GridGain

Best Practices For Disaster Recovery and High Availability

Stan Lukyanov
Customer Solutions, GridGain Systems

Agenda

- What is Disaster Recovery?
- Disaster Recovery Options
- Advanced Topics and FAQ

Agenda

- **What is Disaster Recovery?**
- Disaster Recovery Options
- Advanced Topics and FAQ

What is Disaster Recovery?

What is Disaster Recovery?

What is Disaster Recovery?

What is Disaster Recovery?

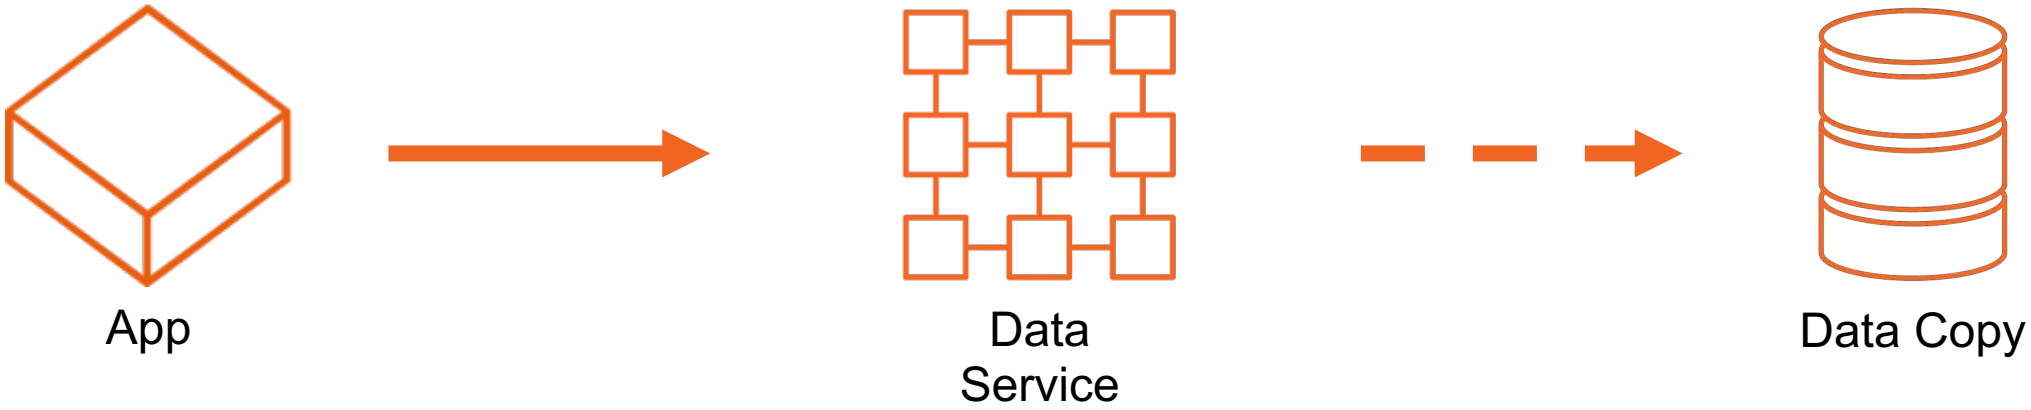

Adding High Availability

Adding High Availability

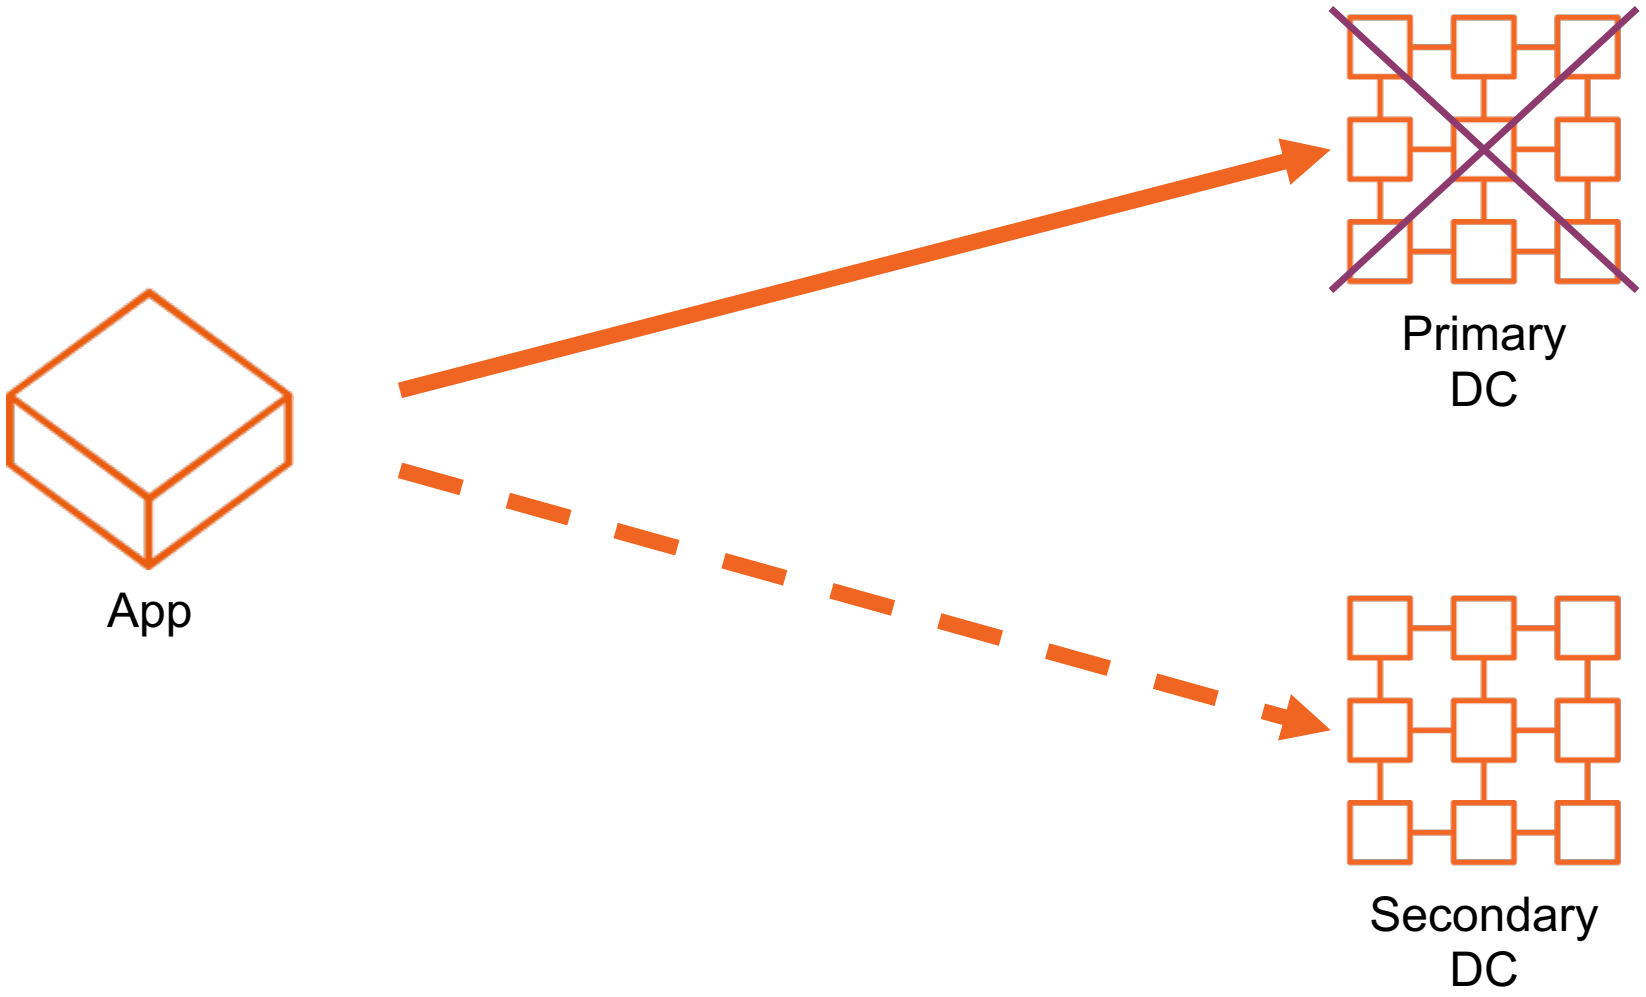

Generic Architecture for Disaster Recovery

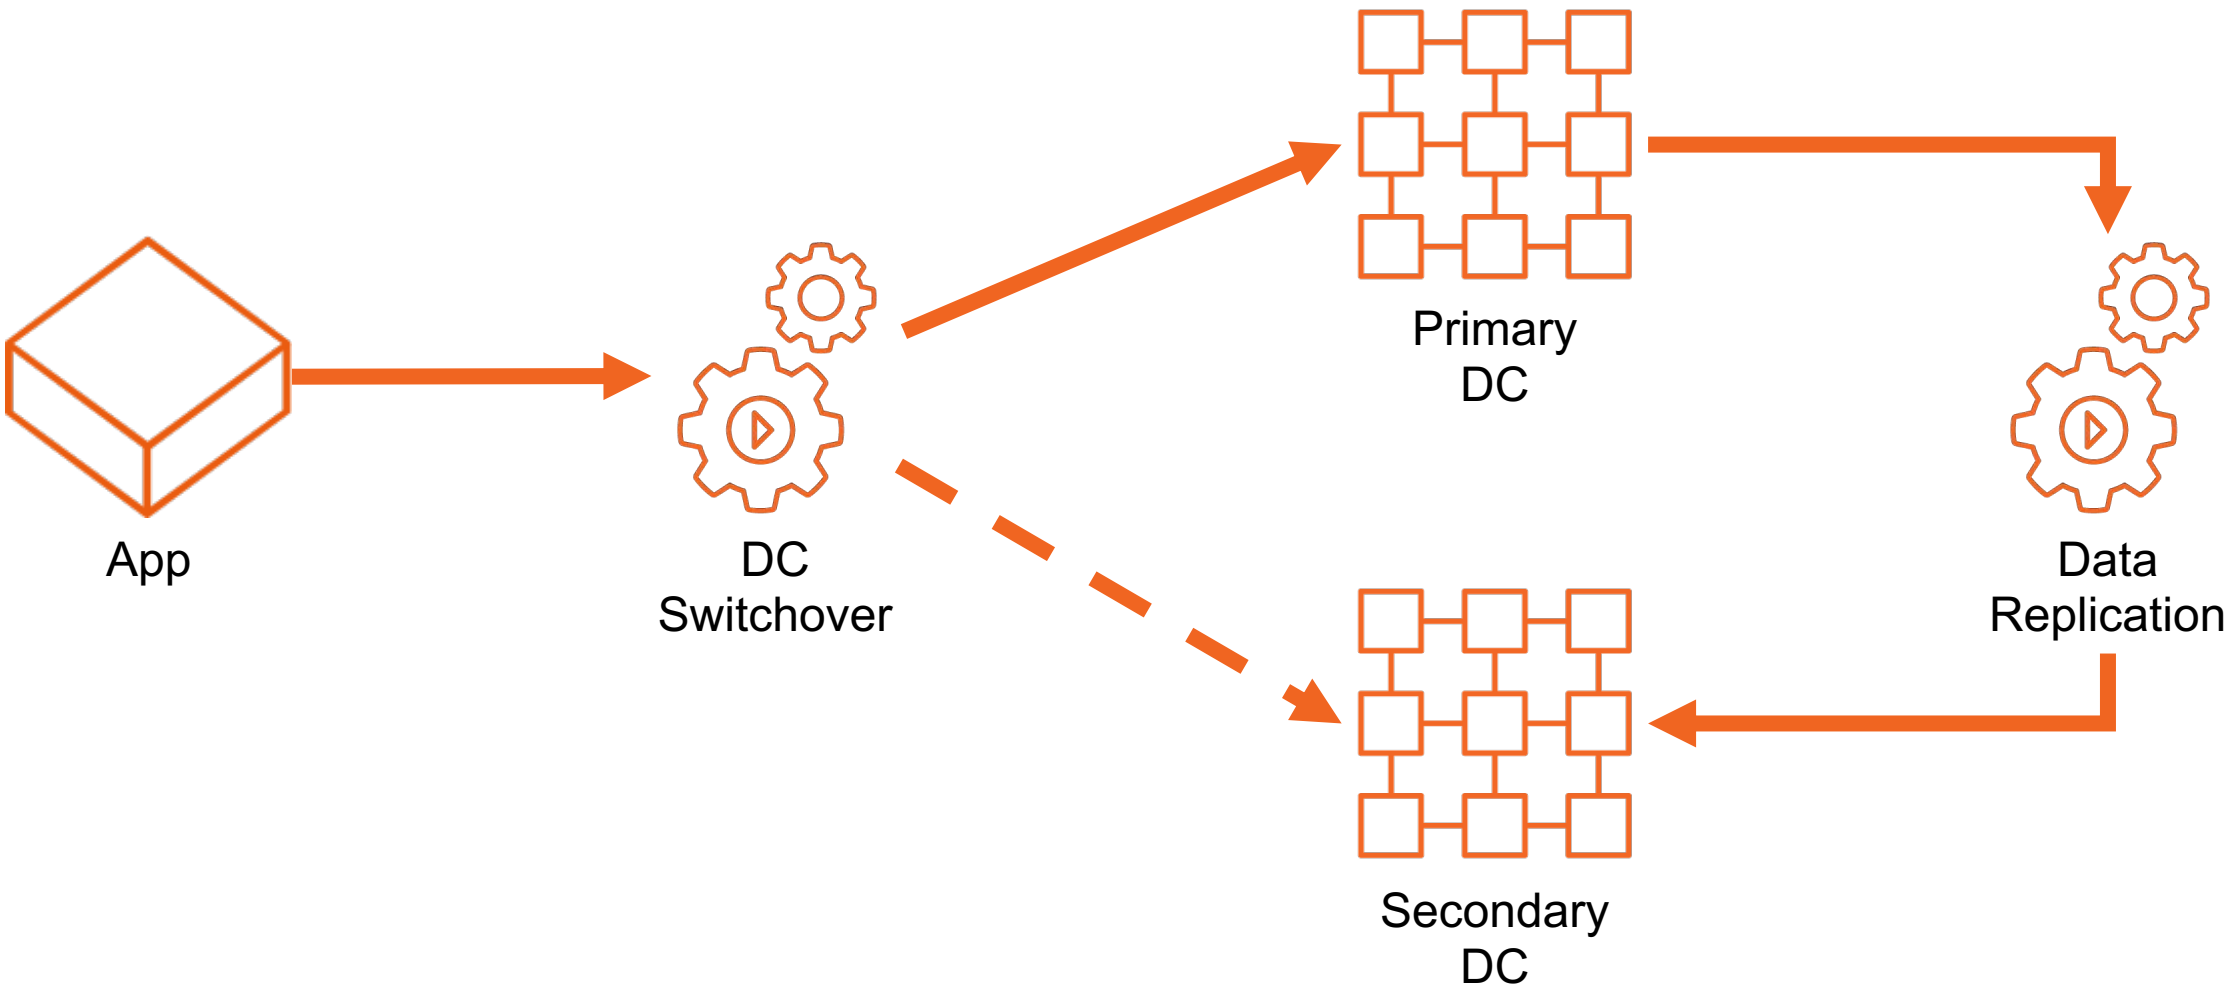

GridGain

How to Compare DR Solutions?

- How long the **service is down** in case of a disaster?
- Is there a **data loss** in case of a disaster?
- How much does the solution **cost**?

Disaster Timeline

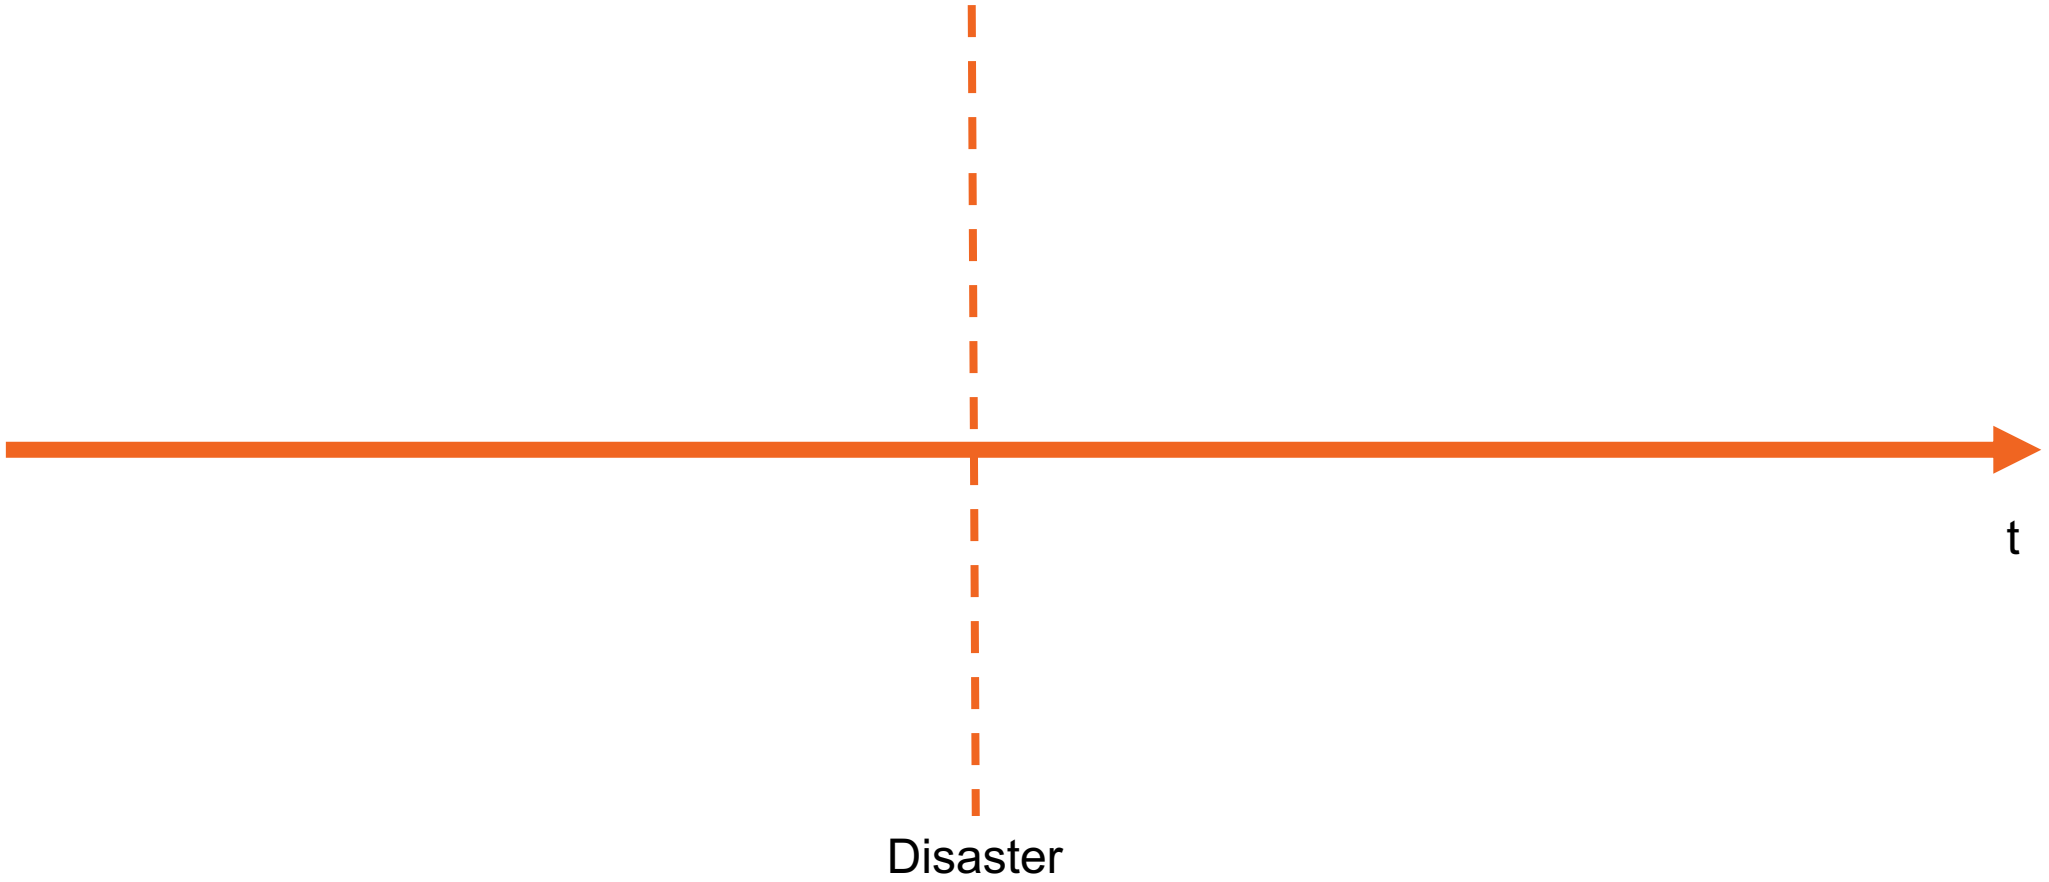

GridGain

Disaster Timeline

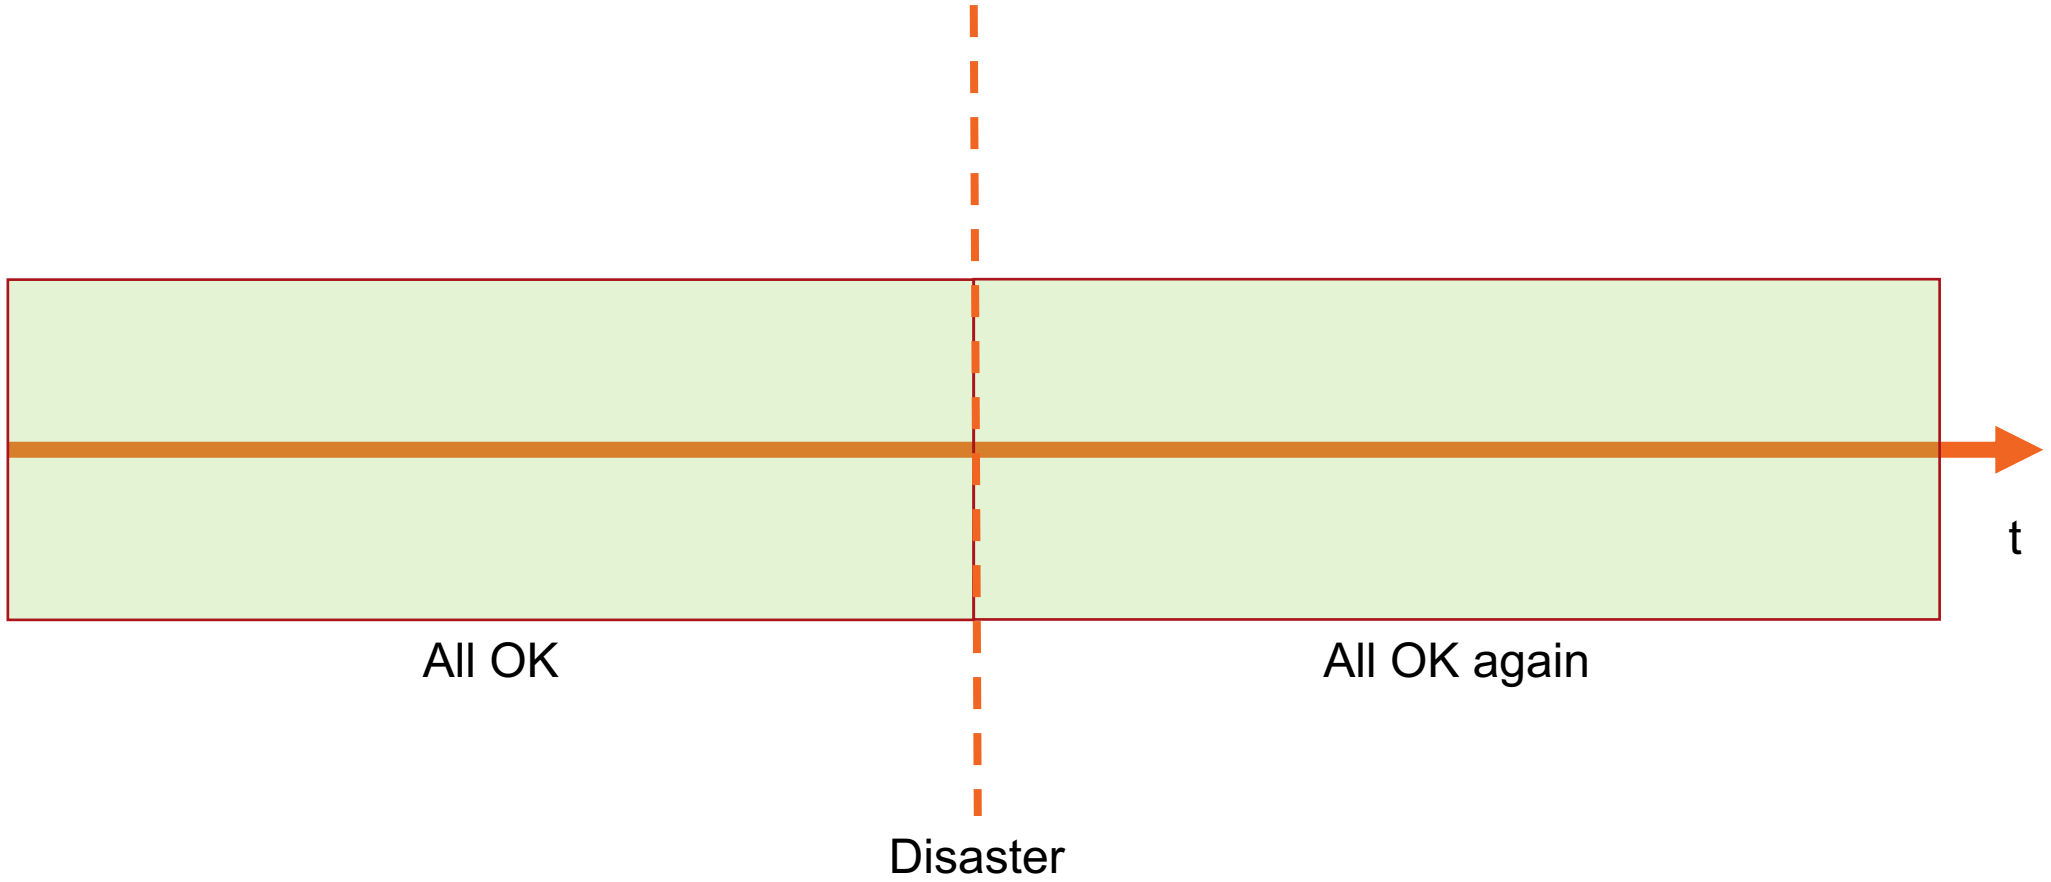

GridGain

Disaster Timeline

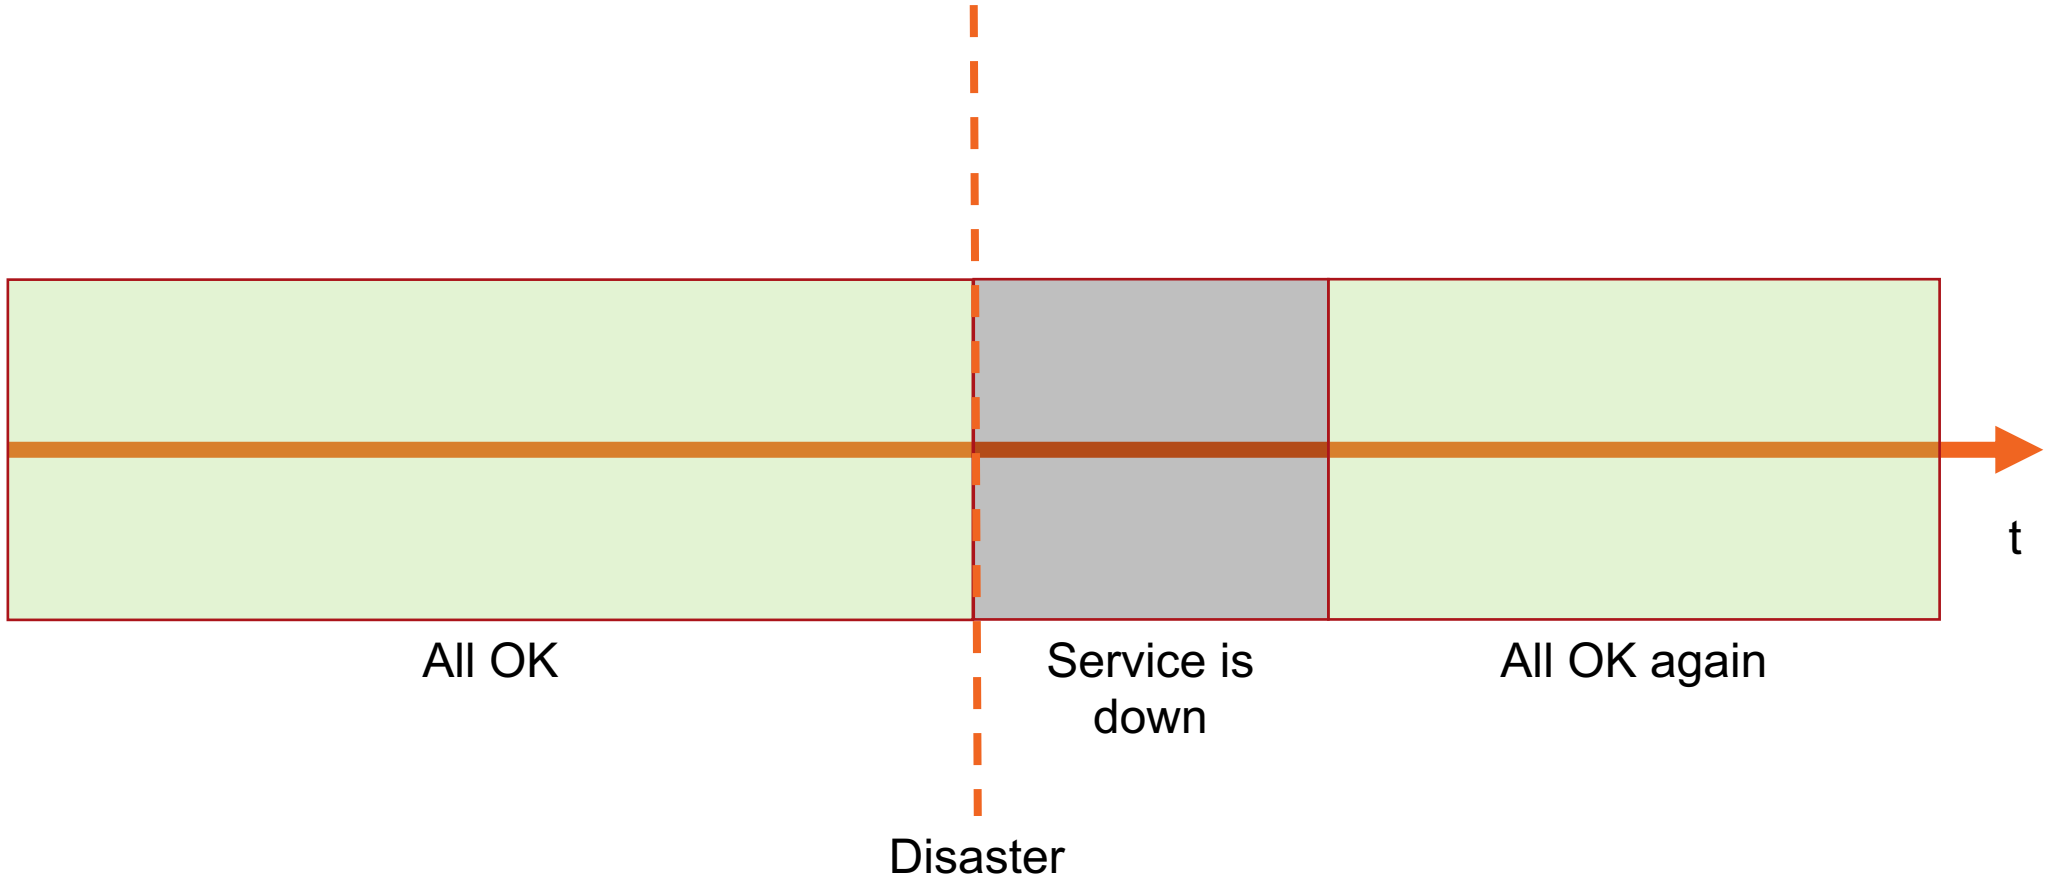

GridGain

Disaster Timeline

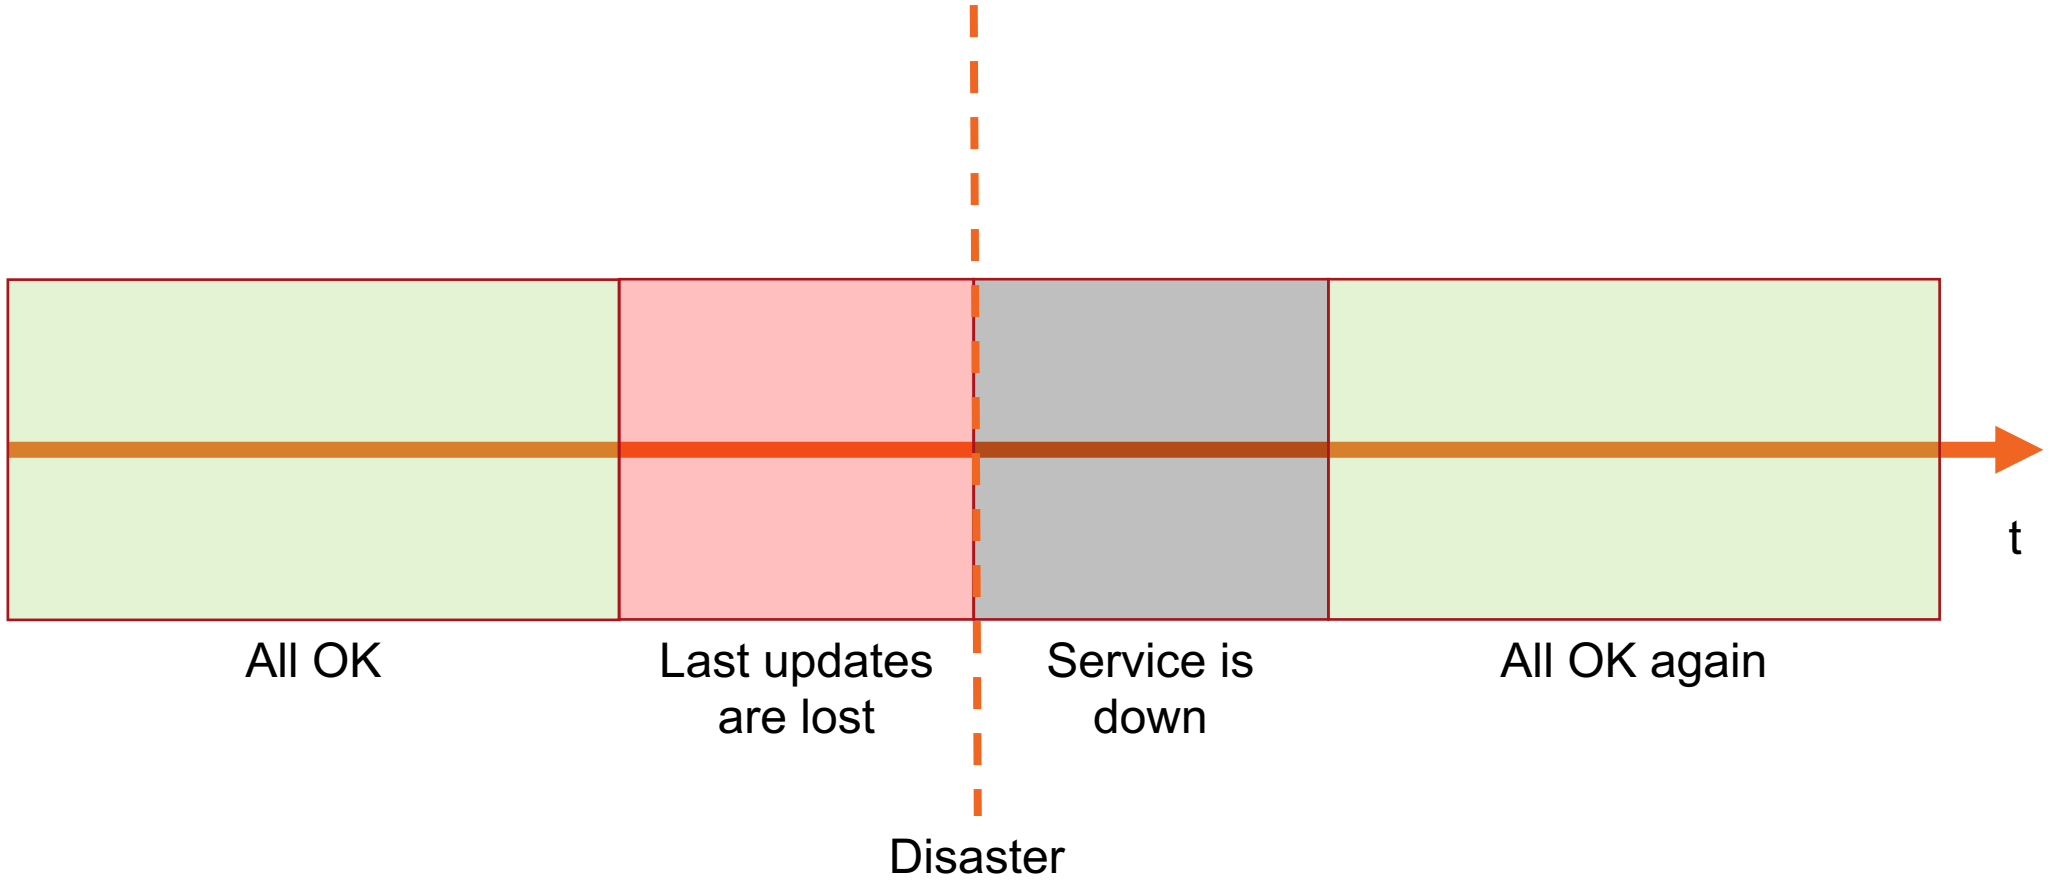

Disaster Timeline

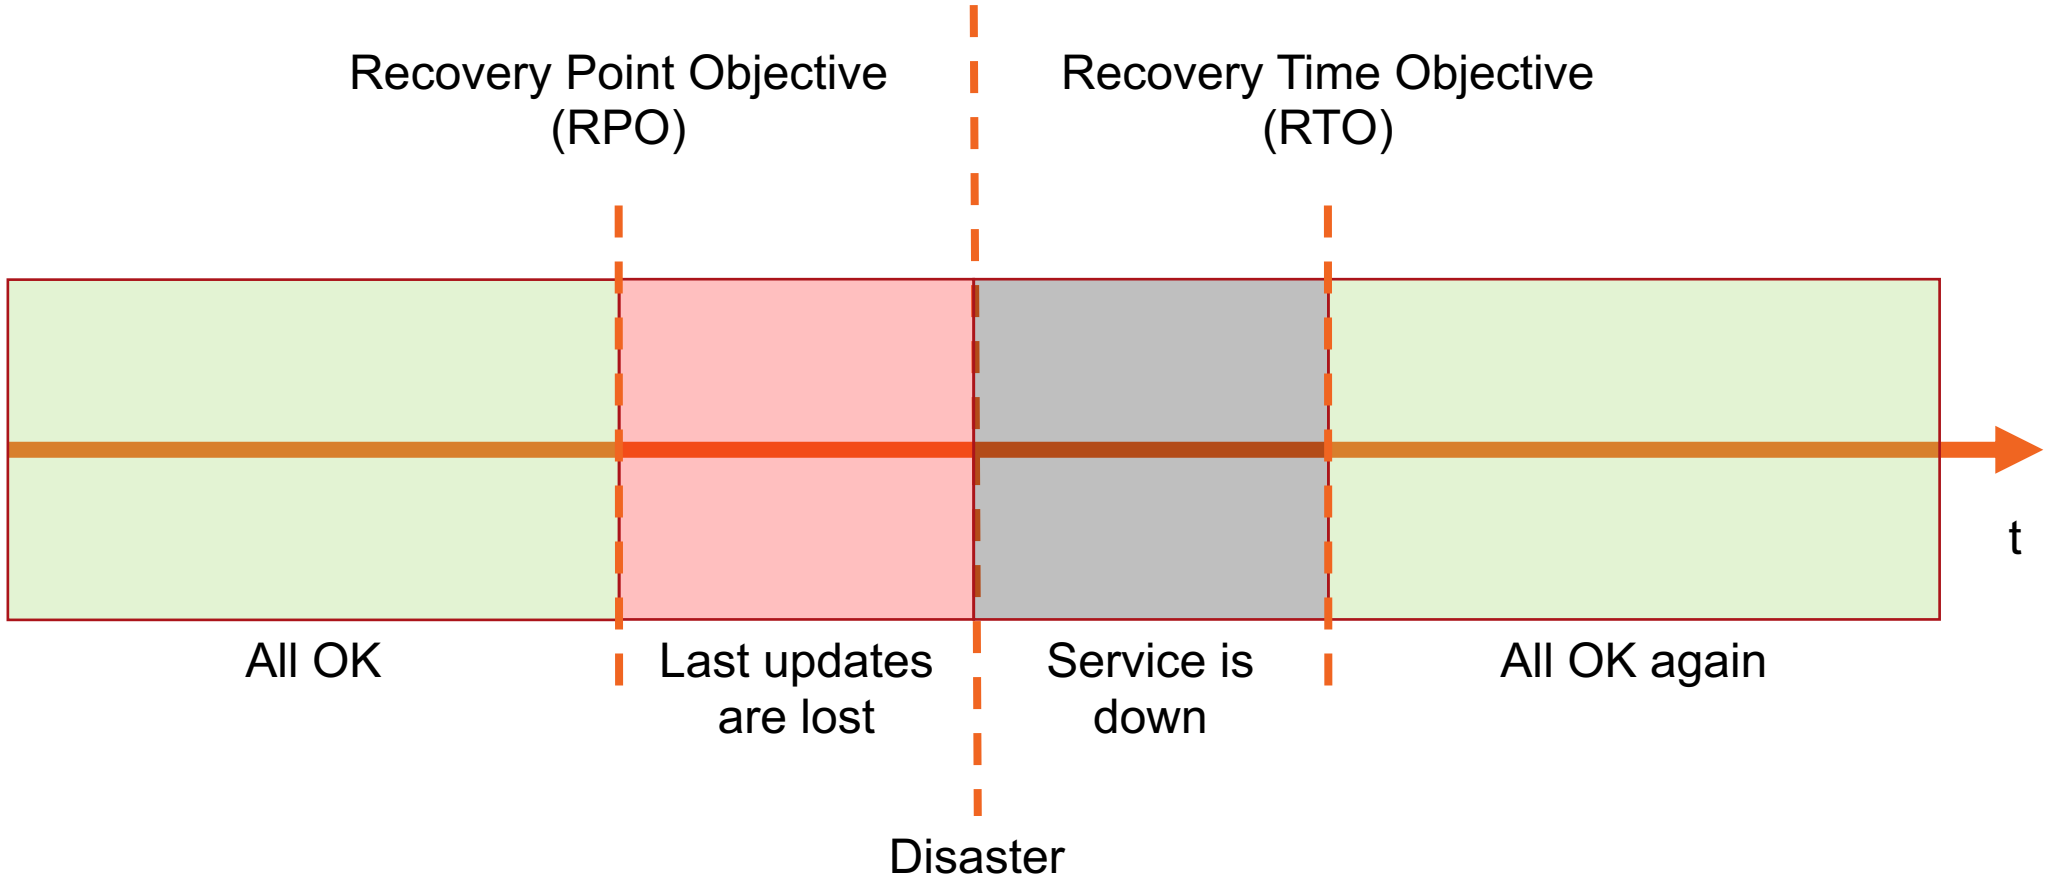

GridGain

Recovery Point Objective (RPO)

- Maximum time for which data is allowed to be lost
- Defined by the replication lag
- RPO = 0 – replication is **synchronous**, i.e. updates happen in both DCs at the same time

Recovery Time Objective (RTO)

- Maximum allowed service interruption time
- RTO = 0 – clients can switch **instantly** AND second DC is in **standby**
- In real life it is either:
  - RTO = client switching time (seconds)
  - RTO = second DC startup time (minutes to hours)

Solutions Comparison

Solution	RPO	RTO	Cost
???	0 to hours	0* to hours	Minimal cost – second DC

Agenda

- What is Disaster Recovery?
- **Disaster Recovery Options**
- Advanced Topics and FAQ

Application-Level Solutions

Application-Level Solutions

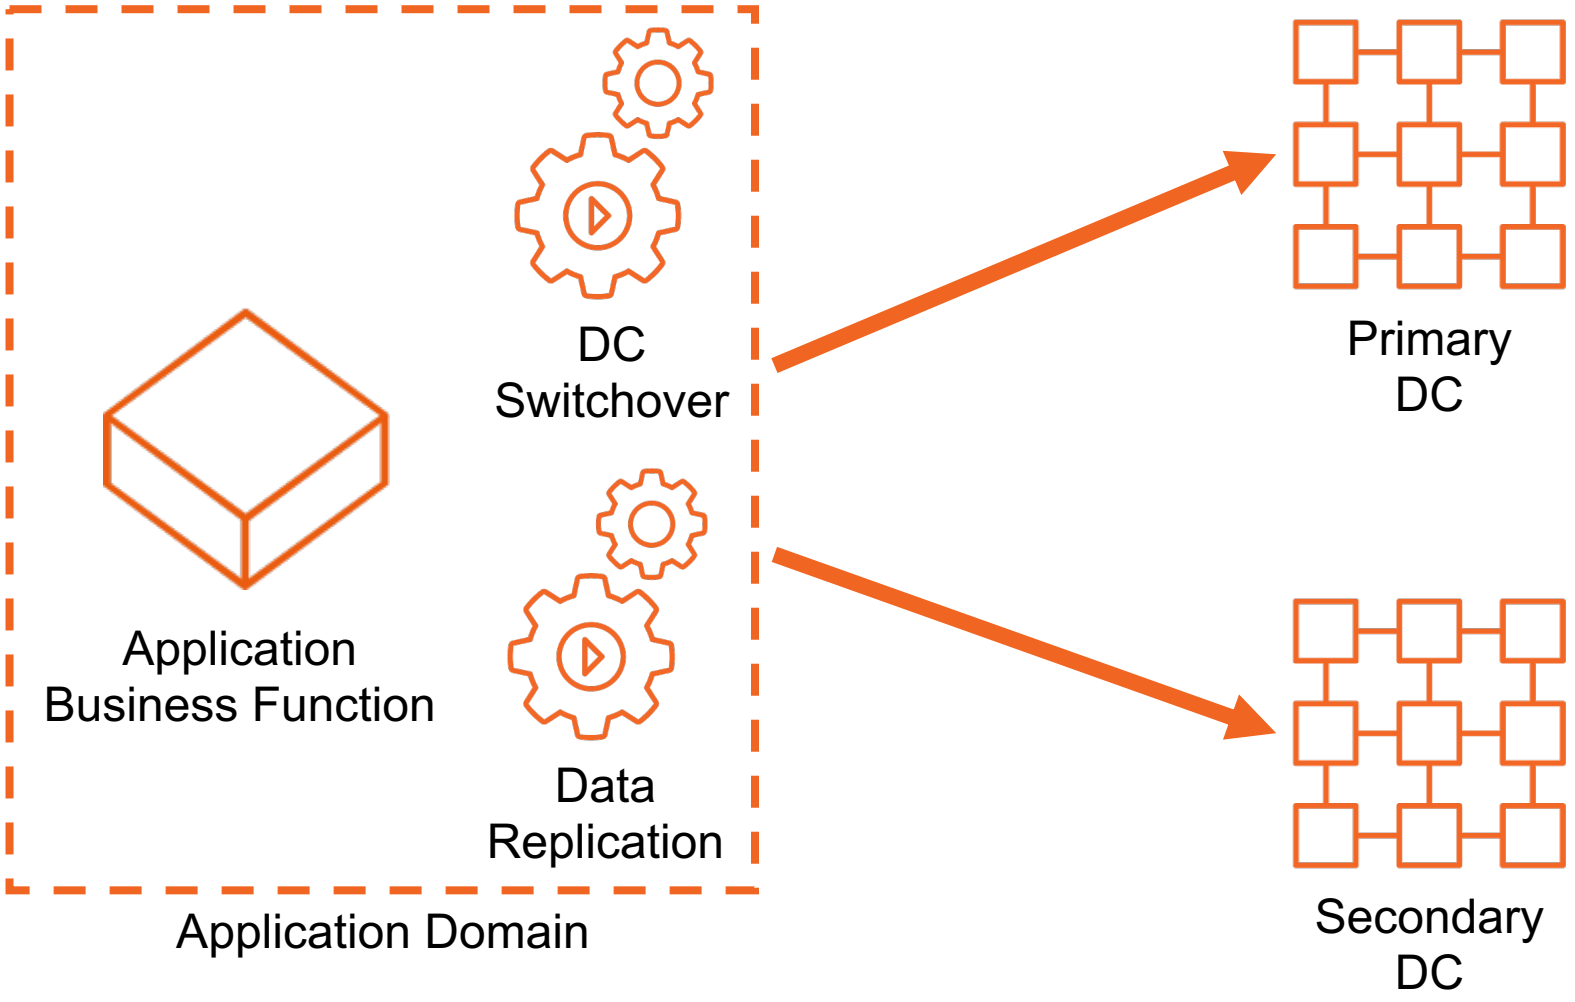

Application-Level Solutions

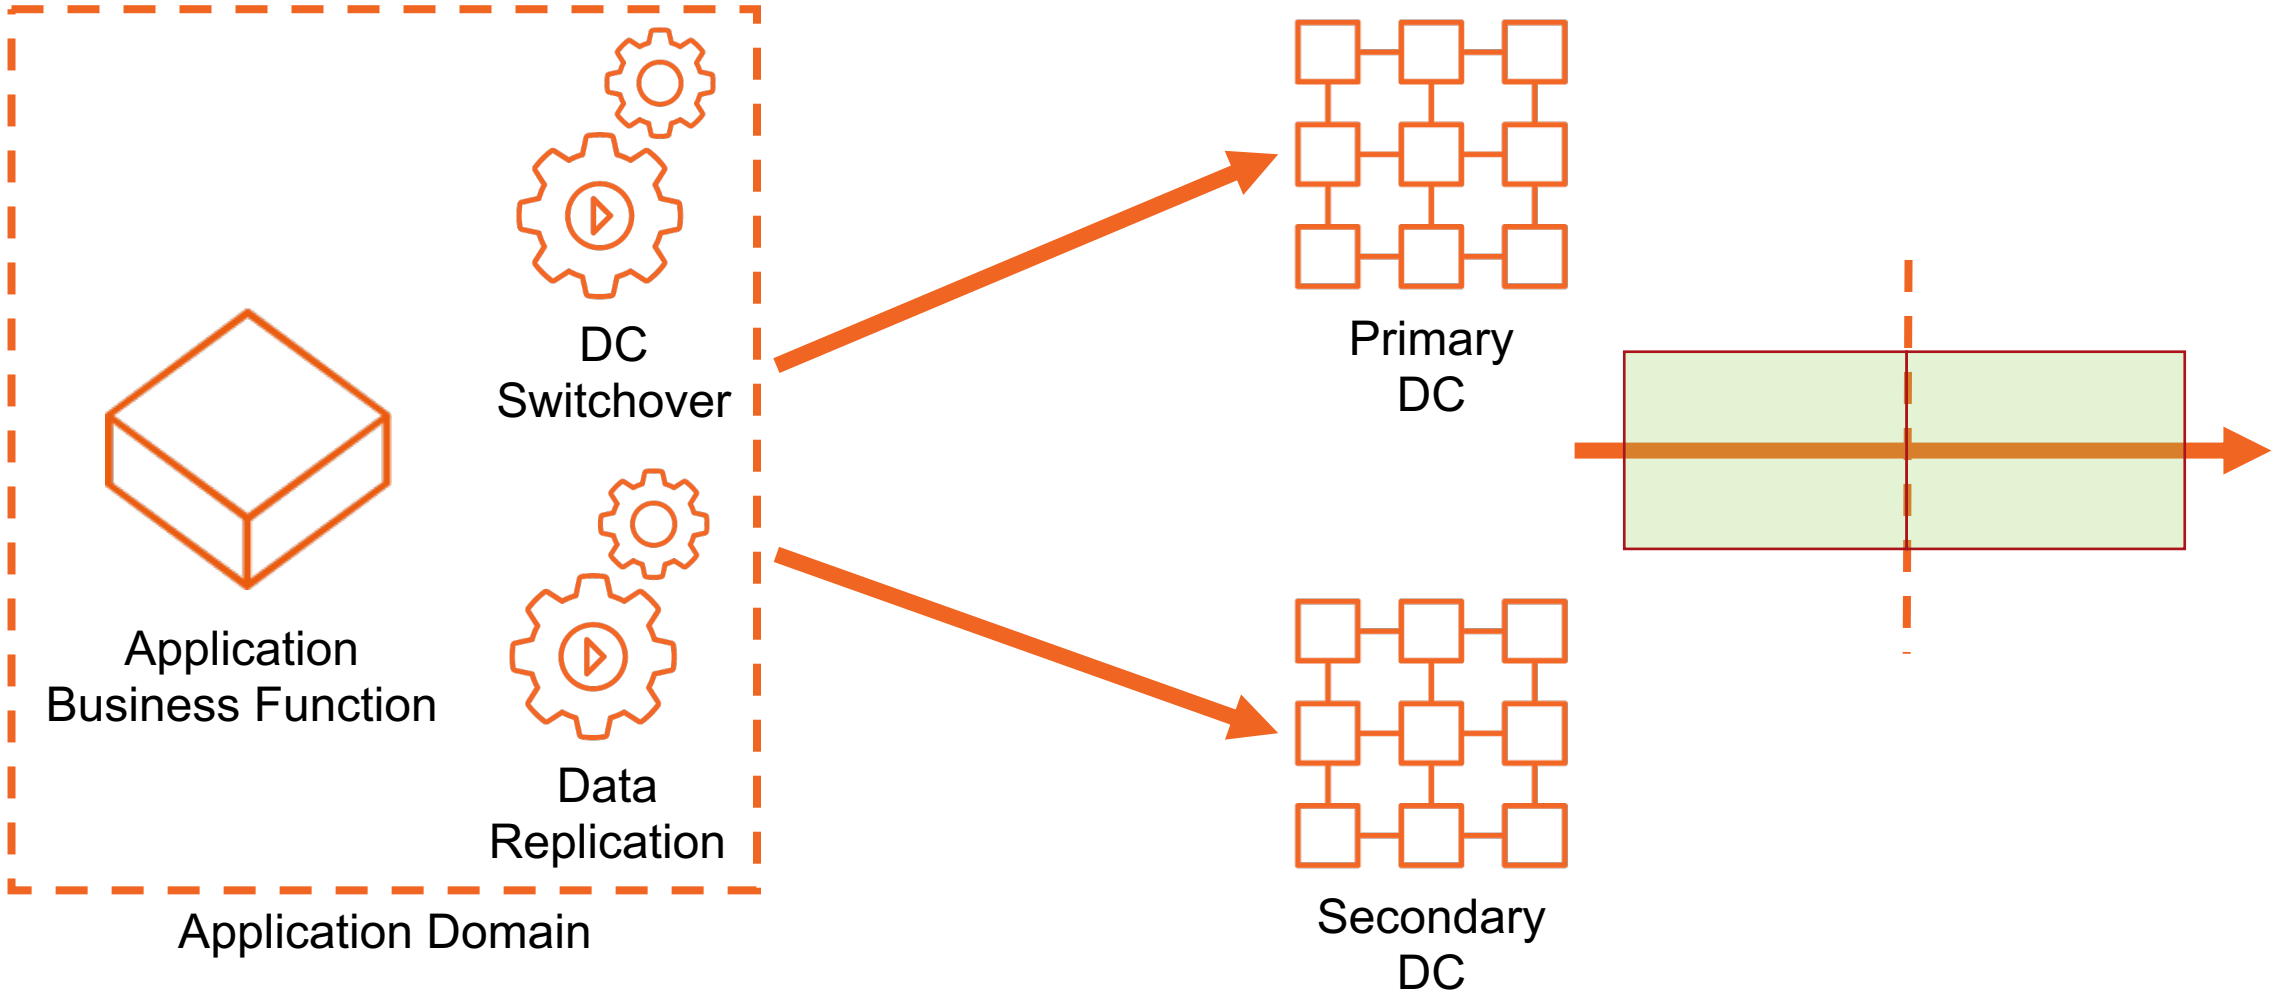

GridGain

Solutions Comparison

Solution	RPO	RTO	Cost
Application-Level	0	0	- Full DIY

System-Level Solutions

System-Level Solutions

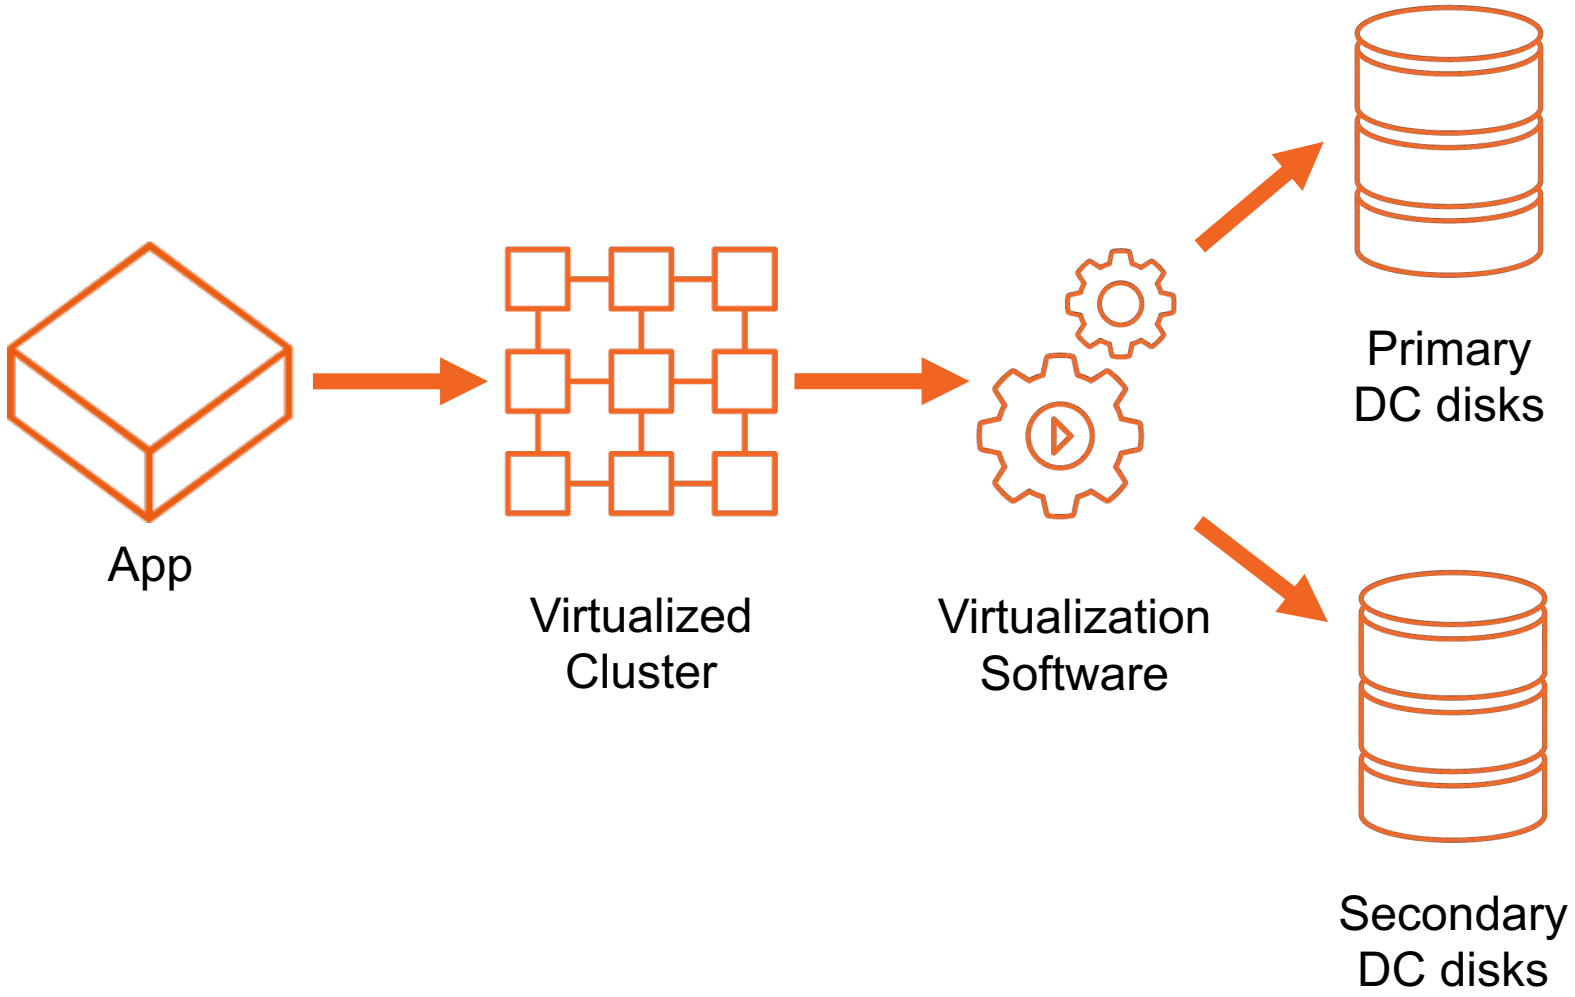

System-Level Solutions

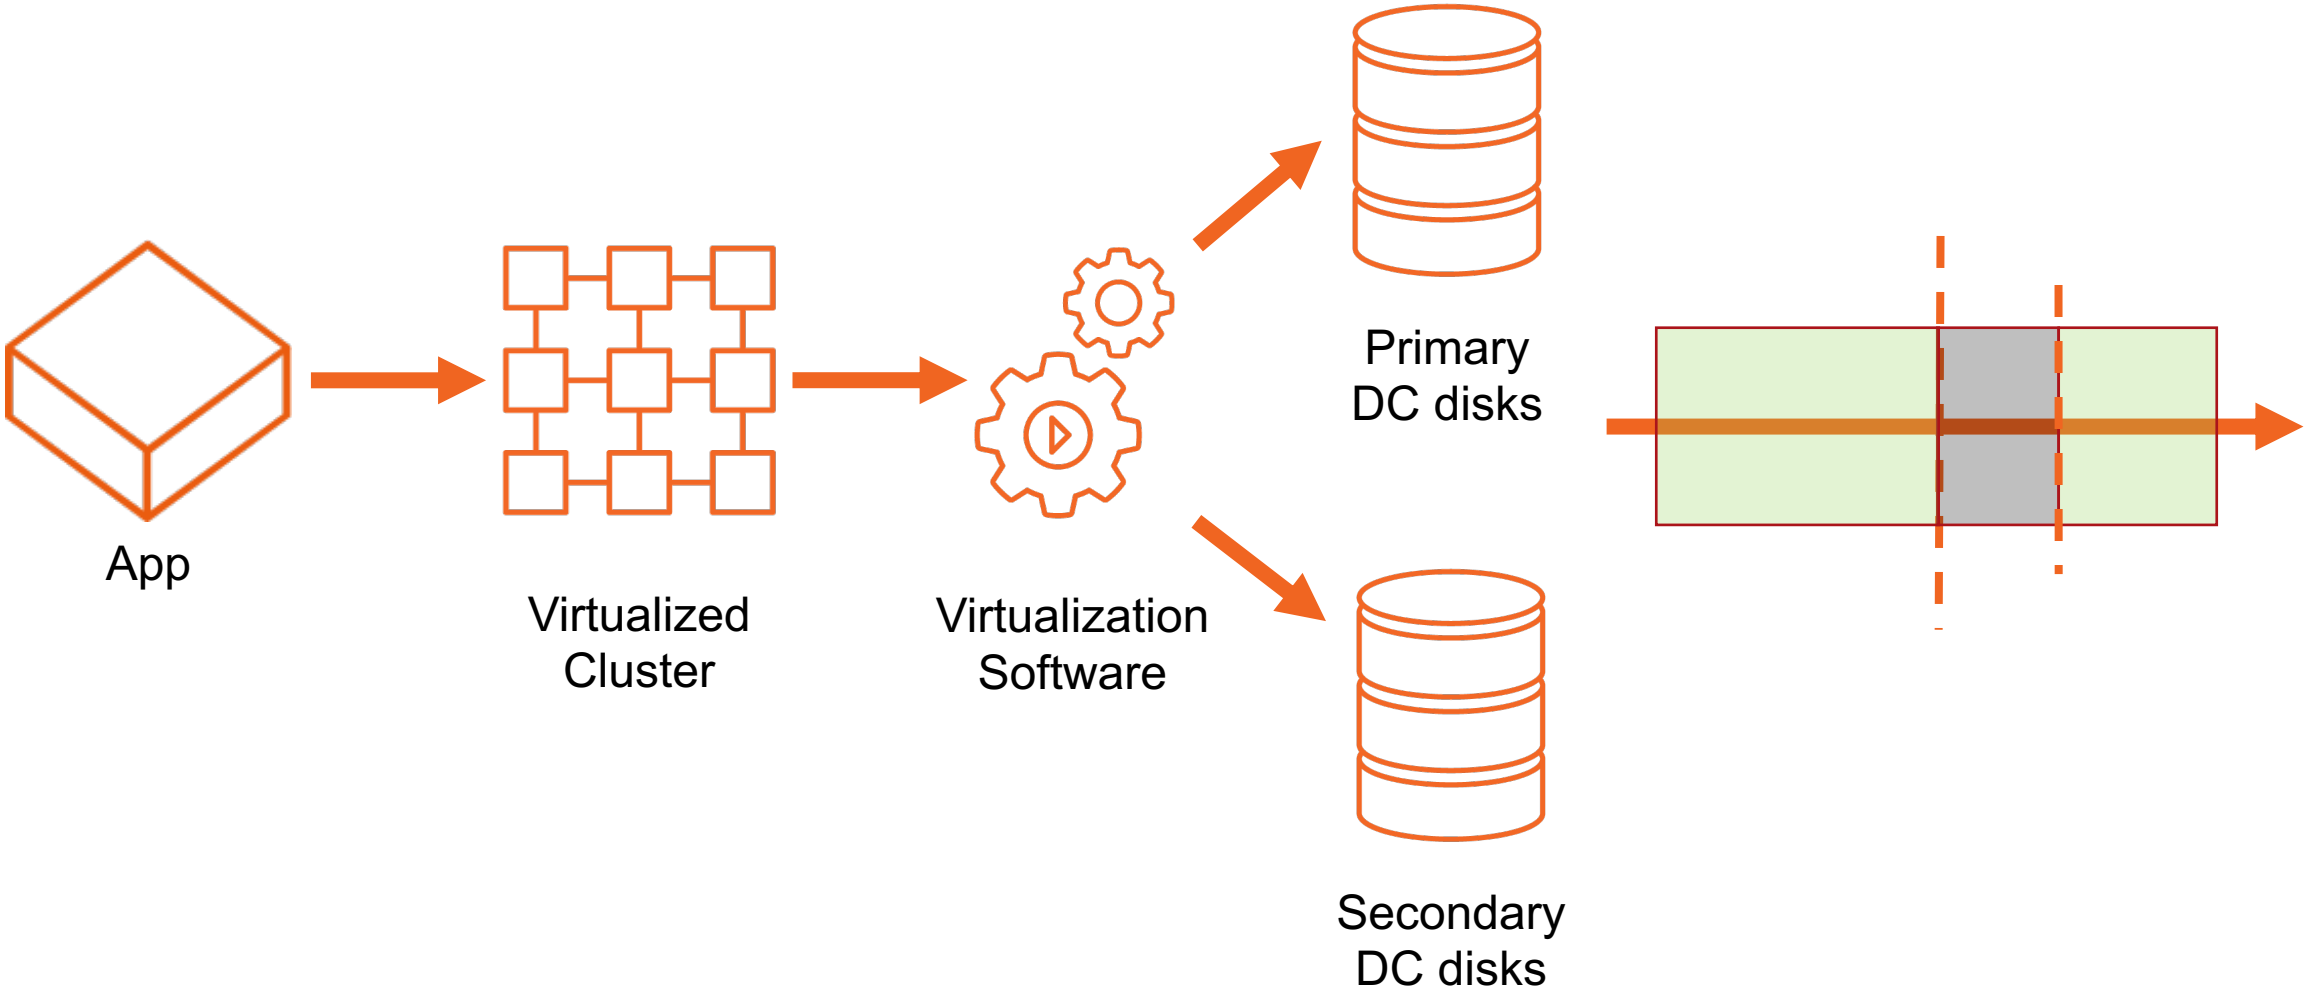

GridGain

Solutions Comparison

Solution	RPO	RTO	Cost
Application-Level	0	0	- Full DIY
System-Level	0 to minutes	Minutes	- VM/Cloud solution - Huge reliance on network for RPO=0

Backup-Based DR

Backup-Based DR

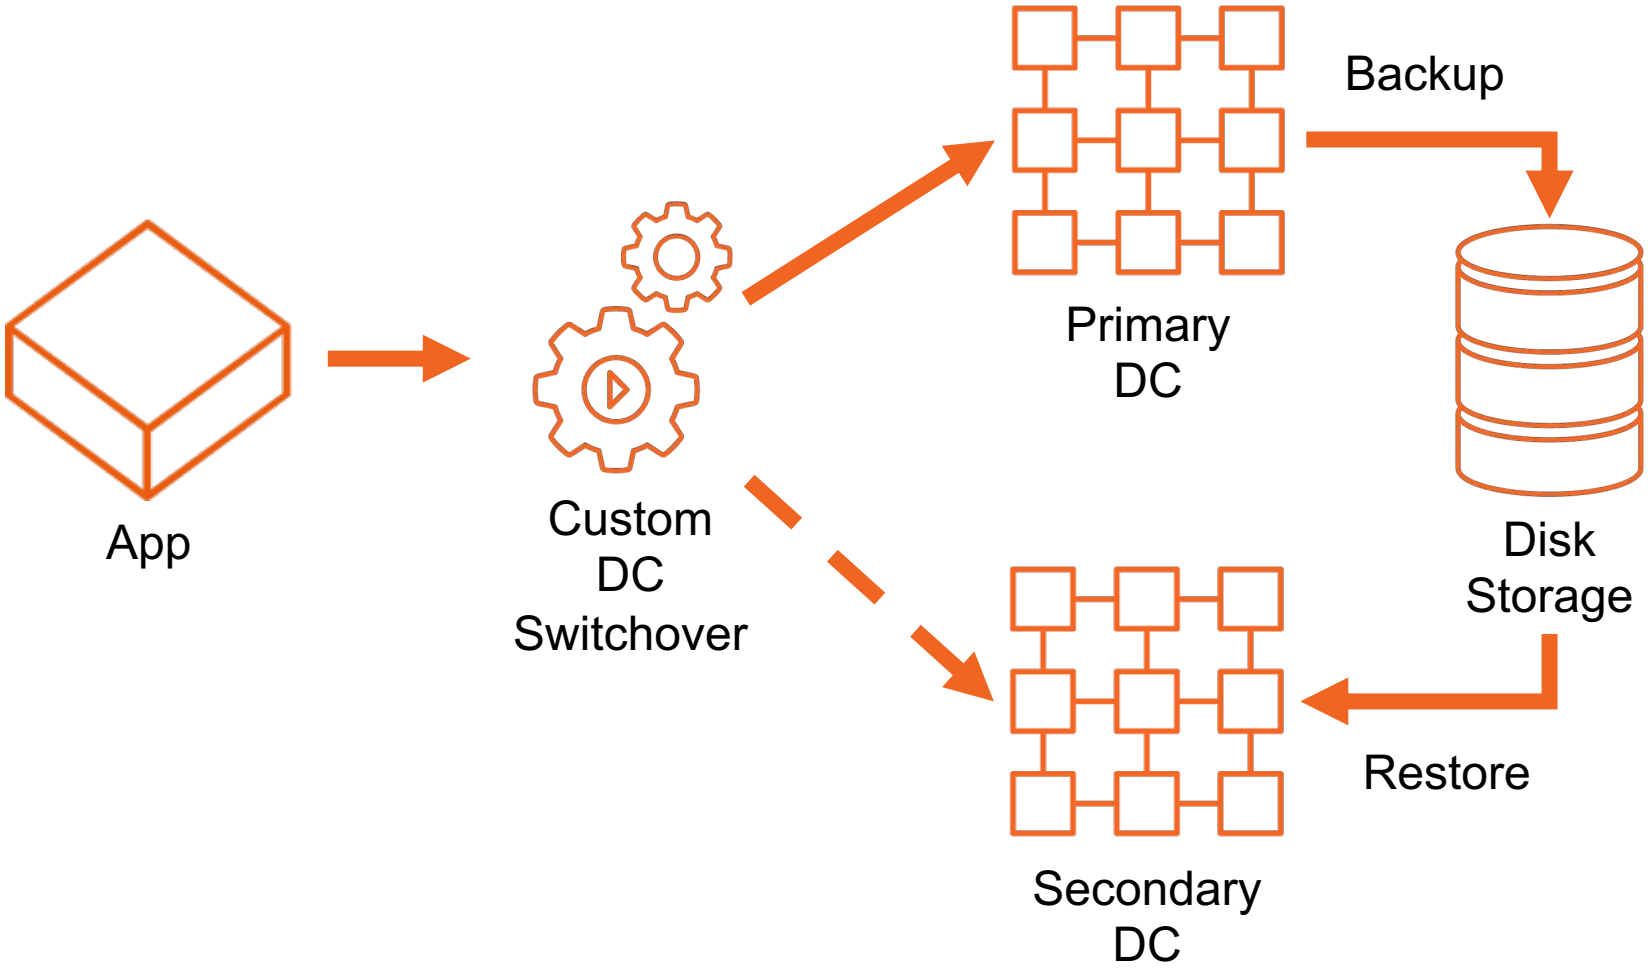

Backup-Based DR

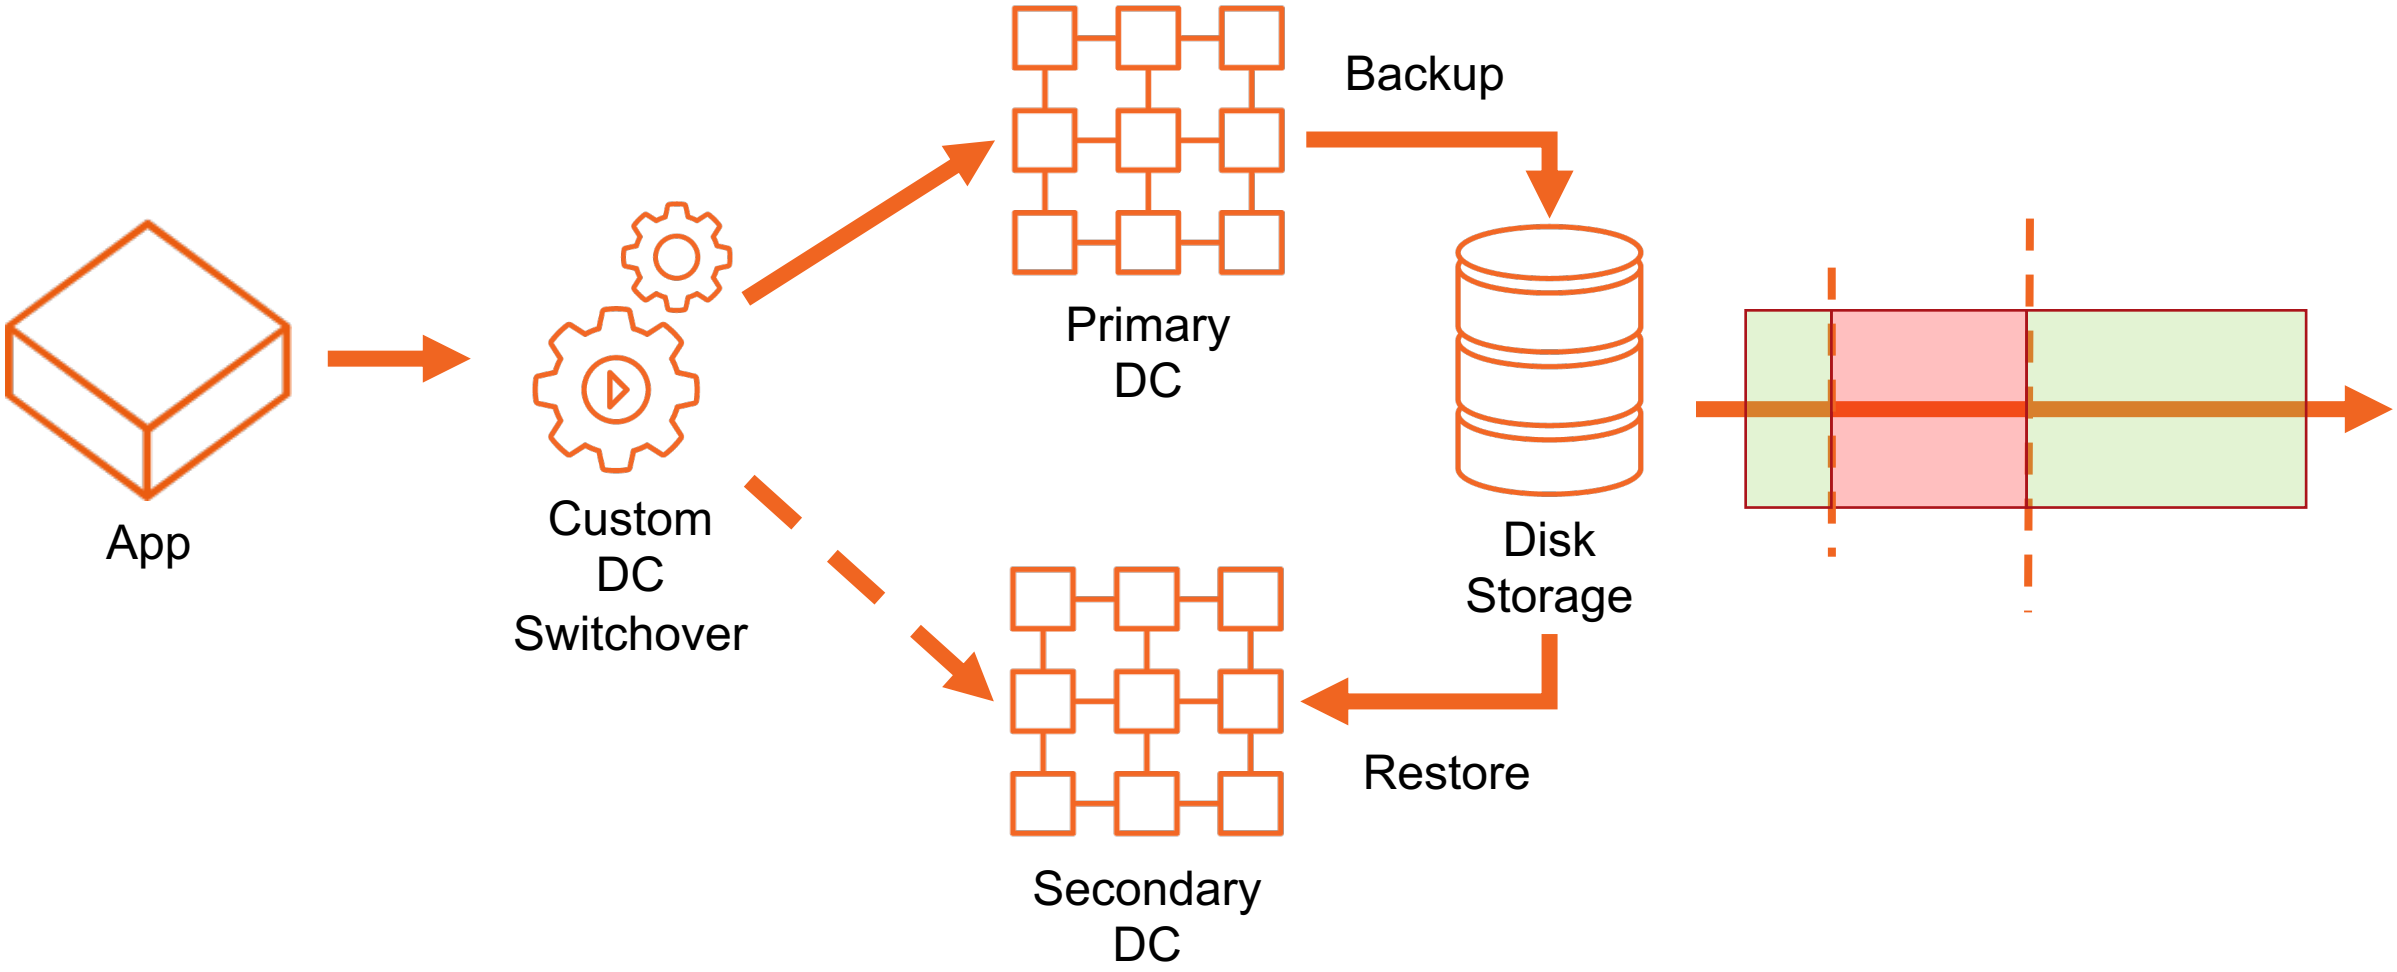

GridGain

Backup-Based DR

GridGain Solution: Snapshots

- Backups of a live cluster – no service disruption
- Incremental backups for more frequent backups – better RPO
- Automatic backup management (scheduling)
- Point-in-Time Recovery
- Only works with Native Persistence
- Available in GridGain Ultimate Edition

Solutions Comparison

Solution	RPO	RTO	Cost
Application-Level	0	0	- Full DIY
System-Level	0 to minutes	Minutes	- VM/Cloud solution - Huge reliance on network for RPO=0
Backup-Based	Hours	0	- Backup solution and disk storage - Custom switchover

Change-Data-Capture

Change-Data-Capture

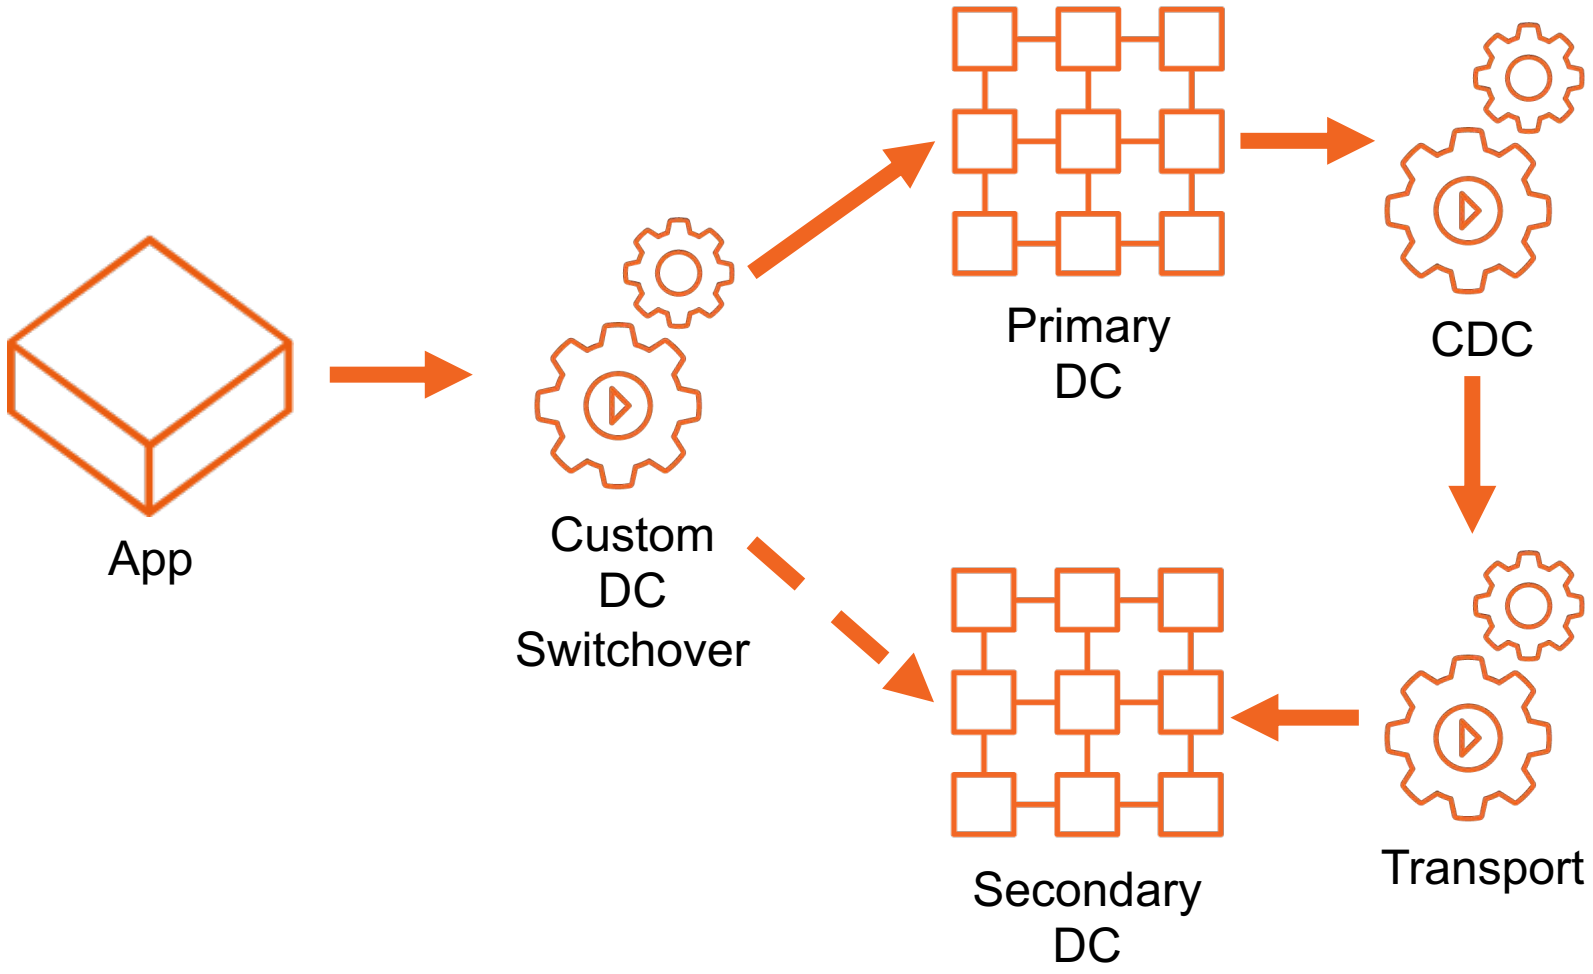

GridGain

Queue-Based CDC

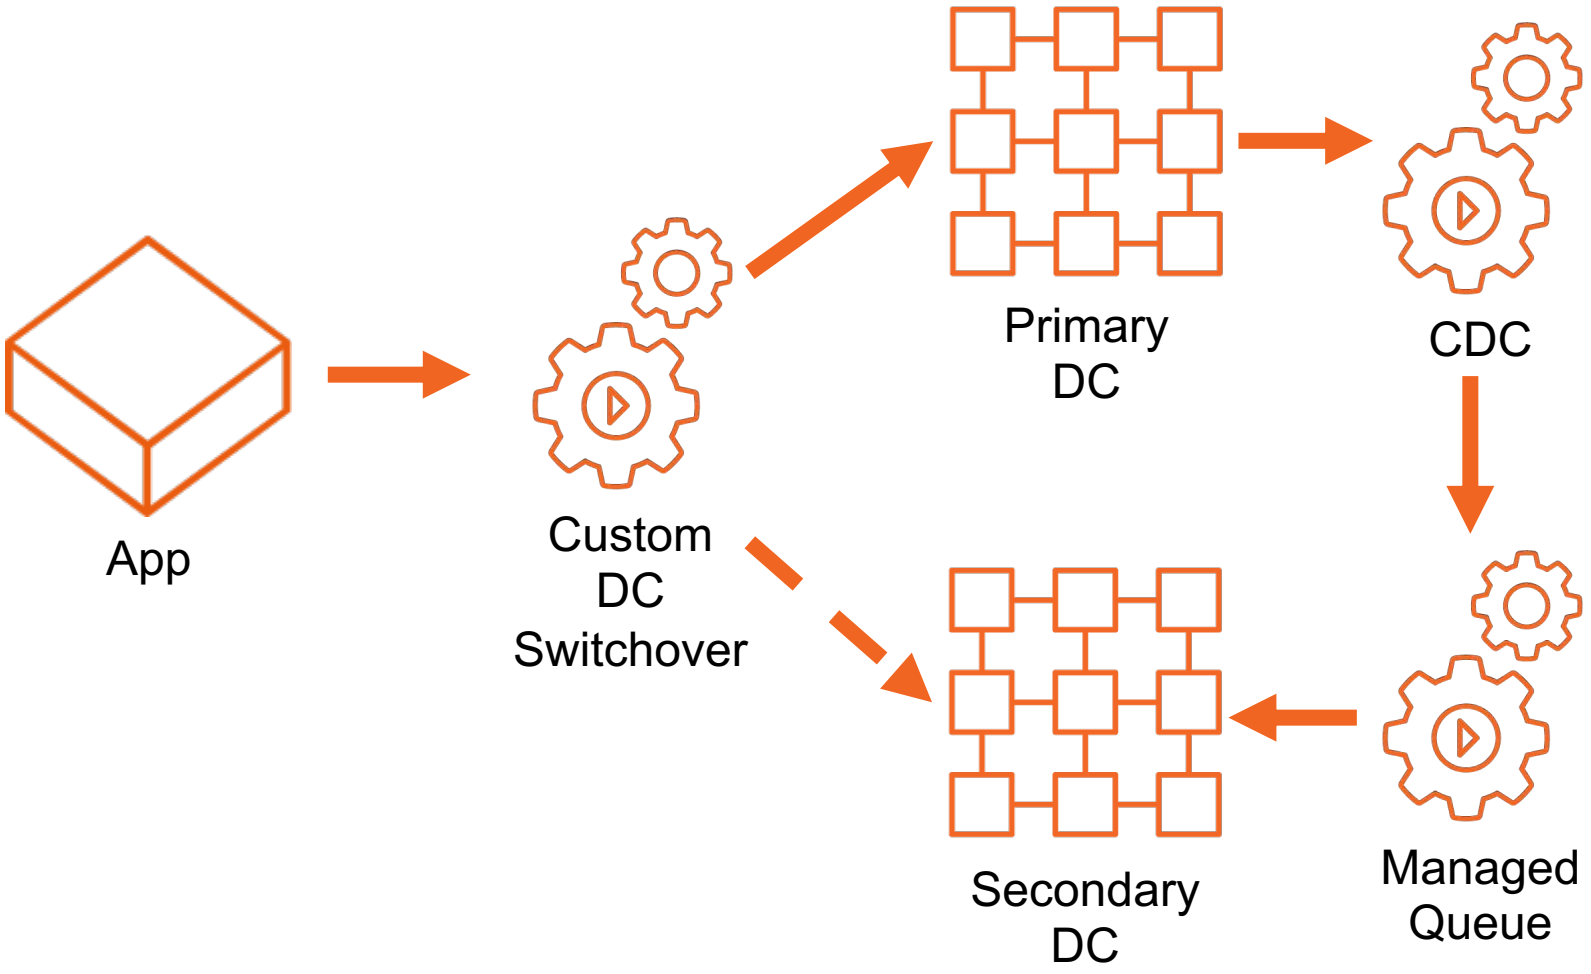

Queue-Based CDC

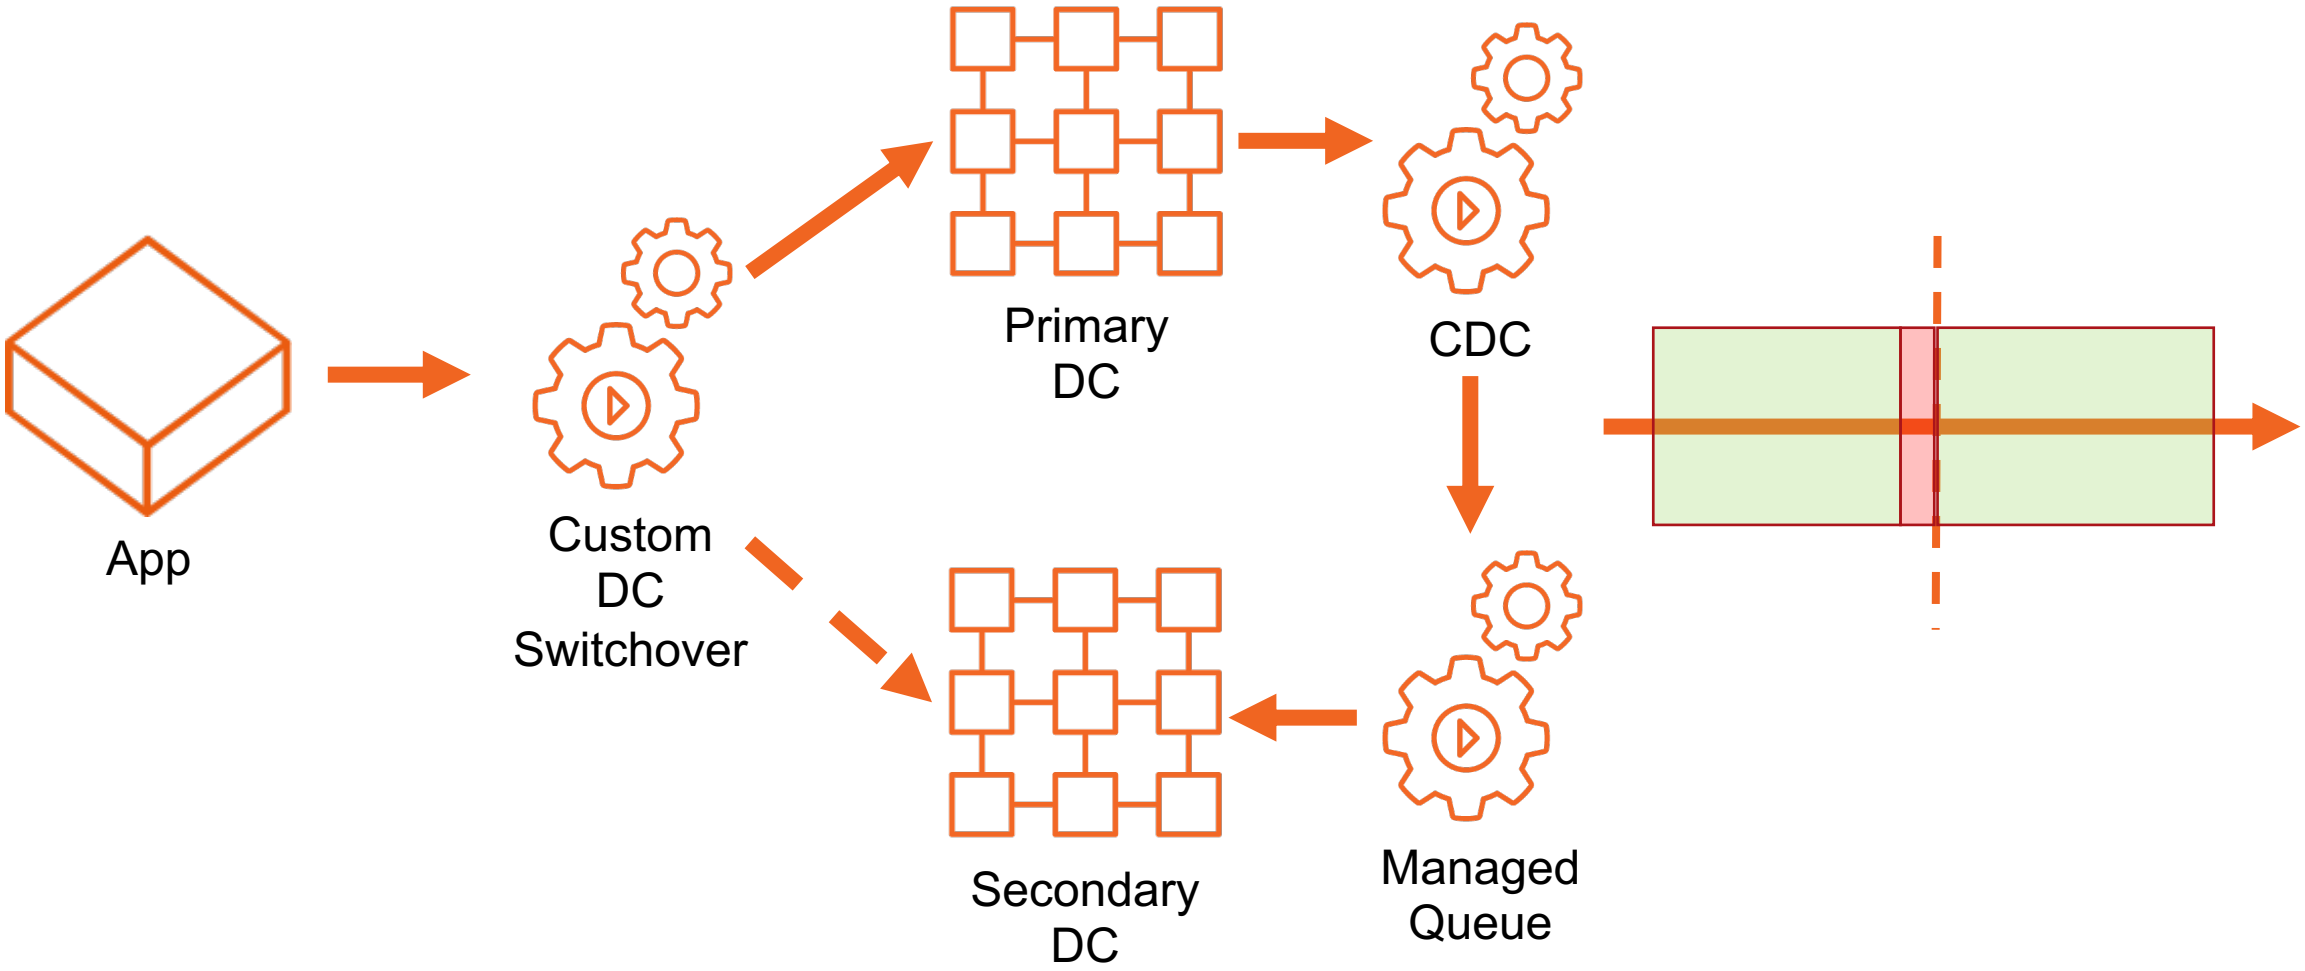

GridGain

Queue-Based CDC

GridGain Solution: Kafka Connector

- Certified by Confluent
- Requires a Kafka instance deployed separately
- Maximum flexibility – allows for heterogenous systems, any topology
- Available in GridGain Enterprise Edition

Solutions Comparison

Solution	RPO	RTO	Cost
Application-Level	0	0	- Full DIY
System-Level	0 to minutes	Minutes	- VM/Cloud solution - Huge reliance on network for RPO=0
Backup-Based	Hours	0	- Backup solution and disk storage - Custom switchover
Queue-based CDC	Seconds	0	- Kafka - Custom switchover

Native CDC

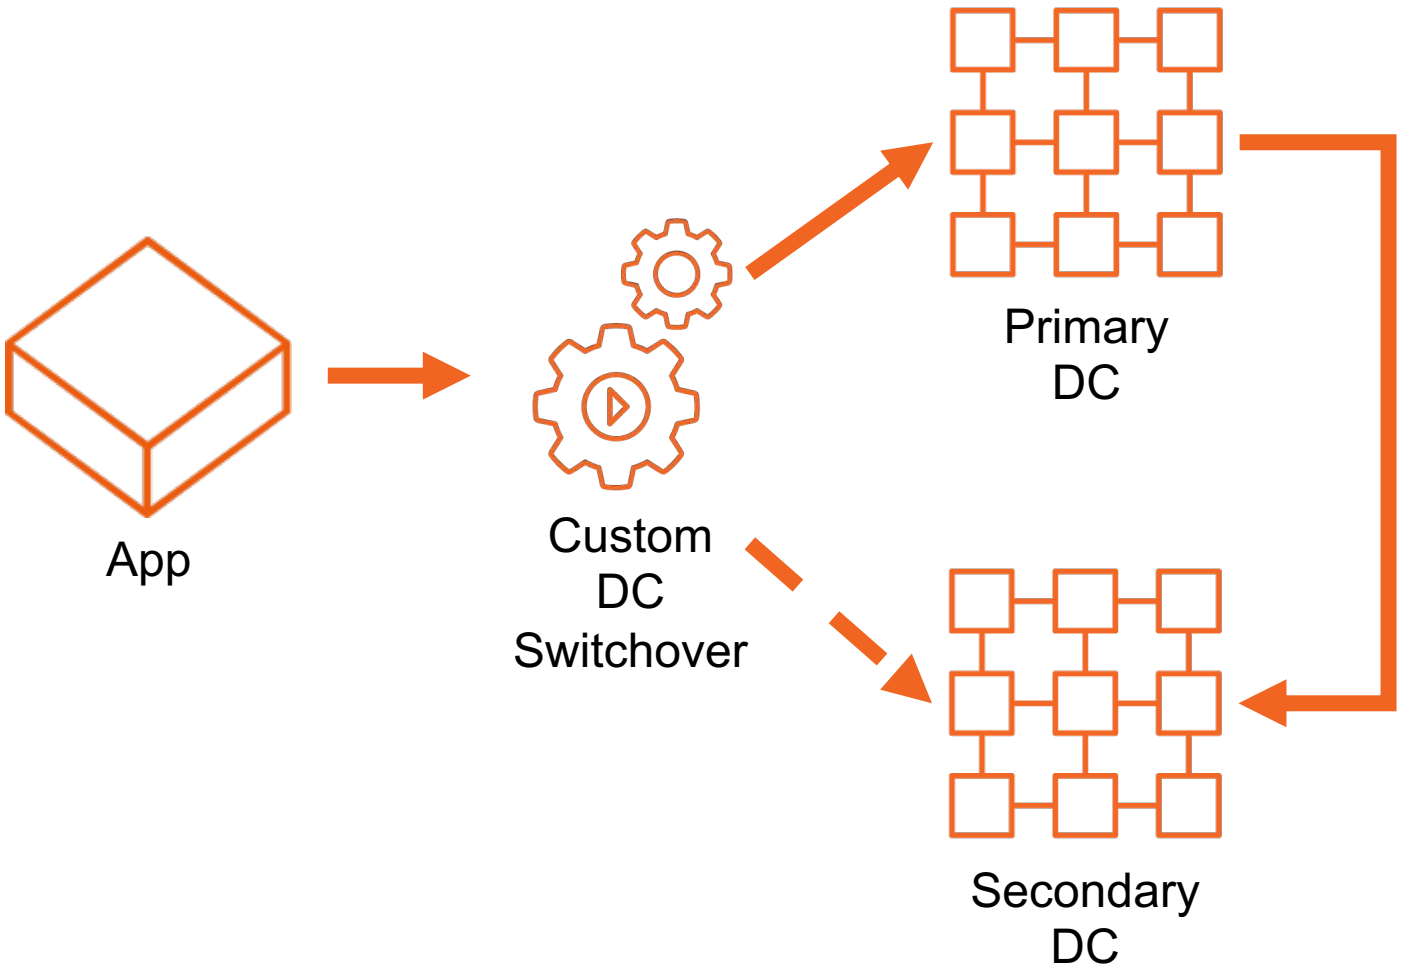

GridGain

Native CDC

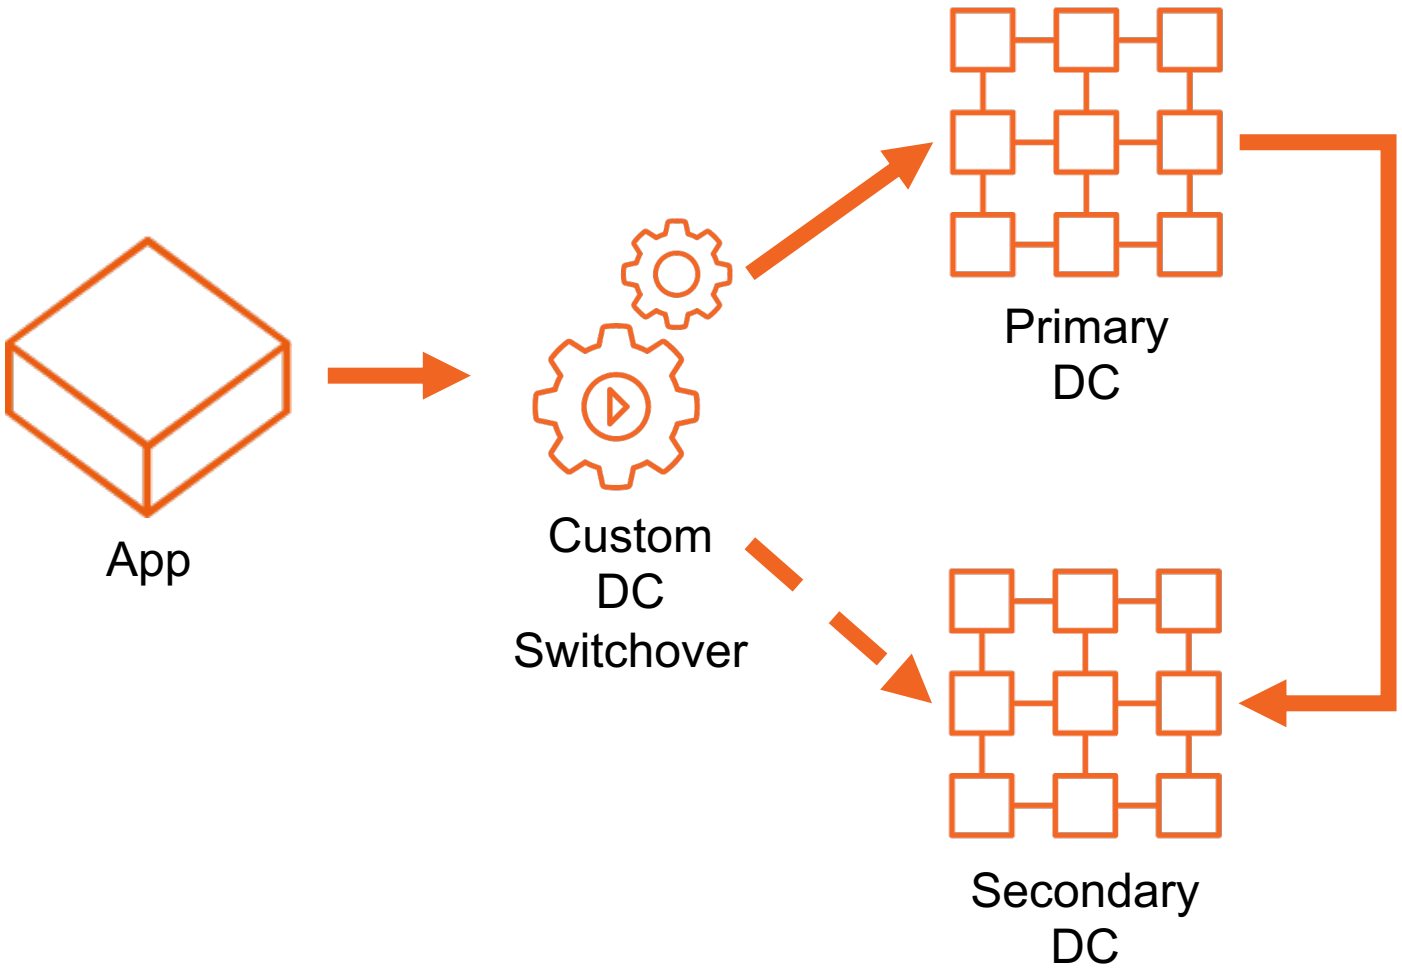

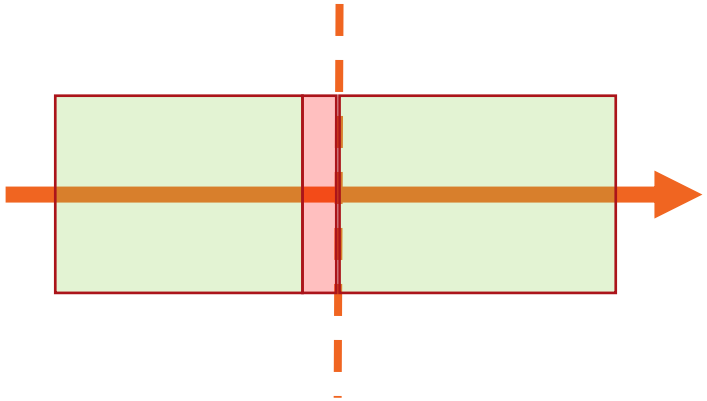

Native CDC

GridGain Solution: Data Center Replication

- No additional software
- Active-Passive or Active-Active
- Allows for complex topologies, up to 32 data centers
- Available in GridGain Enterprise Edition

Solutions Comparison

Solution	RPO	RTO	Cost
Application-Level	0	0	- Full DIY
System-Level	0 to minutes	Minutes	- VM/Cloud solution - Huge reliance on network for RPO=0
Backup-Based	Hours	0	- Backup solution and disk storage - Custom switchover
Queue-based CDC	Seconds	0	- Kafka - Custom switchover
Native CDC	Seconds	0	- Custom switchover

Stretched Cluster

Stretched Cluster

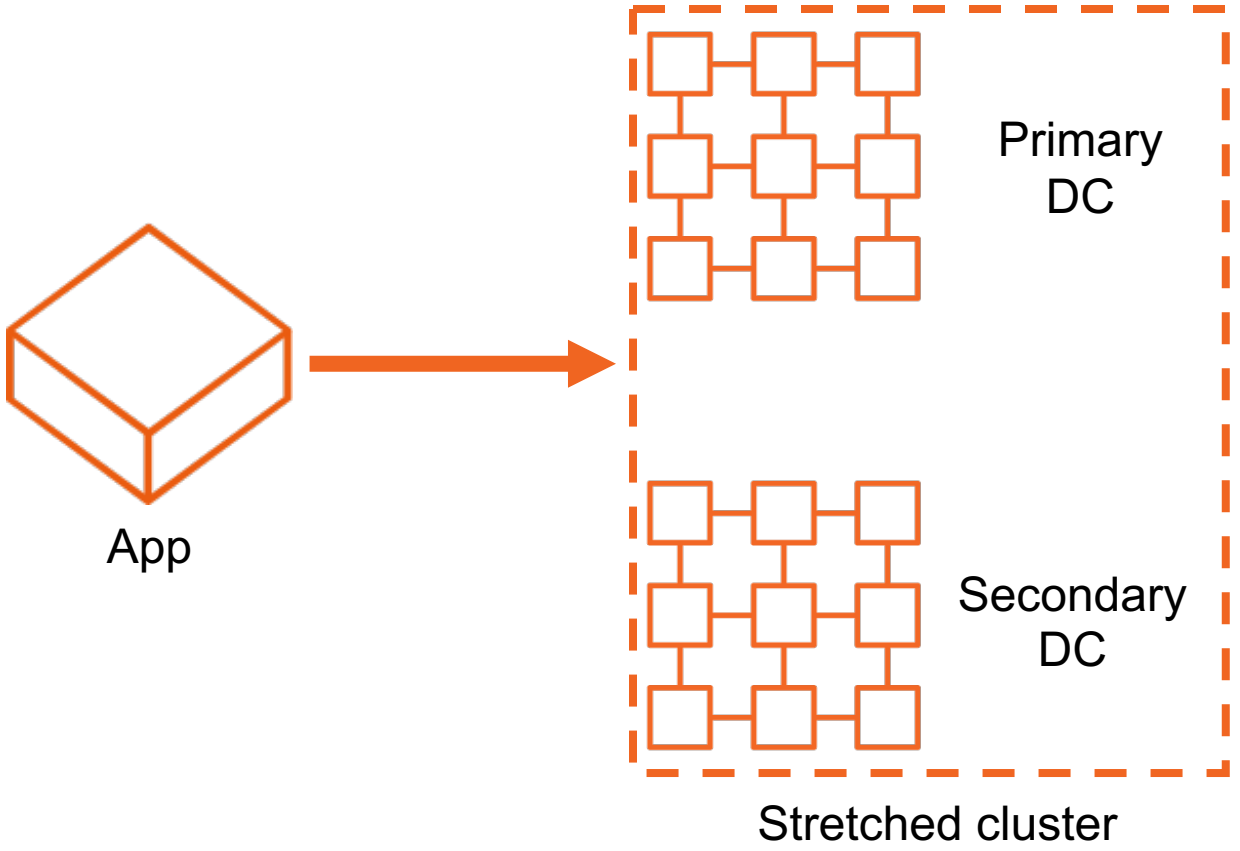

Stretched Cluster

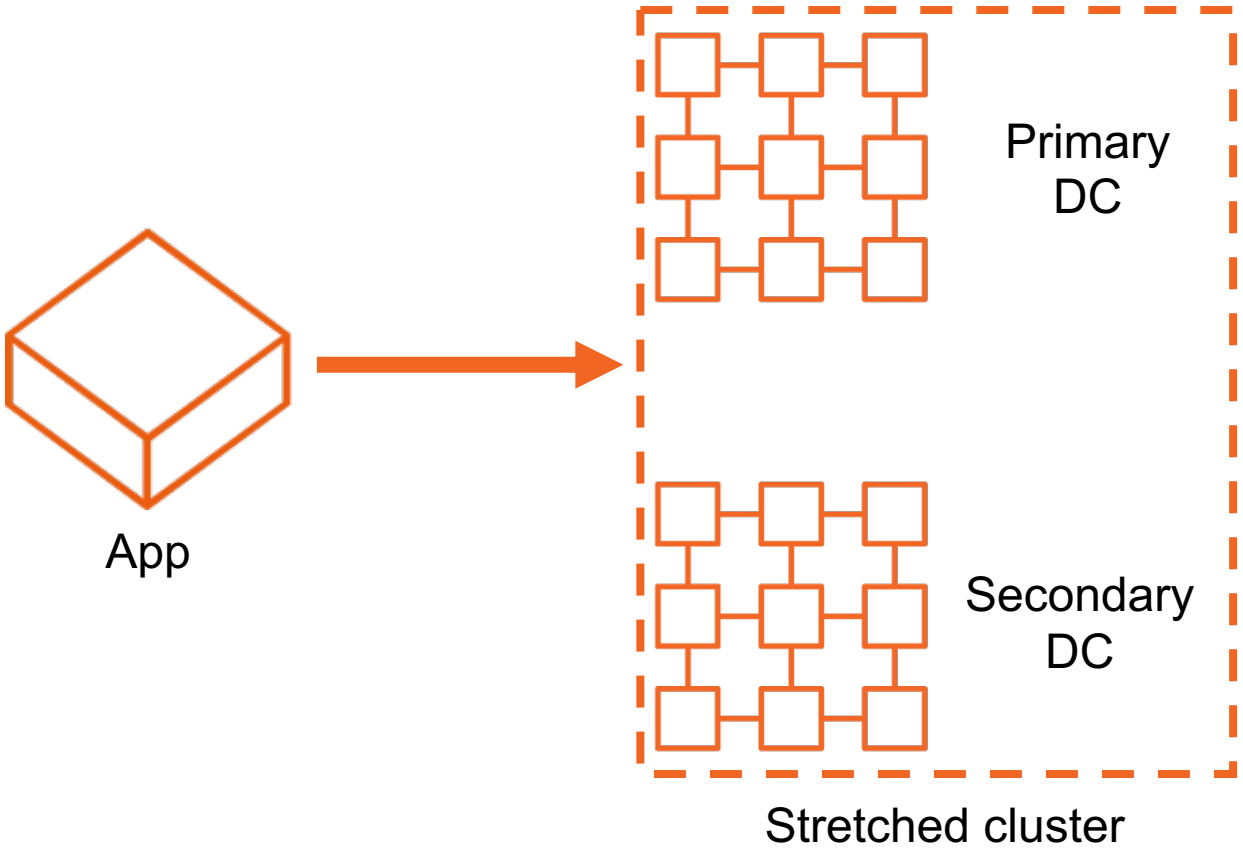

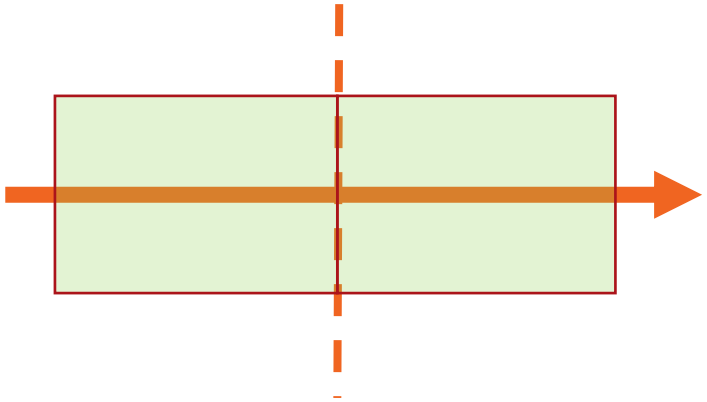

Stretched Cluster

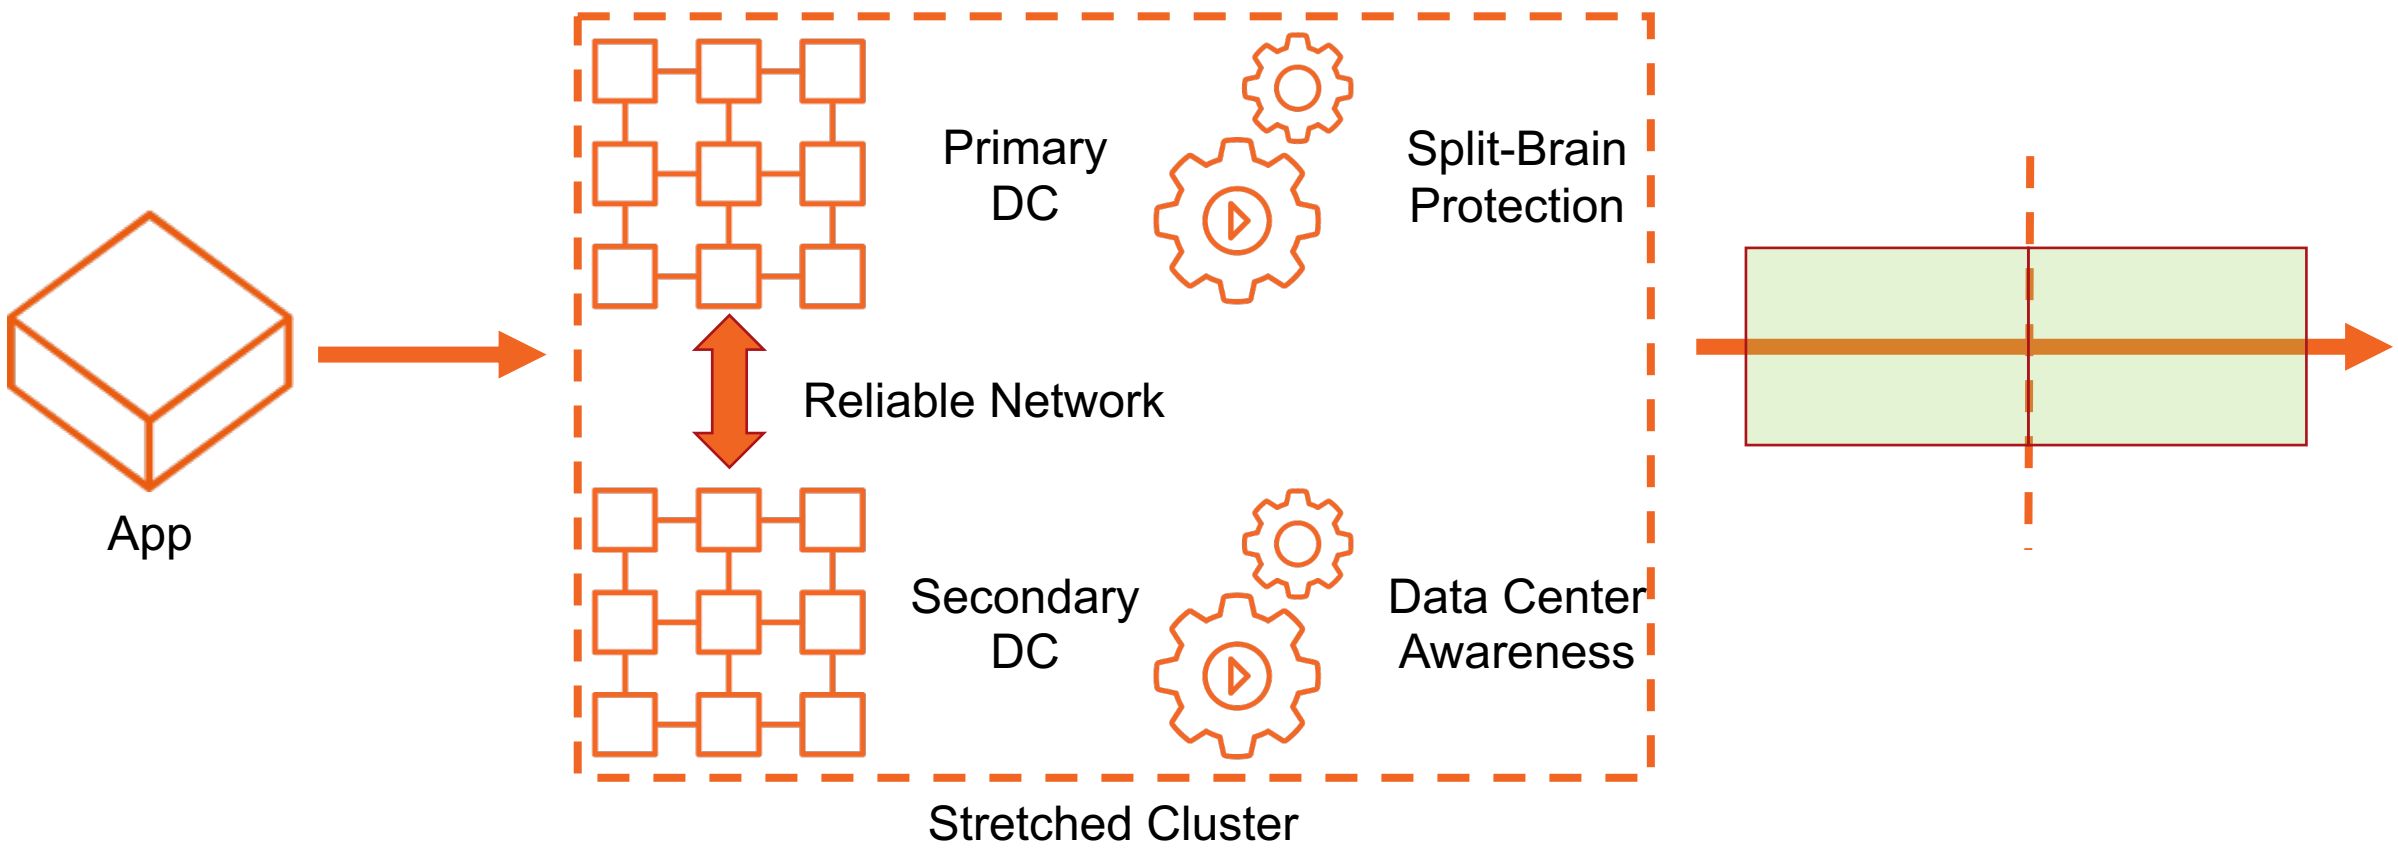

GridGain

Solutions Comparison

Solution	RPO	RTO	Cost
Application-Level	0	0	- Full DIY
System-Level	0 to minutes	Minutes	- VM/Cloud solution - Huge reliance on network for RPO=0
Backup-Based	Hours	0	- Backup solution and disk storage - Custom switchover
Queue-based CDC	Seconds	0	- Kafka - Custom switchover
Native CDC	Seconds	0	- Custom switchover
Stretched Cluster	0	0	- Huge reliance on network - DCs can impact one another - Split-brain protection and data center awareness

Choosing The Solution For GridGain DR

- Default – go with **GridGain Data Center Replication**
  - No additional components needed
- Go for a **stretched cluster** if
  - You need RPO=0
  - AND you have a good network
- Add **GridGain Snapshots** if data is mission-critical – safety first!

Agenda

- What is Disaster Recovery?
- Disaster Recovery Options
- **Advanced Topics and FAQ**

Advanced Topics and FAQ

Split-Brain Explained

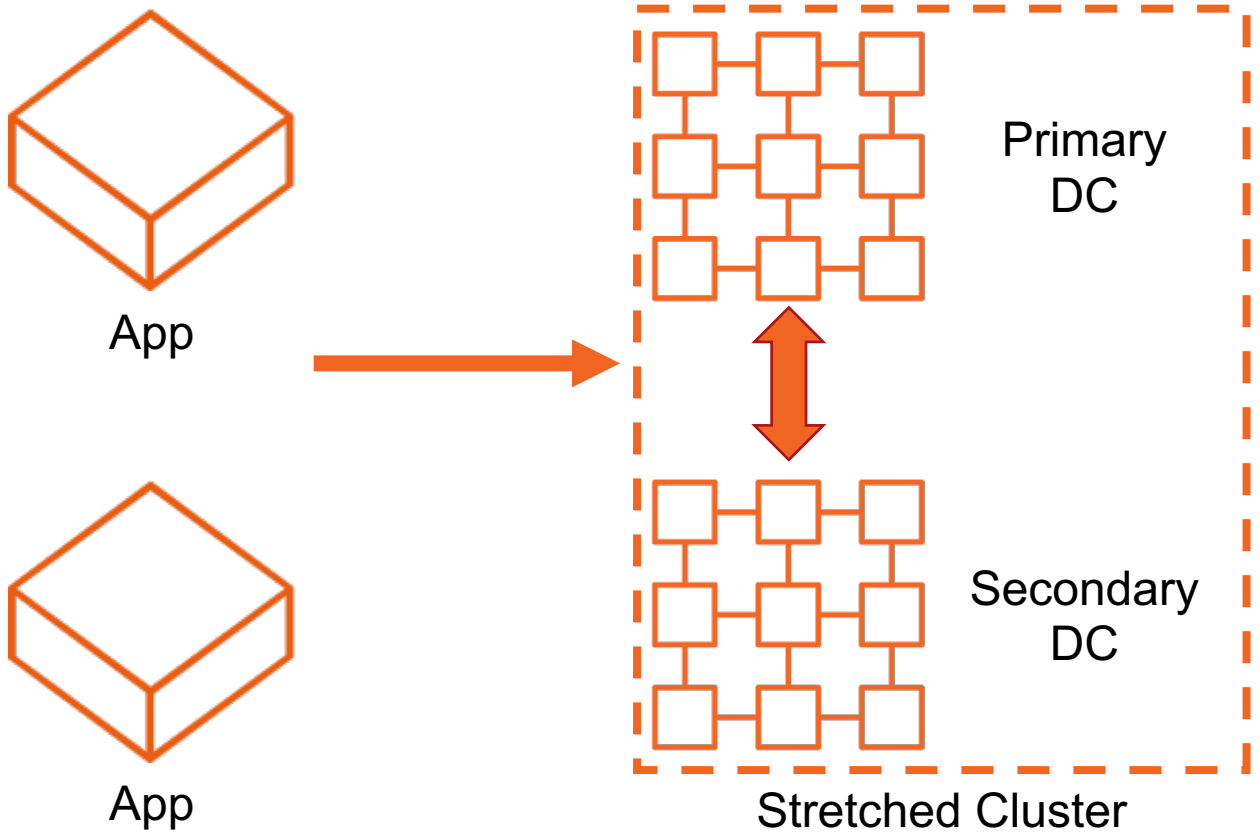

Split-Brain Explained

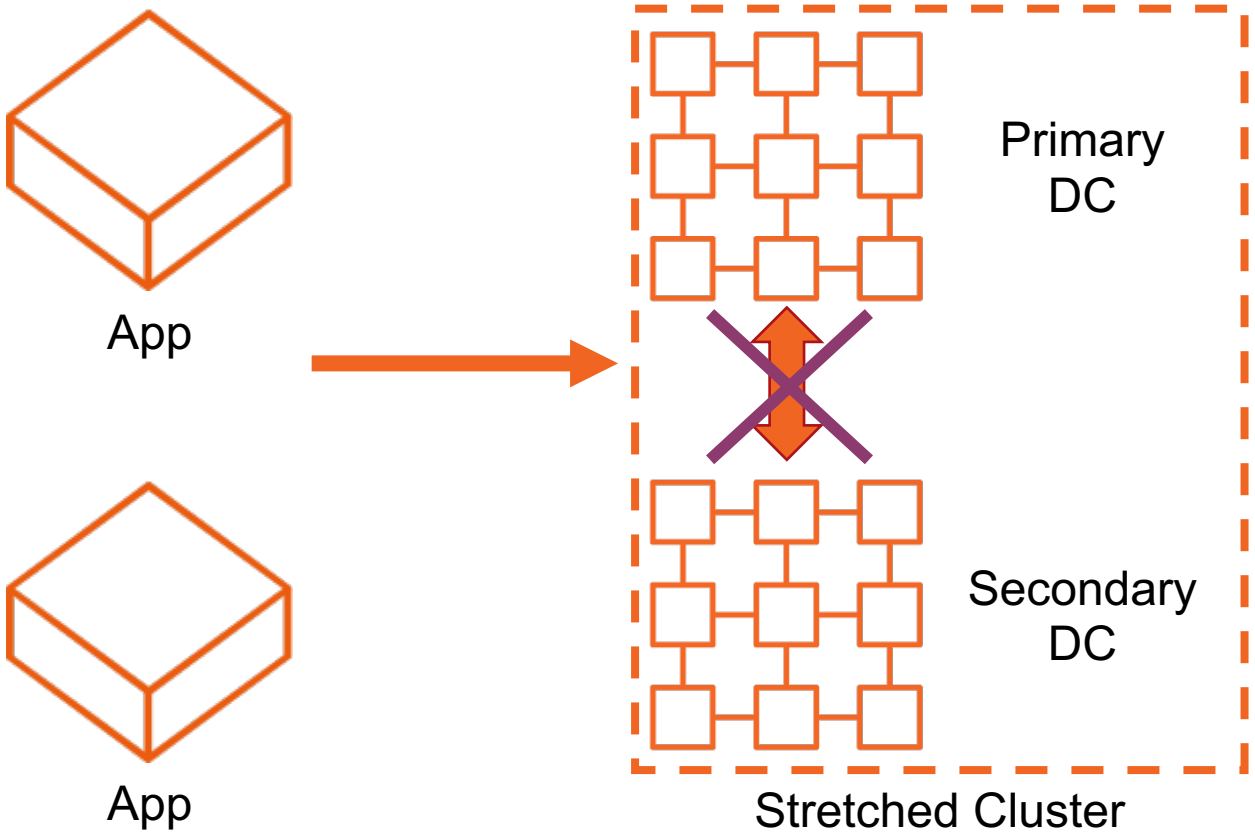

Split-Brain Explained

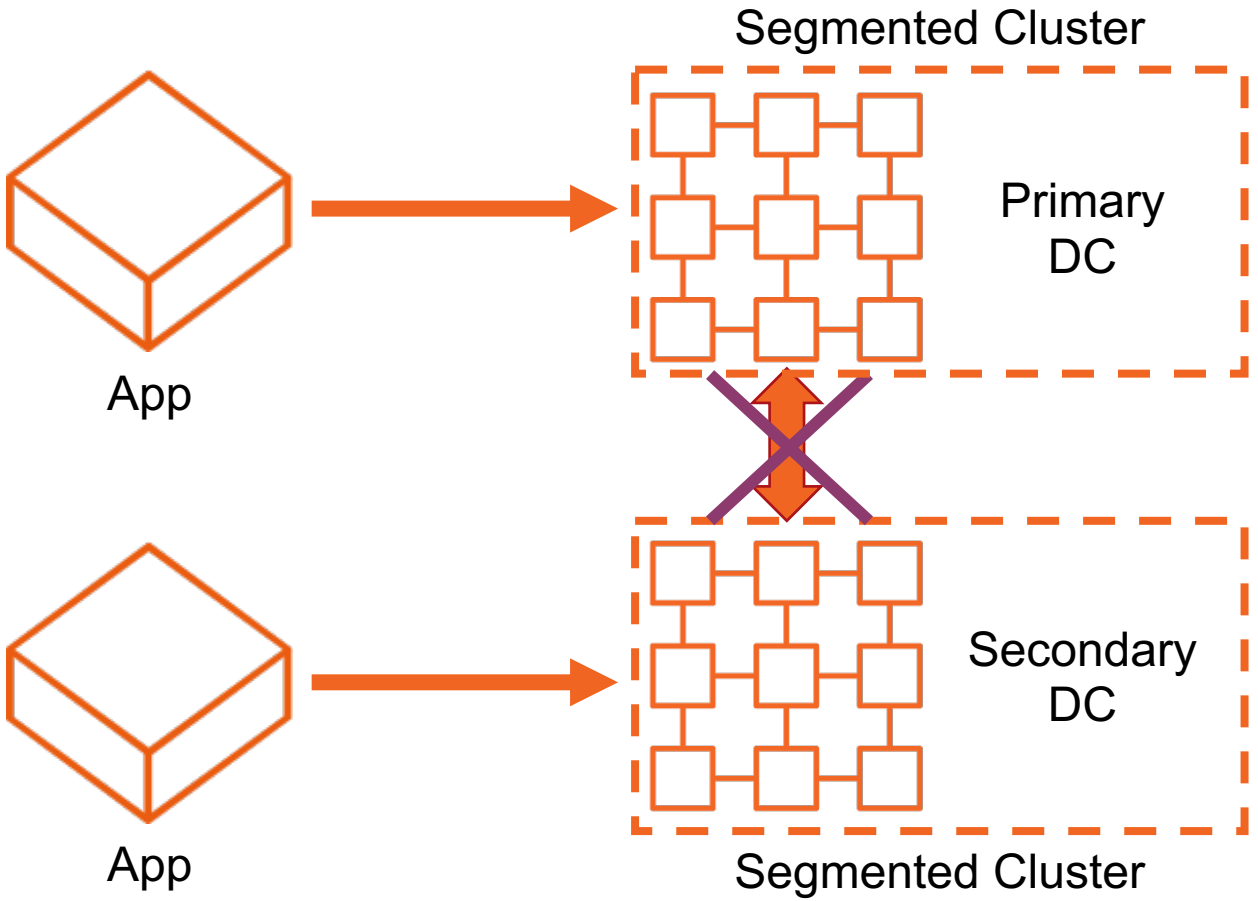

Split-Brain Explained

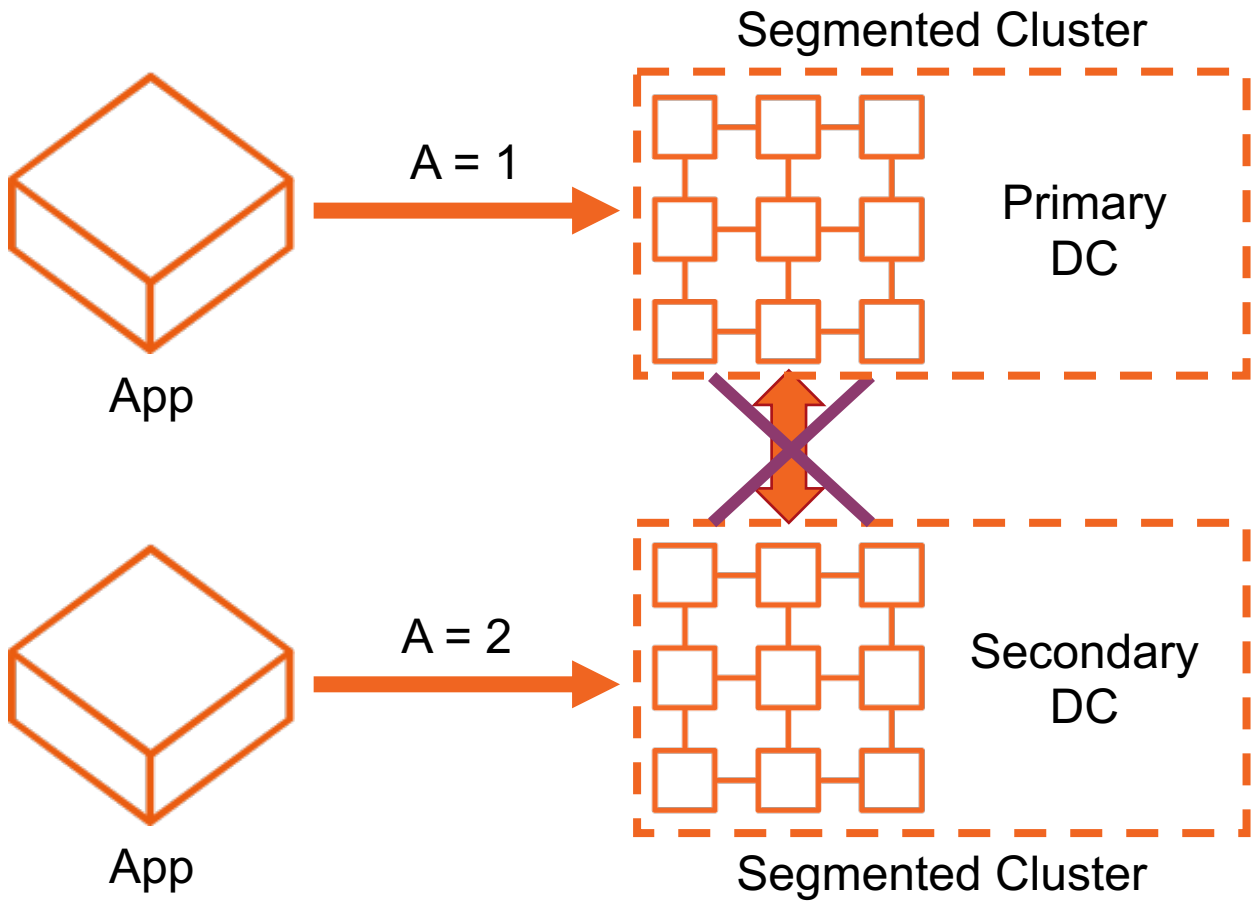

Split-Brain Explained

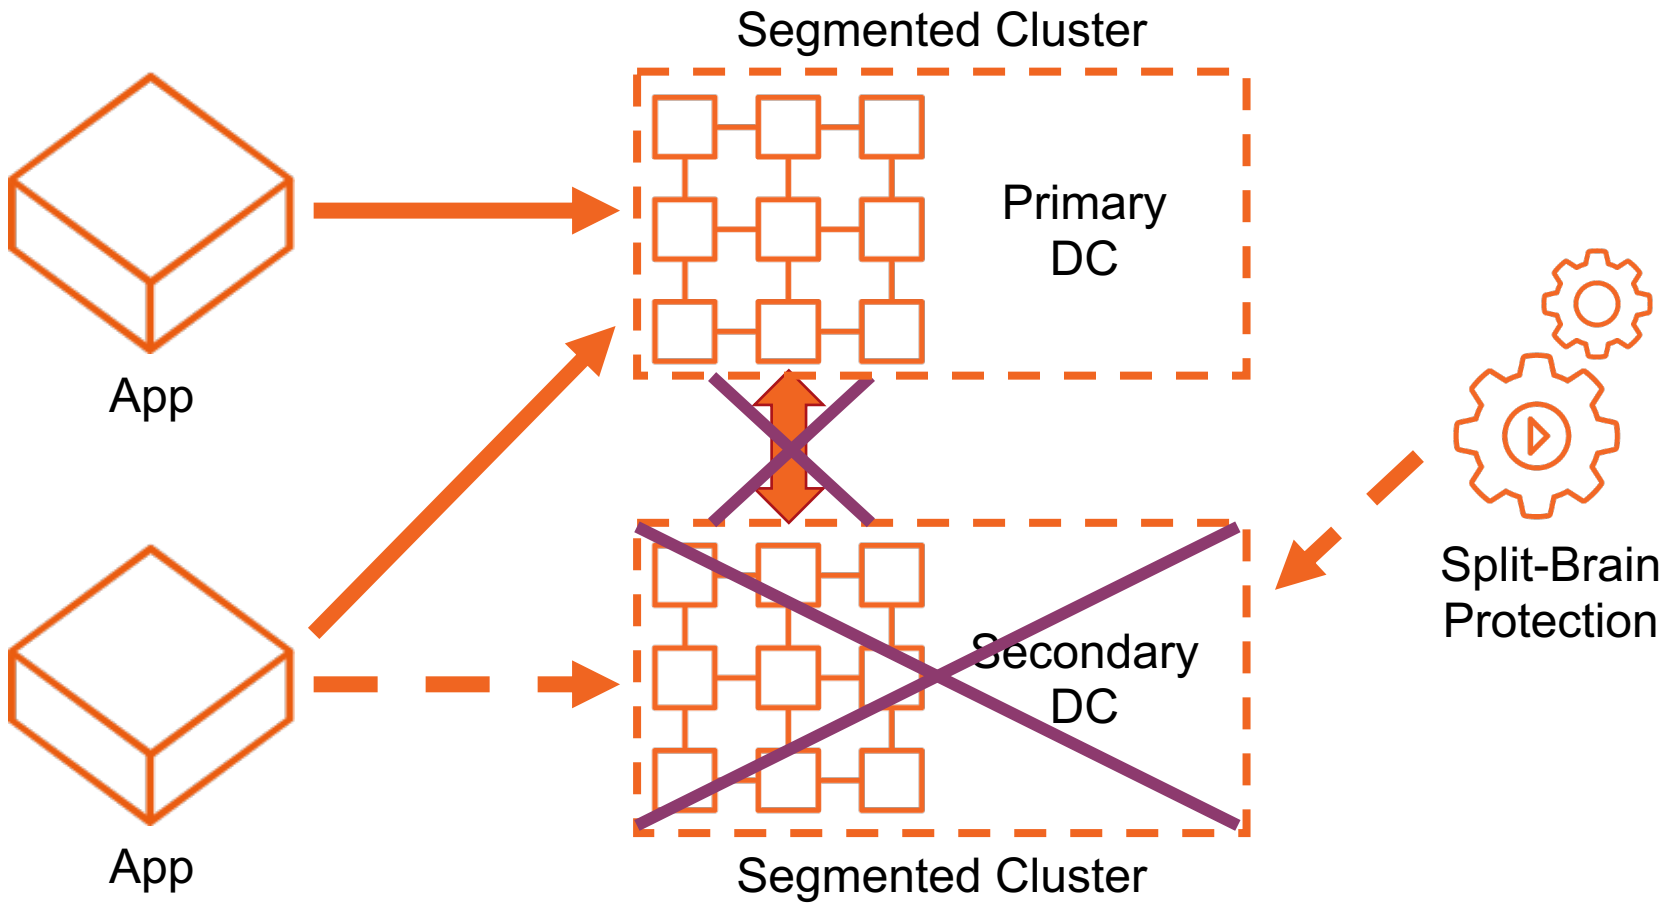

Split-Brain Explained

GridGain Solution For Split-Brain Protection

Solution 1: `TopologyValidator` + `SegmentationResolver`

- `TopologyValidator` prevents updates in the segmented part
- `SegmentationResolver` stops the segmented part
- Available in GridGain Enterprise Edition

Solution 2: Zookeeper Discovery

- Zookeeper is responsible for keeping the cluster together
- Available in Apache Ignite and GridGain Community Edition

Data Center Awareness Explained

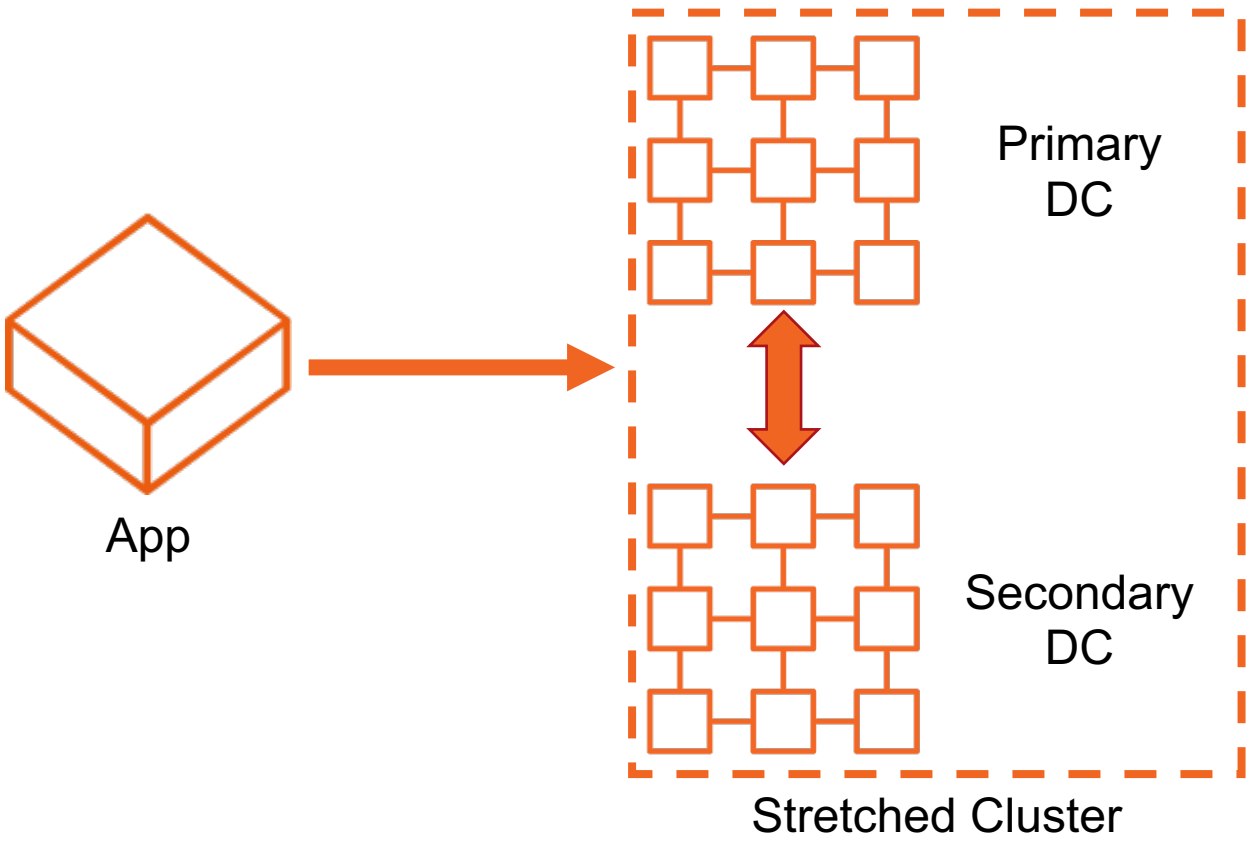

Data Center Awareness Explained

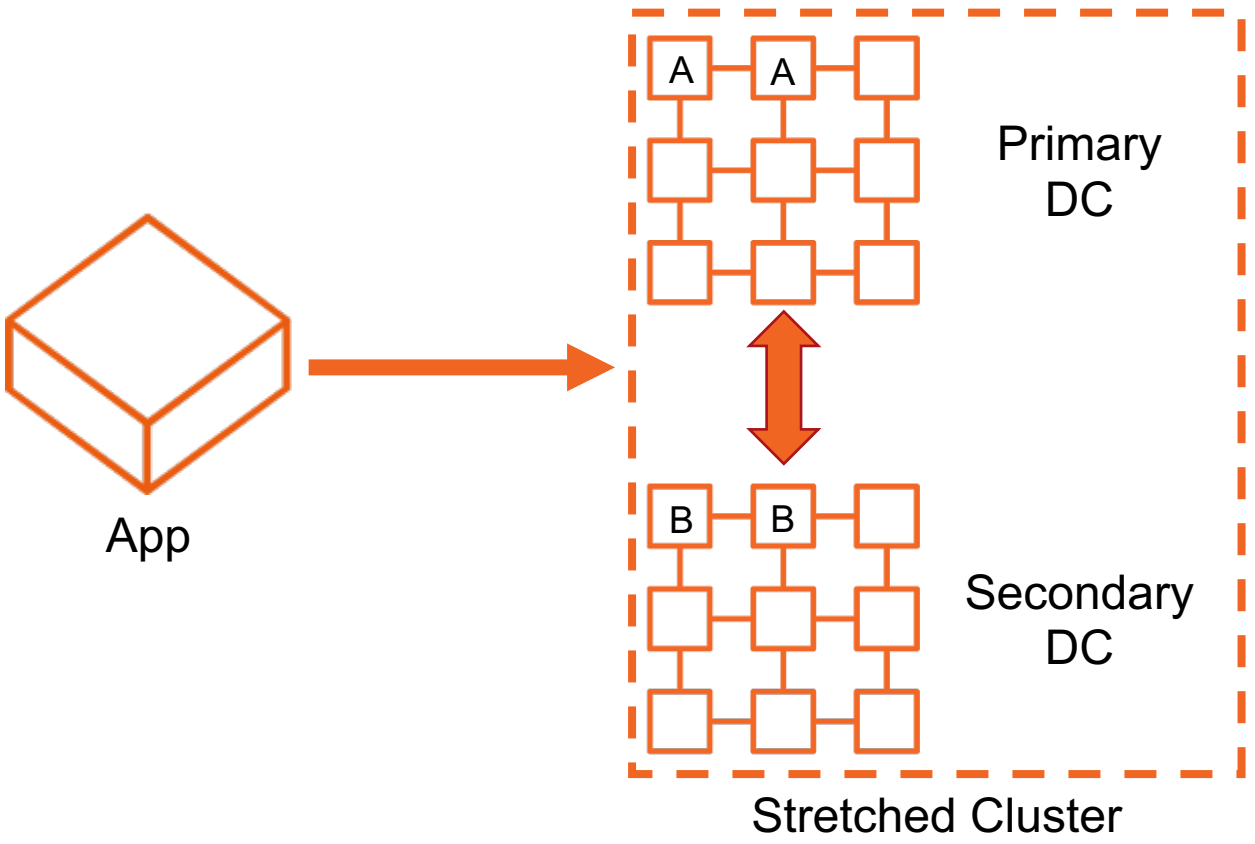

Data Center Awareness Explained

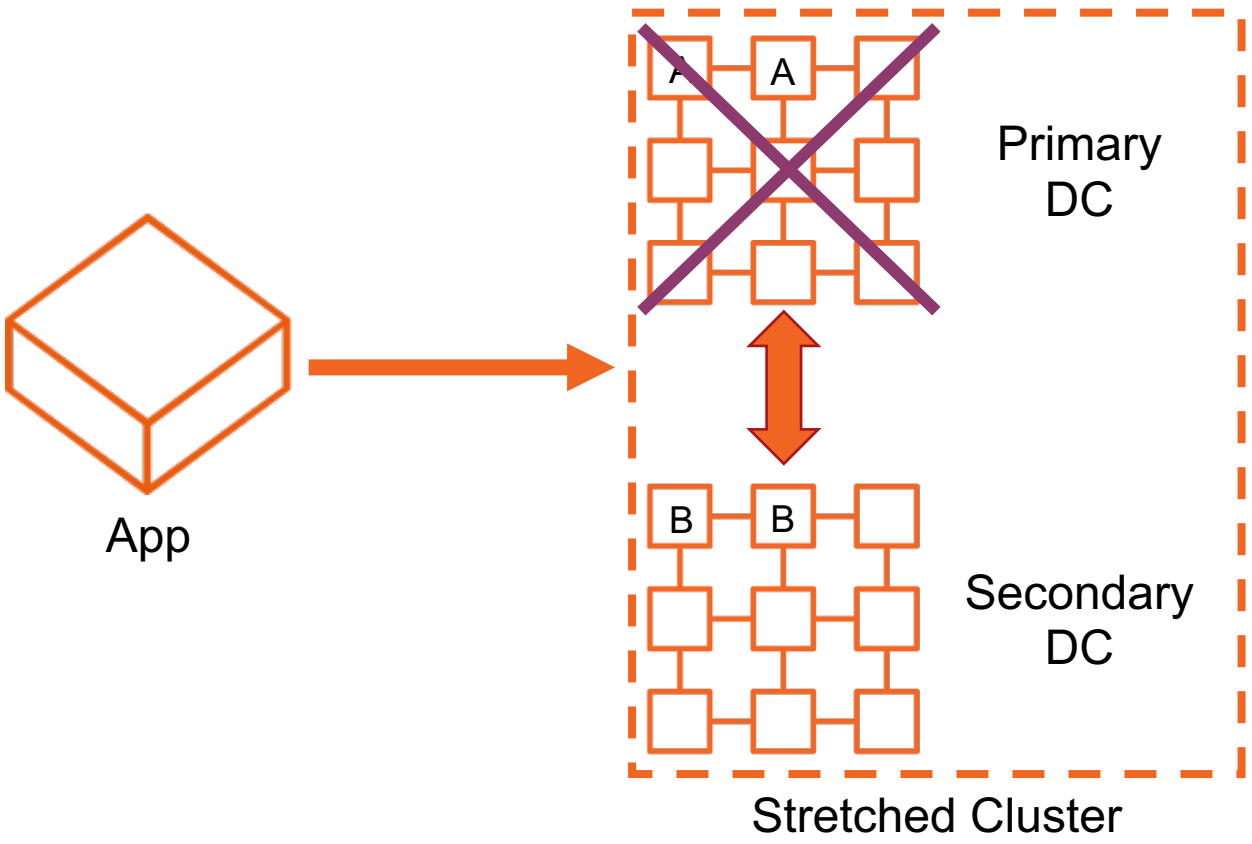

Data Center Awareness Explained

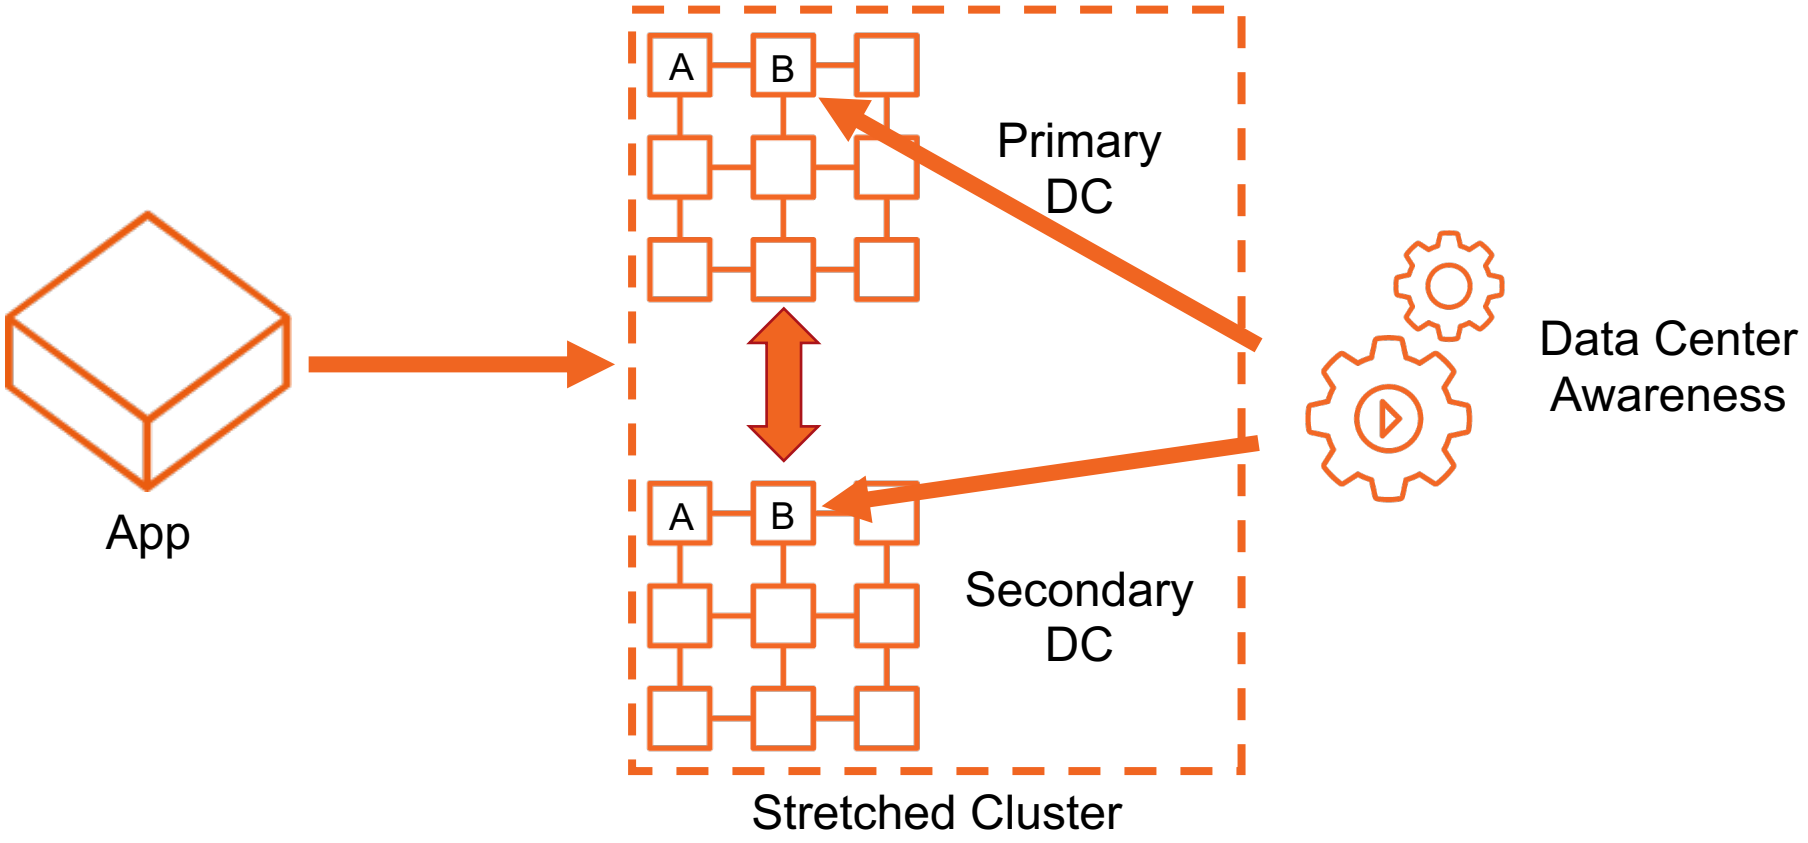

Data Center Awareness Explained

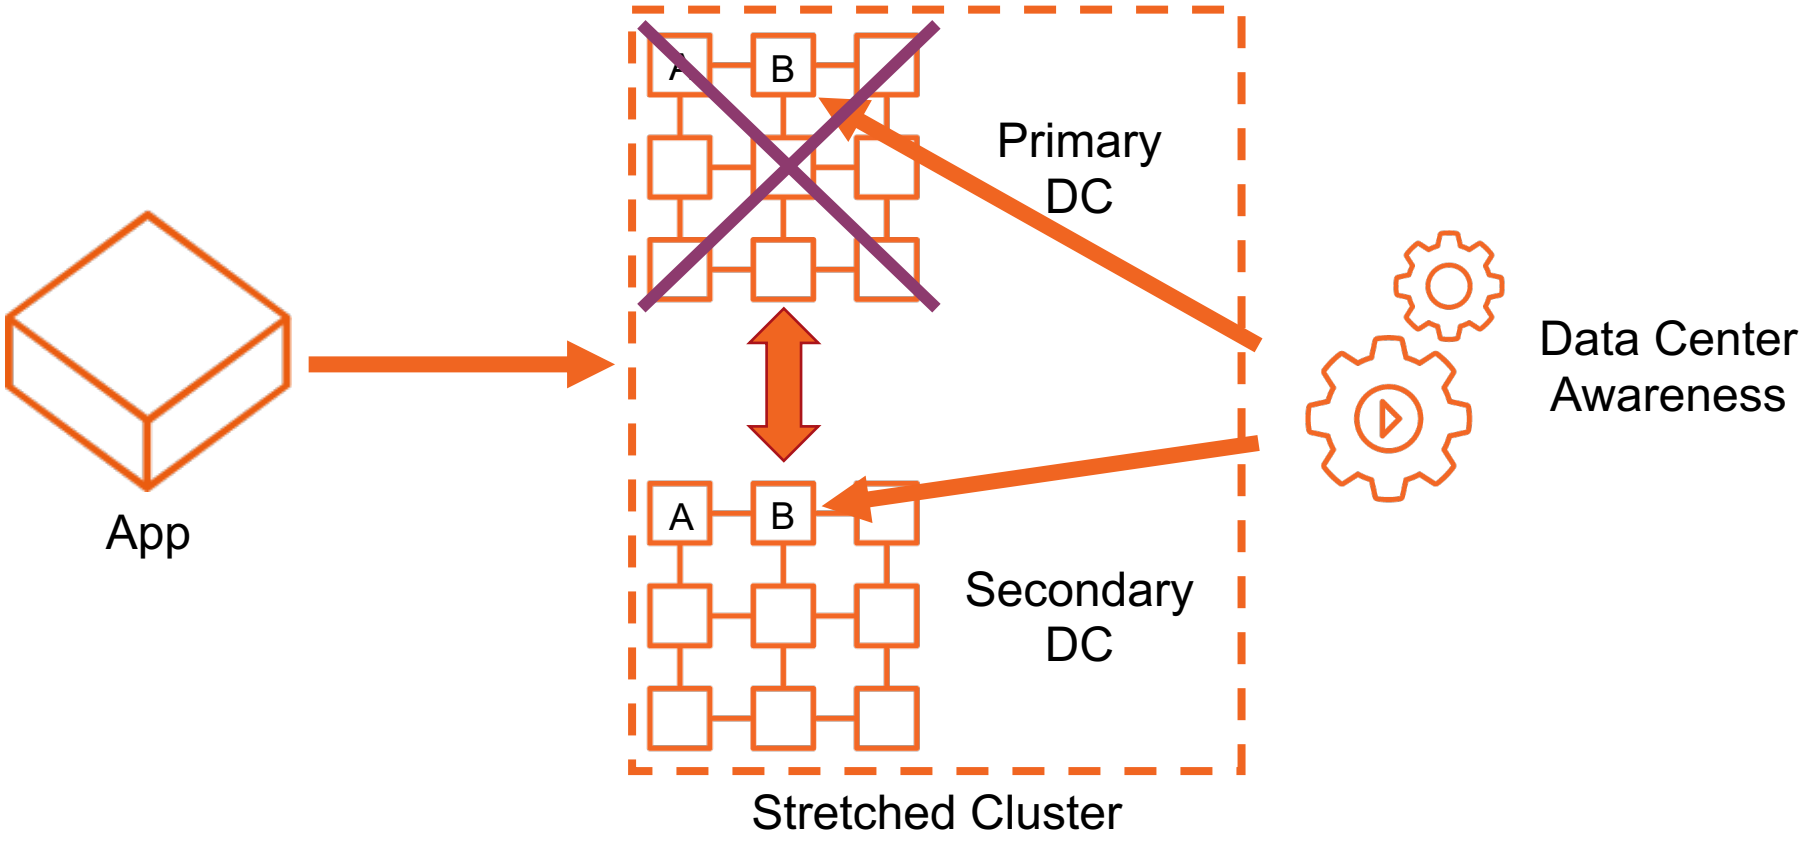

GridGain

Data Center Awareness Explained

GridGain Solution For Data Center Awareness

- Done via cluster settings: `RendezvousAffinityFunction.affinityBackupFilter` controls distribution of backups
- Available in Apache Ignite and GridGain Community Edition

How Many Data Copies Do I Need?

- Depends on
  - Required availability guarantees
  - Estimated hardware failure probability
  - Performance targets
  - etc
- Rule of thumb – at least **two copies per DC**, at lest **two DCs**

Is My Network Good Enough?

- Right answer – it depends
- Rule of thumb for Data Center Replication
  - Throughput matches write activity
  - Latency is mostly irrelevant
- Rule of thumb for stretched cluster
  - No connection interruptions
  - Throughput matches write activity
  - **Latency in millisecond range** (ideally – under 1ms)
  - **Full reachability between all servers and clients**

Active-Active Or Active-Passive?

- Depends on
  - Failover scenarios
  - Write performance requirements
- Rule of thumb
  - Active-passive by default
  - Active-active is for apps sensitive to write latency
  - Active-active requires custom conflict resolution

Summary

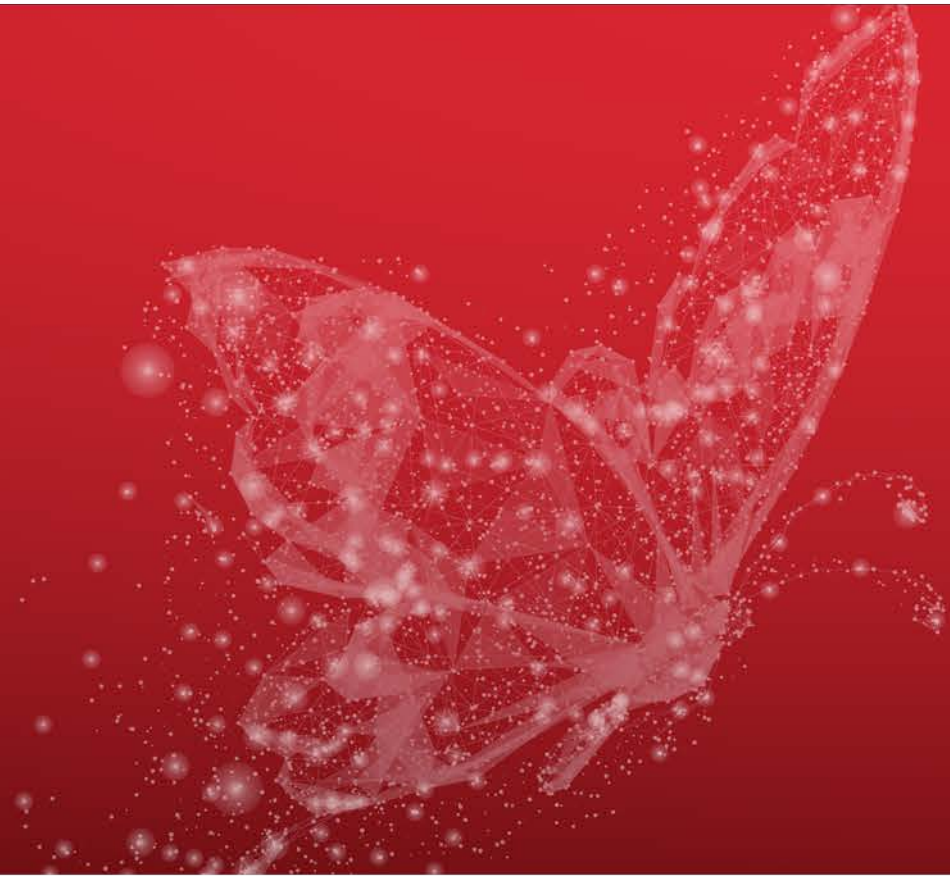

Summary

- Disaster Recovery requires two DCs, a switchover mechanism and a replication mechanism
- Key parameters are RPO and RTO
- Best done on middleware level using middleware tools
- Best DR options for Apache Ignite and GridGain
  - Data Center Replication for RPO > 0
  - Stretched cluster for RPO = 0 – if network allows
  - Add Snapshots if data is critical

Q&A

- Check out Moving Apache® Ignite™ into Production webinars
  - Initial Checklist
  - Best Practices for Native Persistence and Data Recovery
  - Best Practices for Monitoring Distributed In-Memory Computing
  - Best Practices for Deploying Apache Ignite in the Cloud
- GridGain Documentation
  - https://gridgain.com/docs
- Contact me
  - stan@gridgain.com
  - stanlukyanov@gmail.com