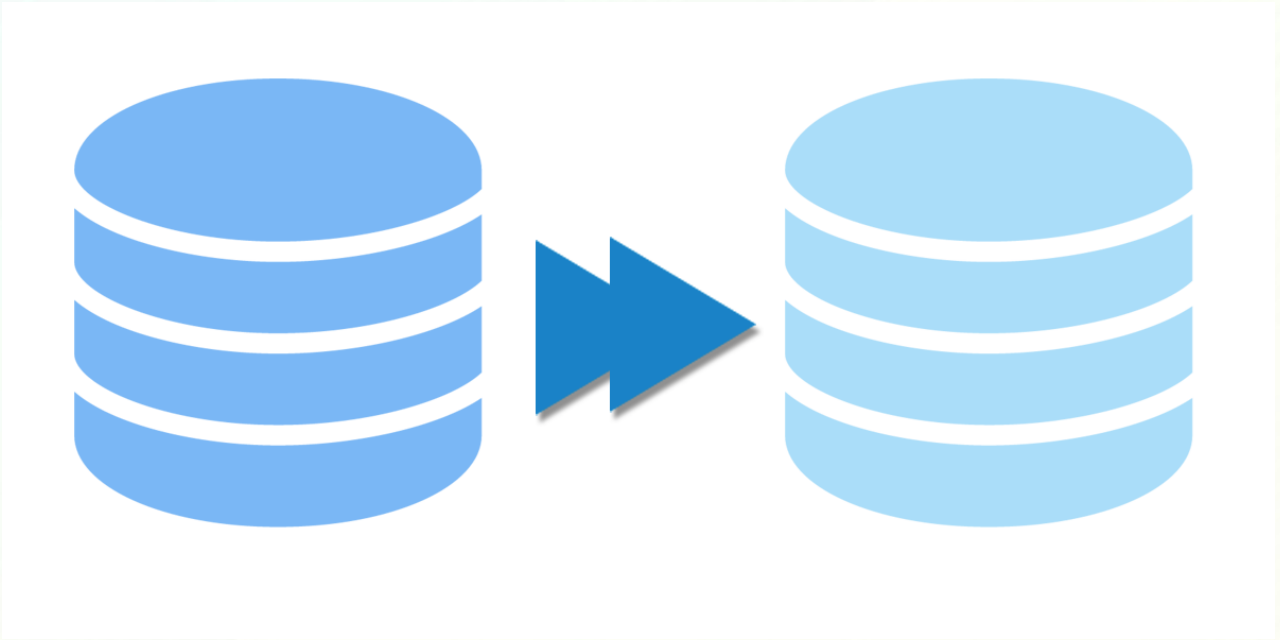

WHY REPLICATE?

- Data security
- Failover
- Data warehousing
- Load balancing
- Increasing system capacity

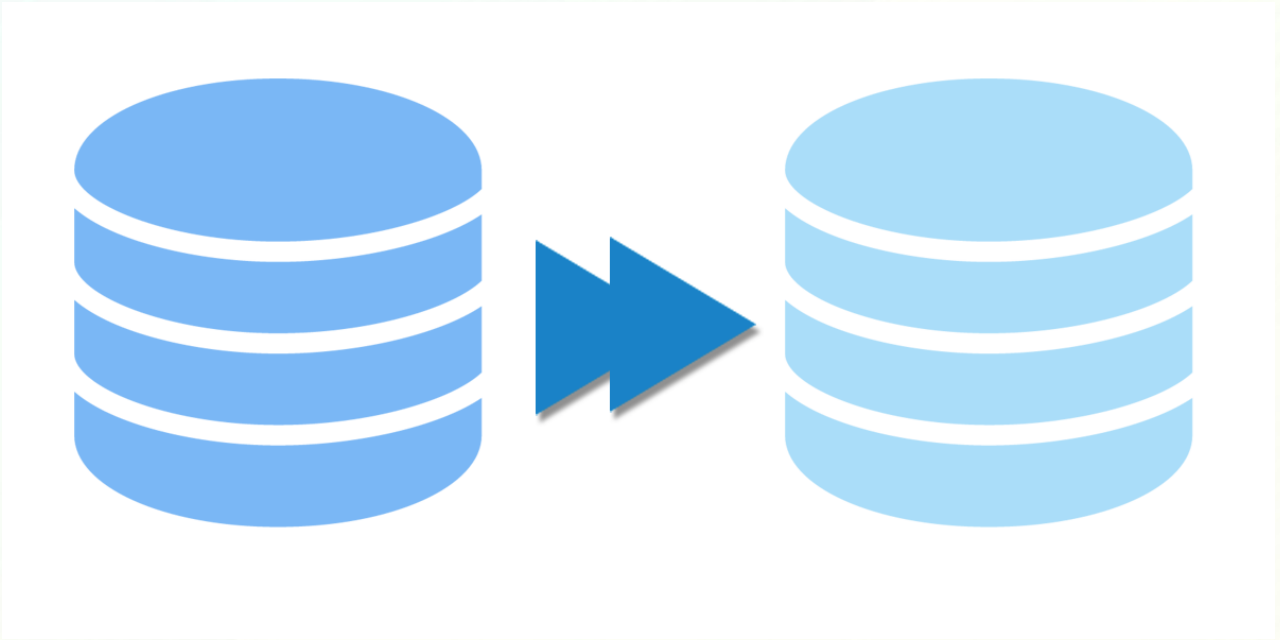

POSTGRESQL REPLICATION

- PostgreSQL is an object-relational database management system (ORDBMS)
- Pioneered many things and concepts
- High maturity level
- Opensource and widely used

POSTGRESQL REPLICATION

- Shared disk storage

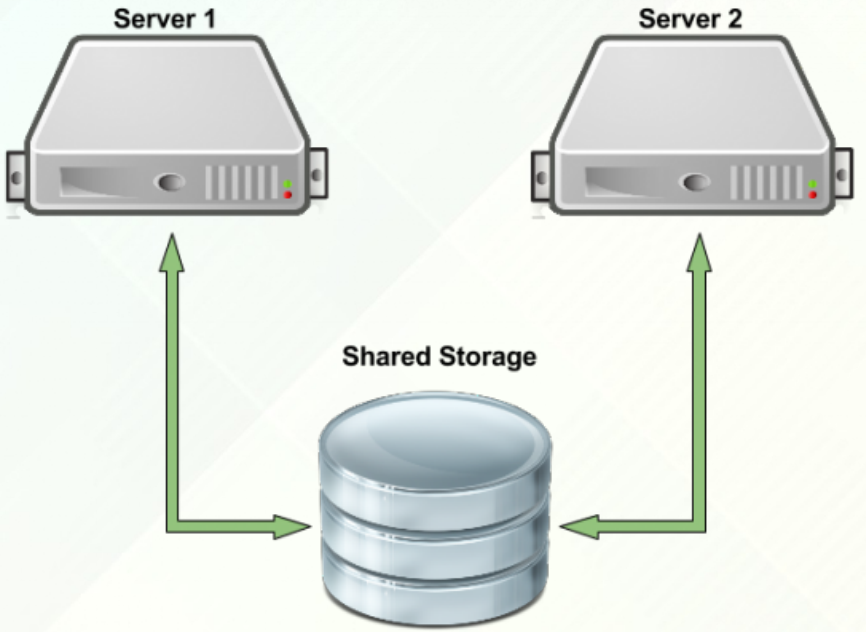

https://www.postgresql.org/docs/10/static/different-replication-solutions.html

POSTGRESQL REPLICATION

- Shared disk storage
- File system replication

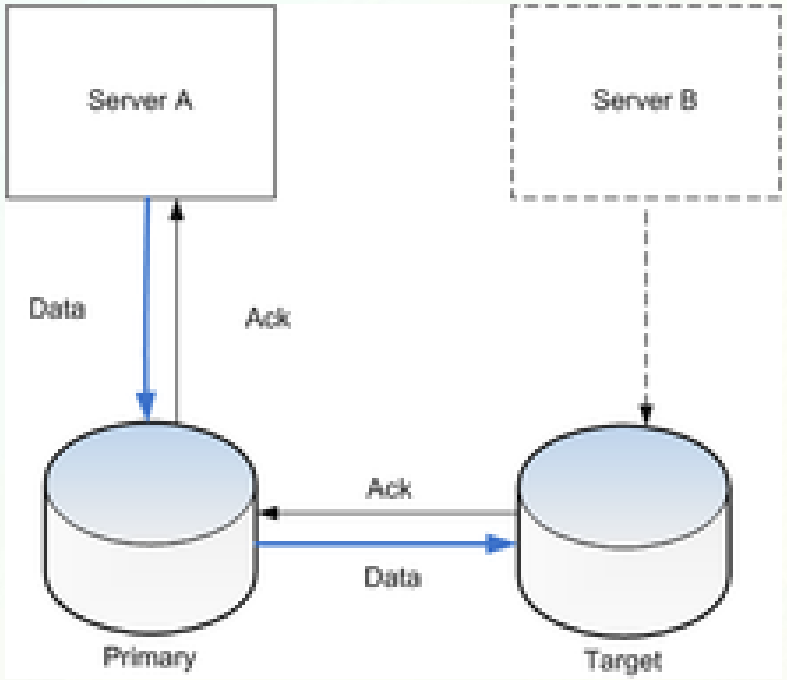

https://www.postgresql.org/docs/10/static/different-replication-solutions.html

POSTGRESQL REPLICATION

- Shared disk storage
- File system replication
- Write-Ahead Log Shipping

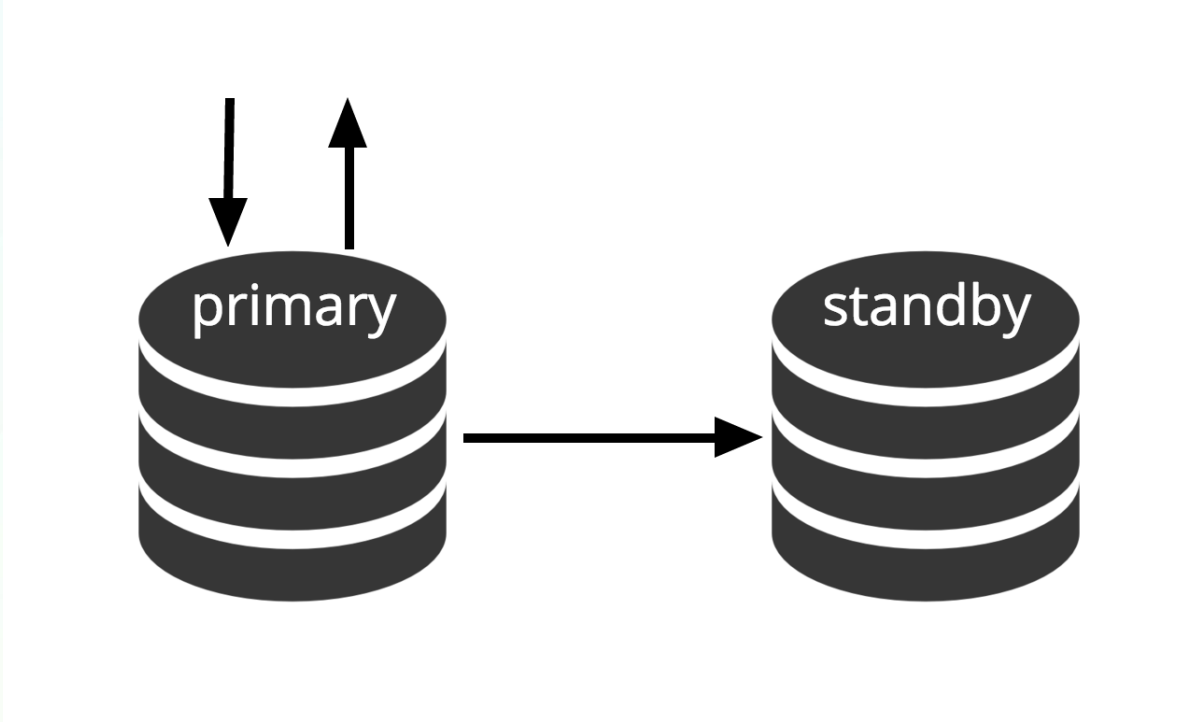

https://www.postgresql.org/docs/10/static/different-replication-solutions.html

POSTGRESQL REPLICATION

- Shared disk storage
- File system replication
- Write-Ahead Log Shipping
- Logical Replication

https://www.postgresql.org/docs/10/static/different-replication-solutions.html

POSTGRESQL REPLICATION

- Shared disk storage
- File system replication
- Write-Ahead Log Shipping
- Logical Replication
- Trigger-Based Master-Standby Replication

https://www.postgresql.org/docs/10/static/different-replication-solutions.html

POSTGRESQL REPLICATION

- Shared disk storage
- File system replication
- Write-Ahead Log Shipping
- Logical Replication
- Trigger-Based Master-Standby Replication
- Statement-Based Replication Middleware

https://www.postgresql.org/docs/10/static/different-replication-solutions.html

POSTGRESQL REPLICATION

- Shared disk storage
- File system replication
- Write-Ahead Log Shipping
- Logical Replication
- Trigger-Based Master-Standby Replication
- Statement-Based Replication Middleware
- Async Multimaster Replication

https://www.postgresql.org/docs/10/static/different-replication-solutions.html

POSTGRESQL REPLICATION

- Shared disk storage
- File system replication
- Write-Ahead Log Shipping
- Logical Replication
- Trigger-Based Master-Standby Replication
- Statement-Based Replication Middleware
- Async Multimaster Replication
- Sync Multimaster Replication

https://www.postgresql.org/docs/10/static/different-replication-solutions.html

REPLICATION IN DISTRIBUTED SYSTEMS

	Monolith	Distributed
Data security	+	?
Failover	+	+
Load balancing	+	?/+
Increasing system capacity	+	?
Data warehousing	+	?

REPLICATION IN GRIDGAIN

- Introduced in 2012/2013
- Completely new feature
- Required a lot of engineering efforts
- Required revisiting of existing logic
- Async KEY/VALUE mode available
- Sync/Async TX replication under development

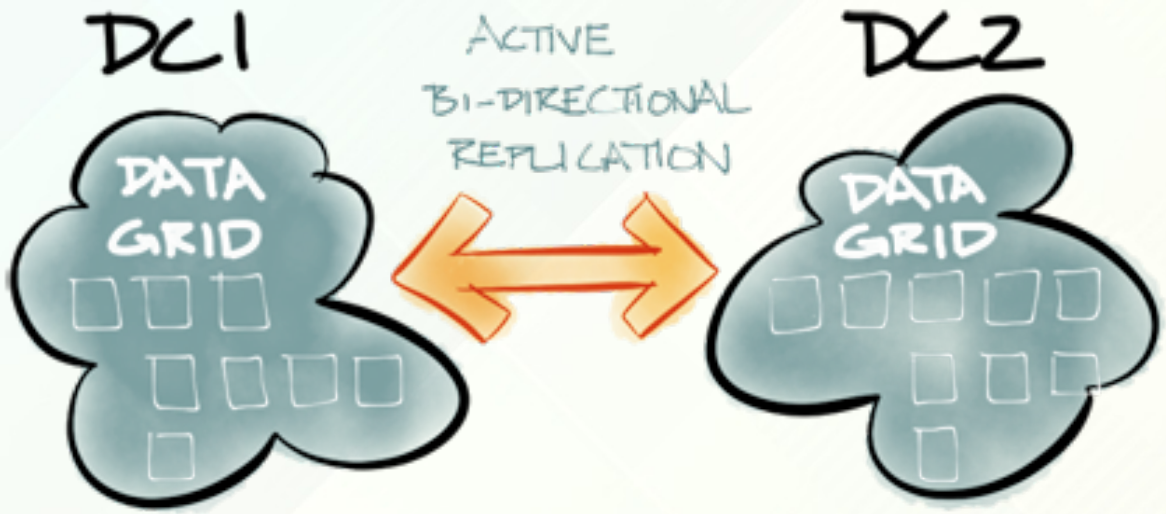

https://docs.gridgain.com/docs/data-center-replication

REPLICATION IN GRIDGAIN: ROLES

- Sender cache
- Sender hub
- Receiver hub
- Receiver cache

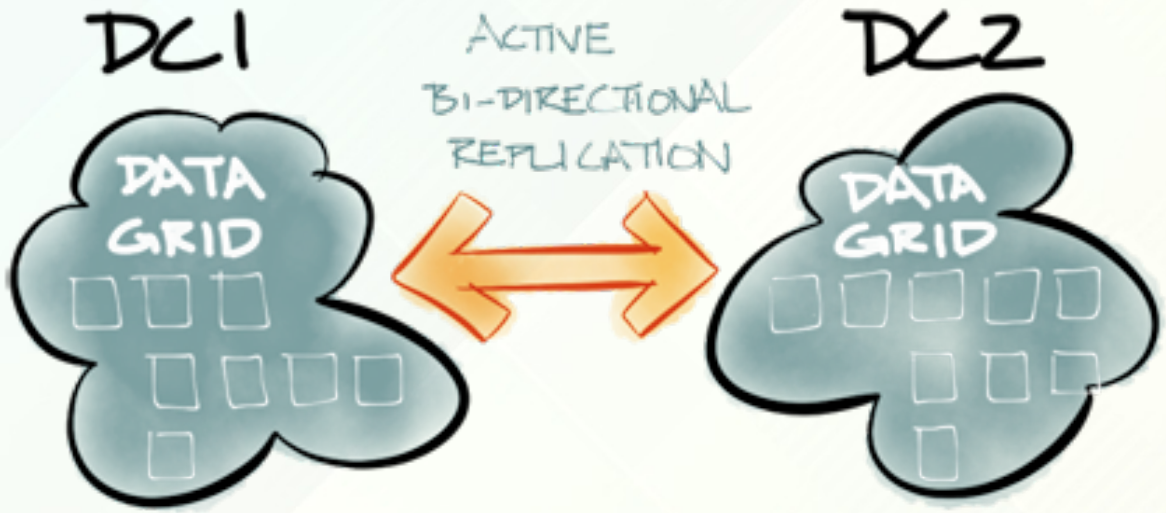

https://docs.gridgain.com/docs/data-center-replication

REPLICATION IN GRIDGAIN: FEATURES

- Complex topologies (up to 32 datacenters)
- Failover
- Pluggable conflict resolution
- Filtering
- Pause/Resume
- Full state transfer

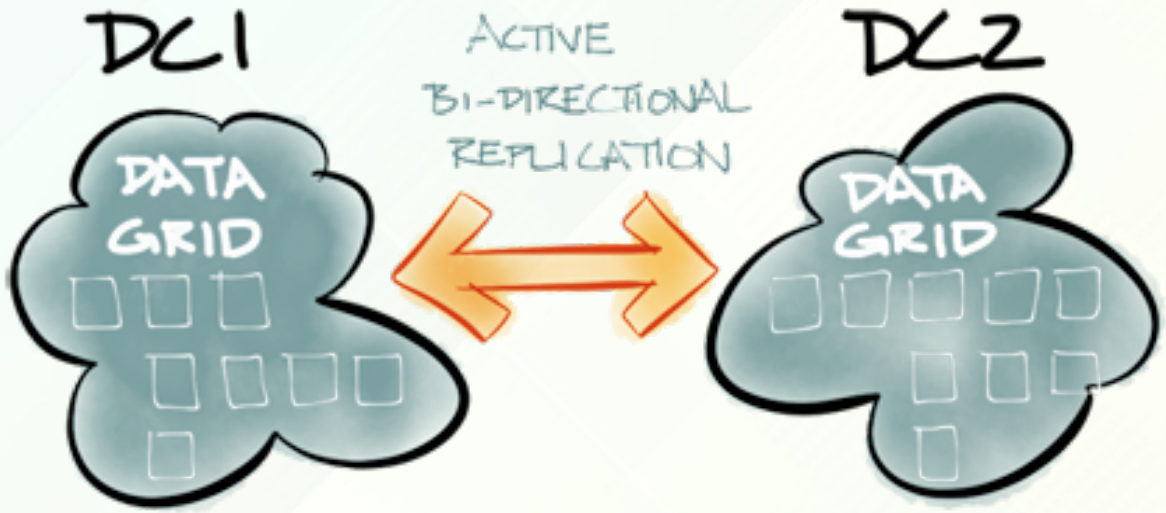

https://docs.gridgain.com/docs/data-center-replication

REPLICATION IN GRIDGAIN: HOW IT WORKS

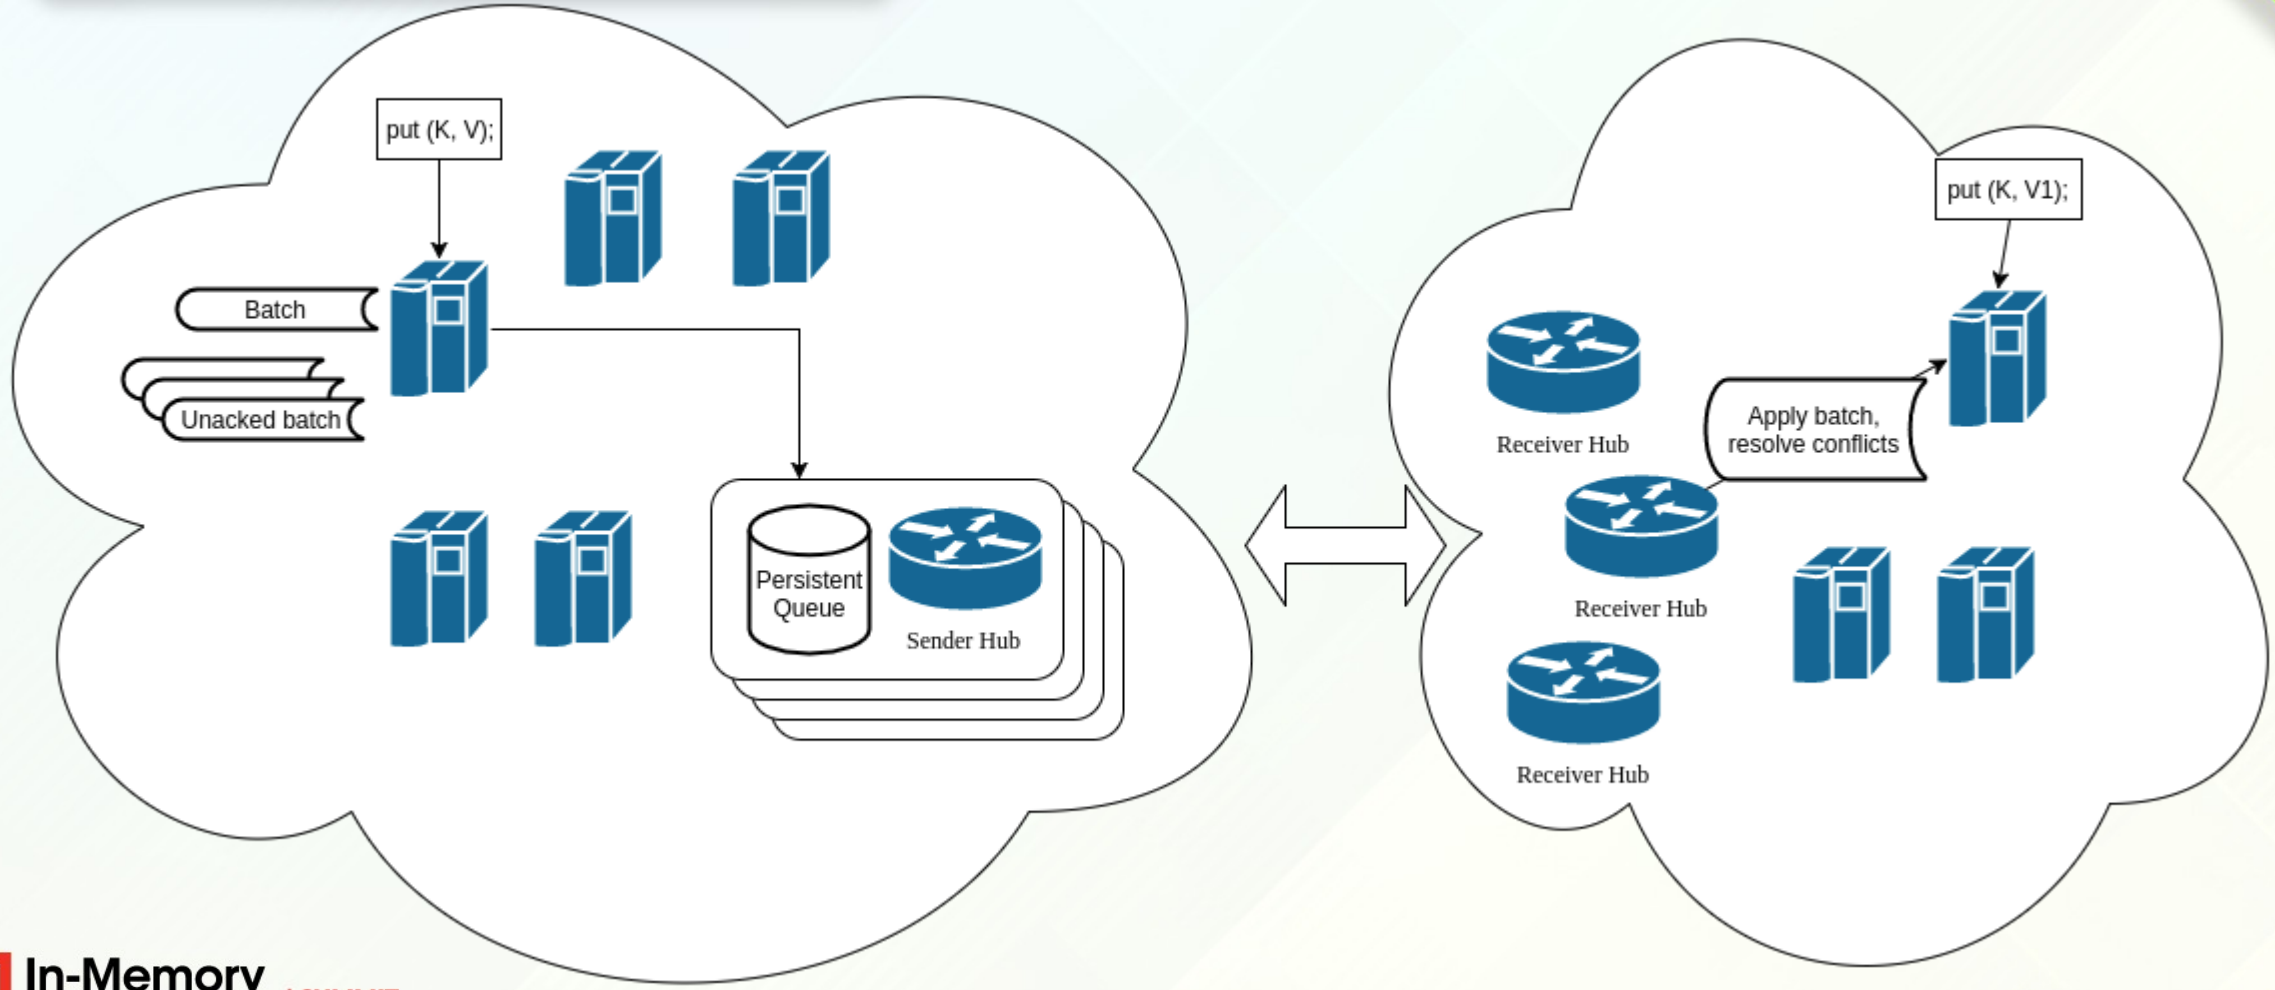

REPLICATION IN GRIDGAIN: COMPLEX TOPOLOGIES

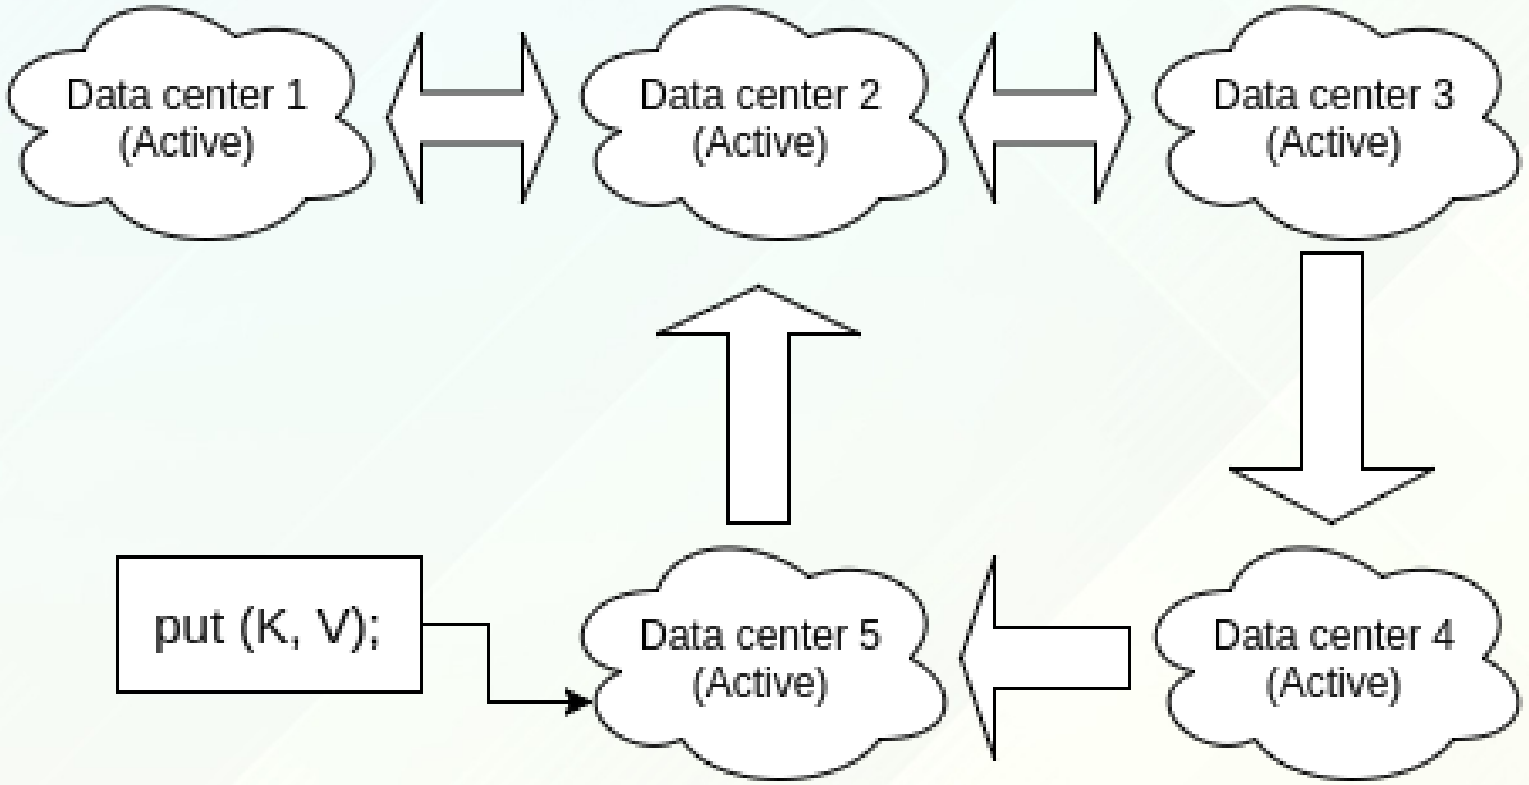

REPLICATION IN GRIDGAIN: WHAT CAN GO WRONG?

Node failure in sending topology

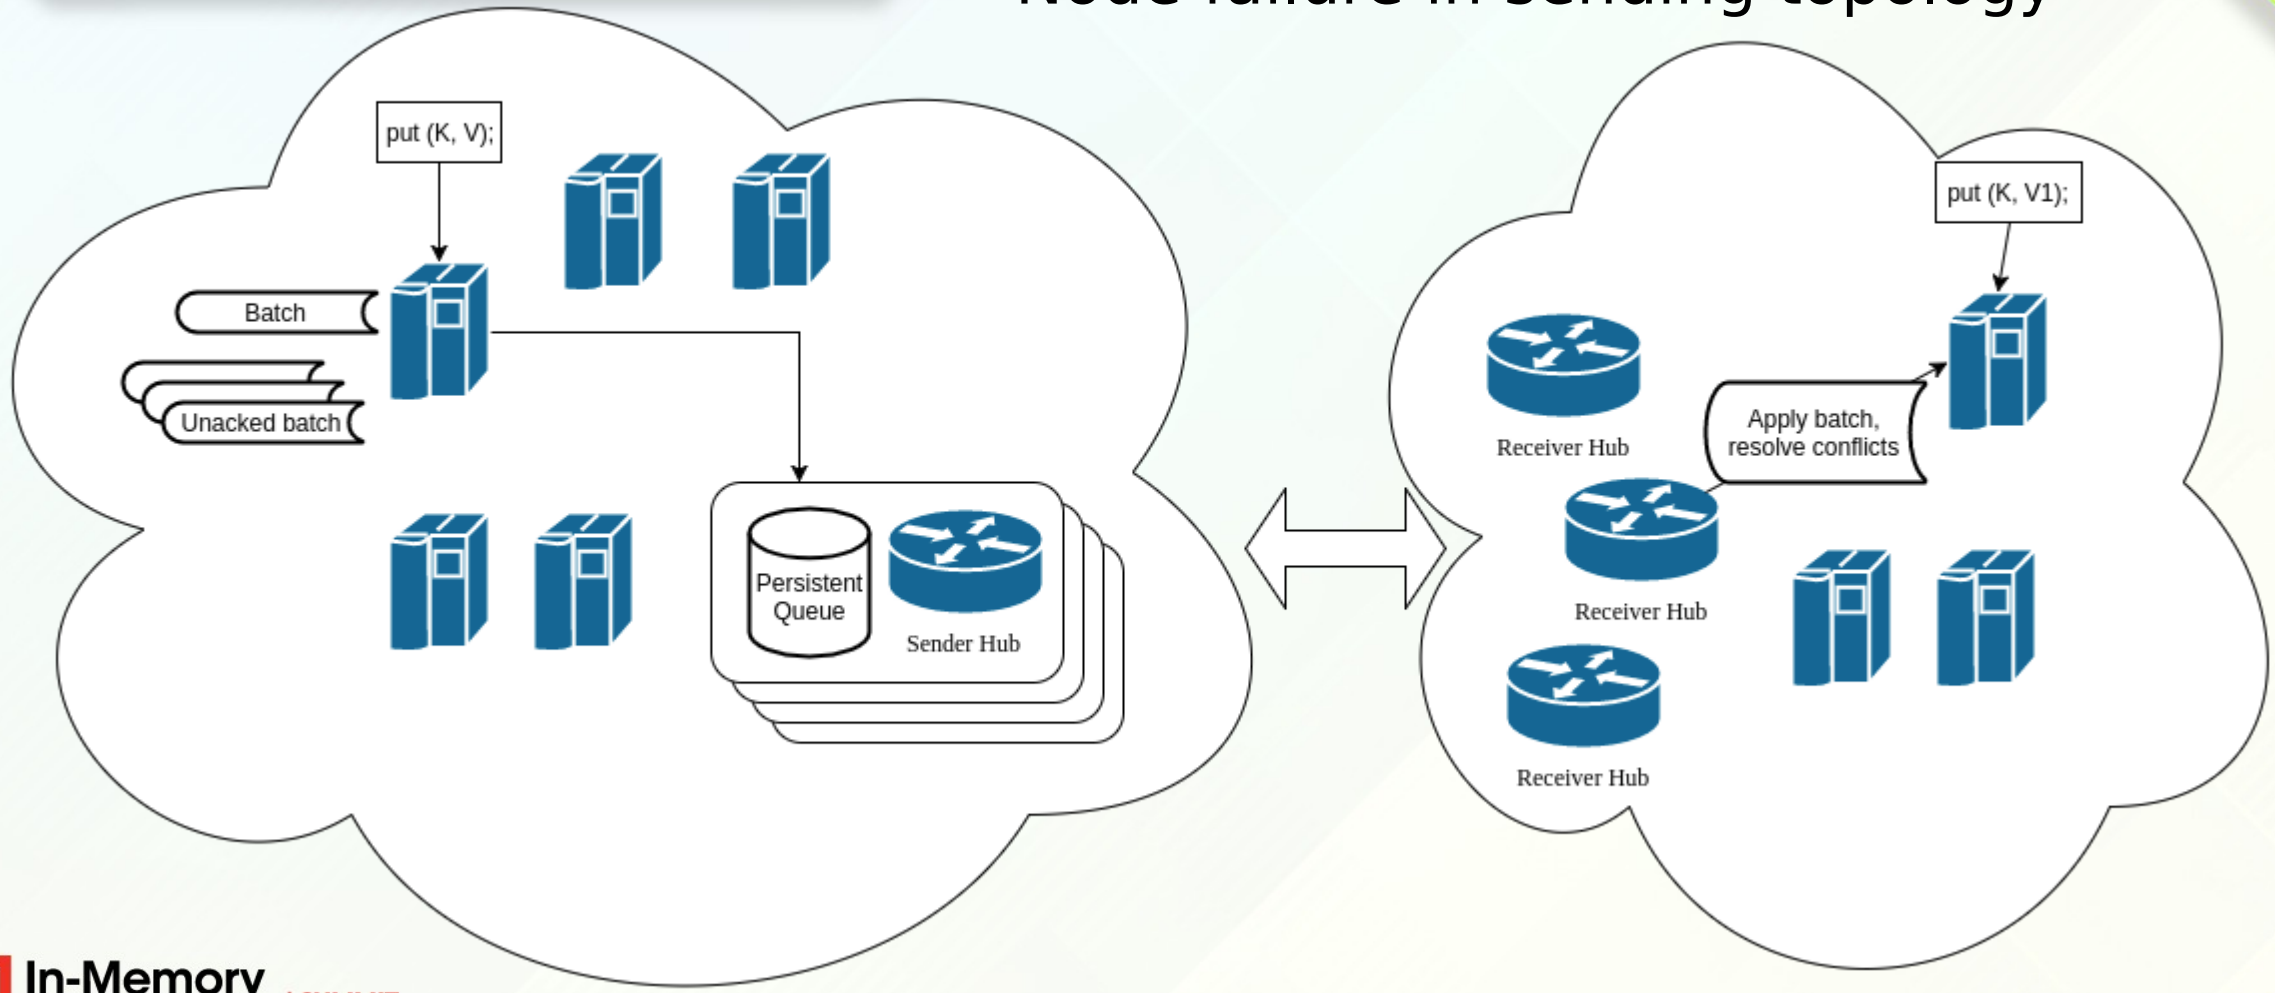

REPLICATION IN GRIDGAIN: WHAT CAN GO WRONG?

Node failure in sending topology

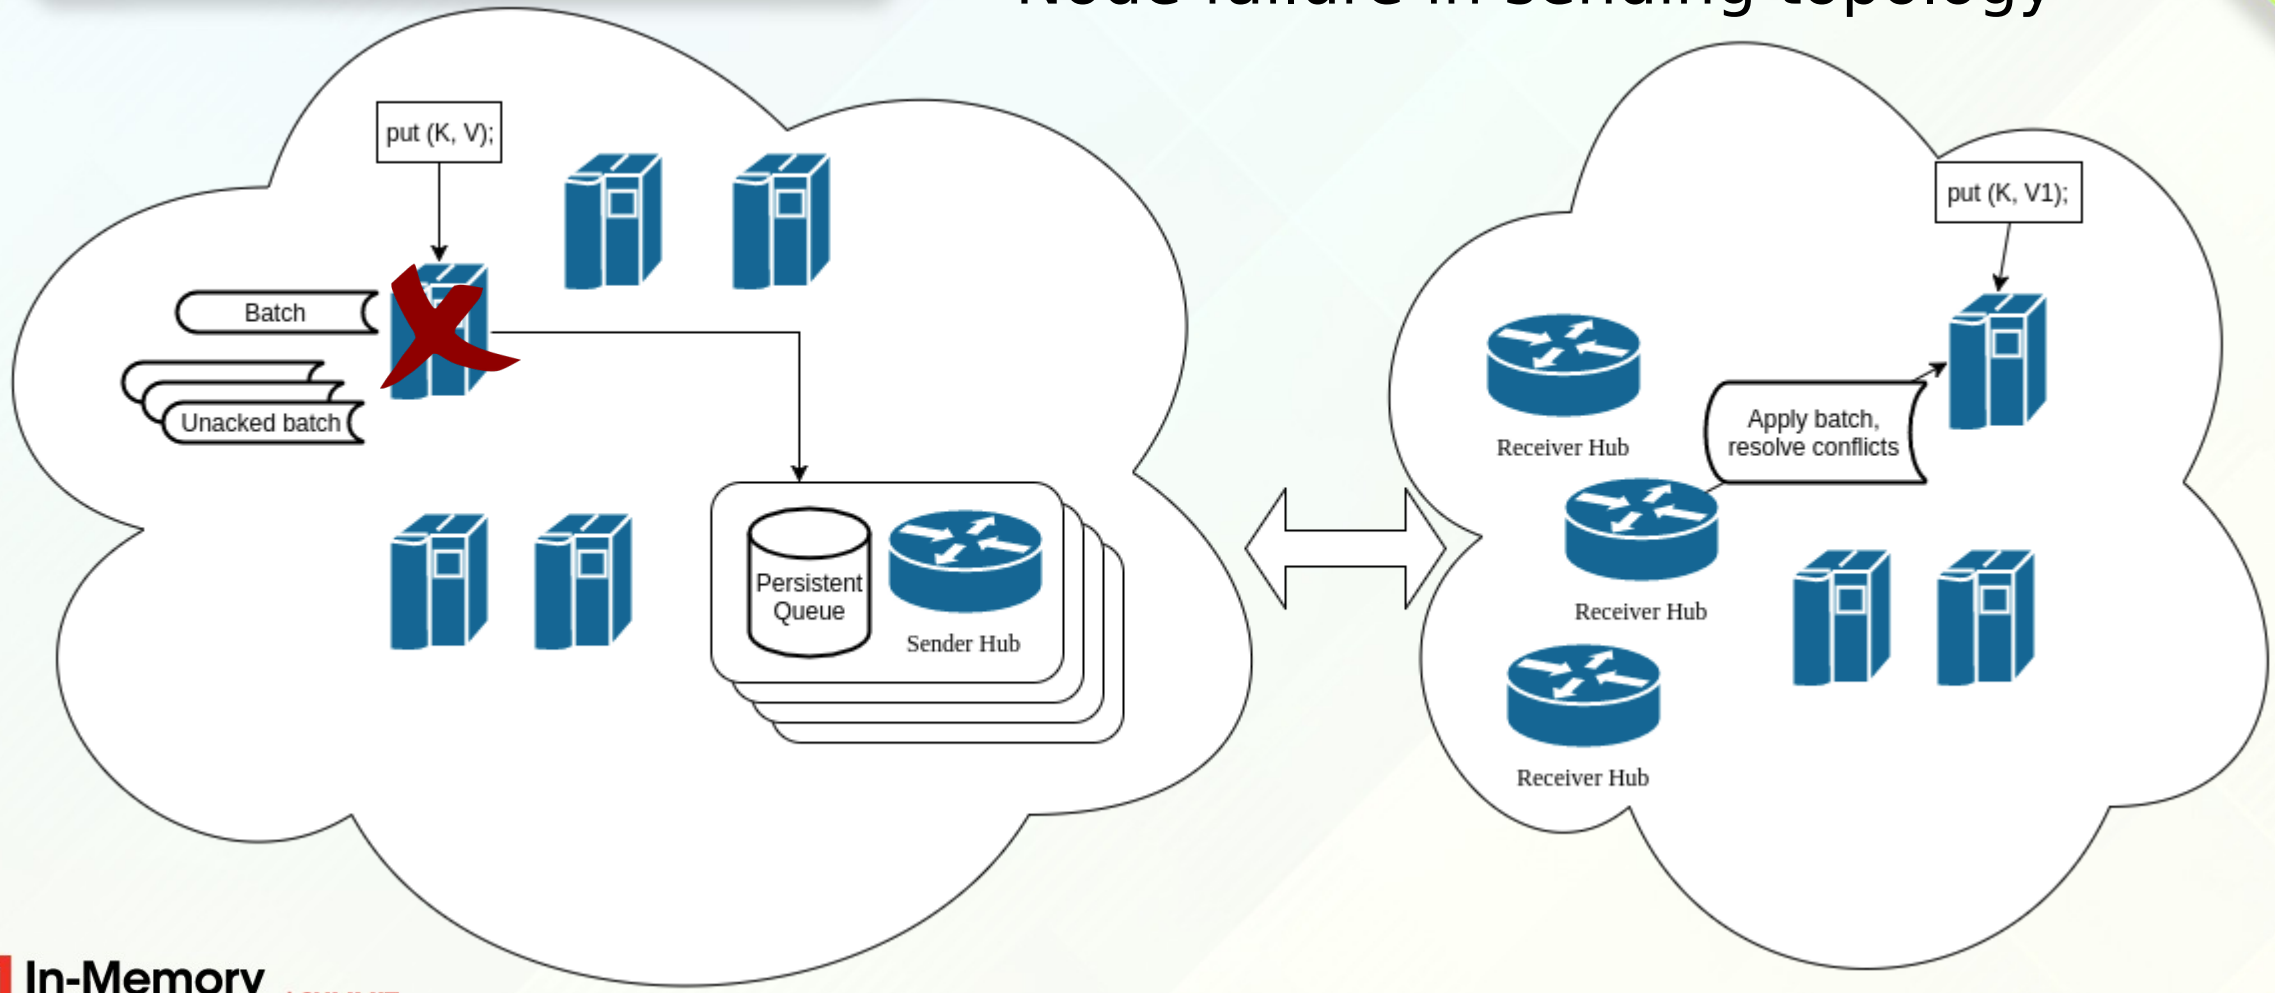

REPLICATION IN GRIDGAIN: WHAT CAN GO WRONG?

Sender hub(s) failure

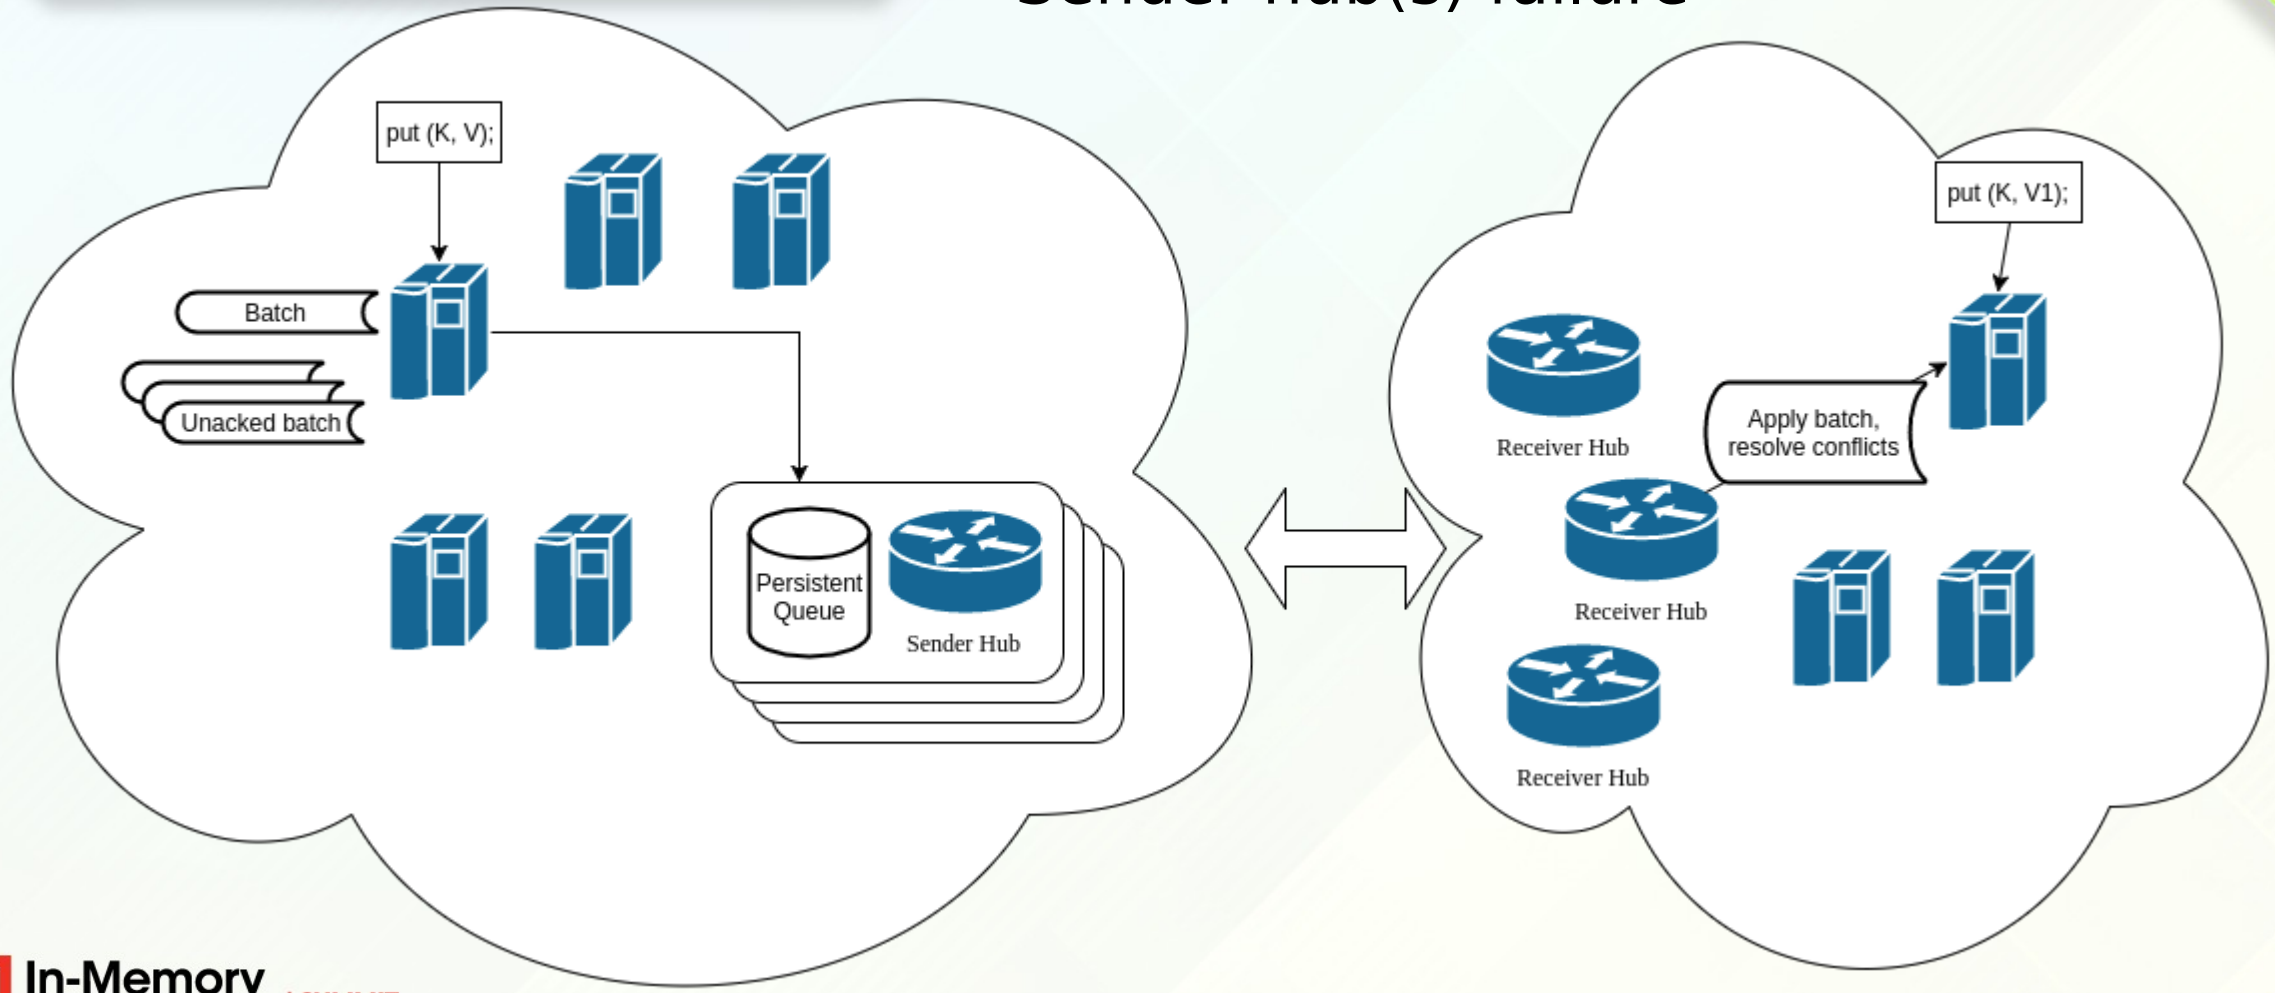

REPLICATION IN GRIDGAIN: WHAT CAN GO WRONG?

Sender hub(s) failure

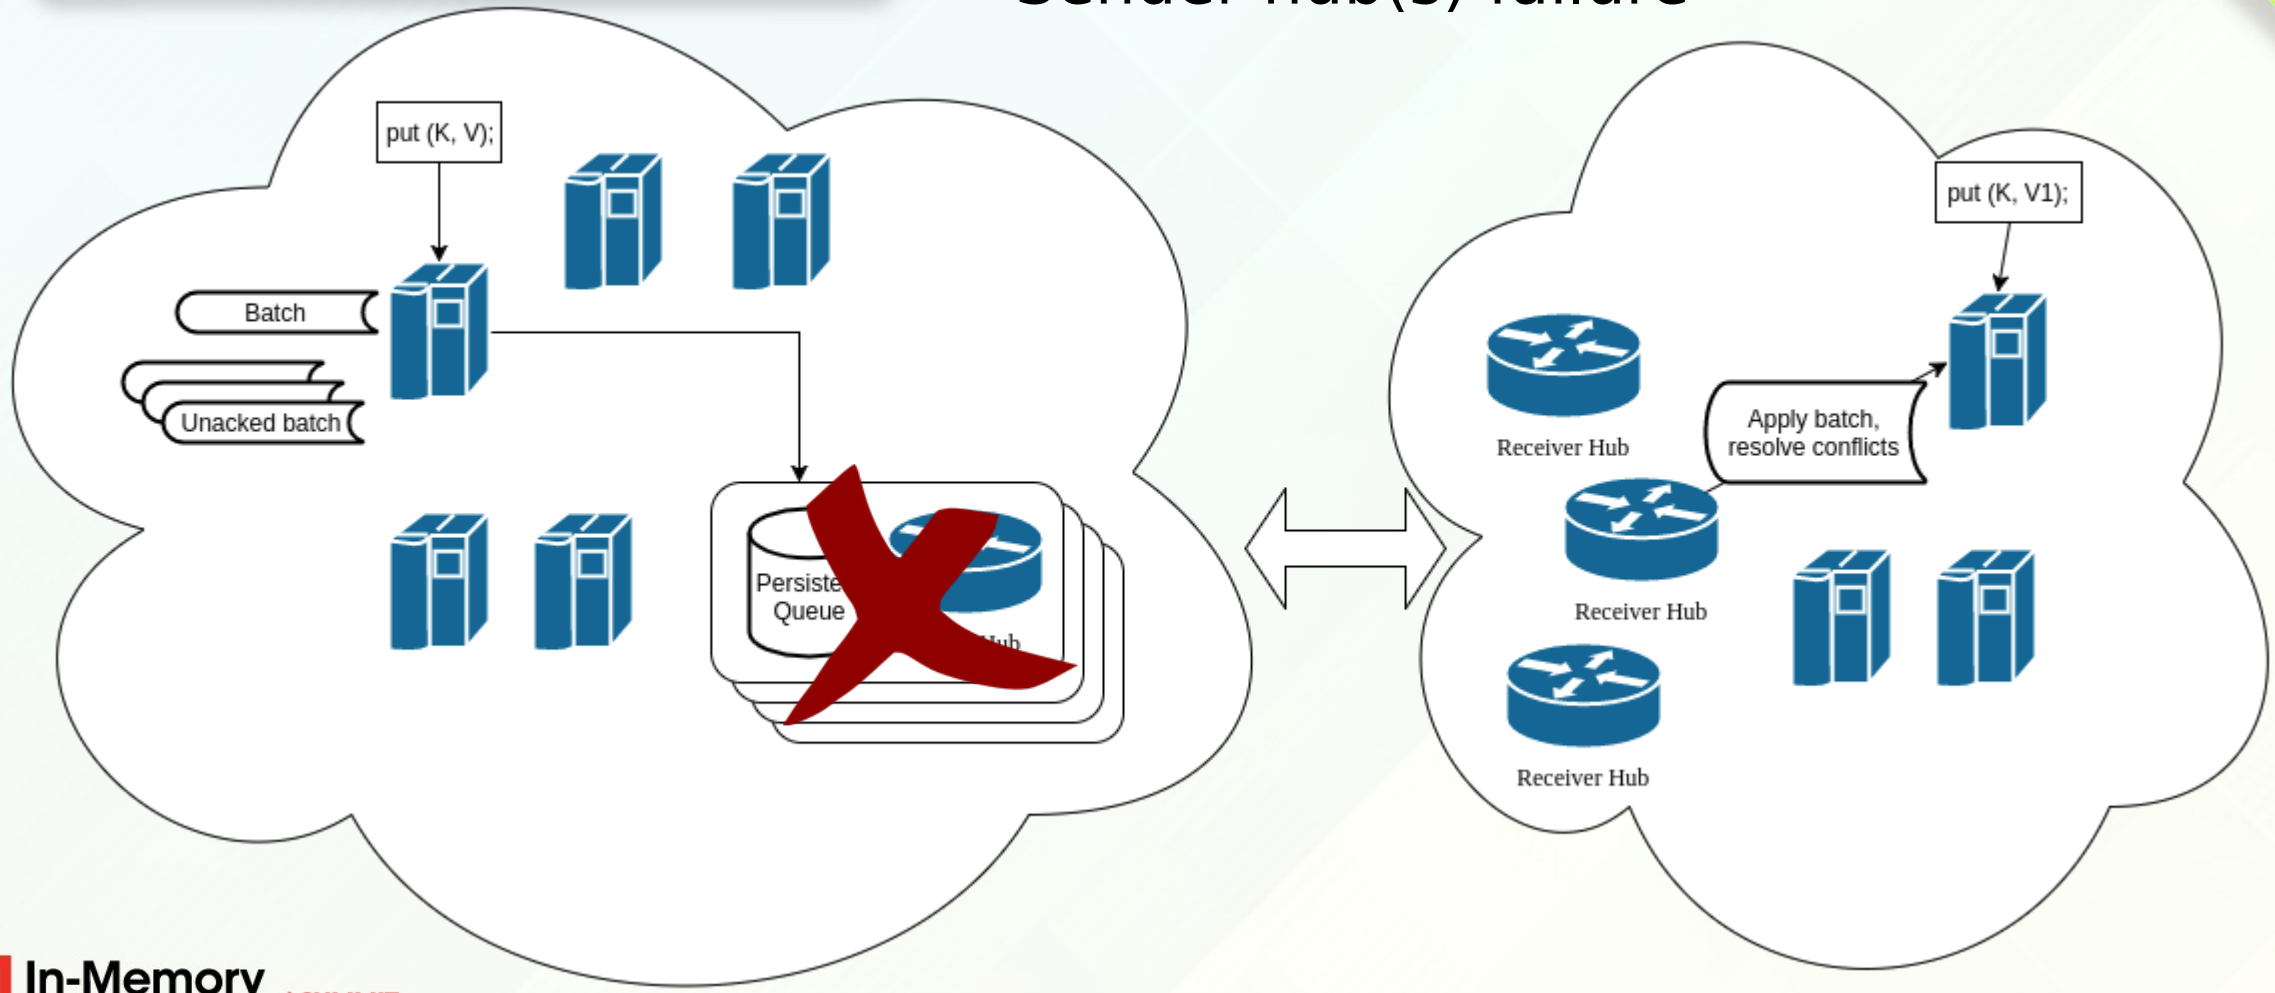

REPLICATION IN GRIDGAIN: WHAT CAN GO WRONG?

Receiver hub(s) failure

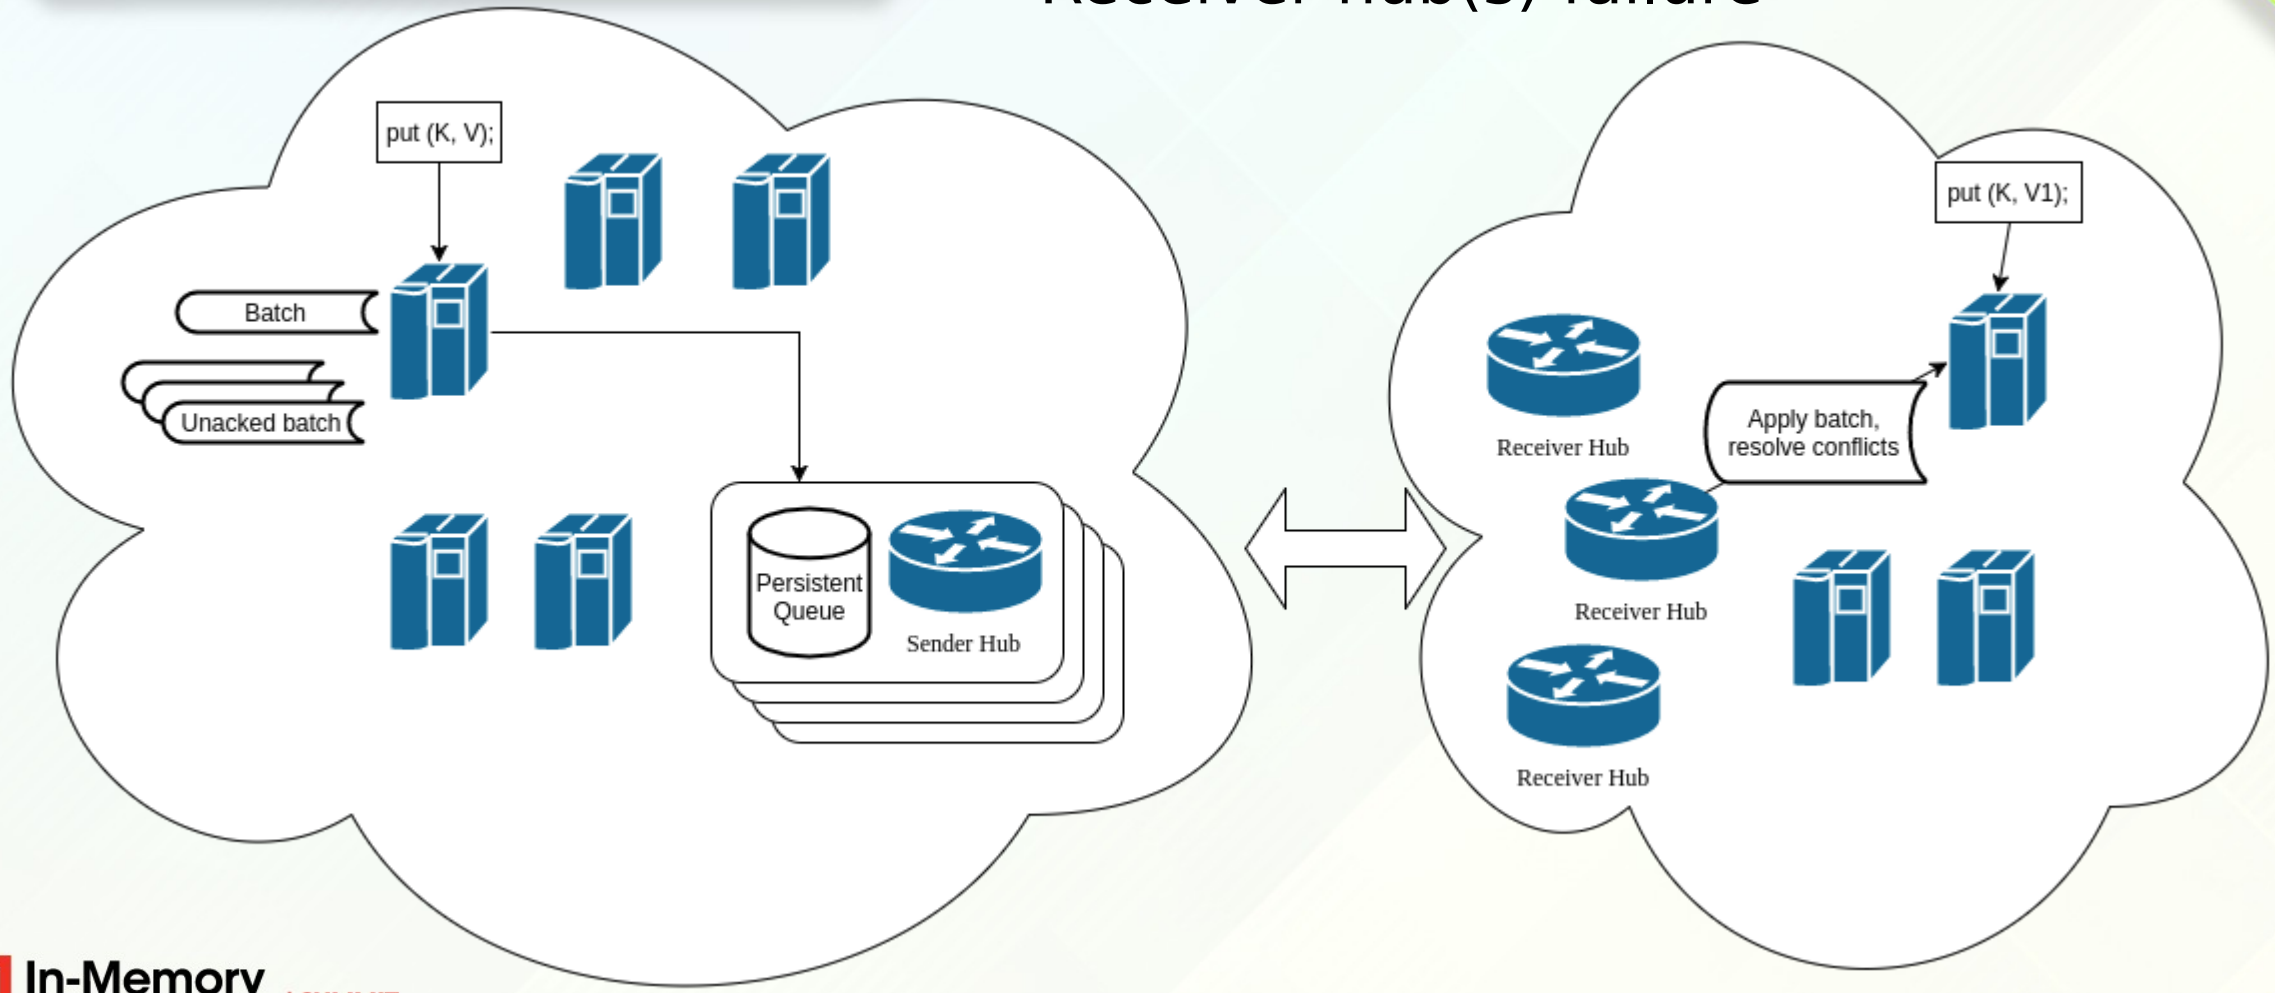

REPLICATION IN GRIDGAIN: WHAT CAN GO WRONG?

Receiver hub(s) failure

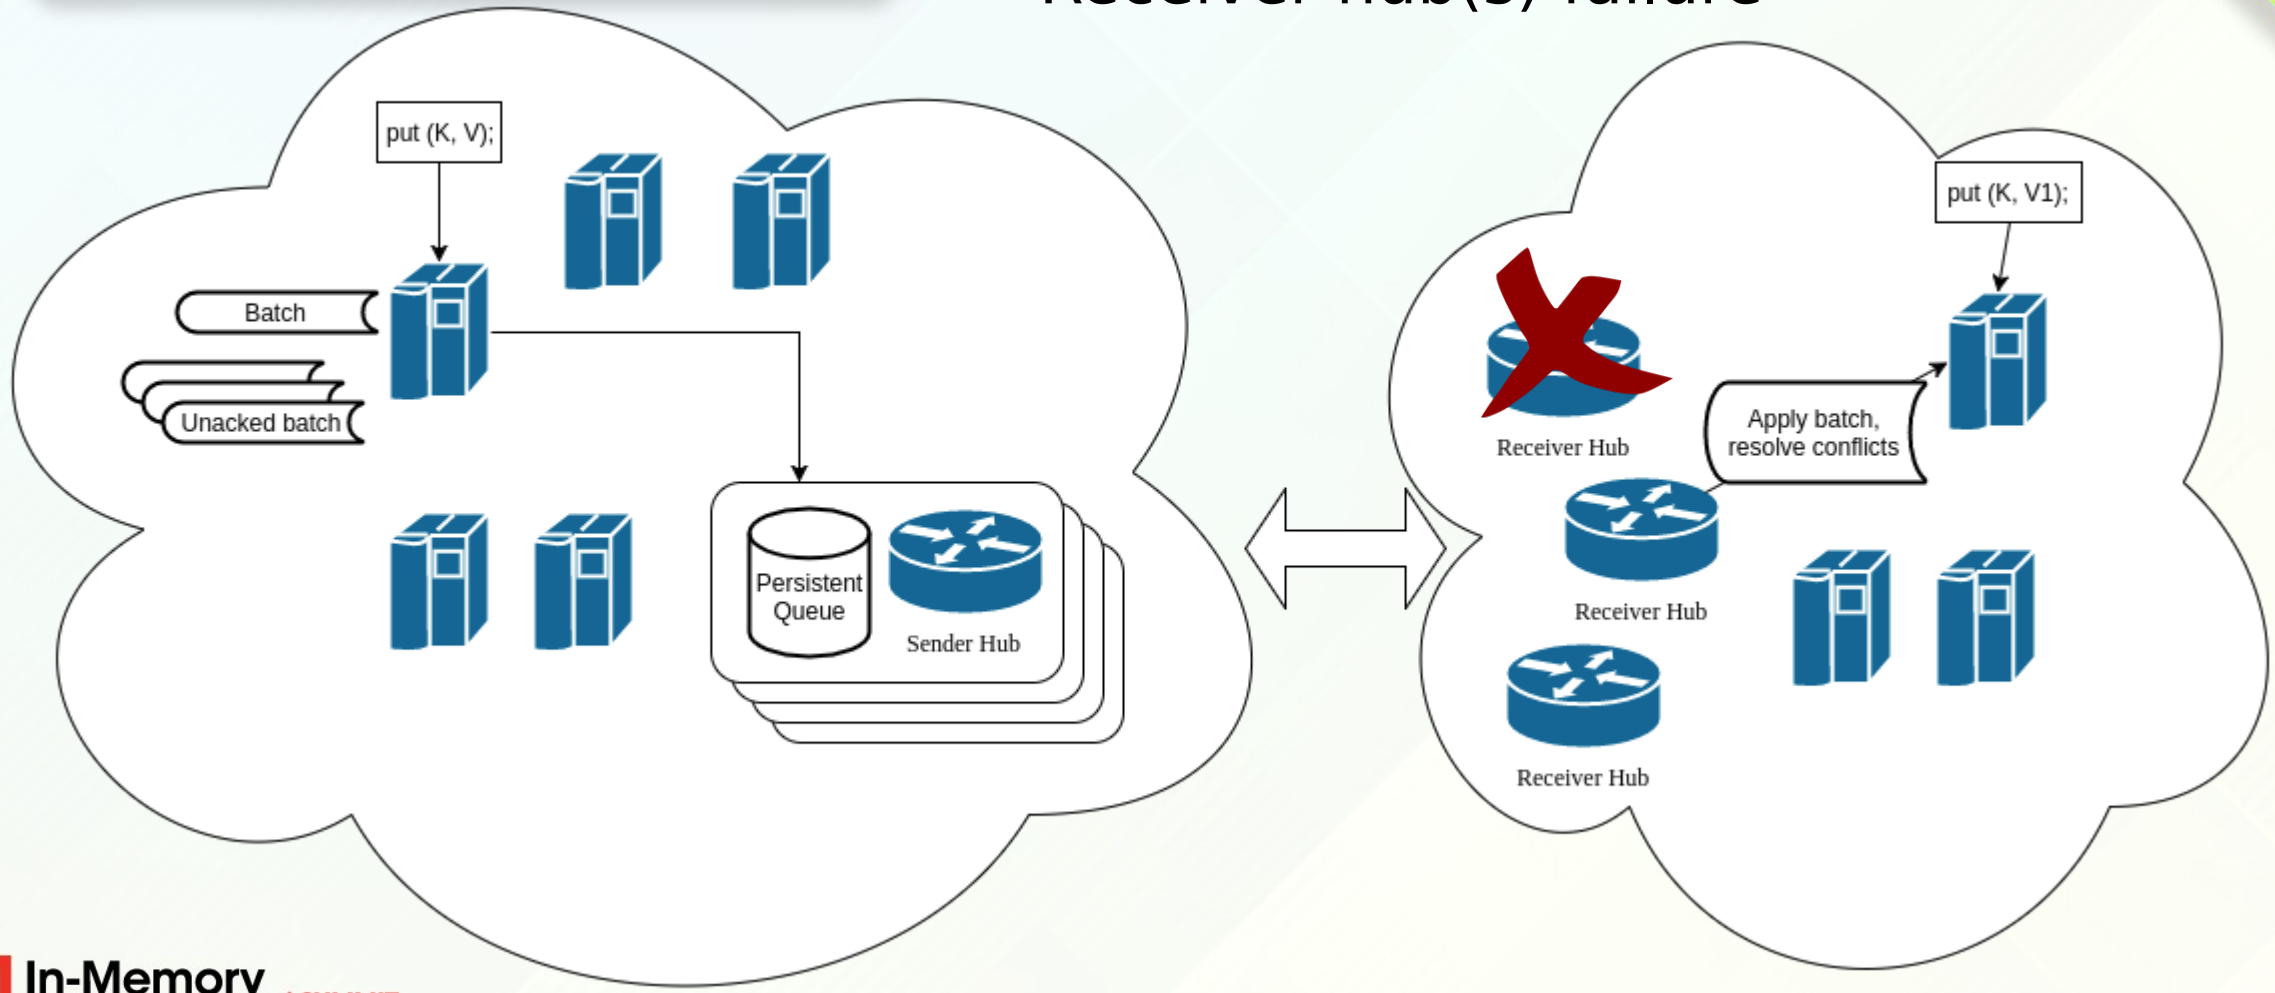

REPLICATION IN GRIDGAIN: WHAT CAN GO WRONG?

Node failure in receiving topology

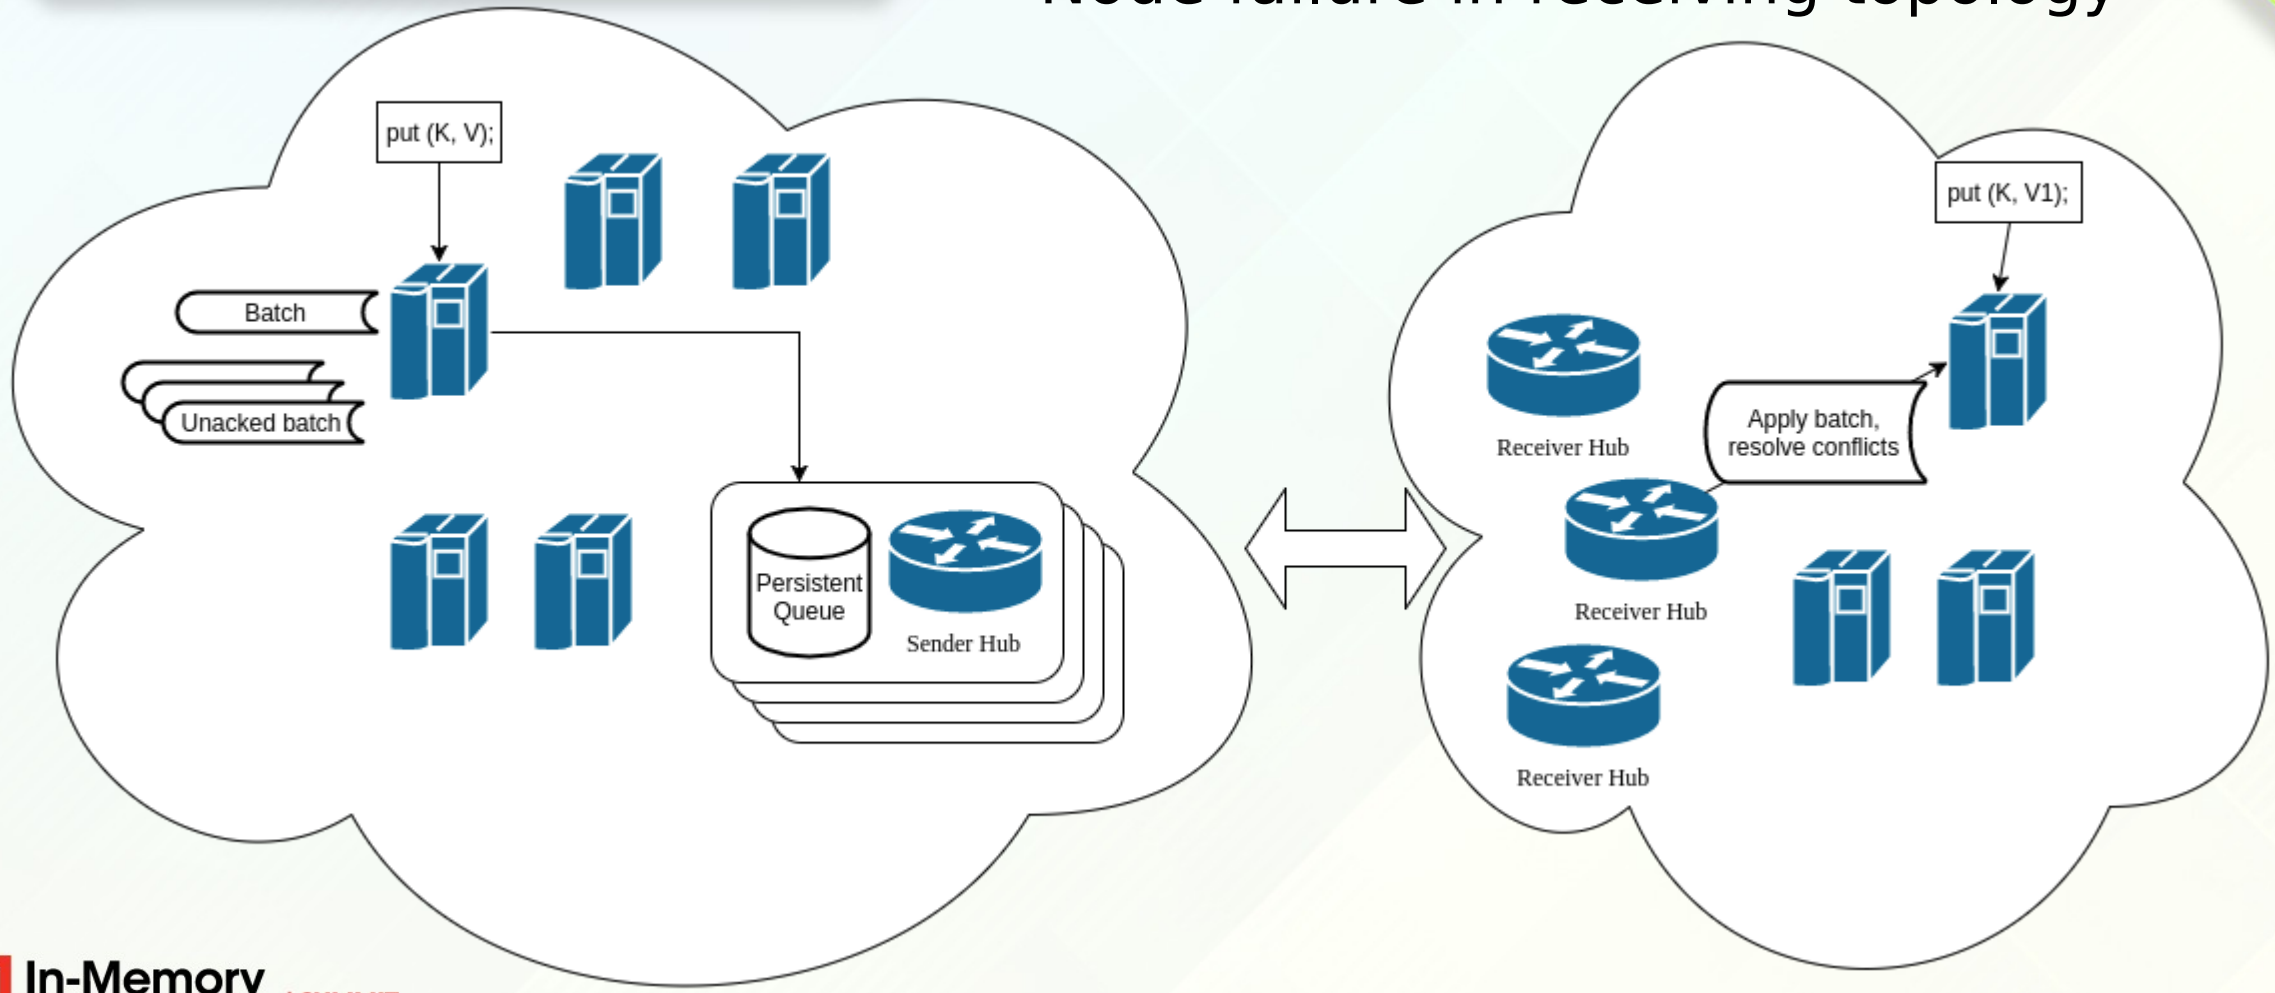

REPLICATION IN GRIDGAIN: WHAT CAN GO WRONG?

Node failure in receiving topology

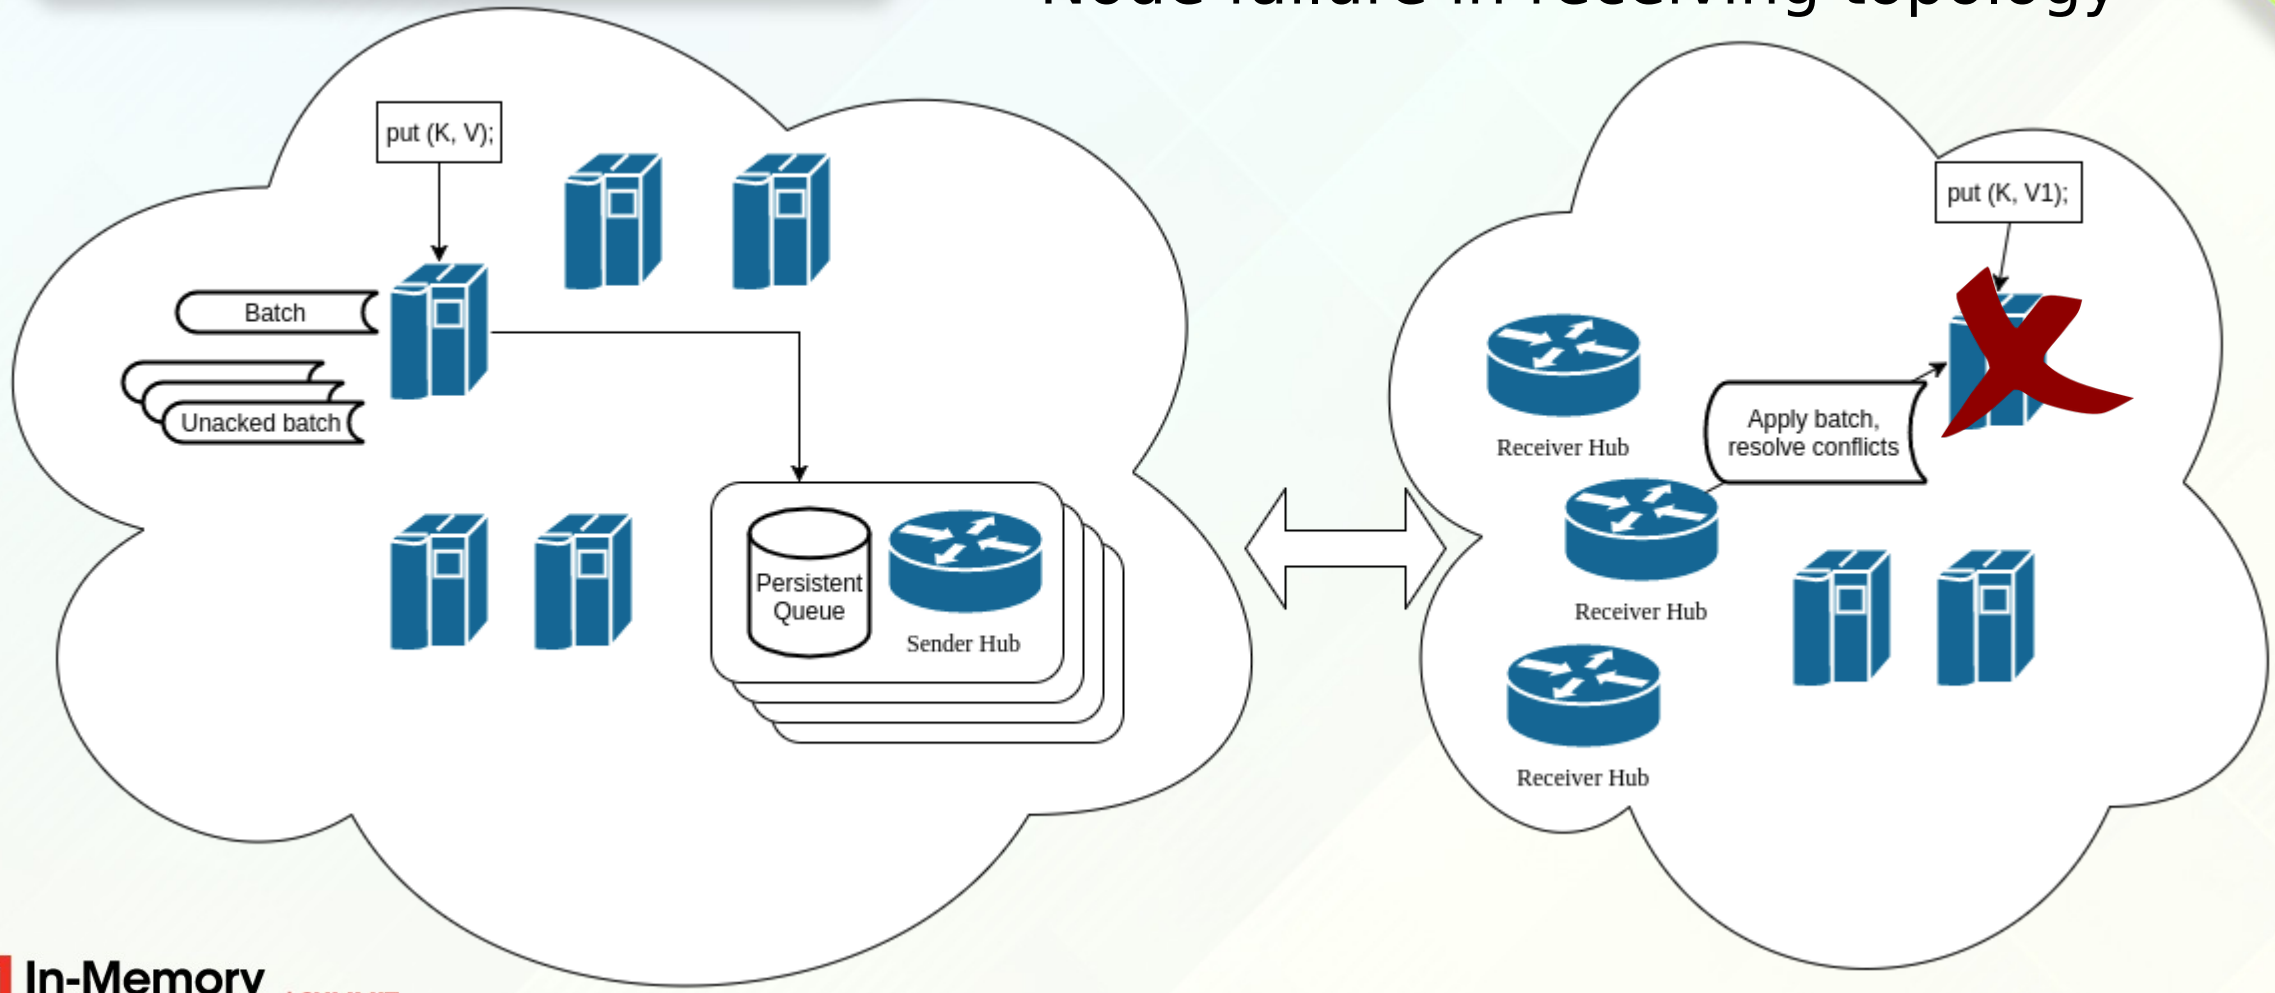

REPLICATION IN GRIDGAIN: WHAT CAN BE BETTER?

Batching on per-node basis vs per-partition basis

Per-node batching

- More efficient from memory standpoint
- Batches collected quickly

Per-partition batching

- No need for additional processing on receiving side
- Less contention – honors thread-per-partition model
- Probably, higher GC pressure
- But still expected to perform better

REPLICATION IN GRIDGAIN: FUTURE PLANS

TX Replication

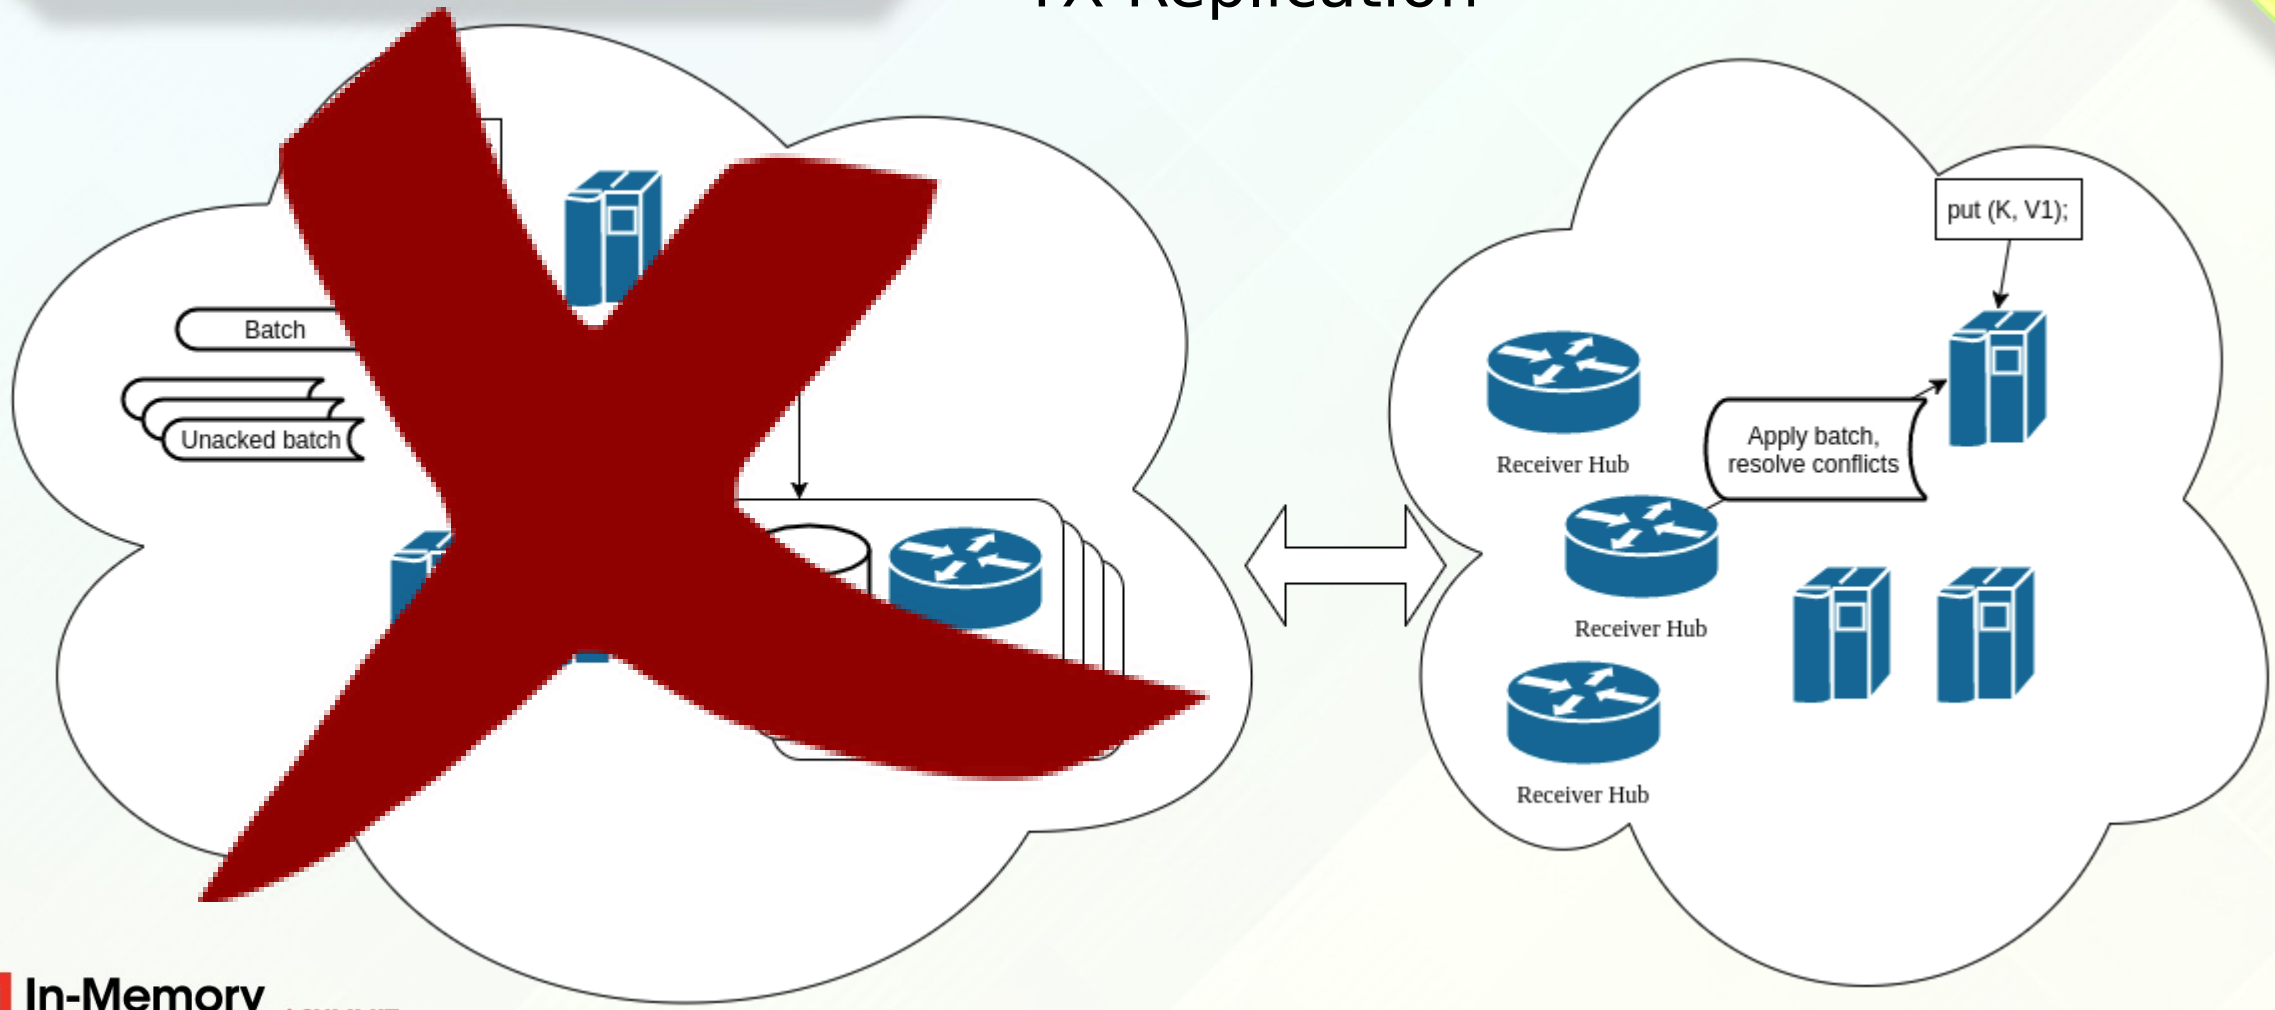

REPLICATION IN GRIDGAIN: FUTURE PLANS

TX Replication – over stretched cluster

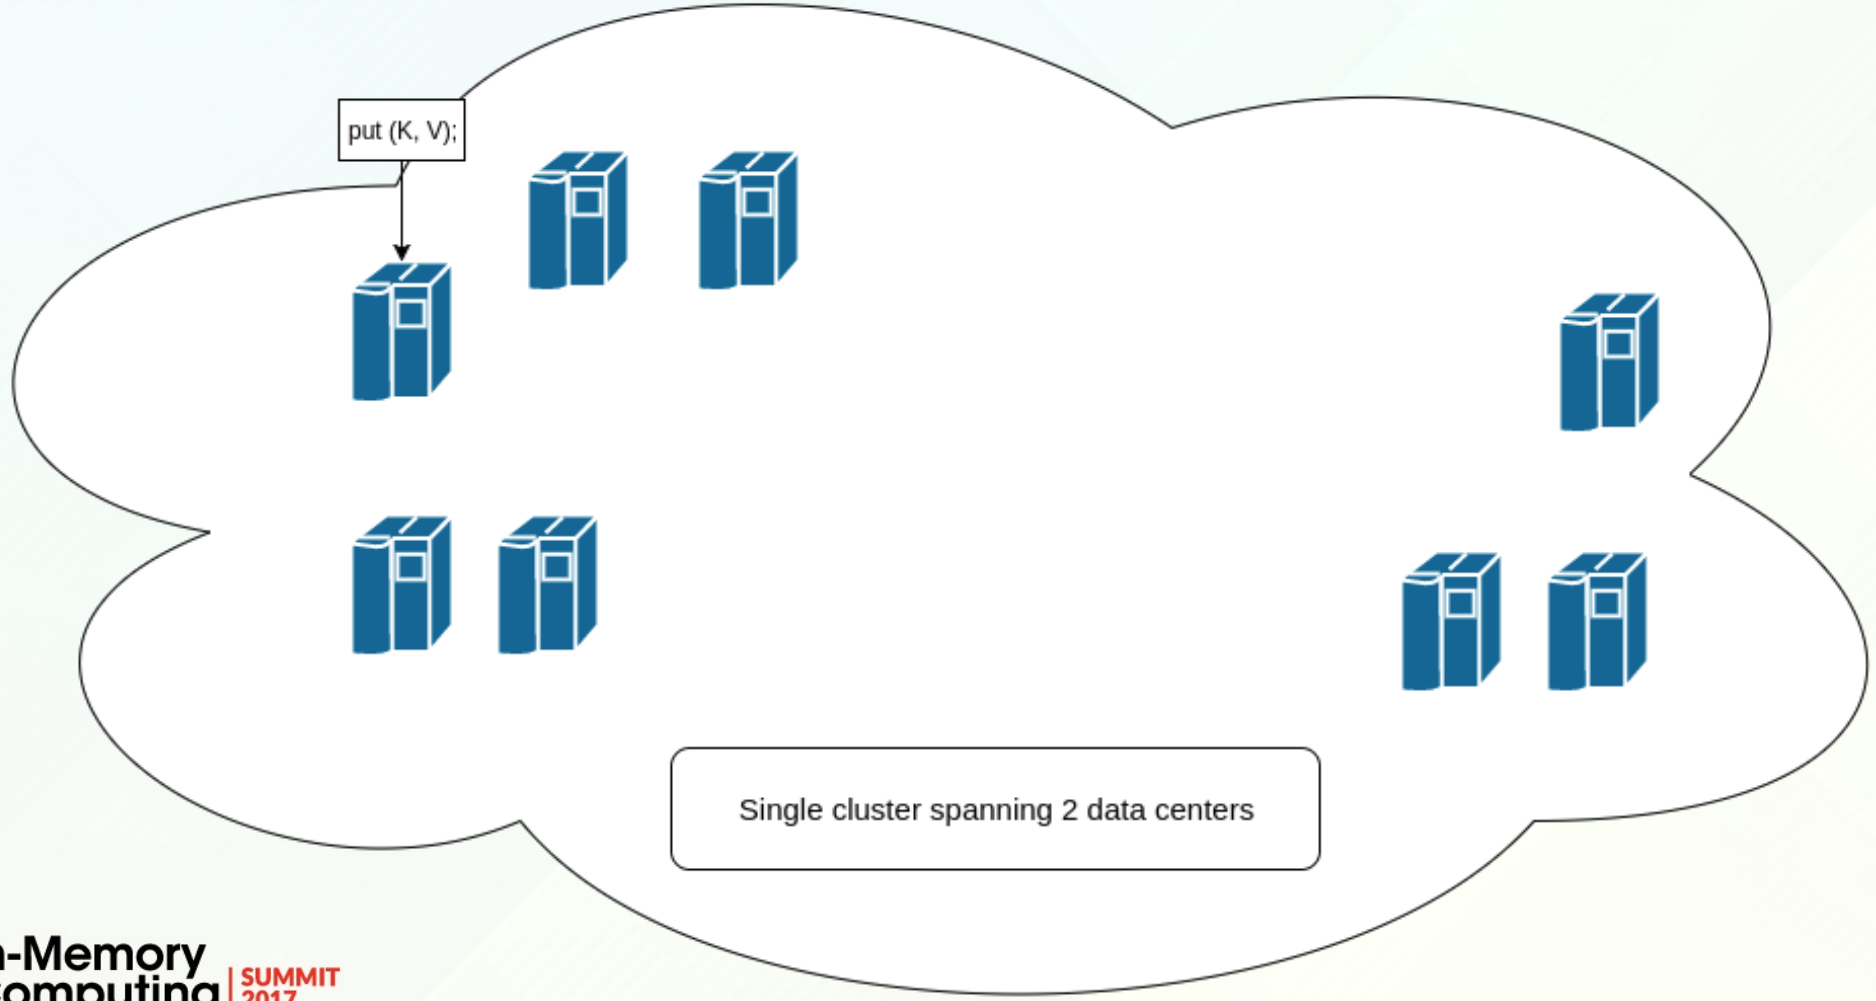

REPLICATION IN GRIDGAIN: FUTURE PLANS

TX replication mechanism

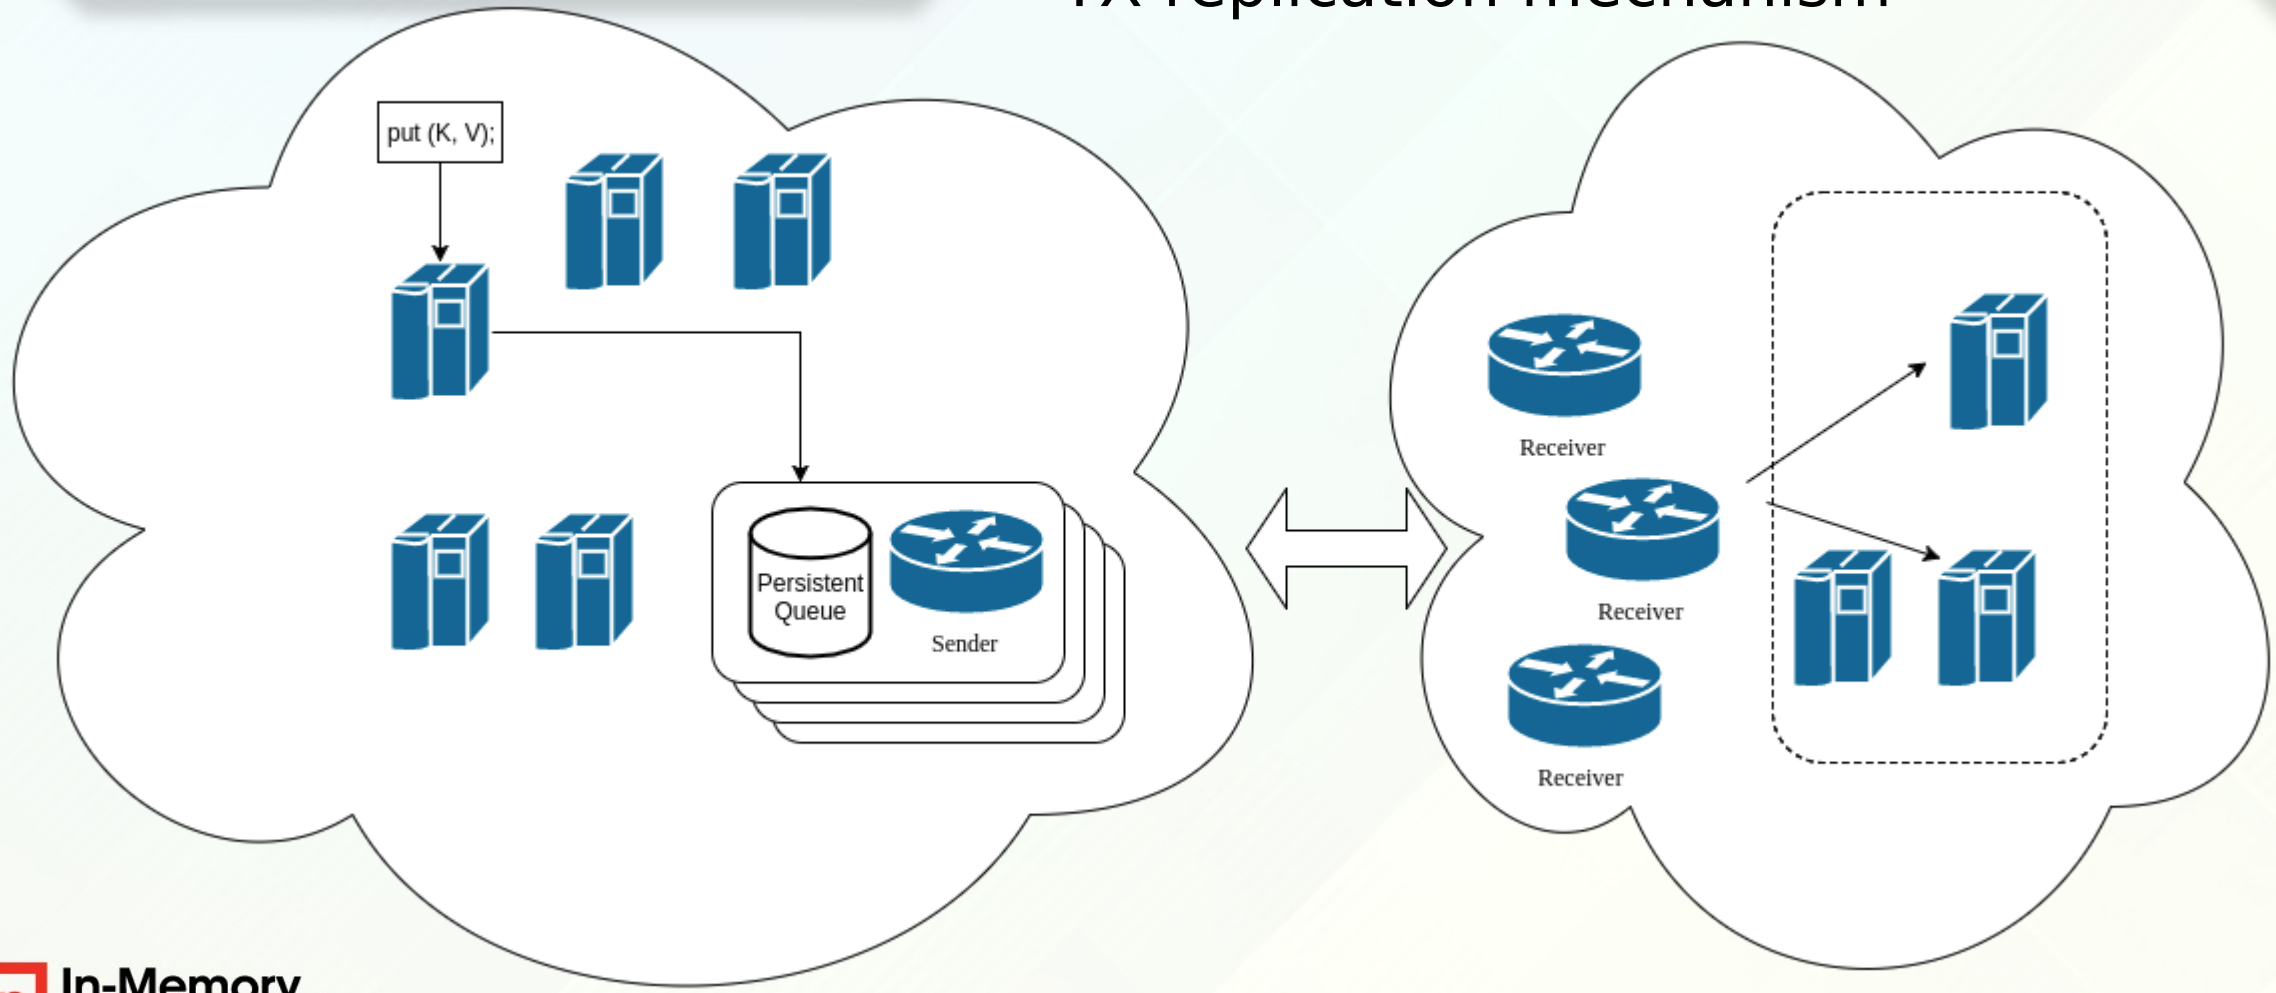

REPLICATION IN GRIDGAIN: FUTURE PLANS

TX Replication – sync modes

- Strict SYNC mode

Main primary node (PN) → Sender → Receiver → Stand-In PN →Receiver →Sender → Main PN

- Merciful SYNC mode

Main PN → Sender → Receiver → Receiver's WAL → Sender → Main PN

- ASYNC mode

Main PN->Sender->Sender WAL->Main PN

LESSONS LEARNED

- Is replication able to solve your problem?

LESSONS LEARNED

- Is replication able to solve your problem?
- Pickup proper settings: sync/async, physical vs logical changes.

LESSONS LEARNED

- Is replication able to solve your problem?
- Pickup proper settings: sync/async, physical vs logical changes.
- Be aware of internals – know what makes it work.

LESSONS LEARNED

- Is replication able to solve your problem?
- Pickup proper settings: sync/async, physical vs logical changes.
- Be aware of internals – know what makes it work.
- Make sure to test, tune and monitor.

CONTACTS

yzhdanov@gridgain.com

http://ignite.apache.org

dev@ignite.apache.org
user@ignite.apache.org

QUESTIONS?

ANY QUESTIONS?