THE GLUE

EMPOWERING BUSINESS INNOVATION

Dynamically Re-Configurable Platform for µServices

Sven Beauprez

Visualise. Create. Innovate. Repeat.

? Who?

- The Glue
  - FinTech startup founded in 2015
  - Platform to speed up innovation in the financial sector
  - Open up backend systems to 3rd parties

- Sven Beauprez
  - Lead Architect at The Glue with a focus on innovation and long-term product strategy
  - Co-Founder of IoTBE & Long-time steering member of Devoxx in the early days
  - https://be.linkedin.com/in/svenbeauprez


EVENT DRIVEN JOURNEYS
CONTAINER BASED
THE GLUE
THE GLUE
THE GLUE
IN MEMORY DATA GRID

Traditional Monolith Application

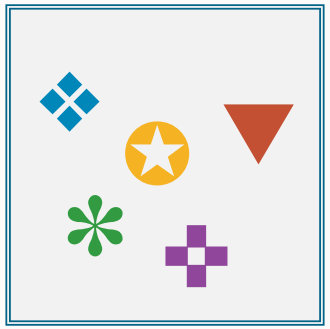

Monolith vs µServices

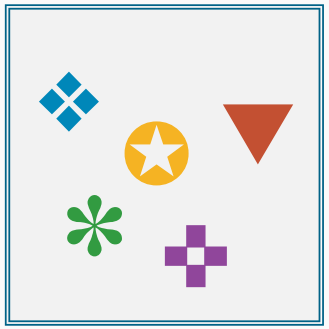

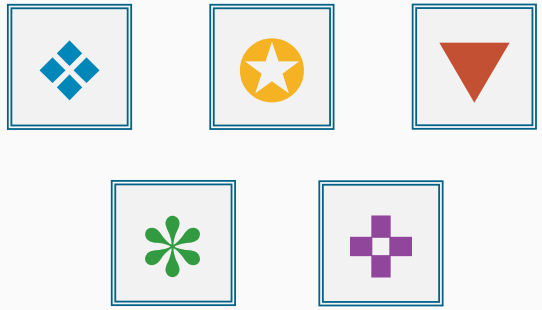

Monolith vs µServices: Deployment

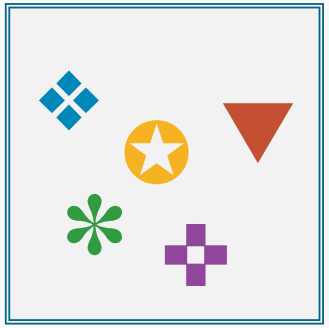

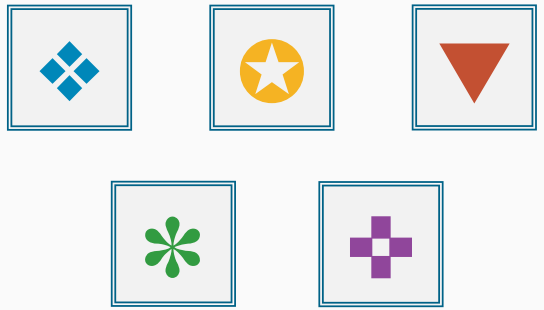

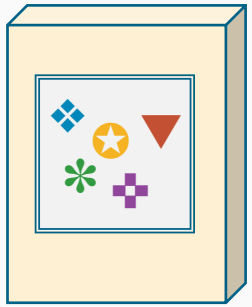

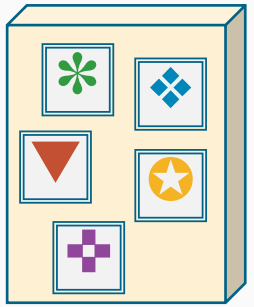

Monolith vs µServices: Scaling

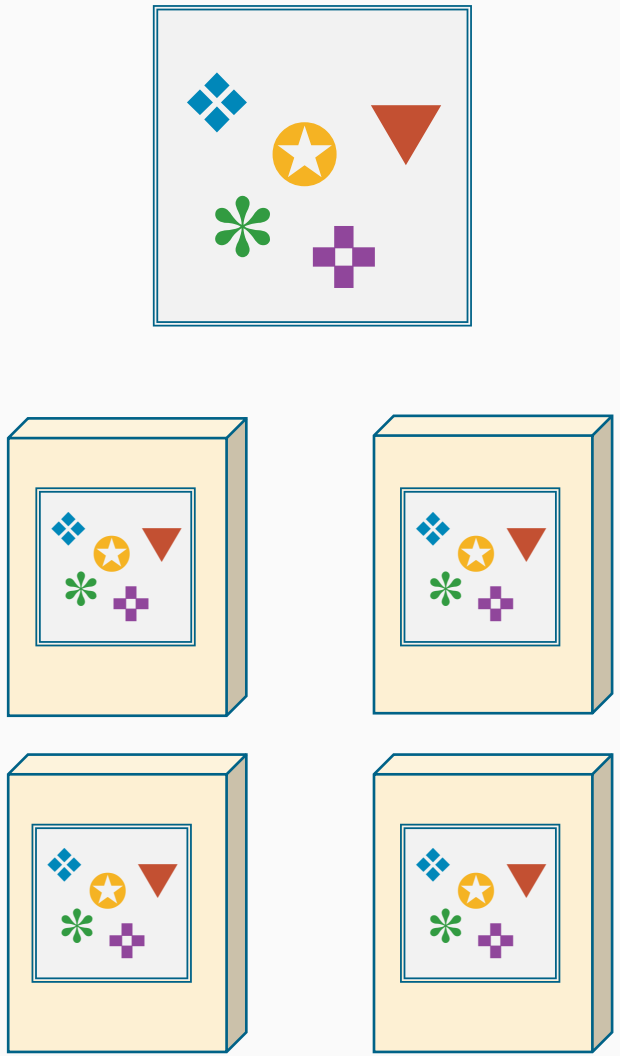

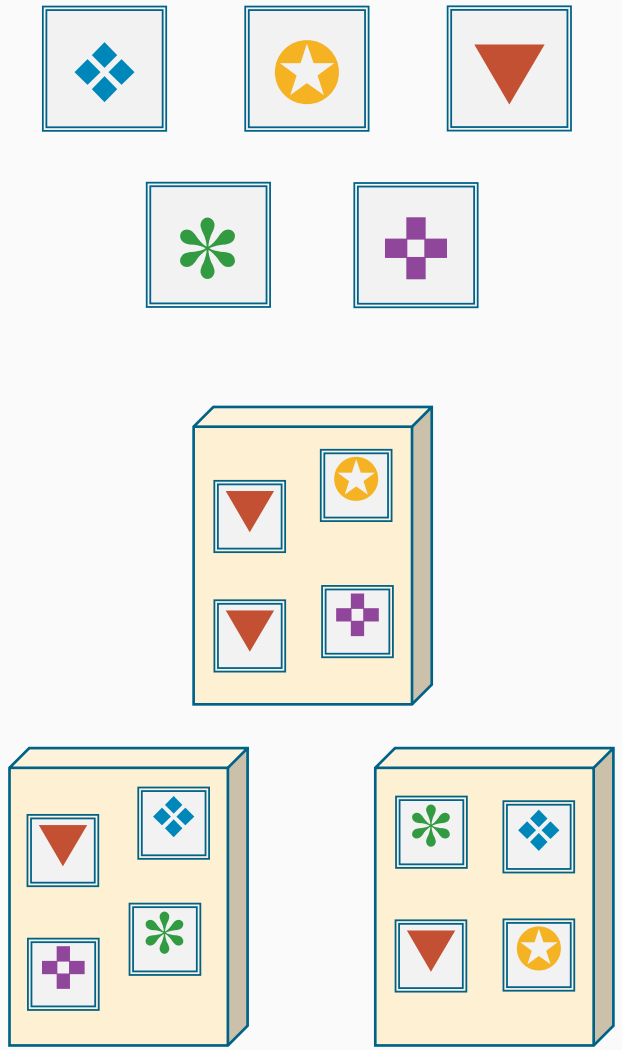

Characteristics of µServices

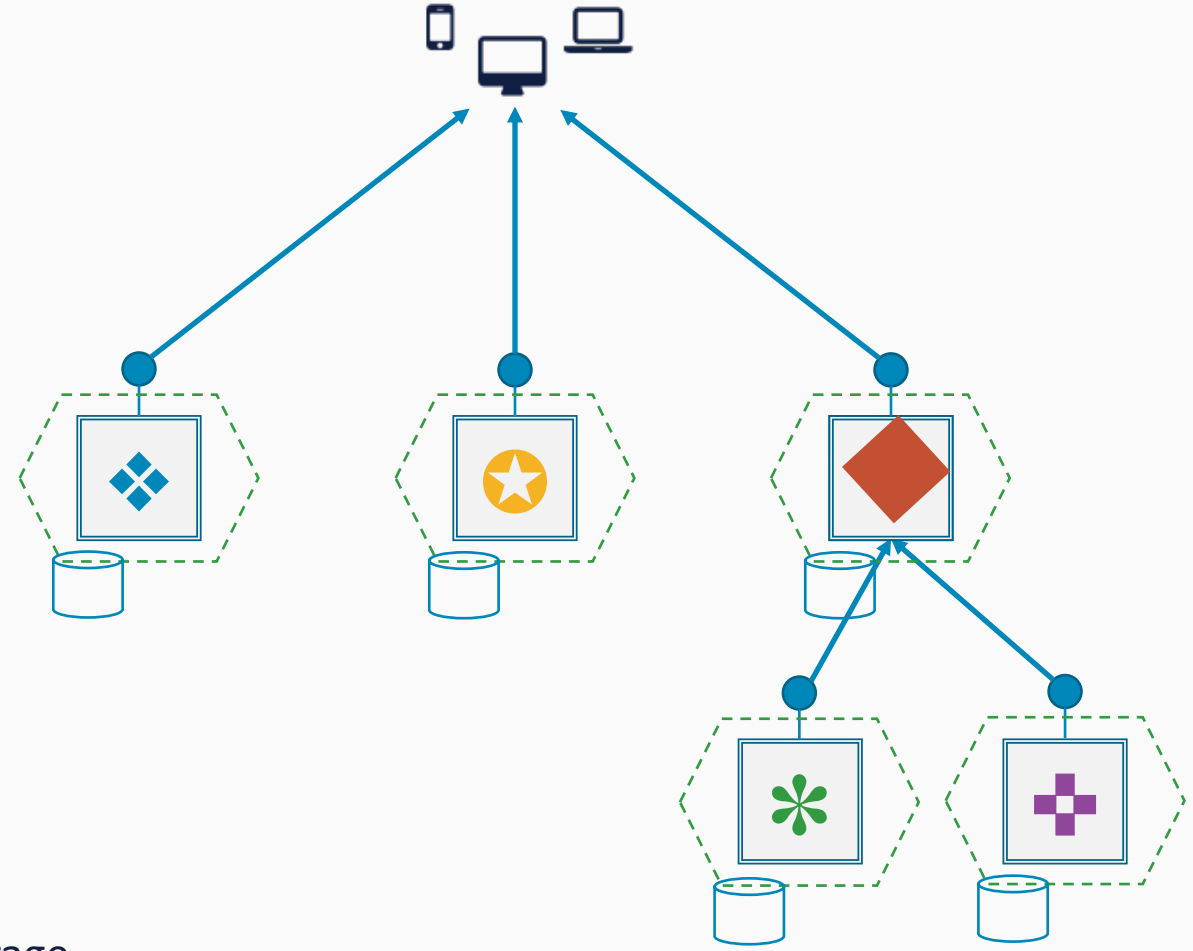

Each service
- runs in isolation
- has a bounded context
- has decentralised data storage
- embraces evolution

Pitfalls & Challenges

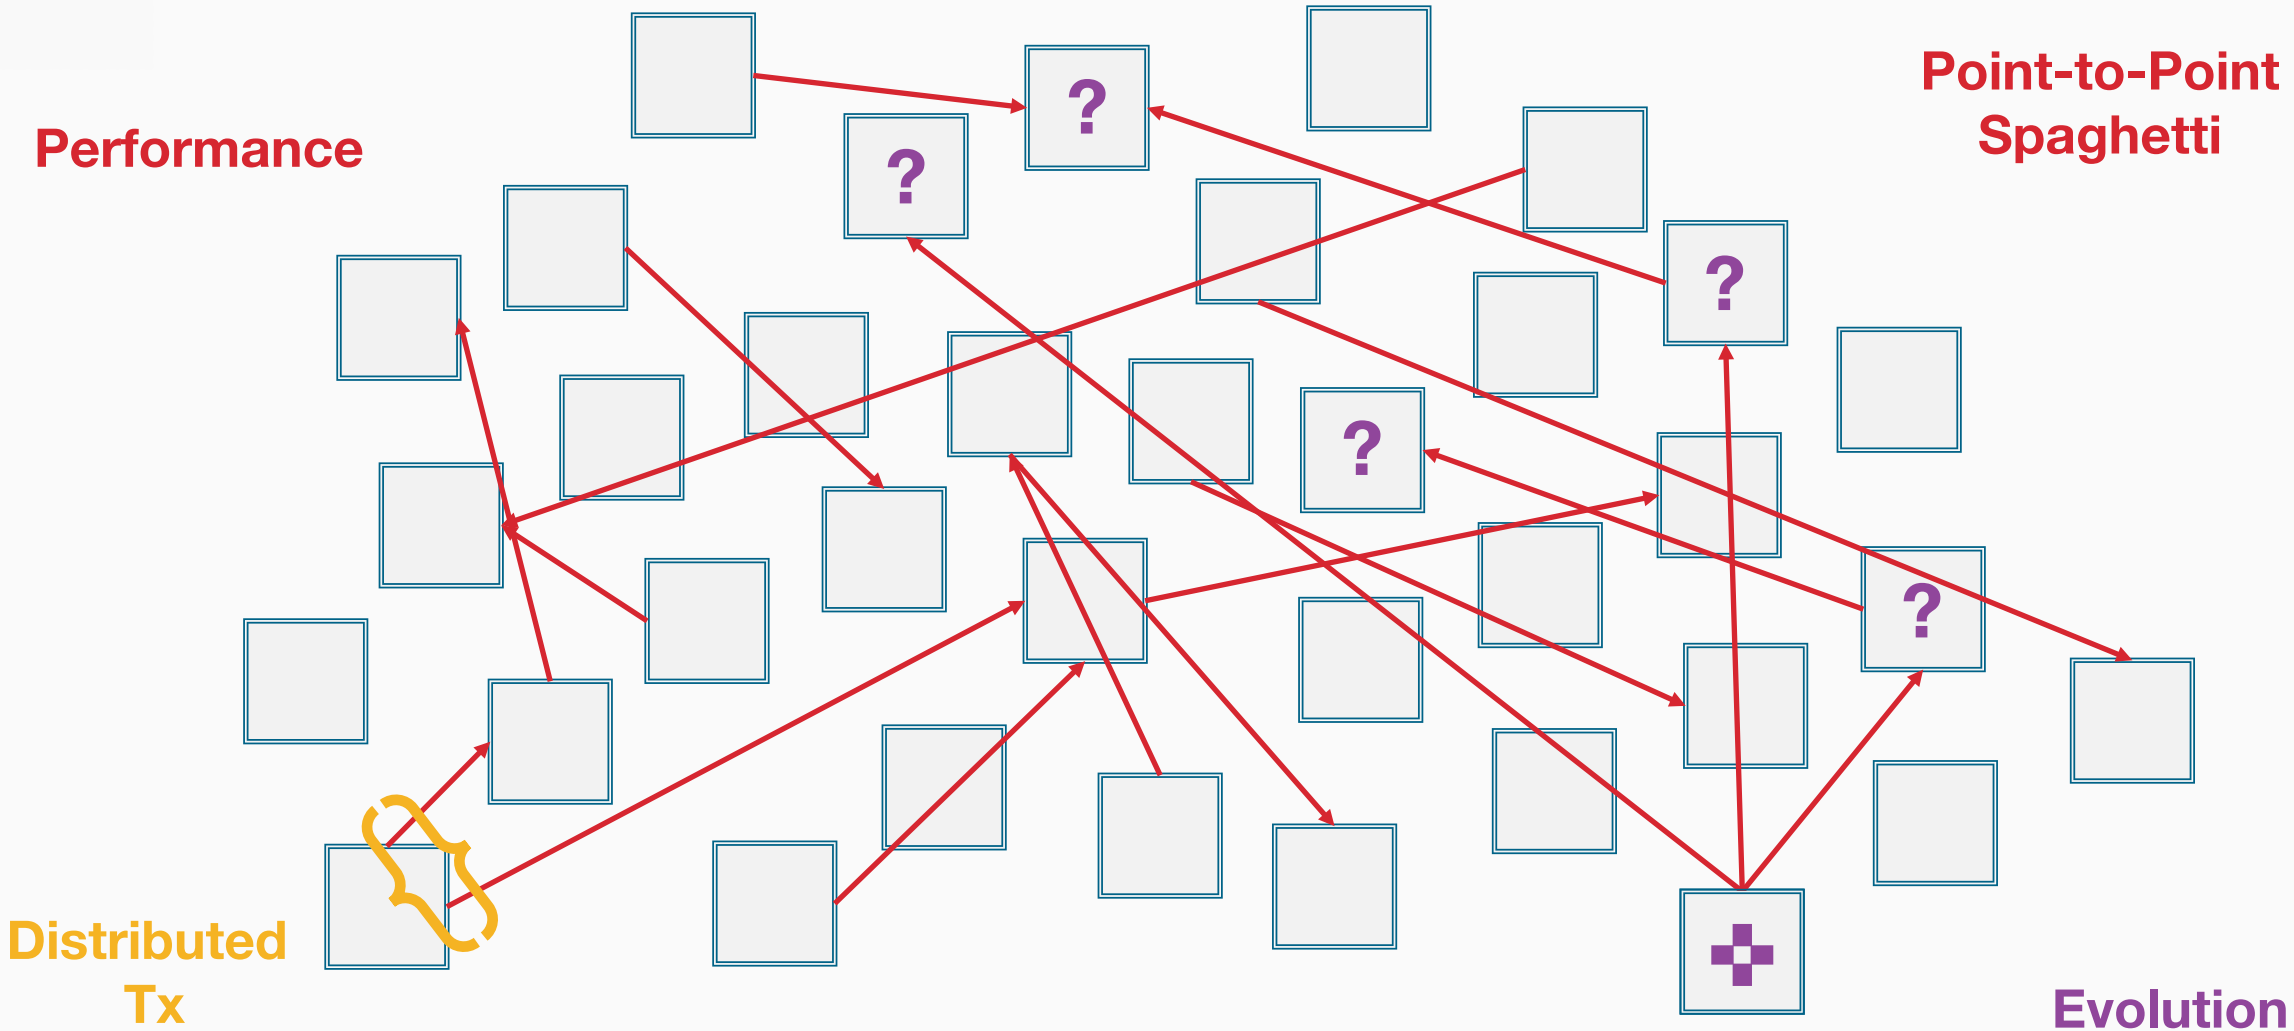

Local Caching

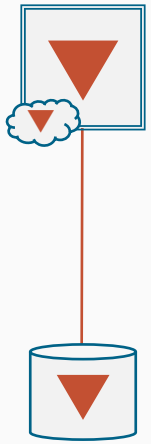

Local Caching

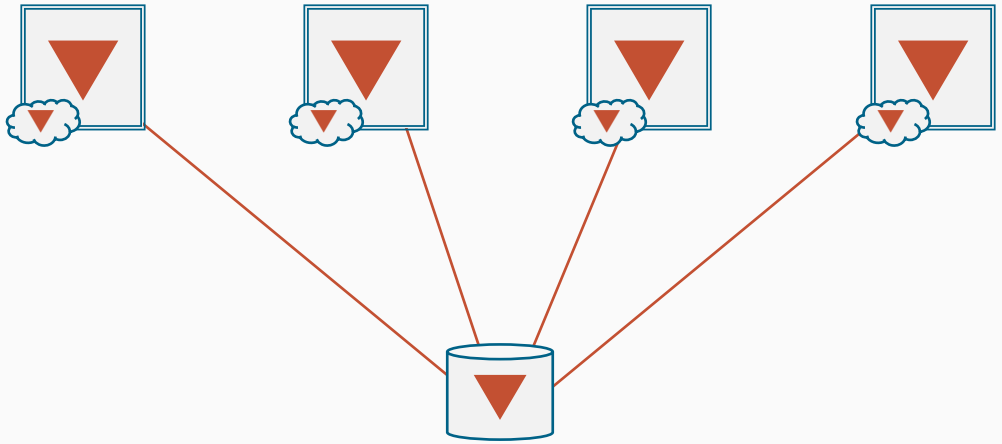

Remote Caching

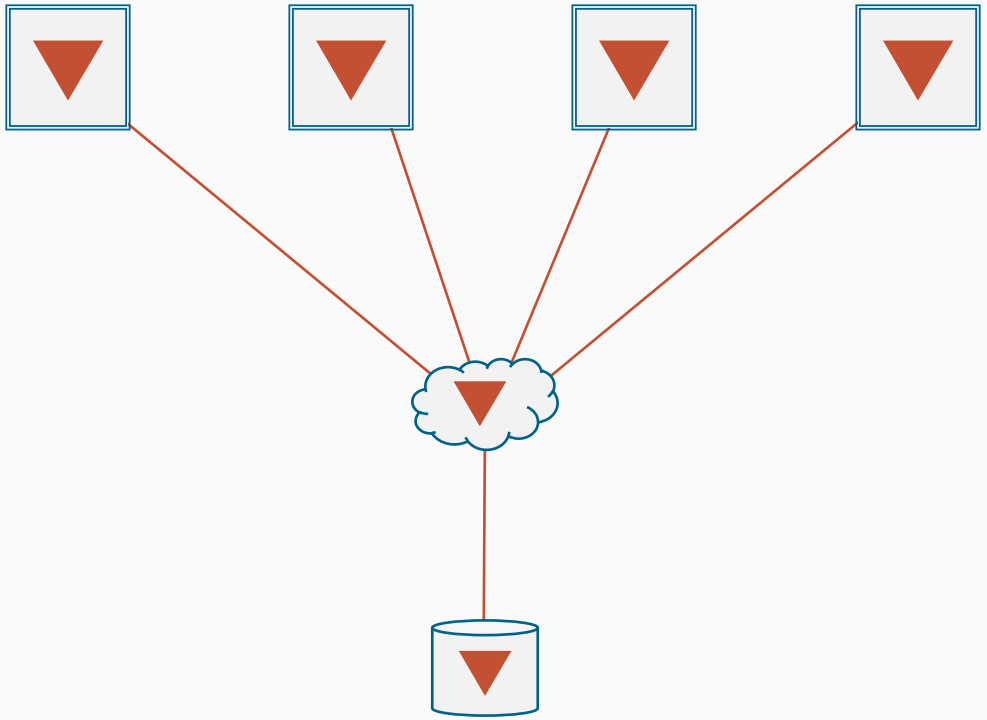

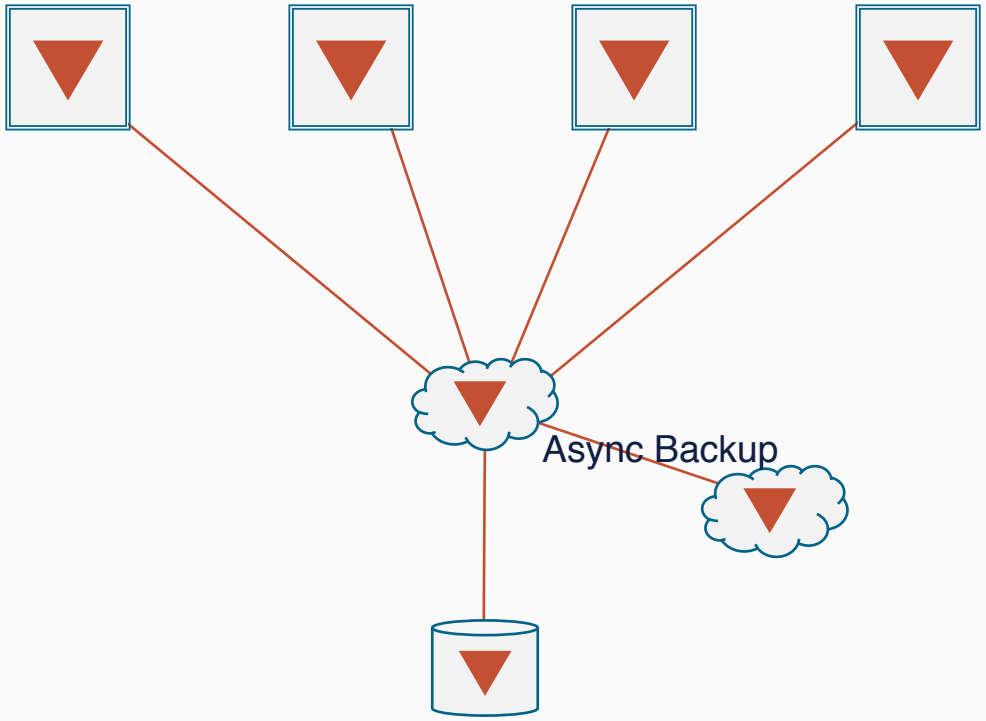

Replicated Distributed Caching

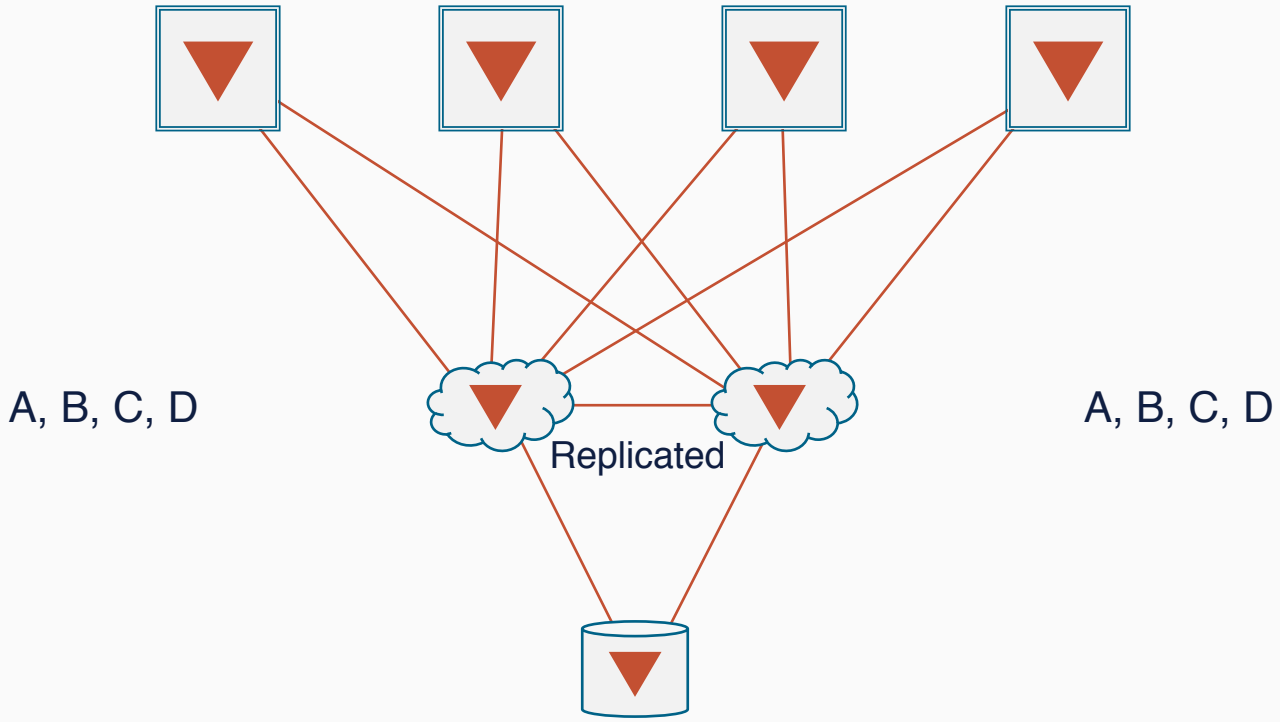

Replica = Backup

Partitioned Distributed Caching

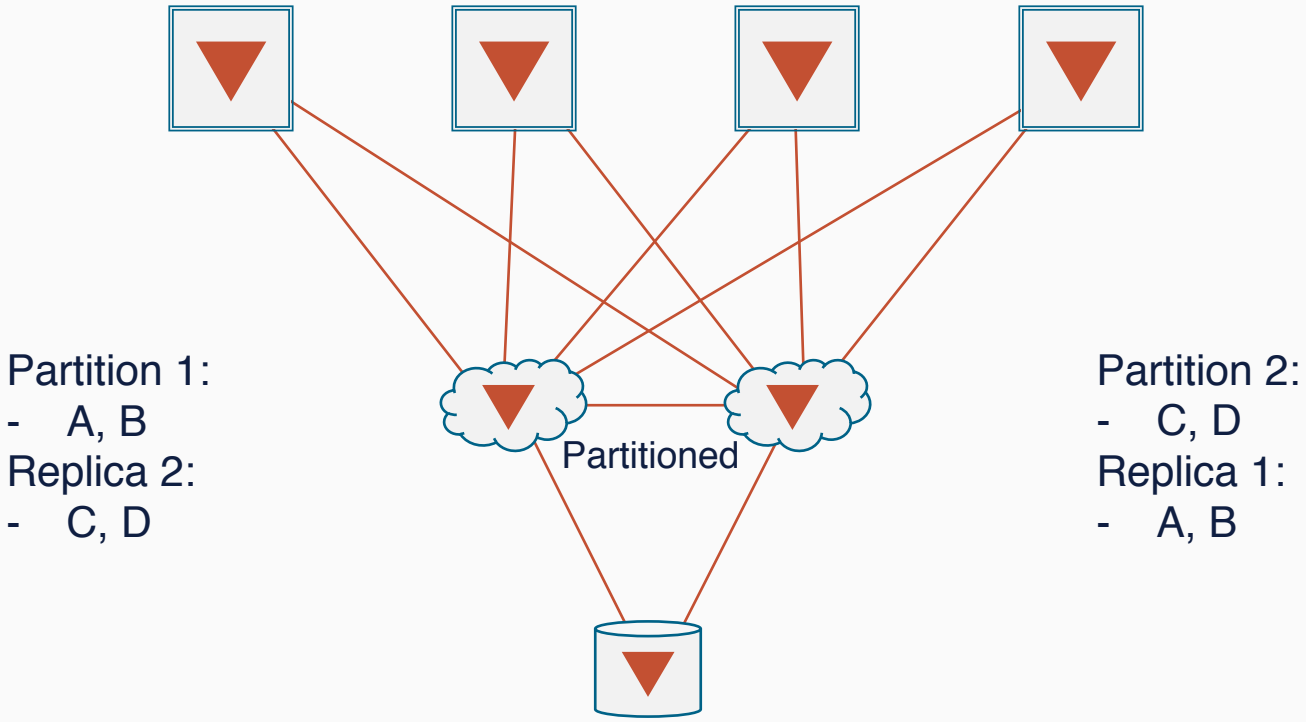

Replica = Backup

Partitioned Distributed Caching: Scaling

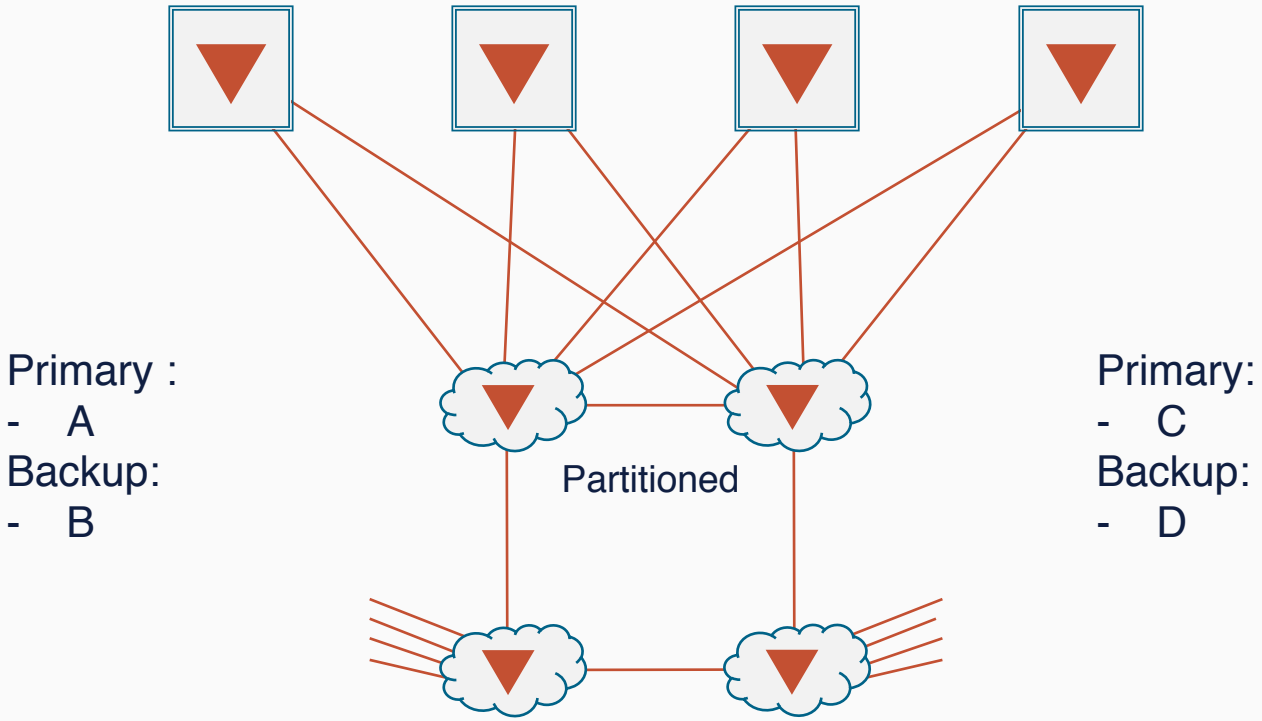

Replica = Backup

Adding Data

G

H

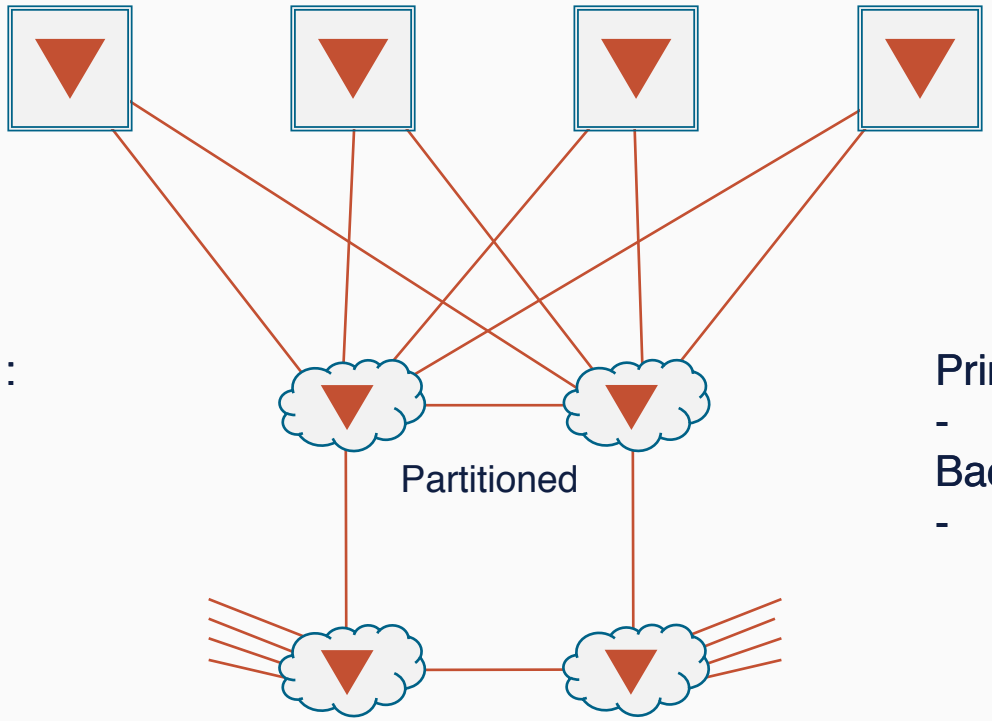

Primary :
- A, G

Backup:
- B

Primary:
- C, H

Backup:
- D

Primary:
- D

Backup:
- A, G

Primary:
- B

Backup:
- C, H

Replica = Backup

Grid Systems: Affinity Collocation

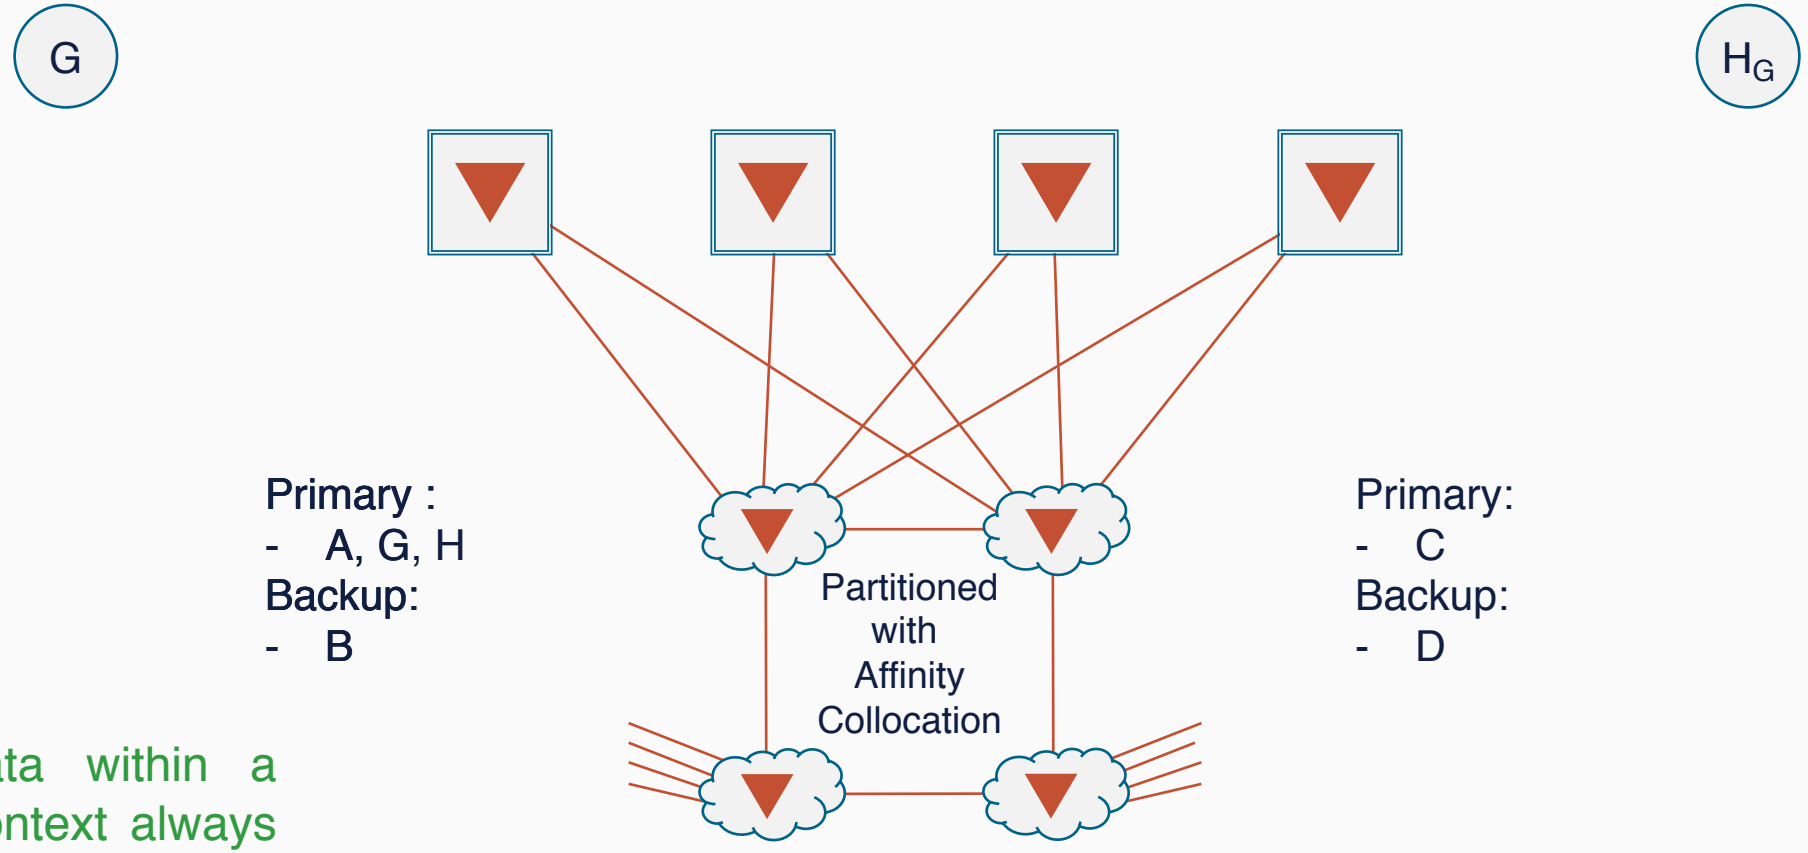

Related data within a bounded context always go to the same partition

Replica = Backup

Grid Systems: Affinity Collocation

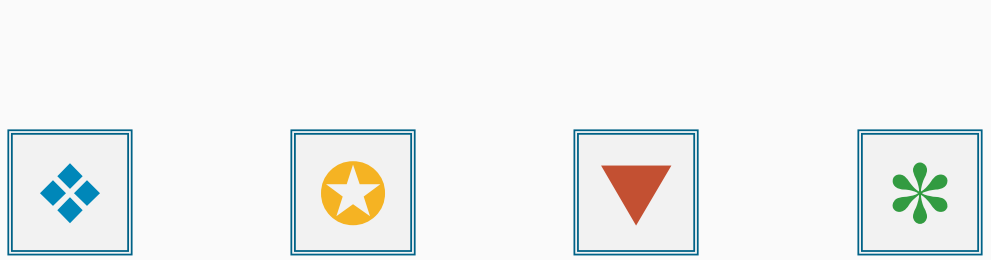

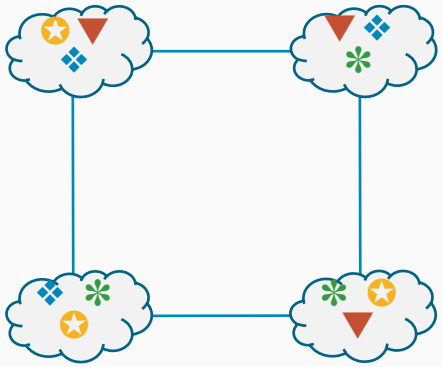

Primary:
- ✻ C

Backup:
- ✻ A, D

Primary:
- ✻A, D

Backup:
- ✻B

Primary:
- ✻ B

Backup:
- ✻ C

Replica = Backup

Enter “The Glue”

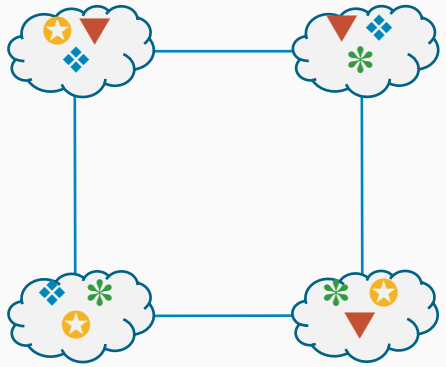

Enter “The Glue”

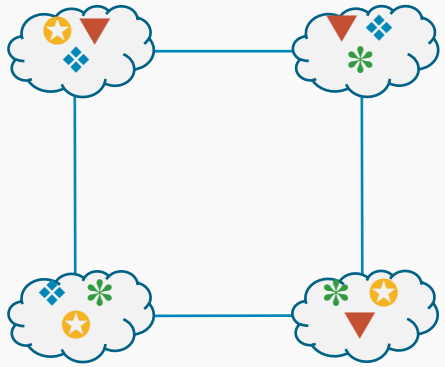

Data Centric

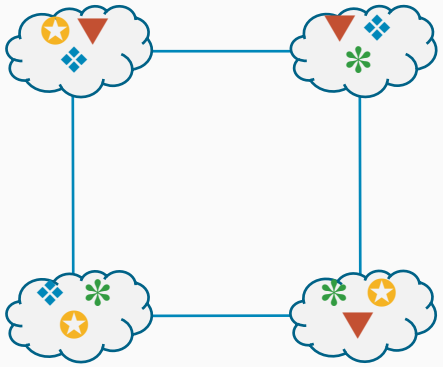

Data Centric

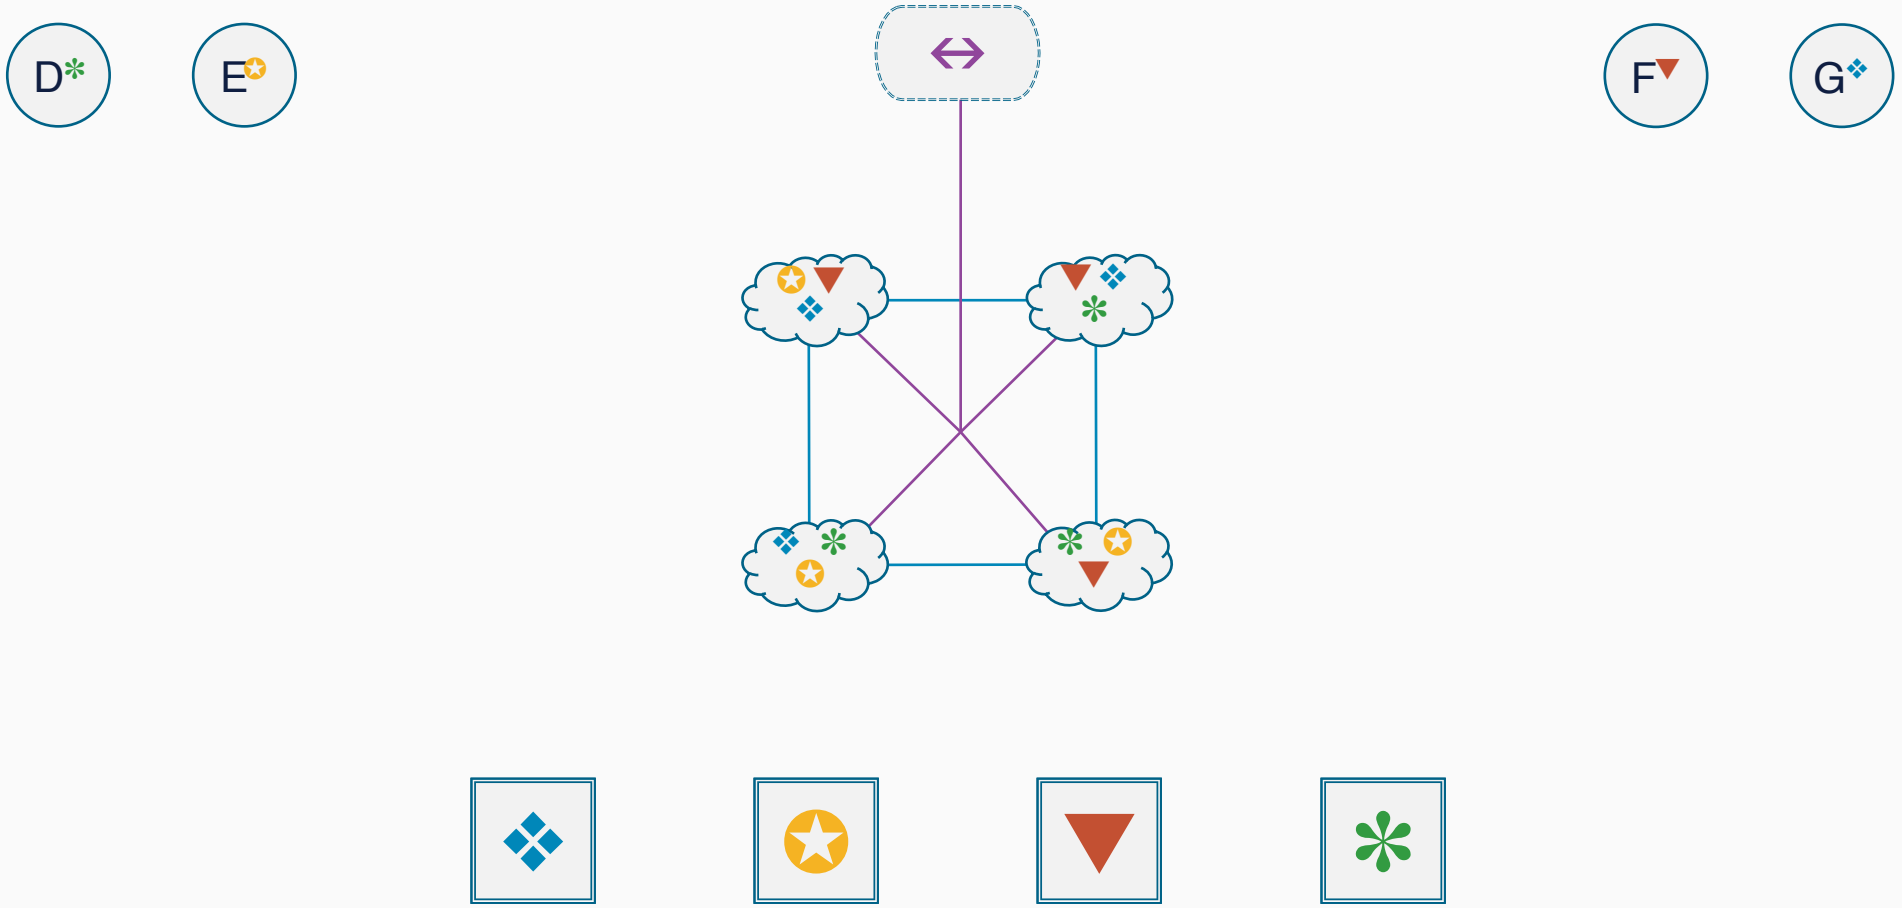

Compute Collocation

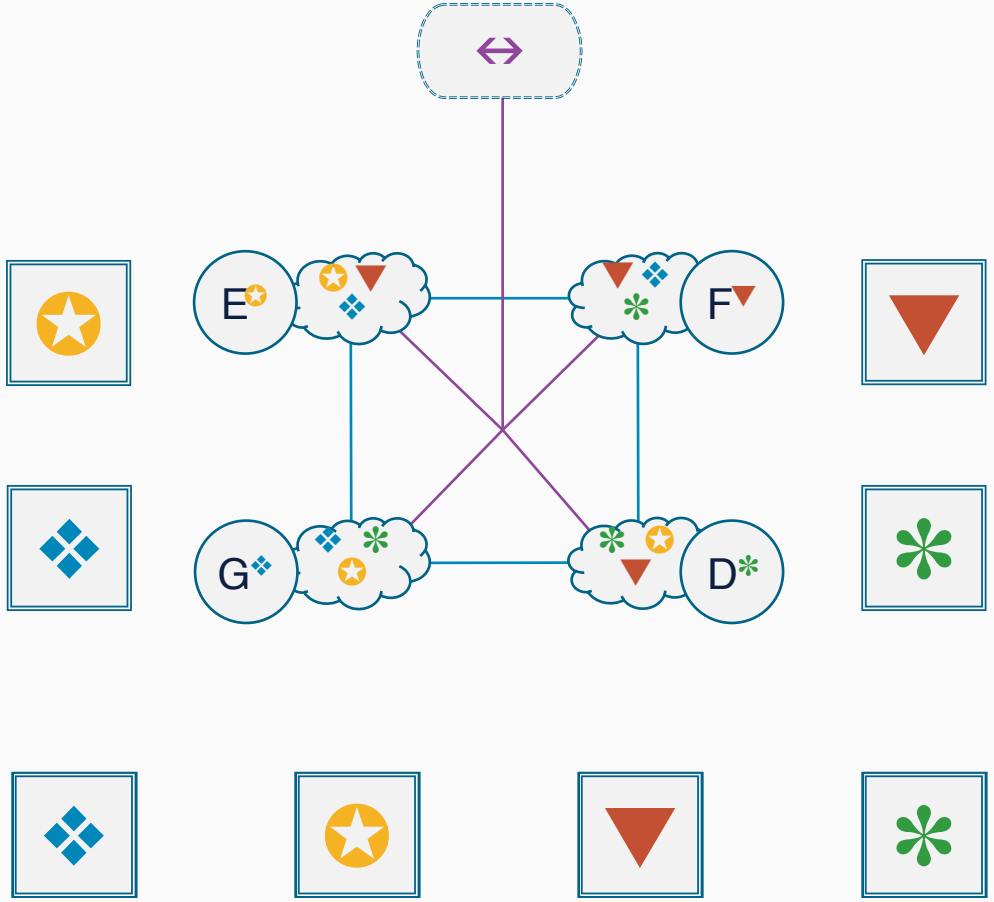

Compute Collocation

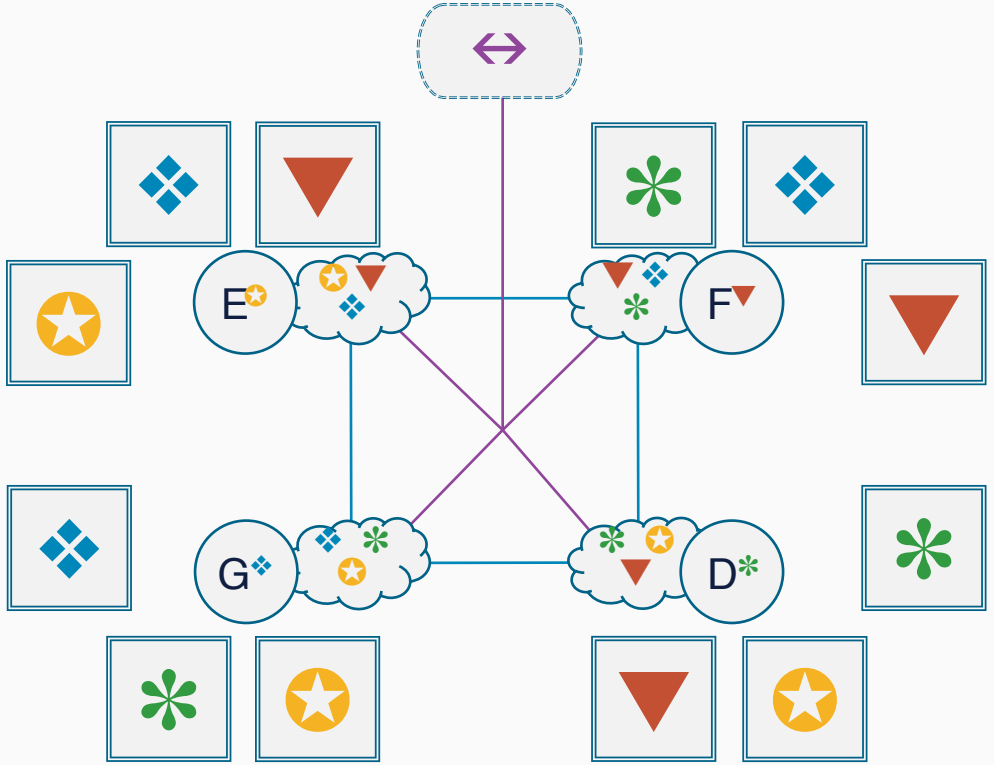

Embedding Grid Nodes

Embedding Grid Nodes

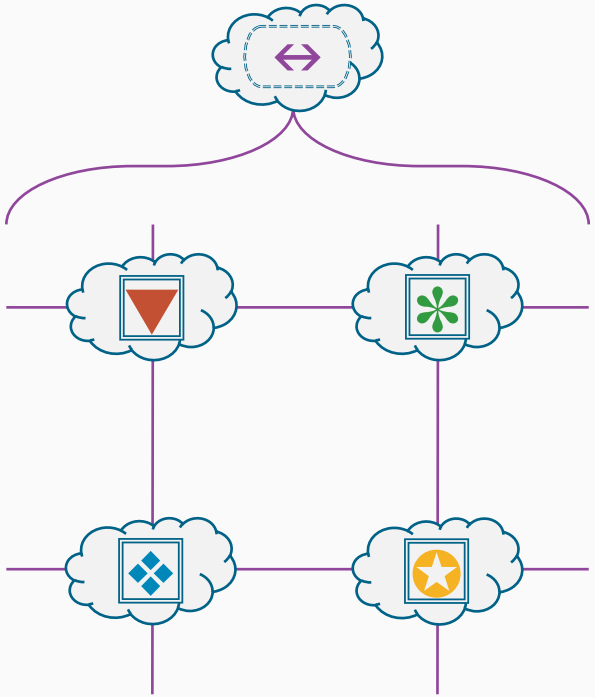

Scaling

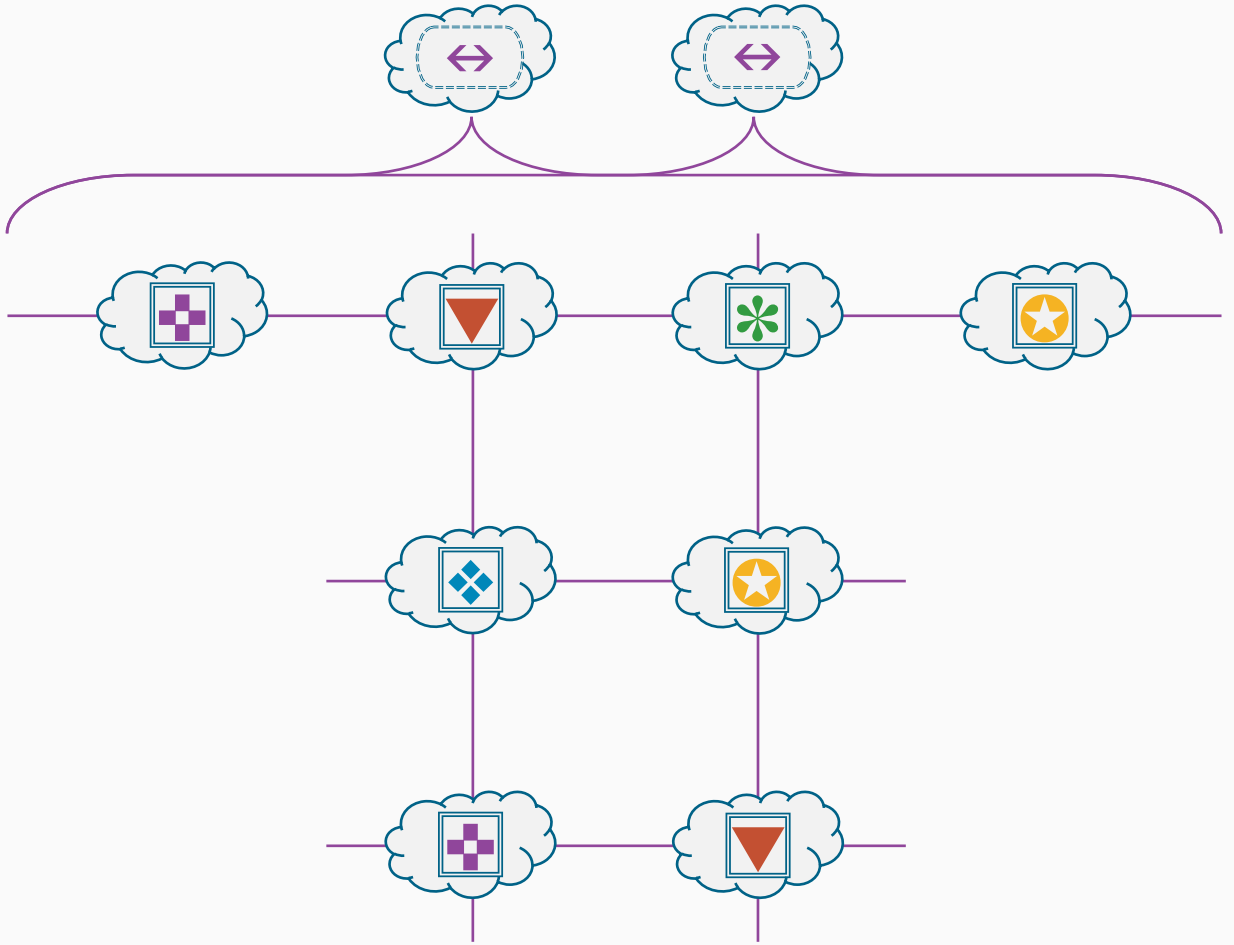

Scaling

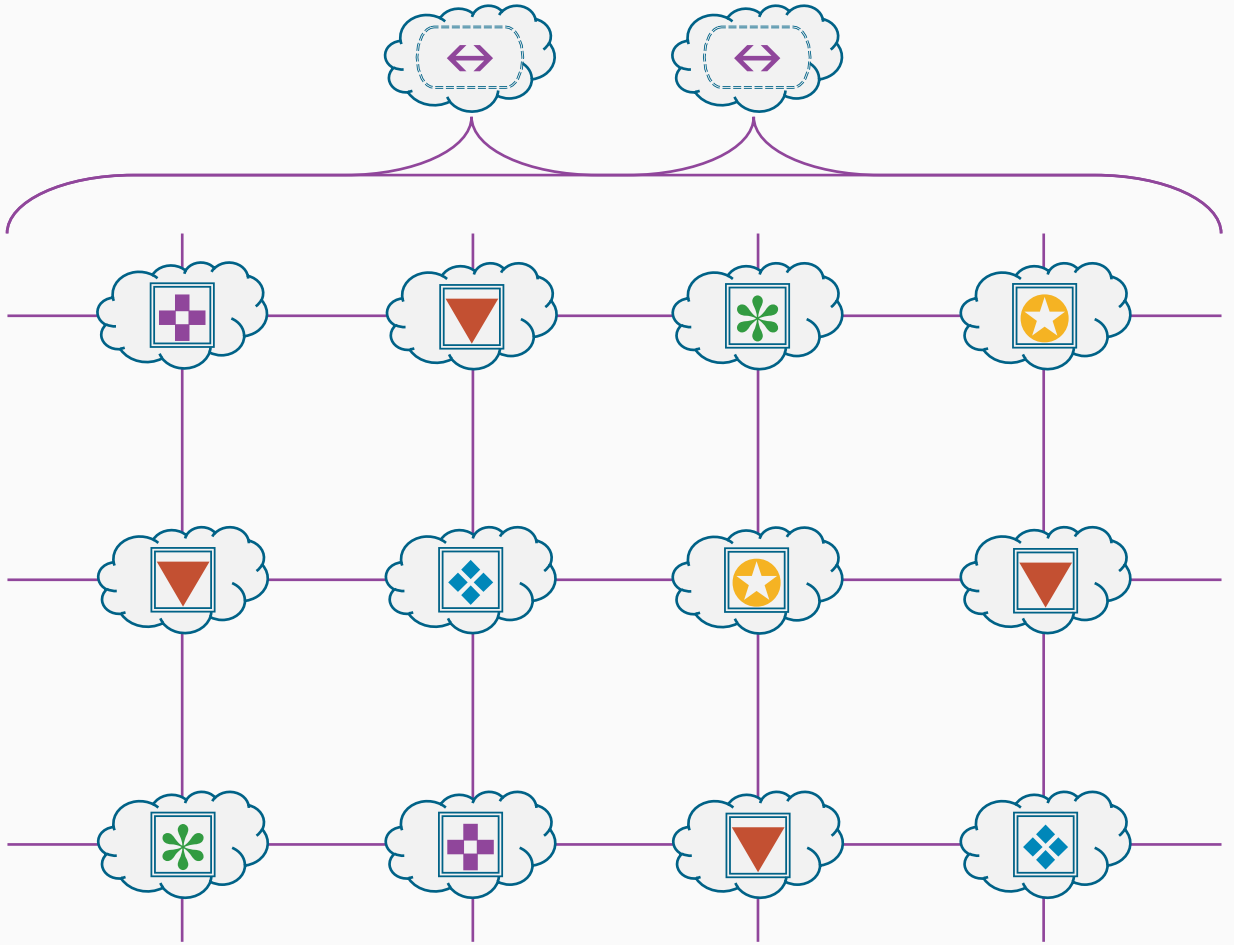

Event Driven / Reactive Processing

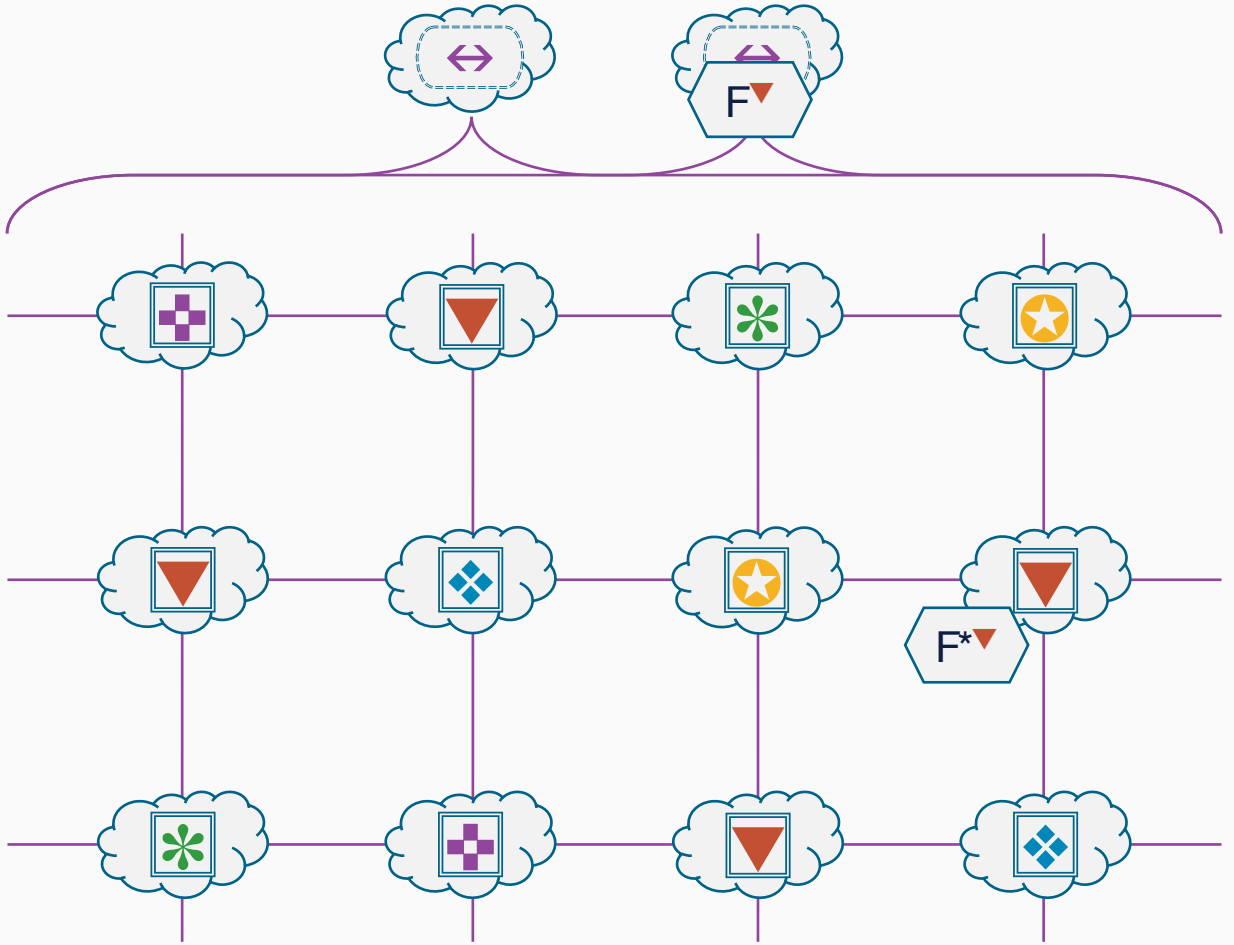

Event Driven / Reactive Processing

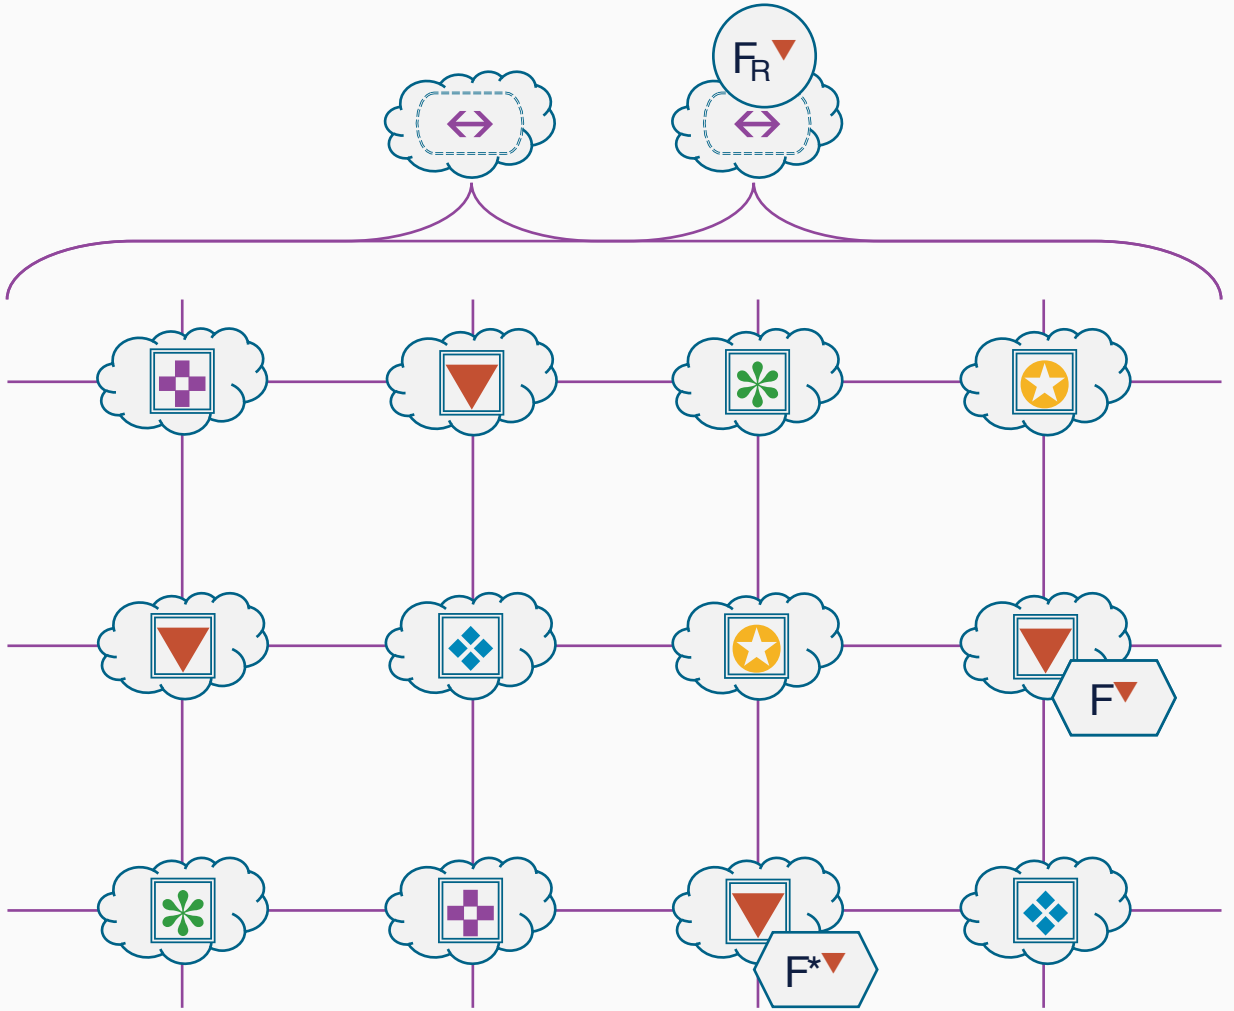

Exactly Once Processing

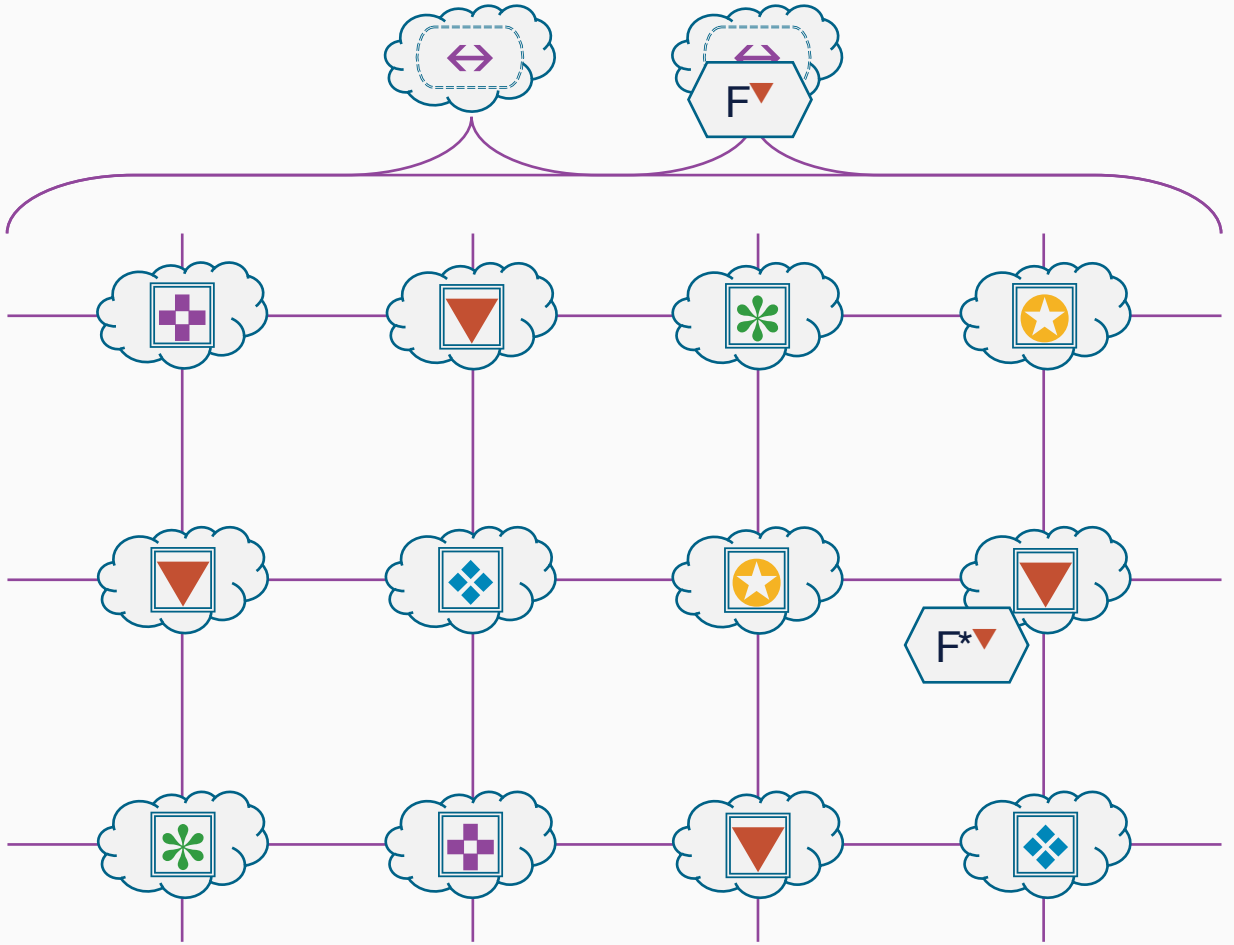

Exactly Once Processing

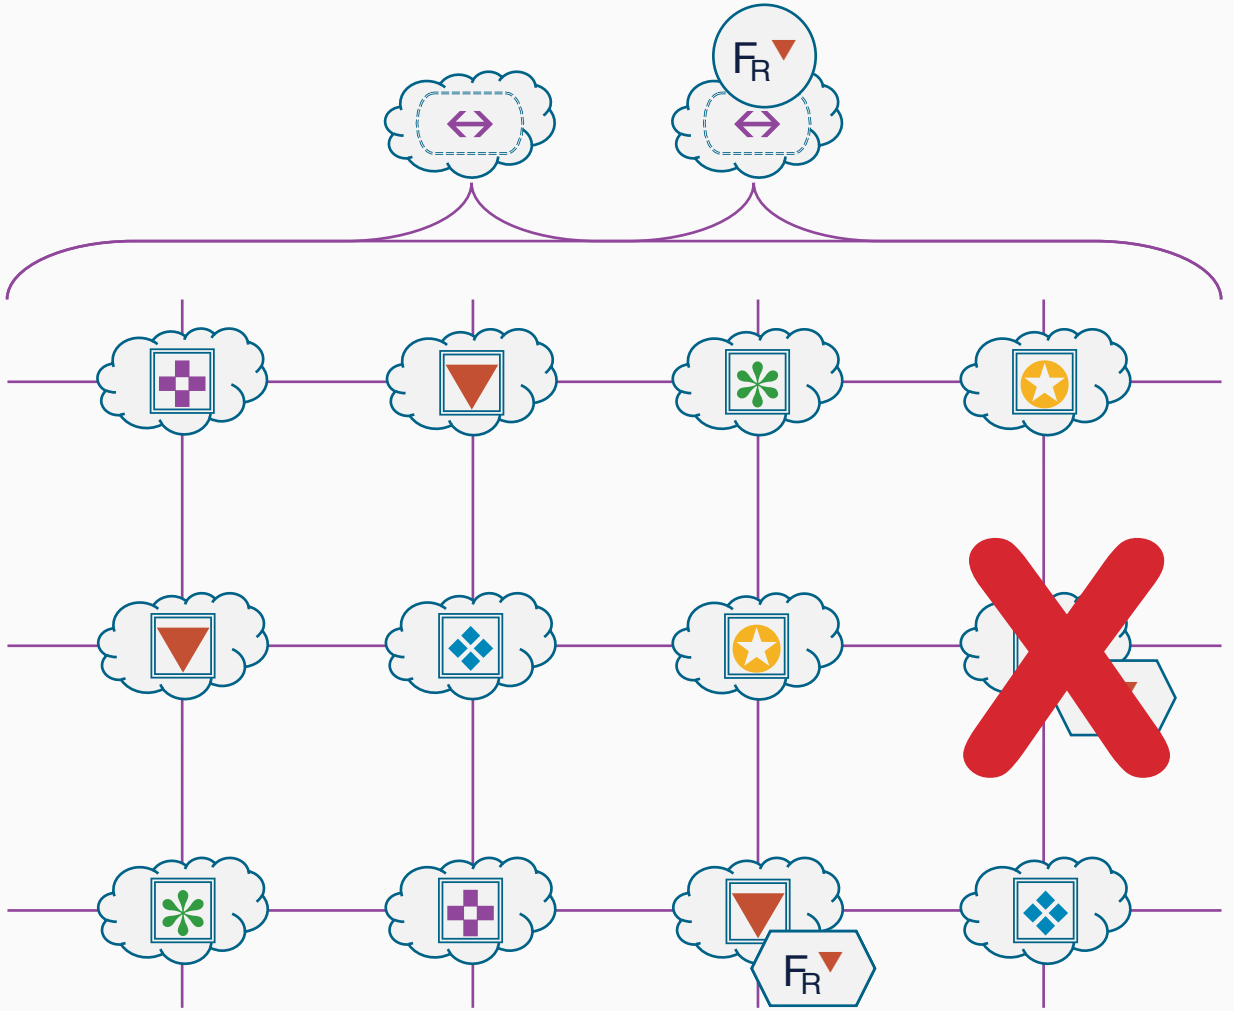

Distributed Transactions

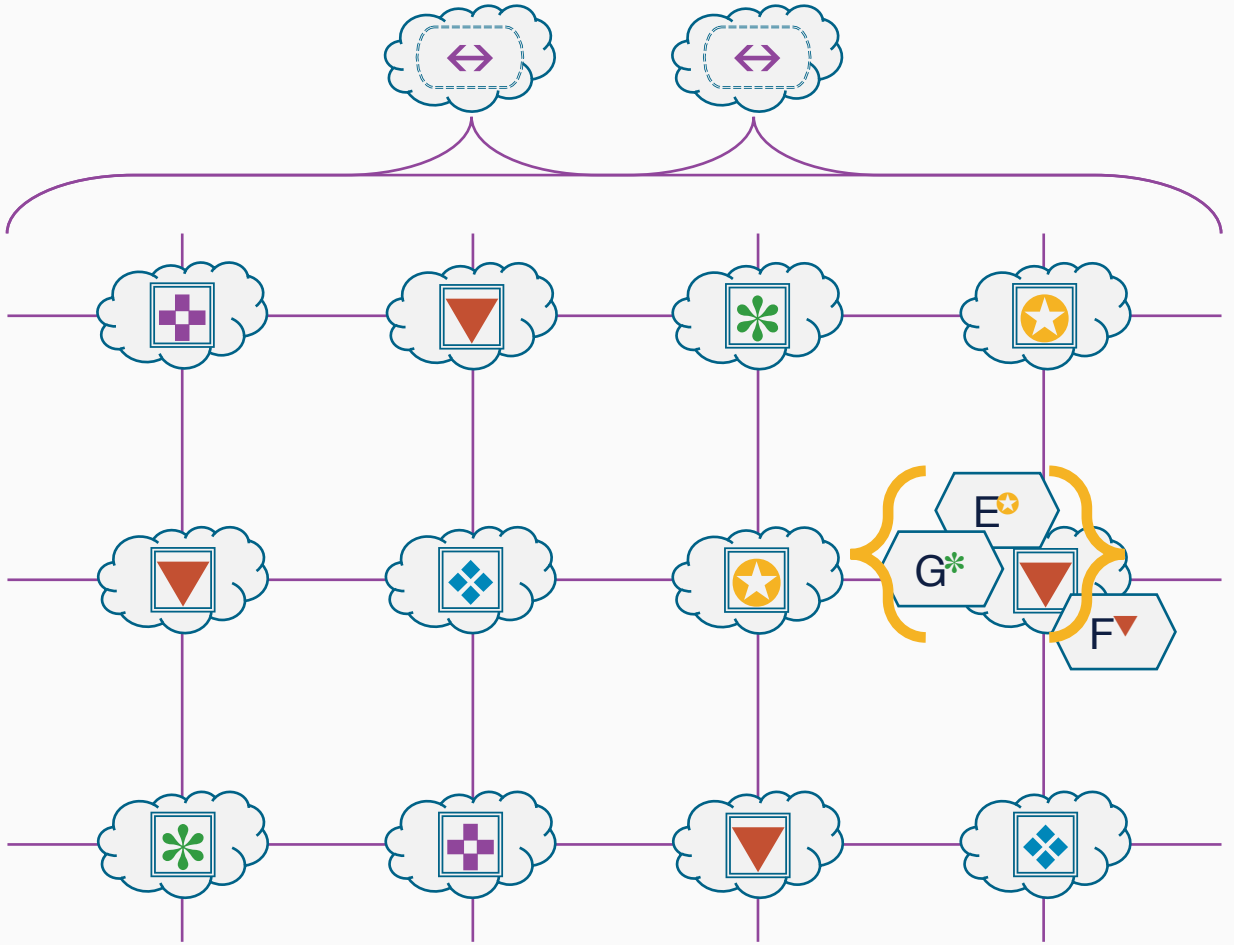

Adding Communication Services

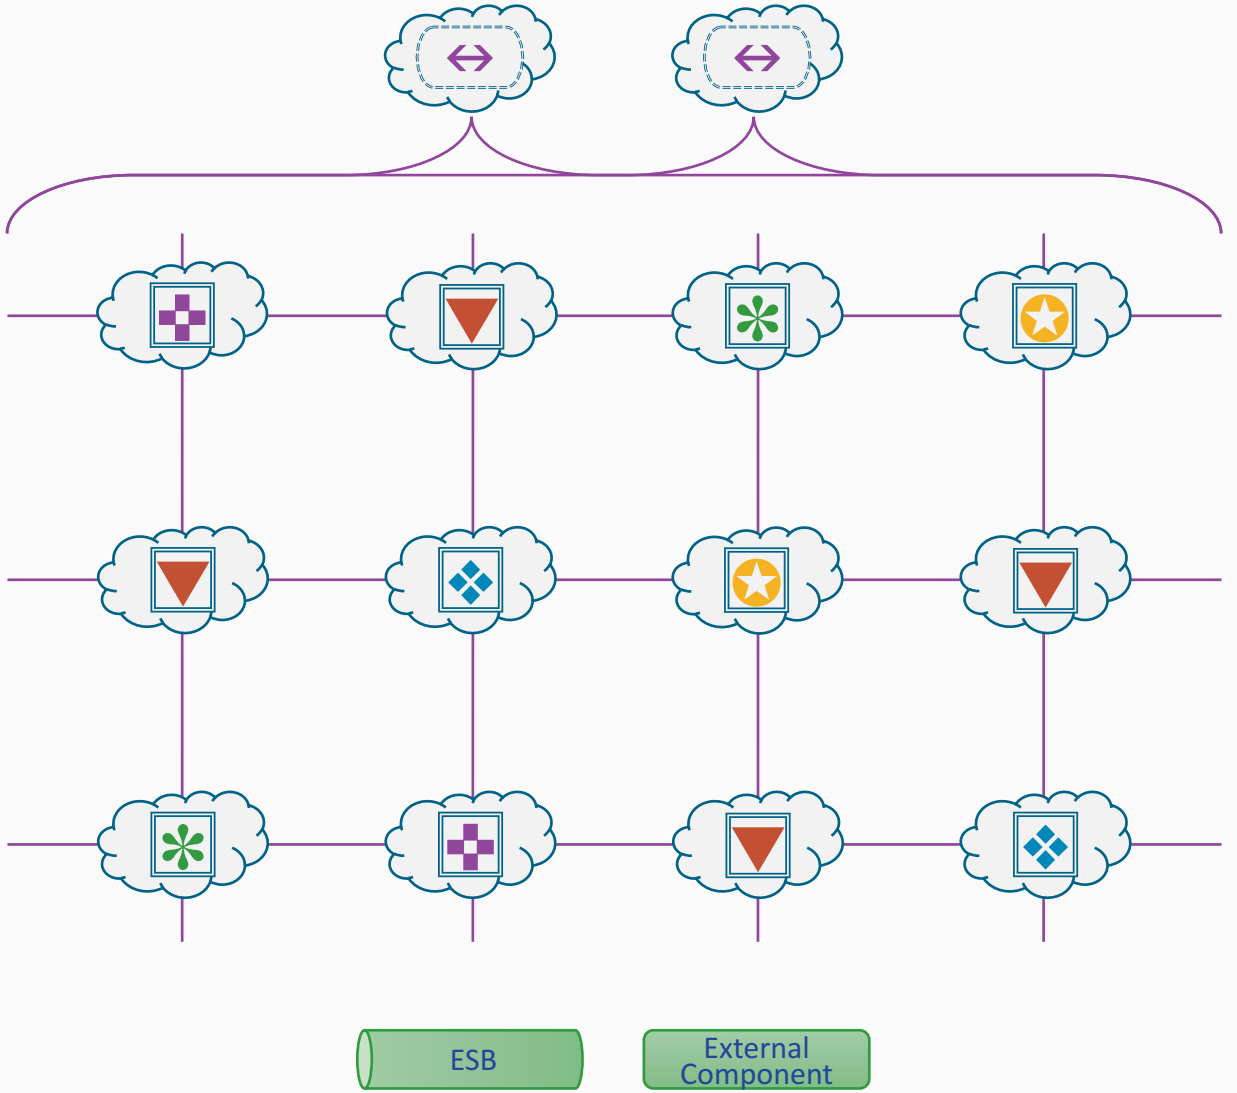

Adding Communication Services

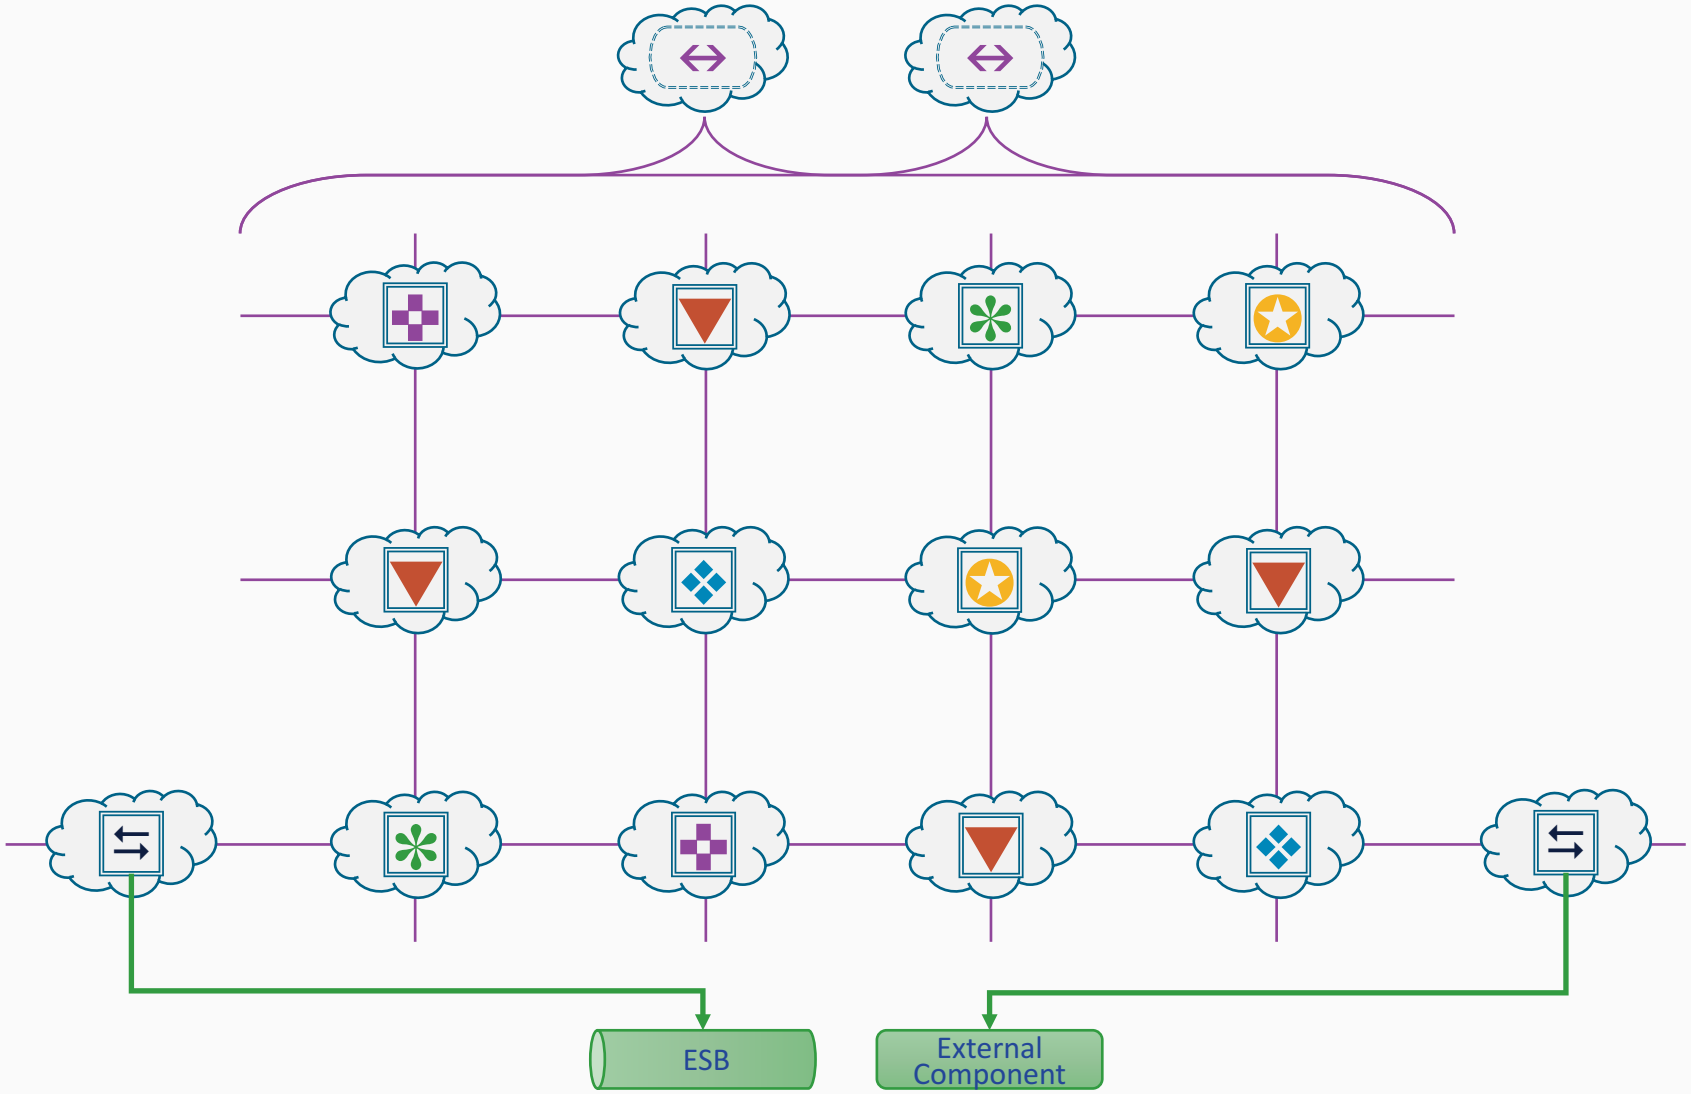

Adding Entity Caches

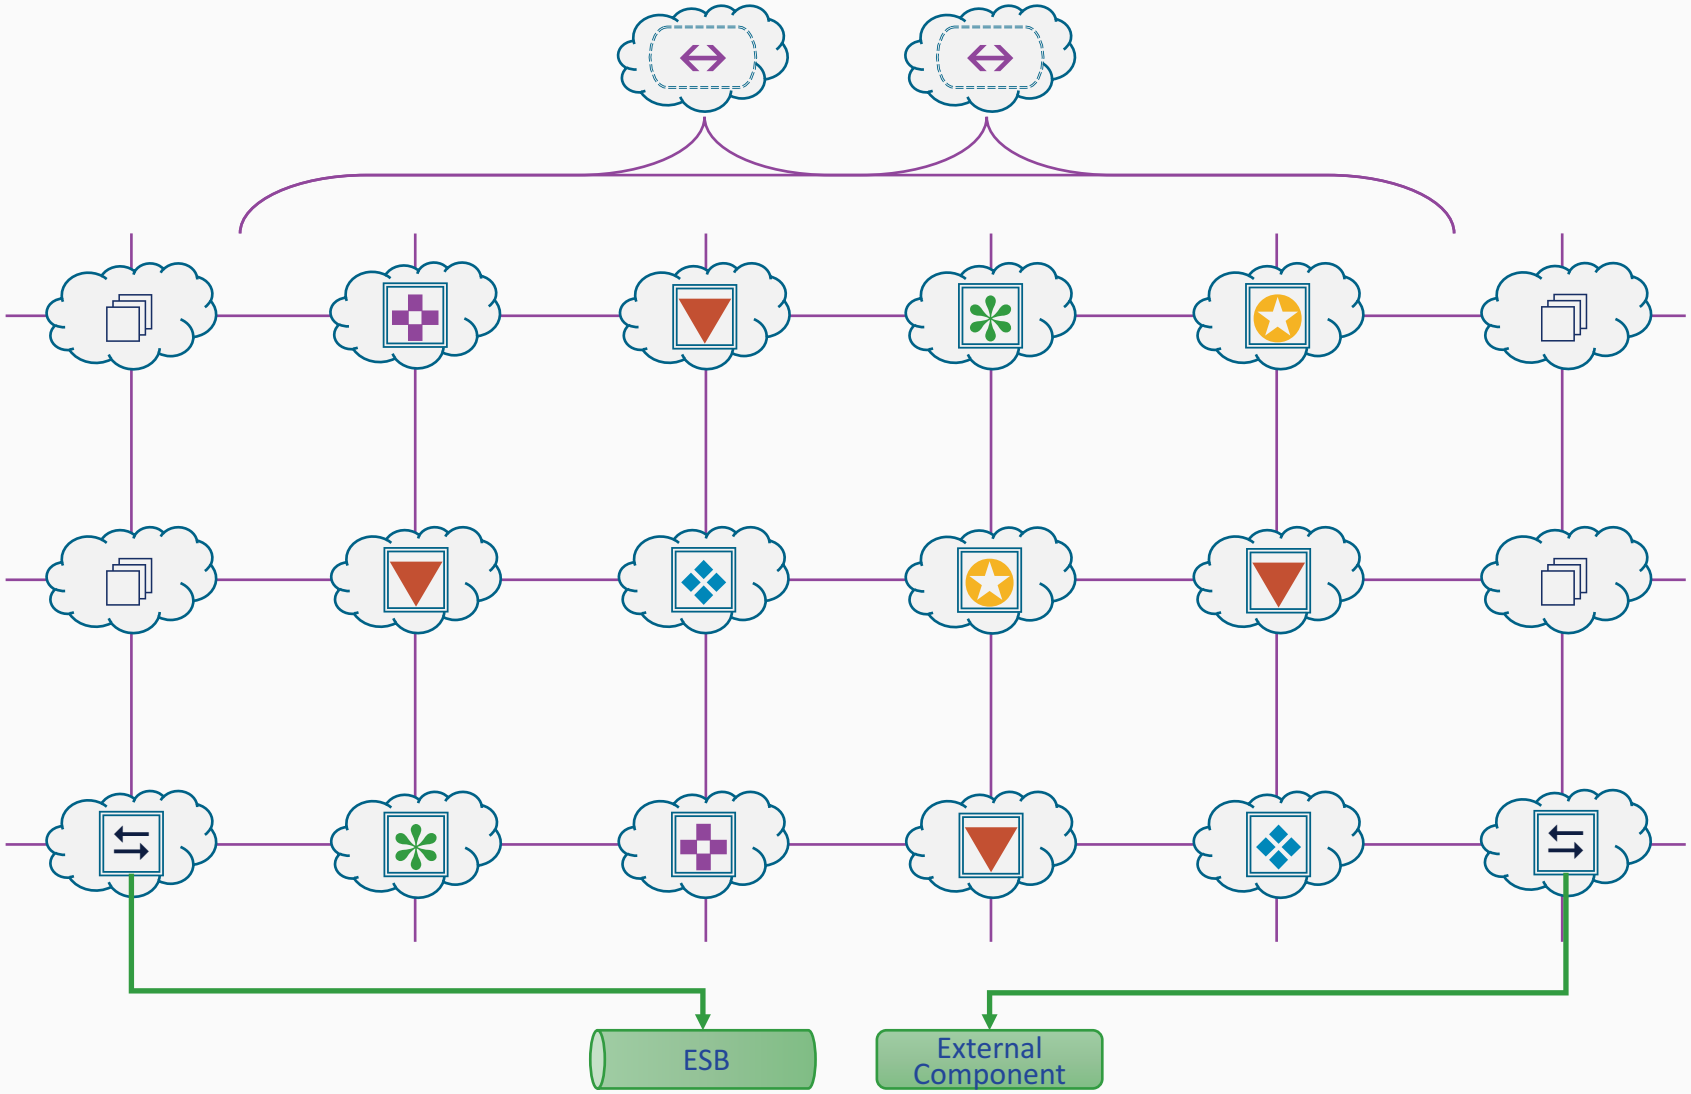

Containerization

Containerization

Containerized Grid

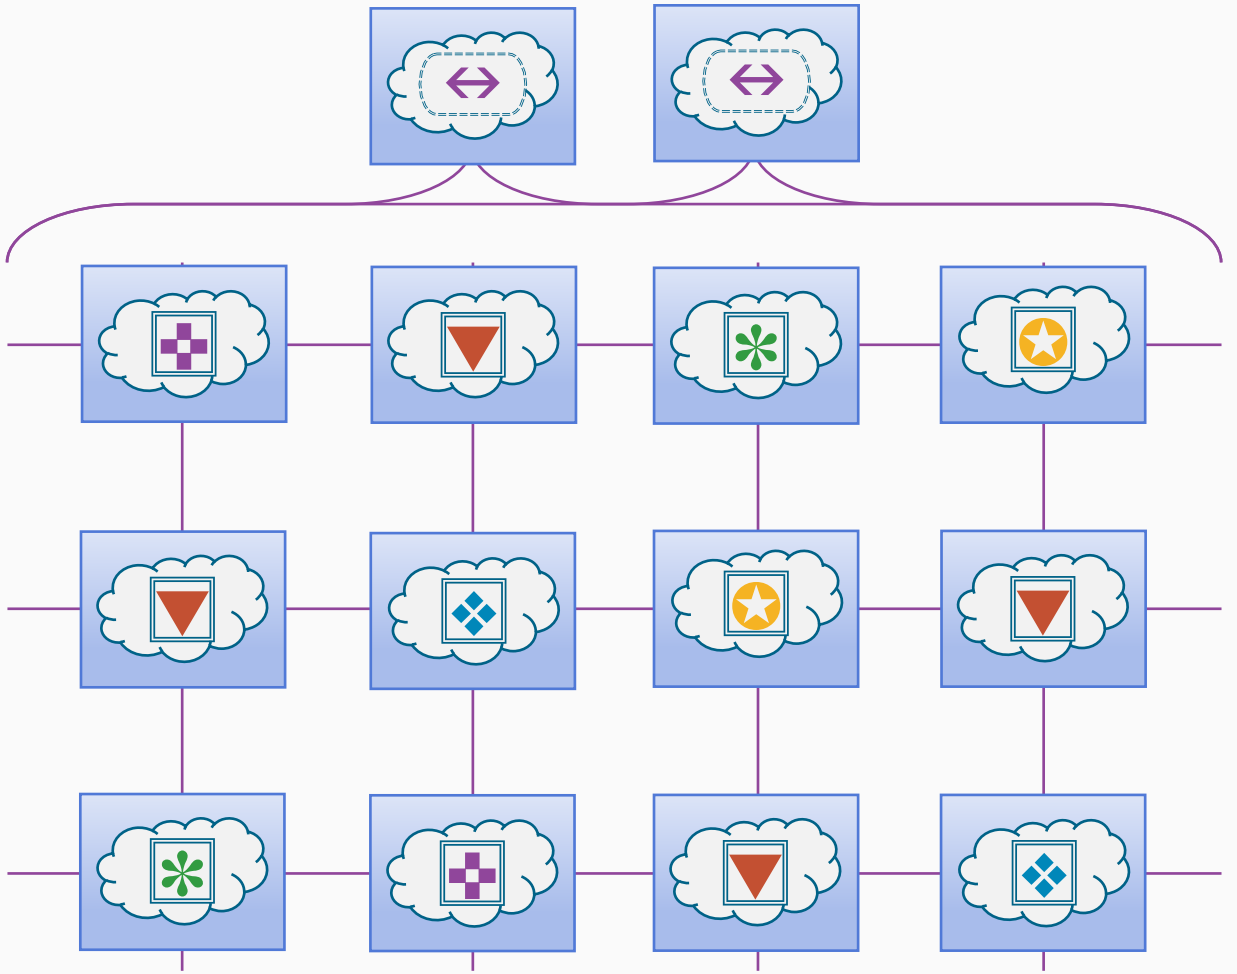

Deployment

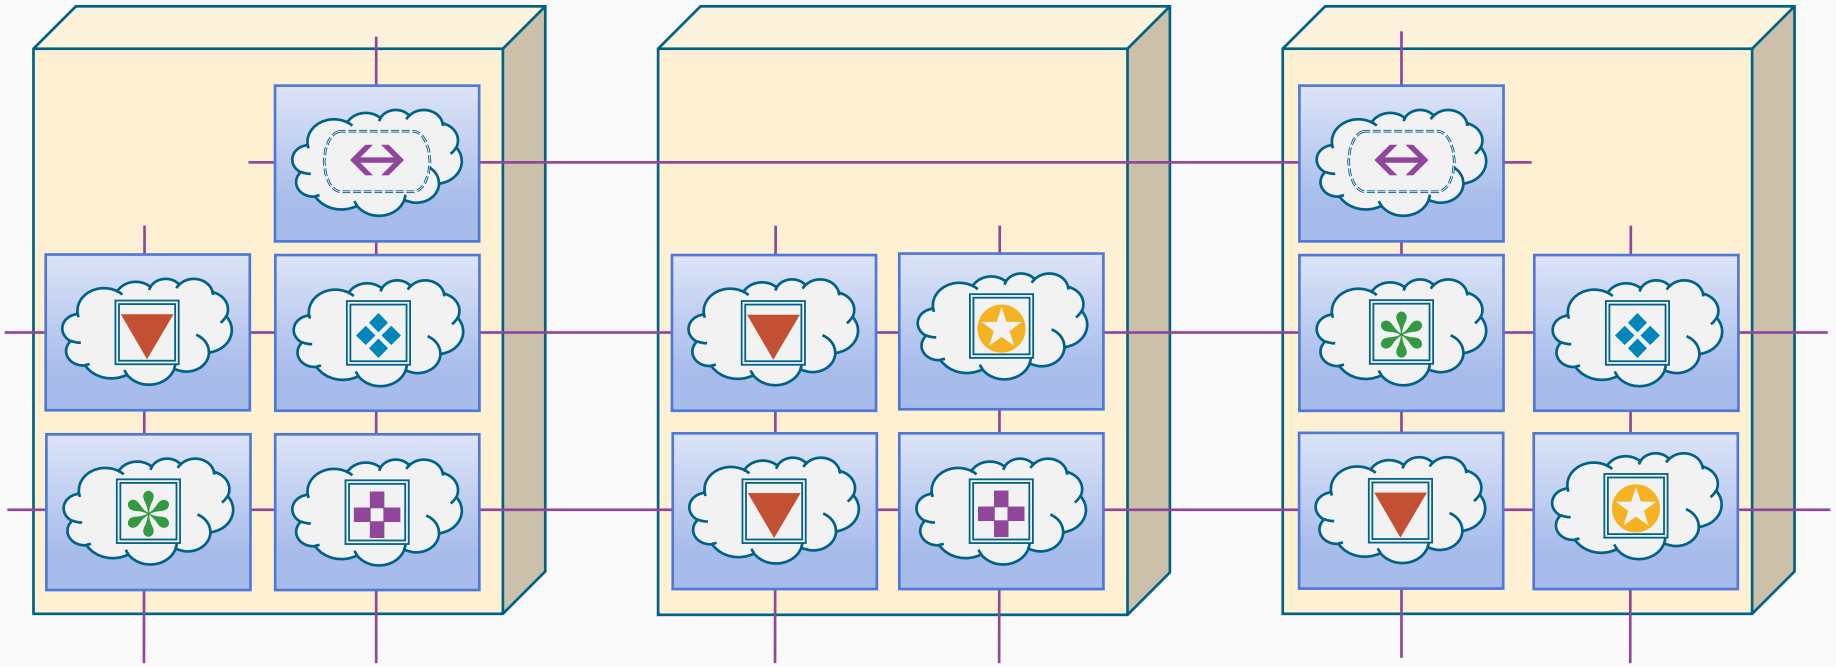

Software Defined Network

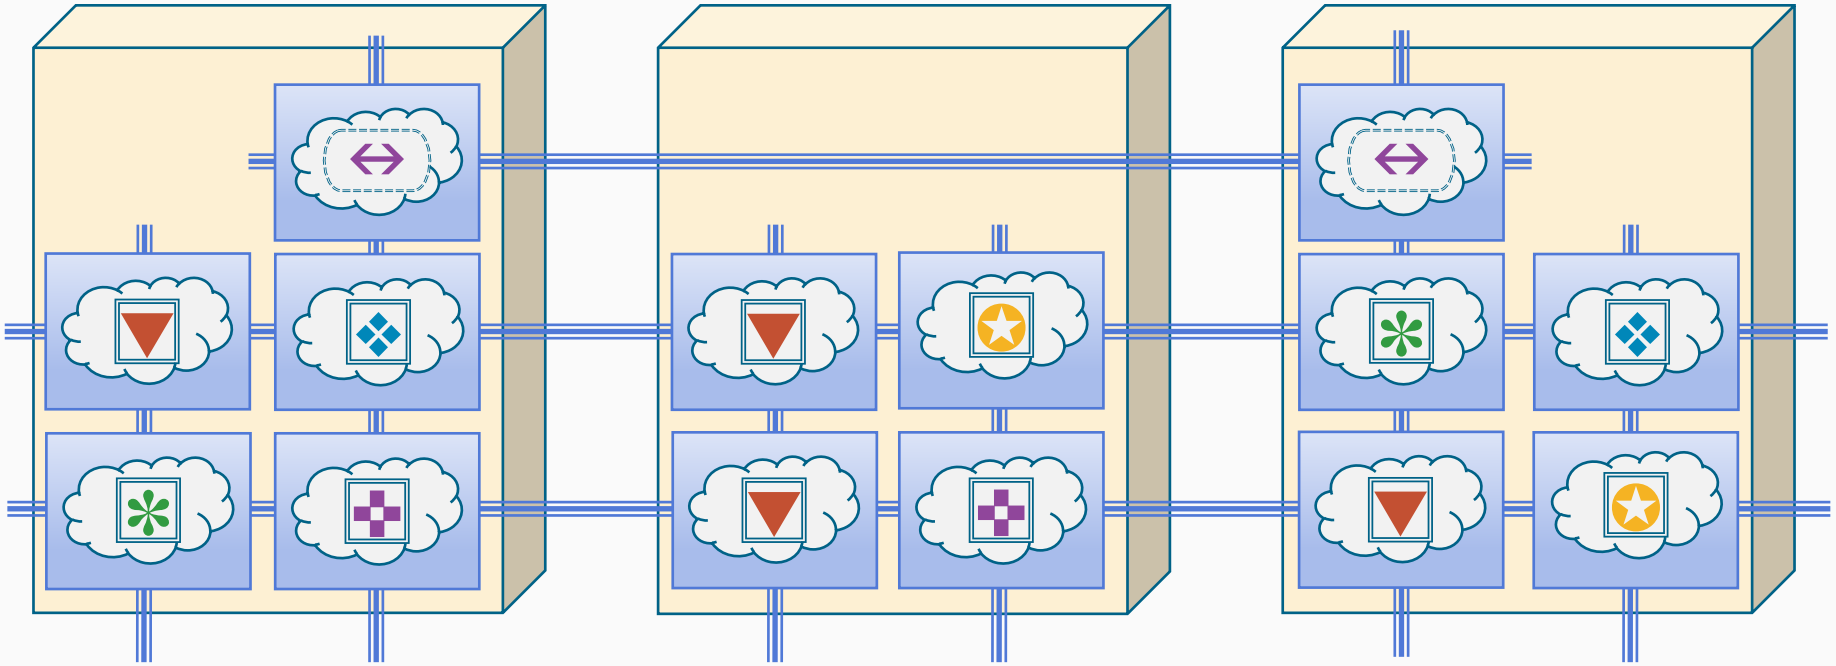

Orchestration with Single Master

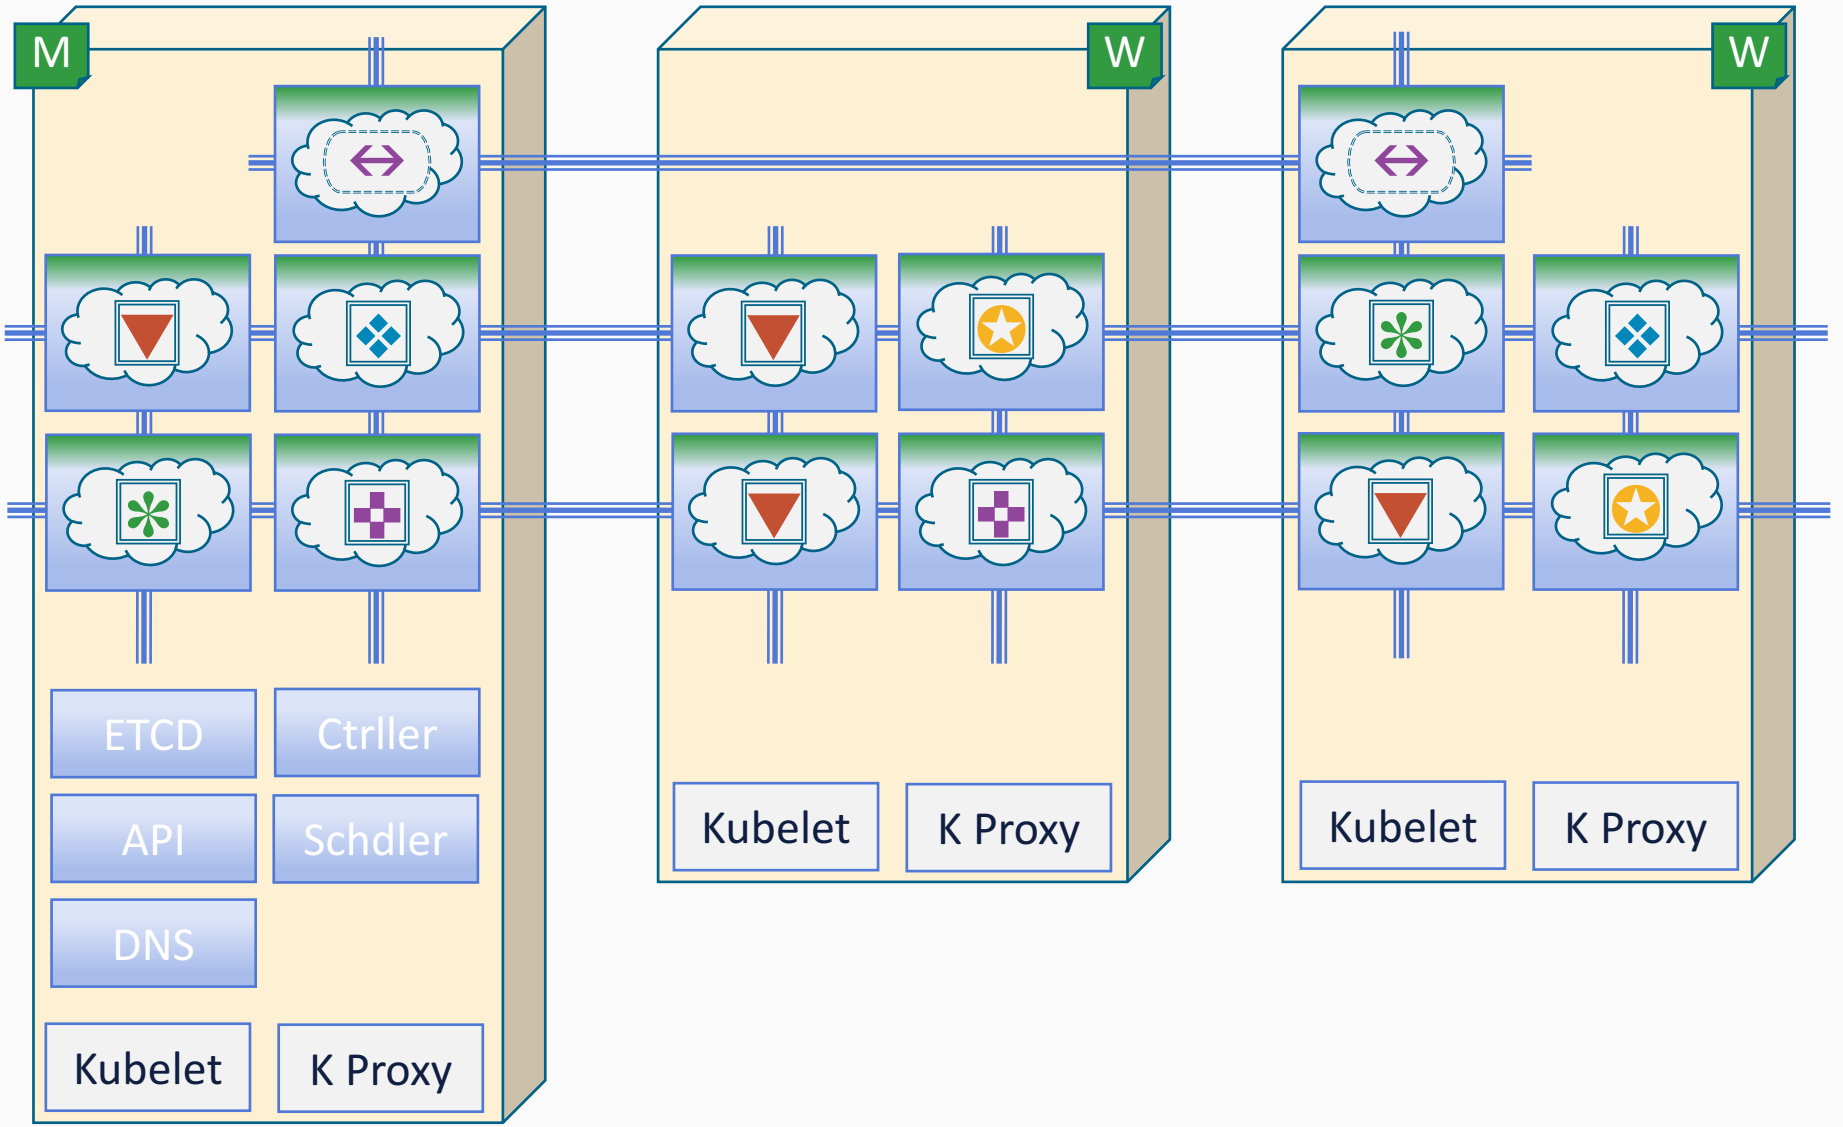

Orchestration with Single Master

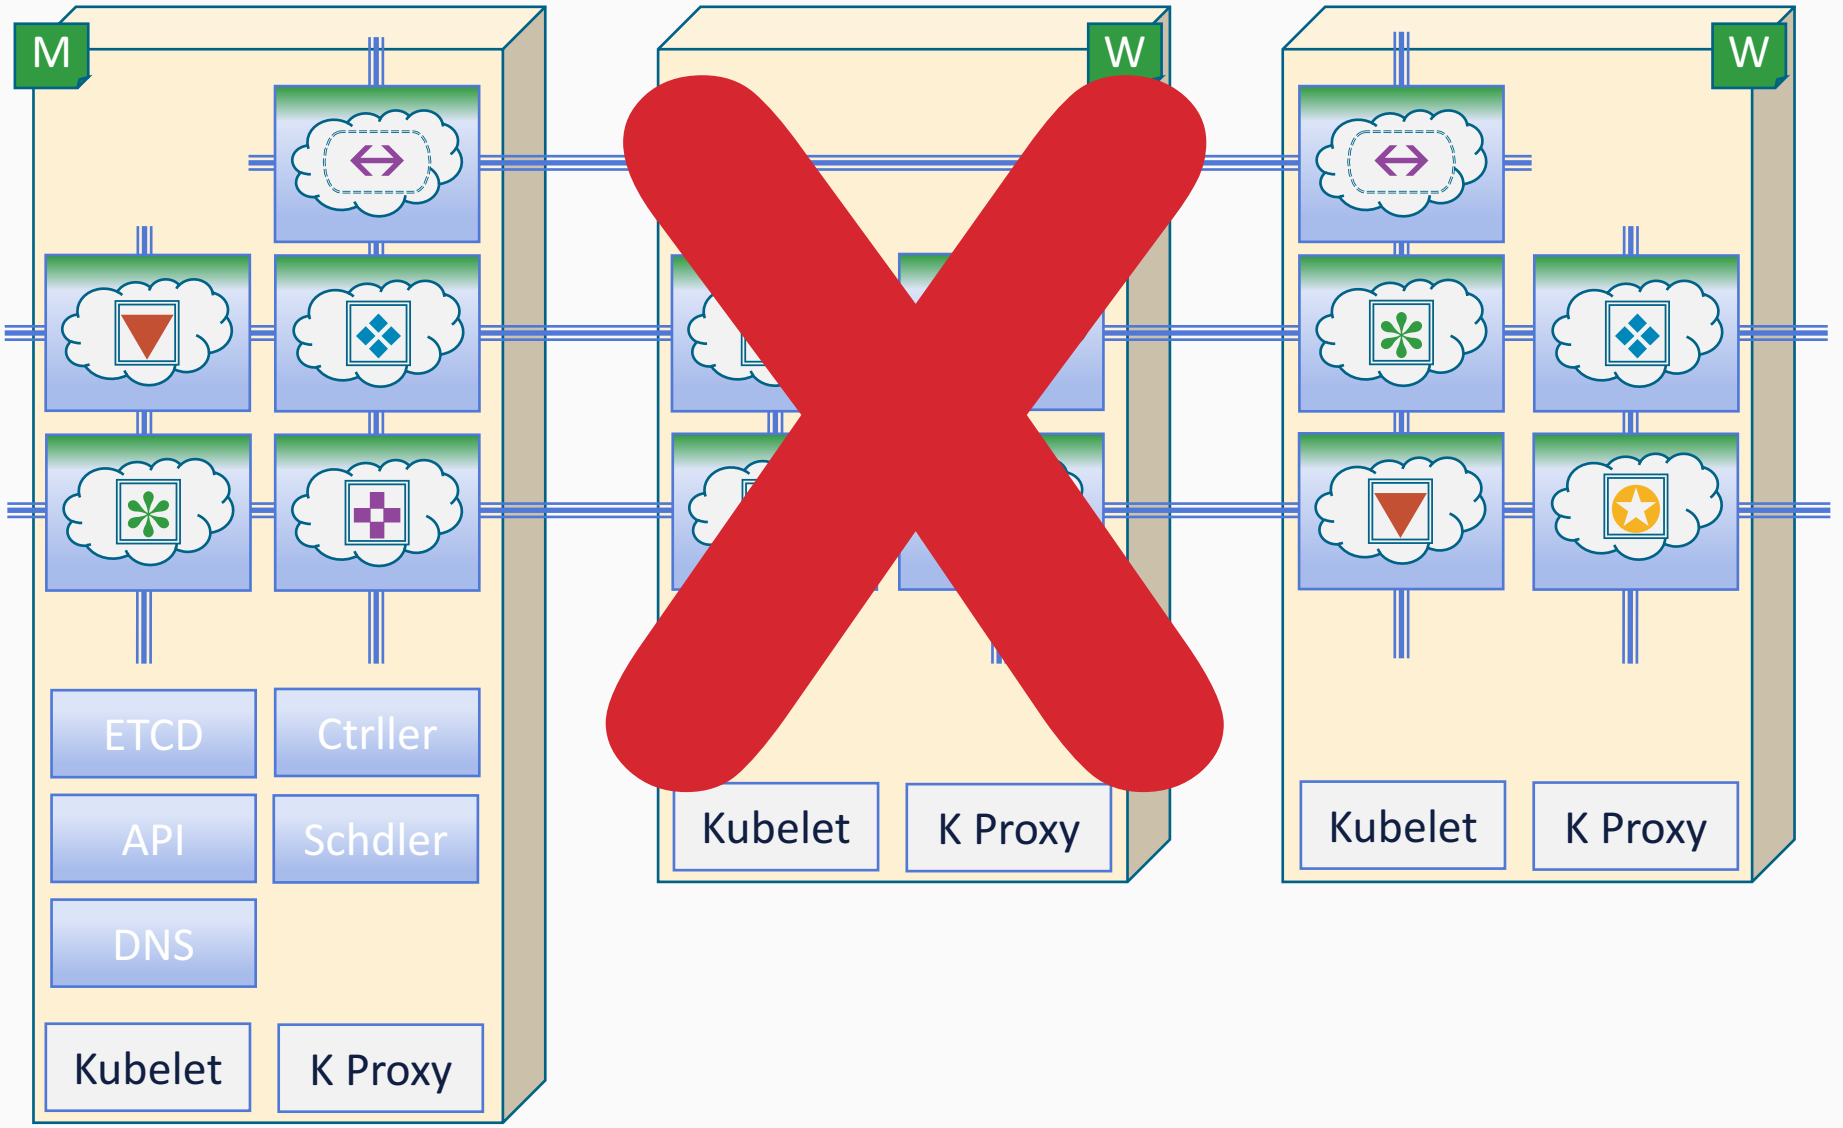

Grid Rebalancing

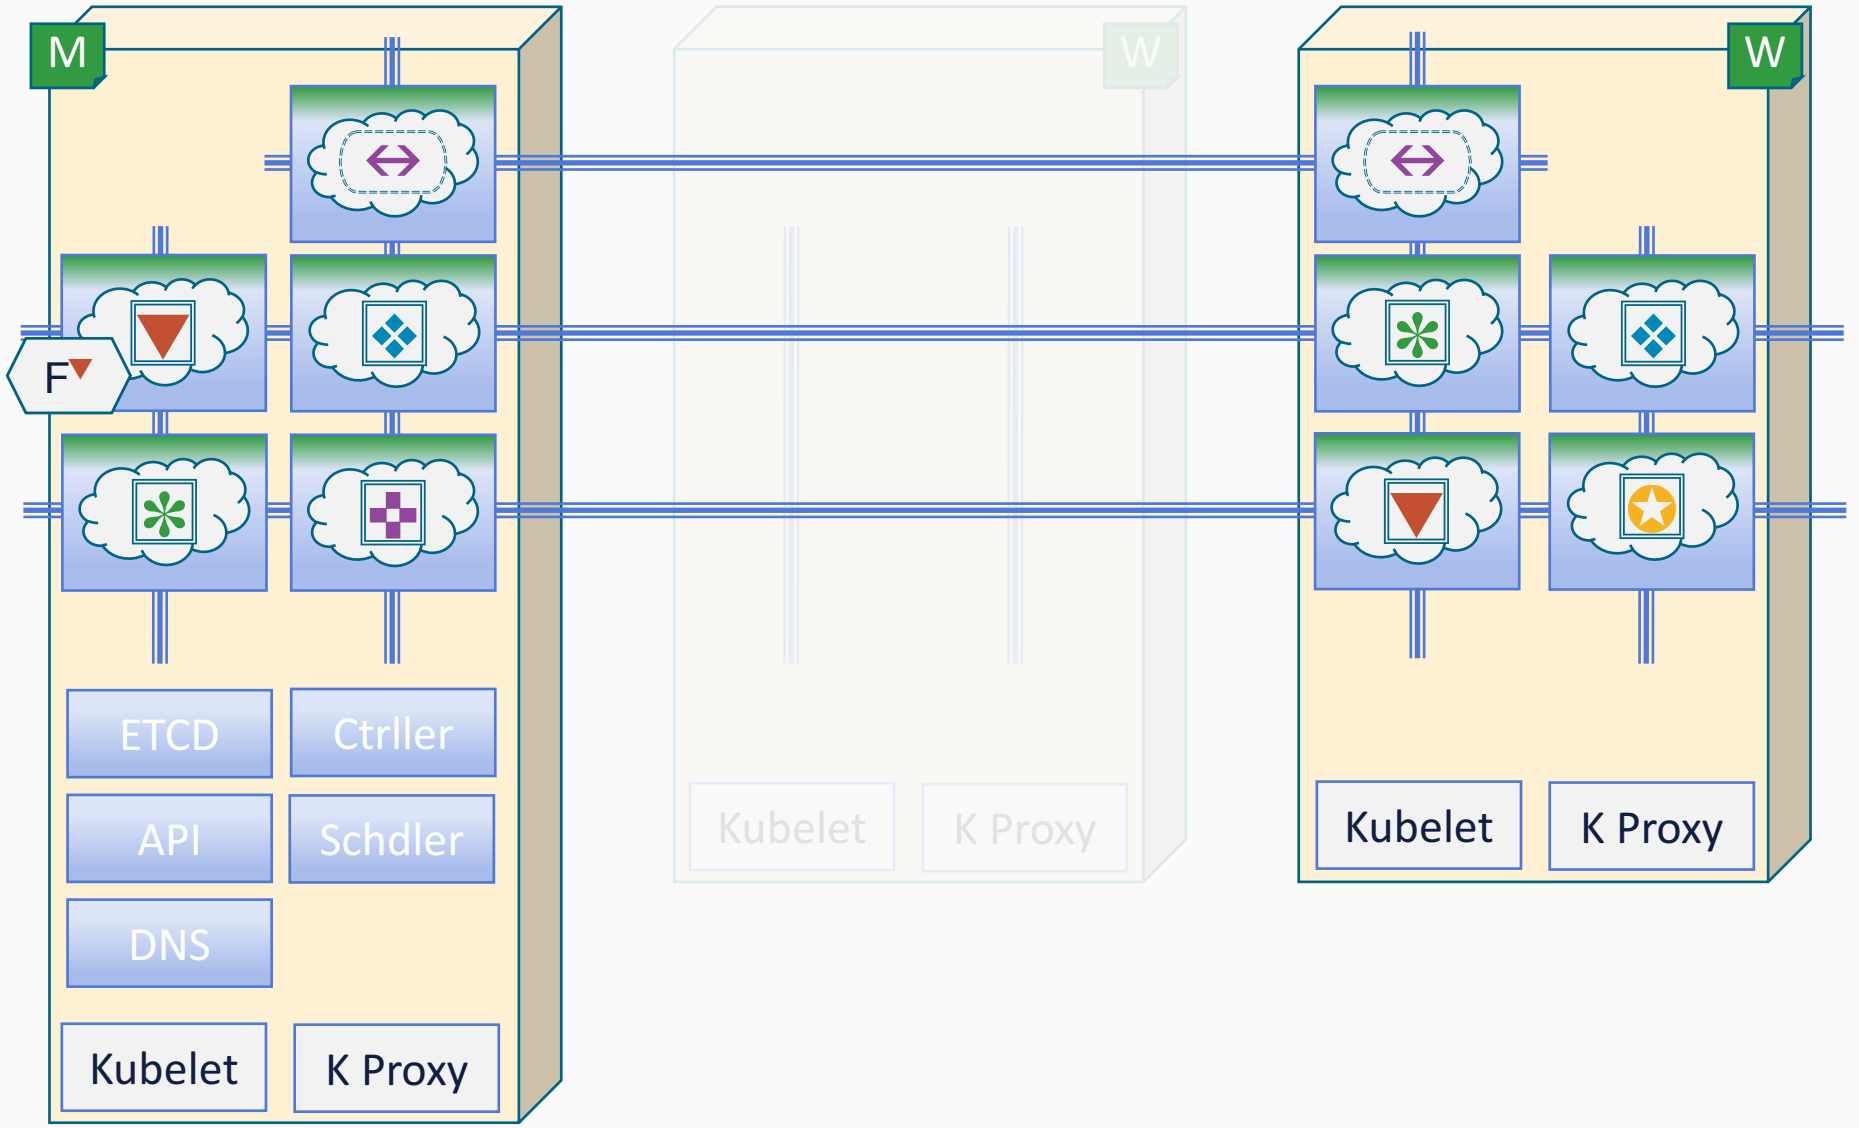

Dynamic Re-Configuration

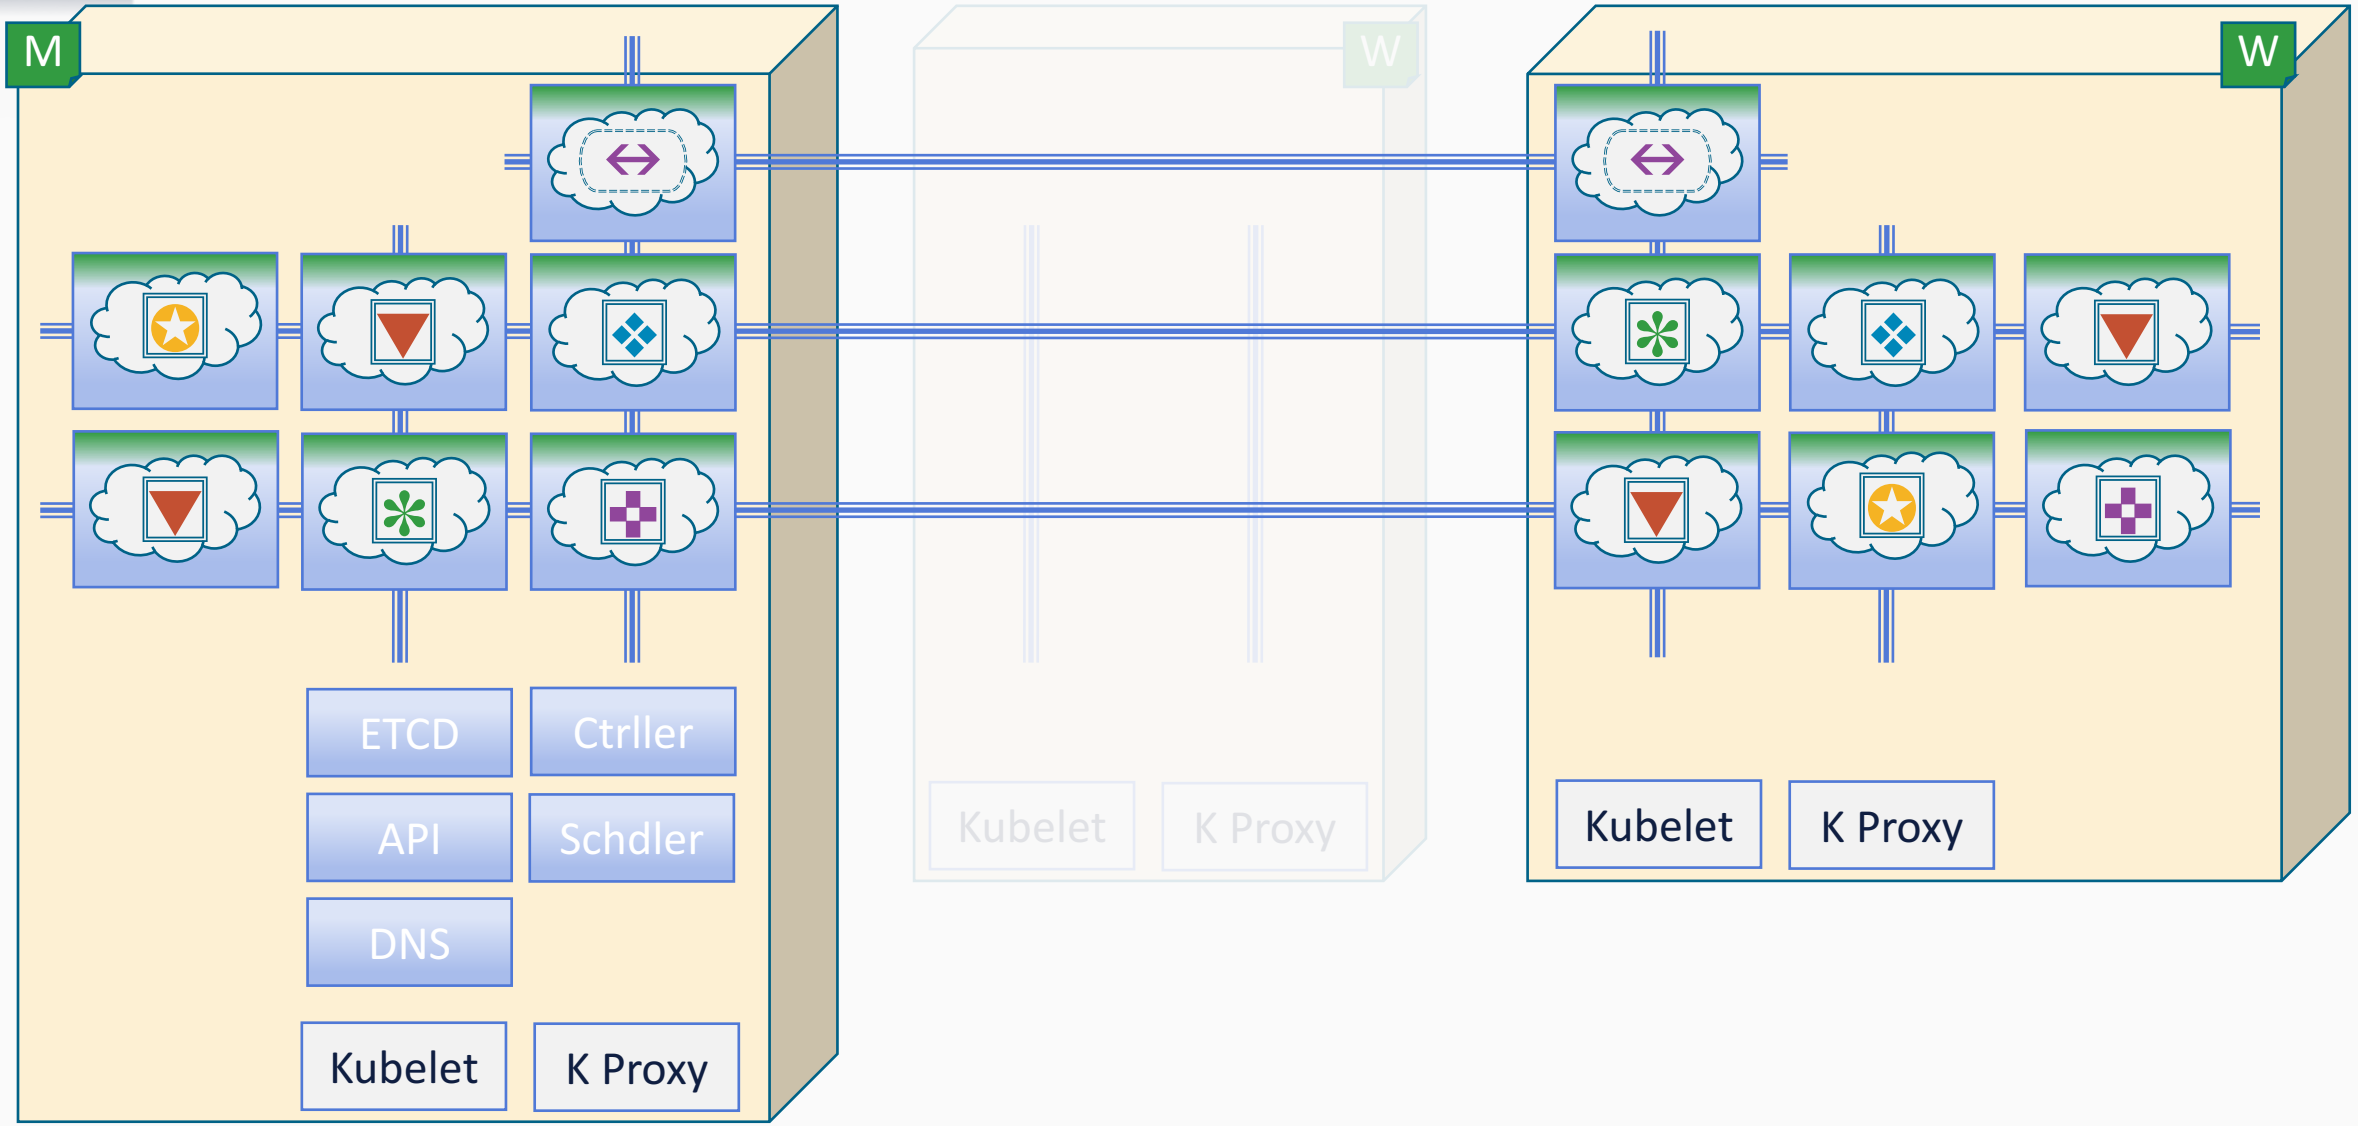

High Available Setup

The central foundation of a highly available solution is a redundant, reliable storage layer. The number one rule of high-availability is to protect the data.

Whatever else happens, whatever catches on fire, if you have the data, you can rebuild. If you lose the data, you're done.

High Available Setup

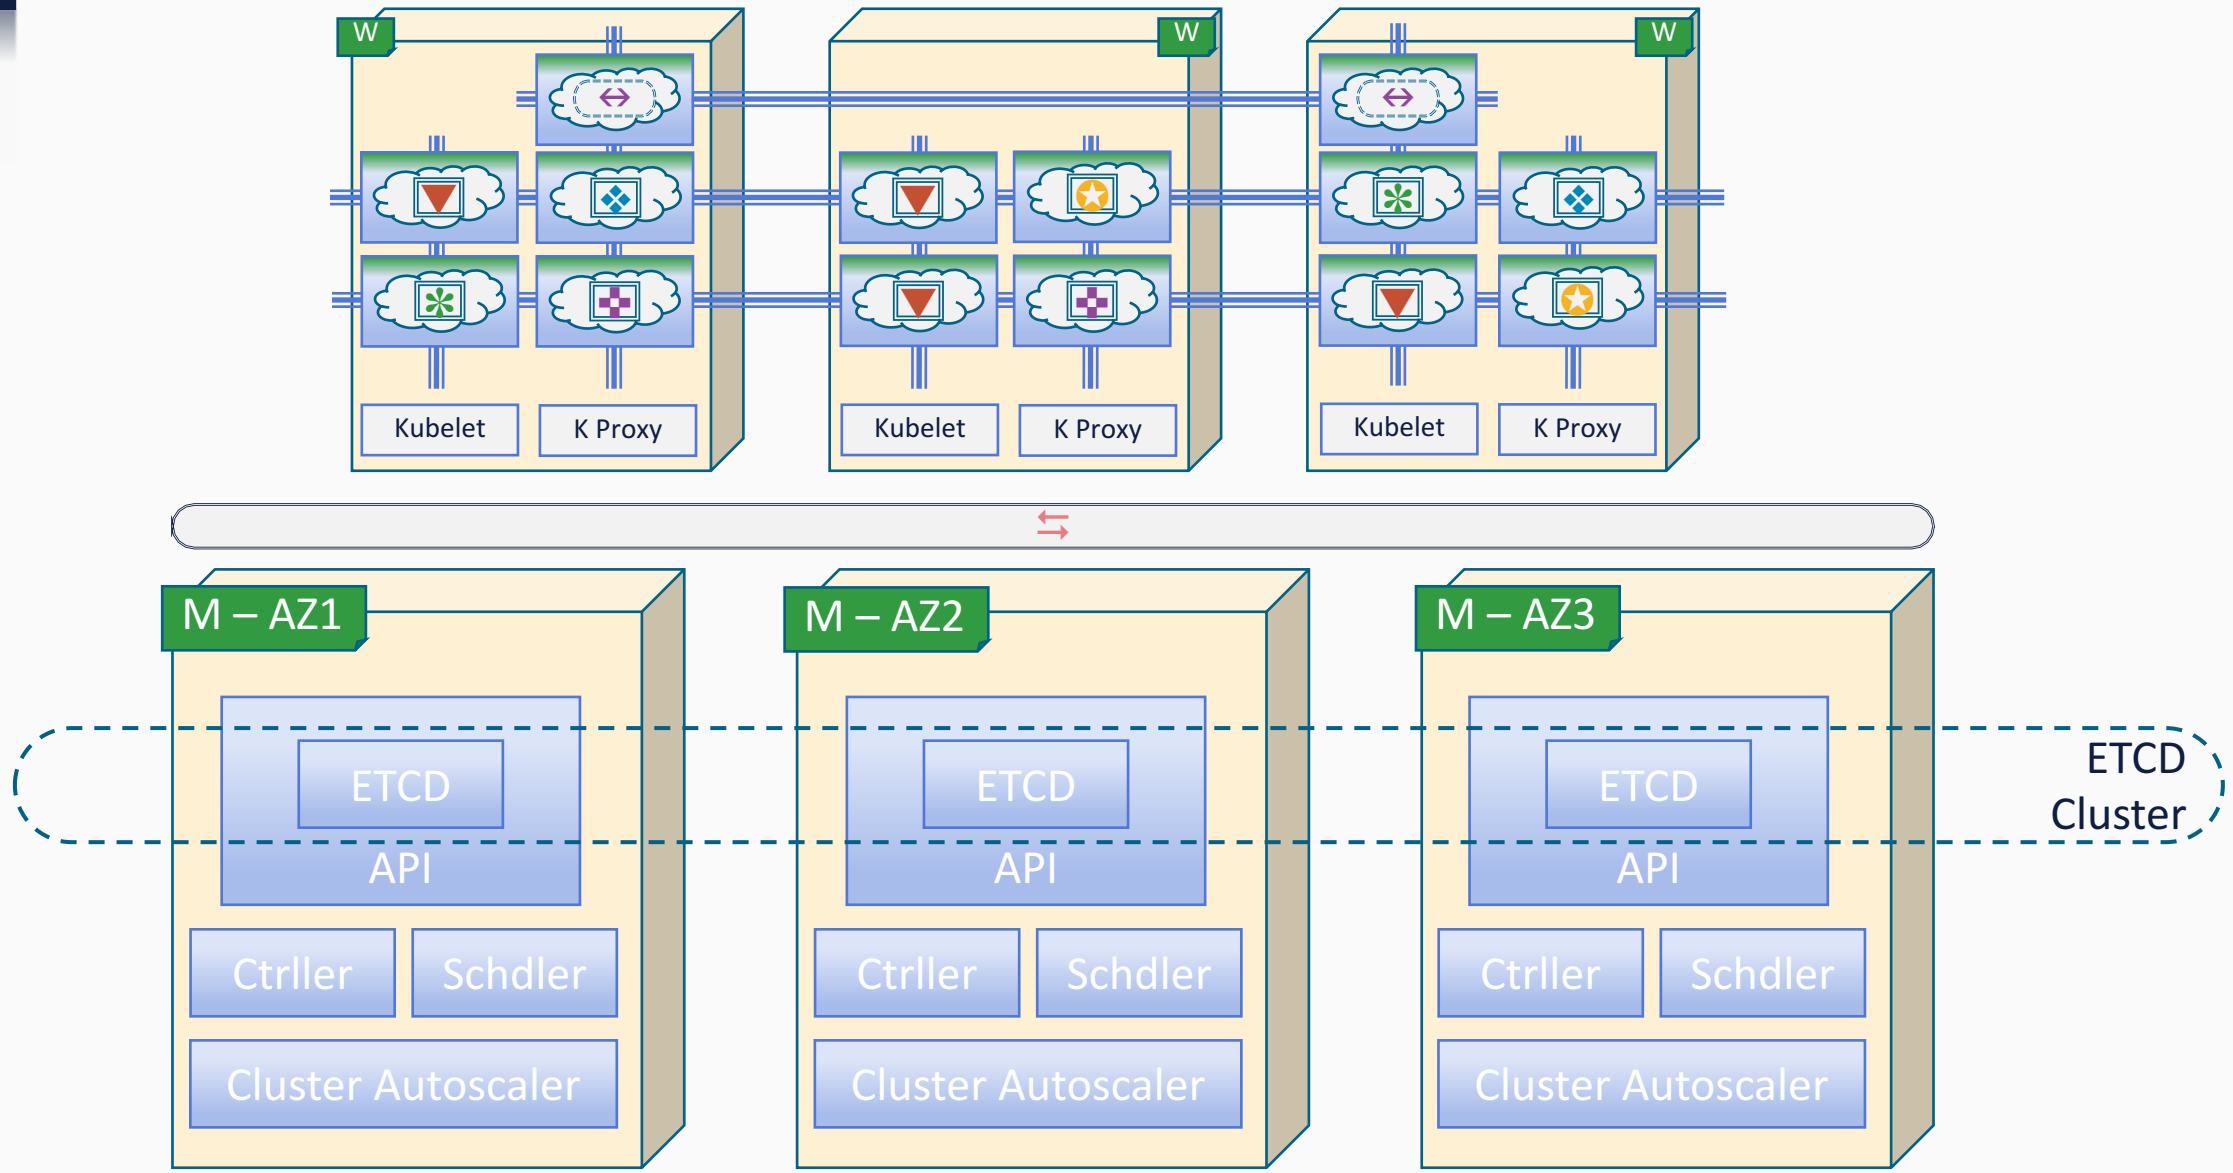

Exactly Once Processing

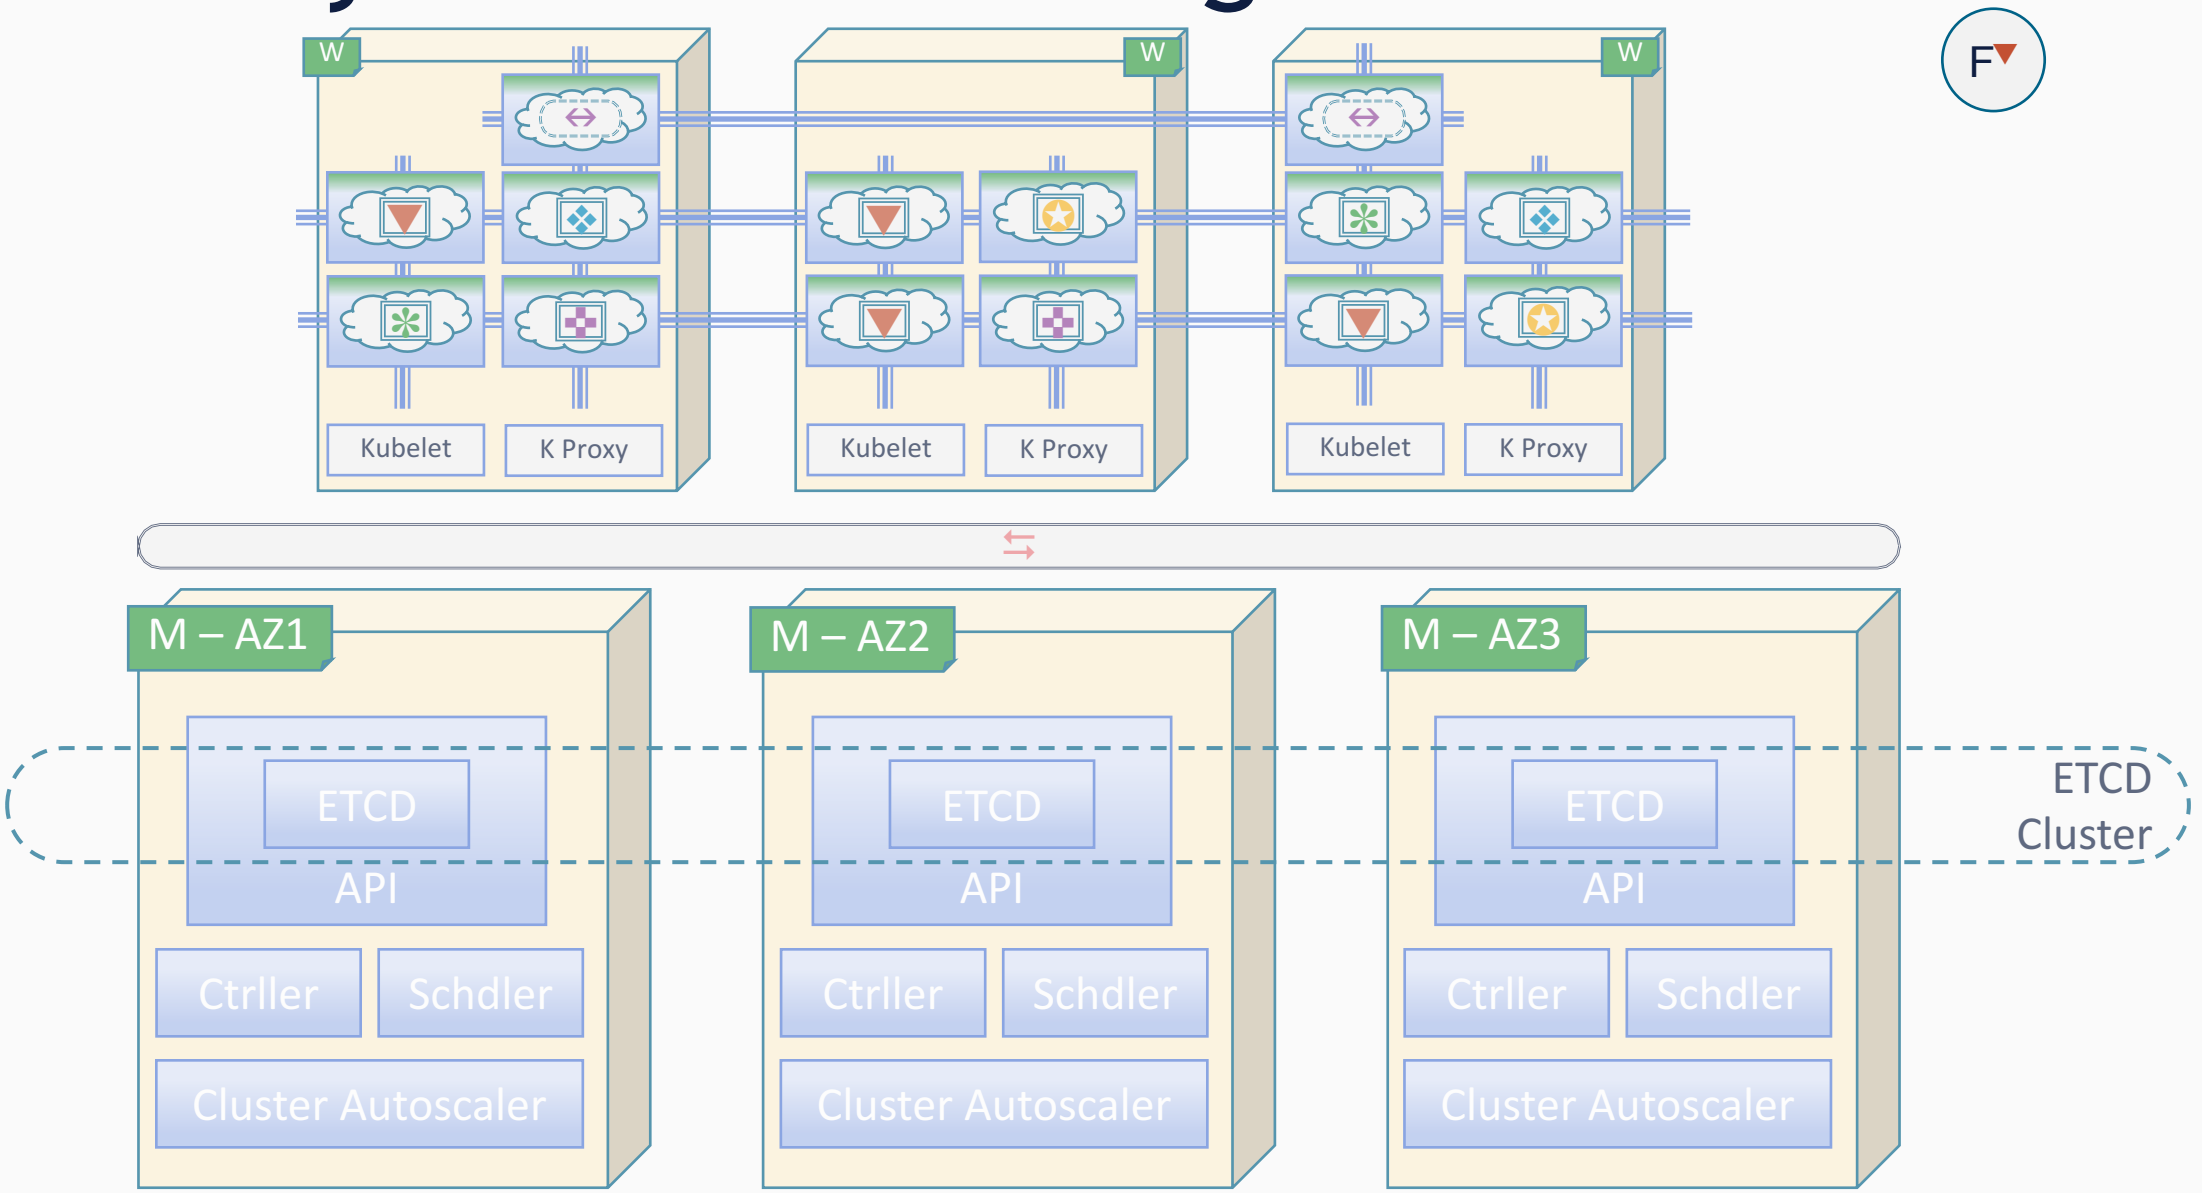

Exactly Once Processing

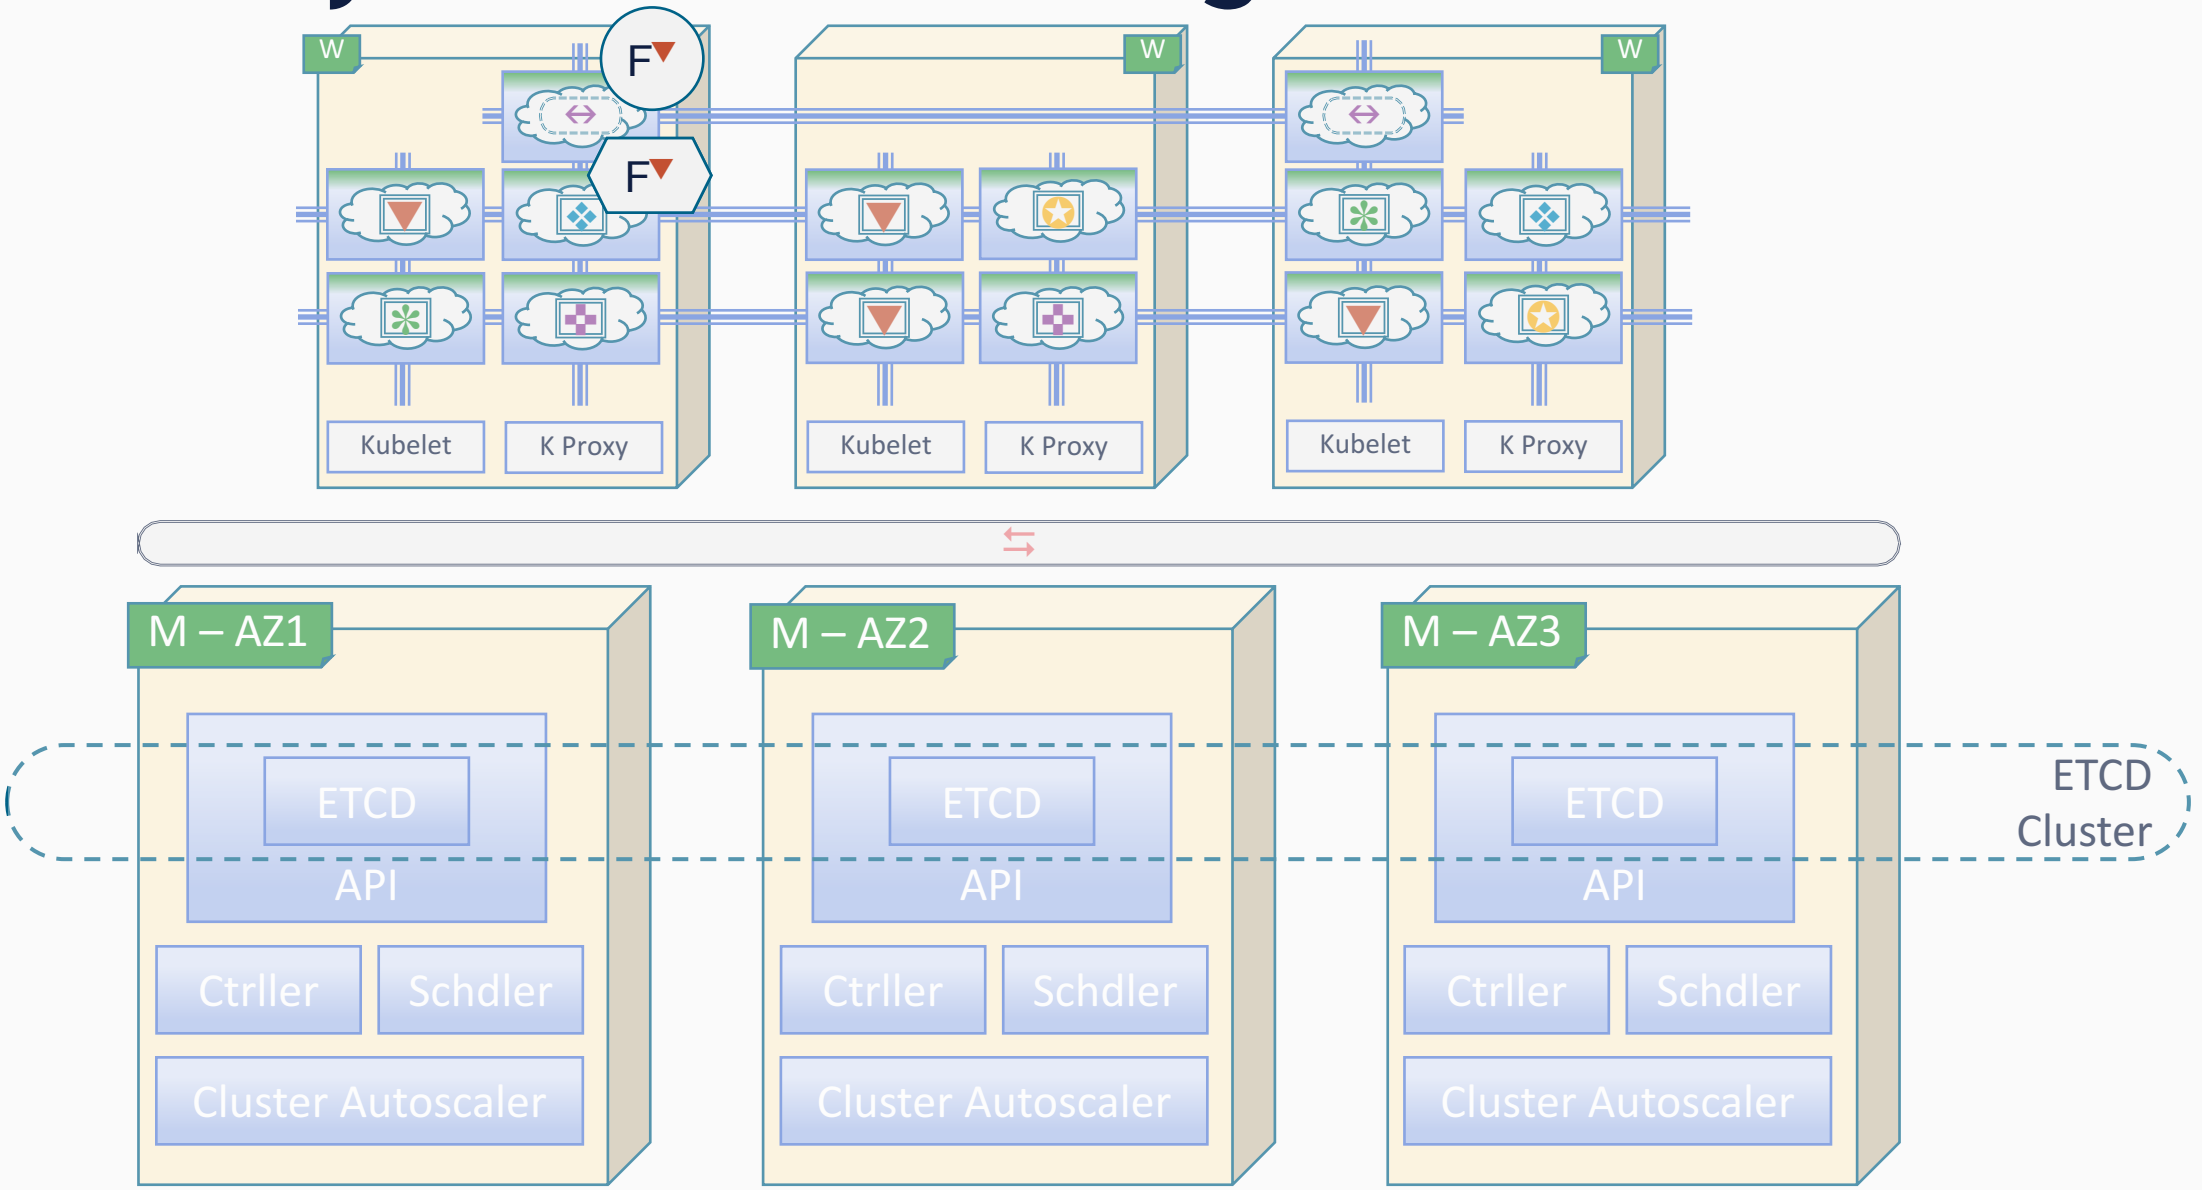

Exactly Once Processing

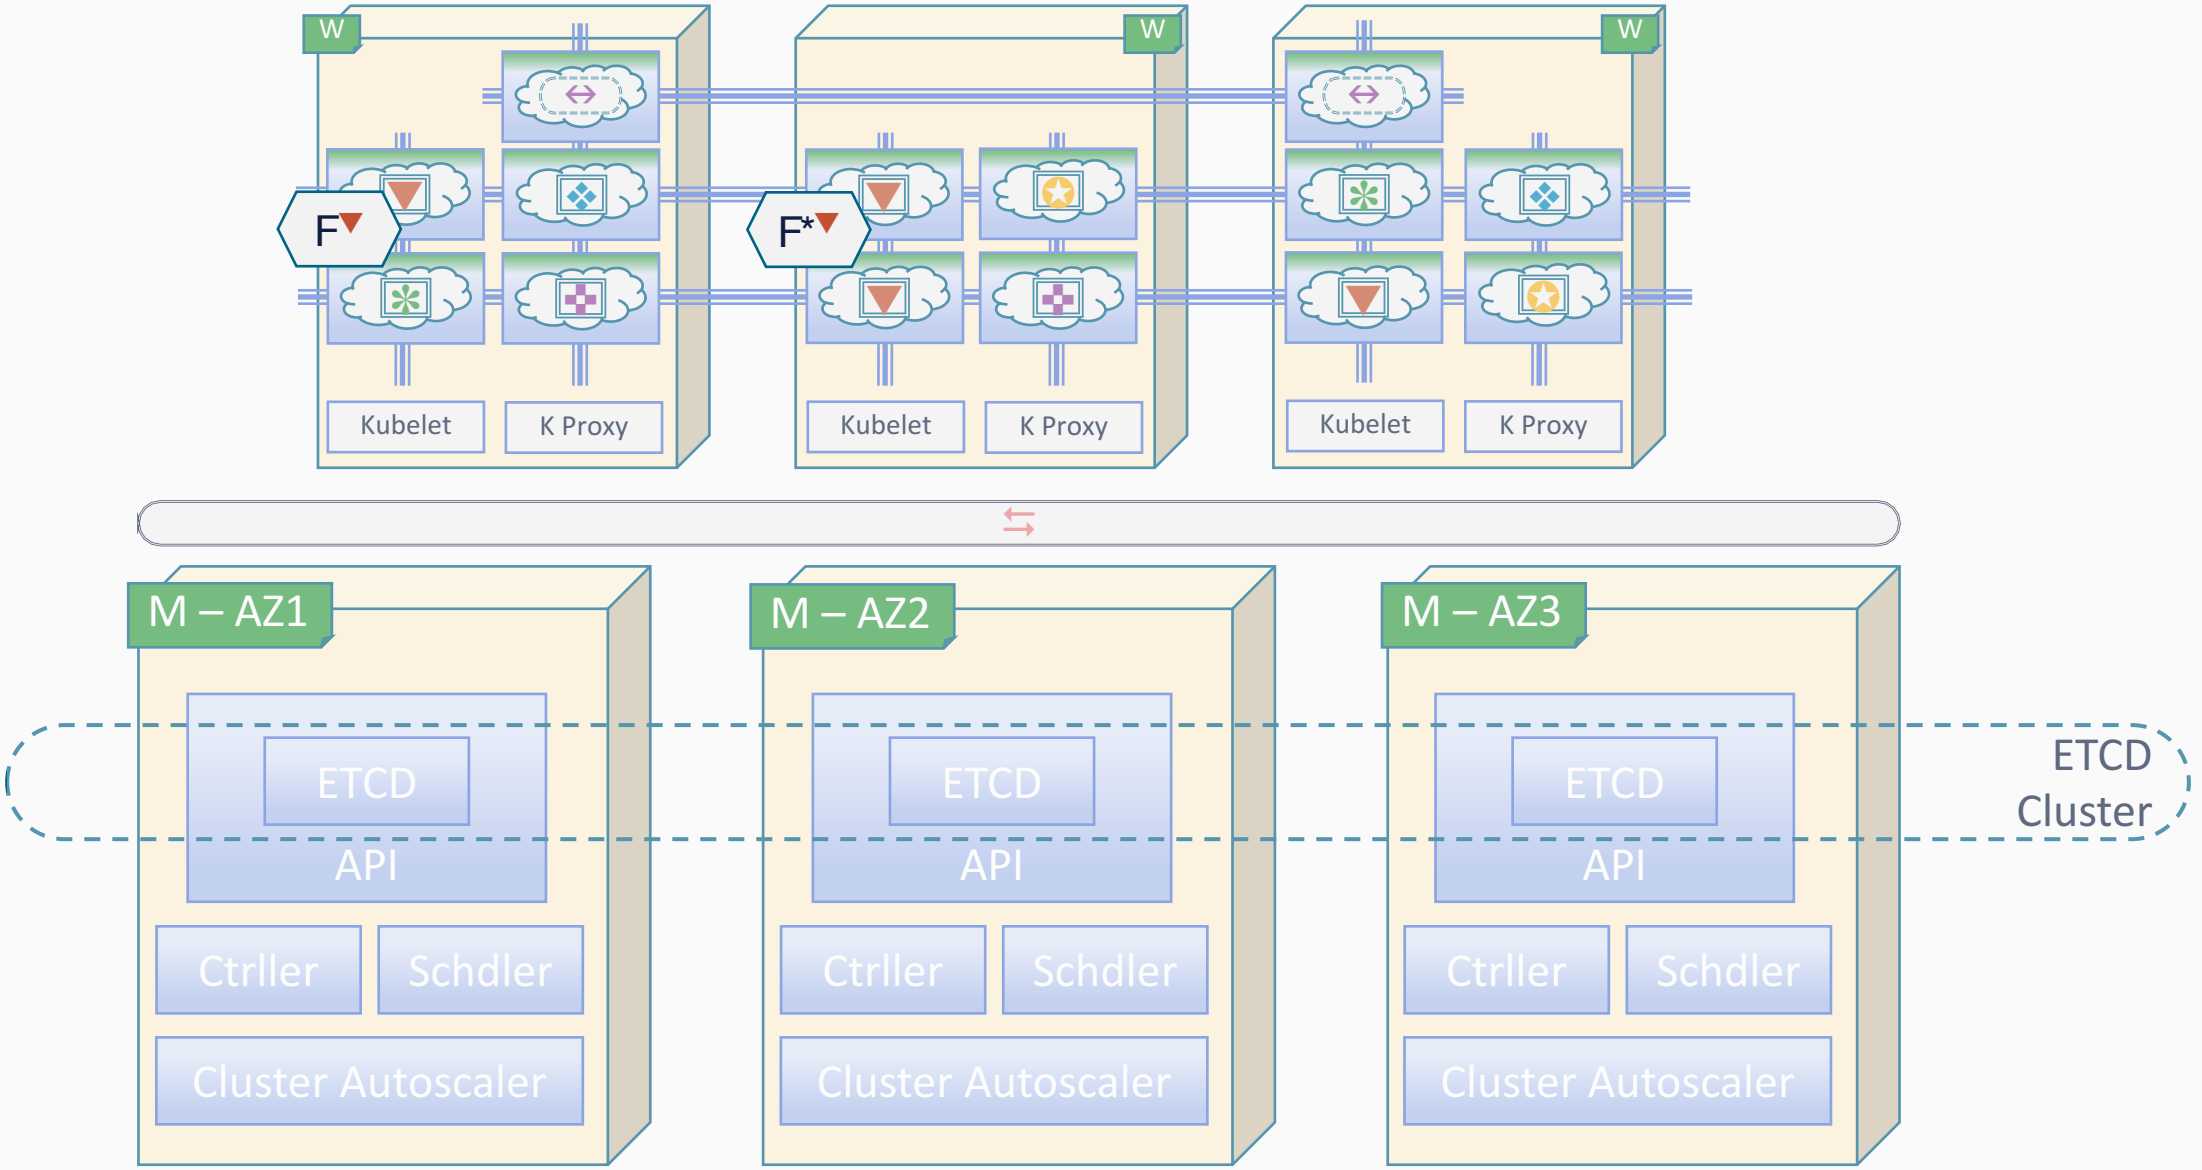

Exactly Once Processing

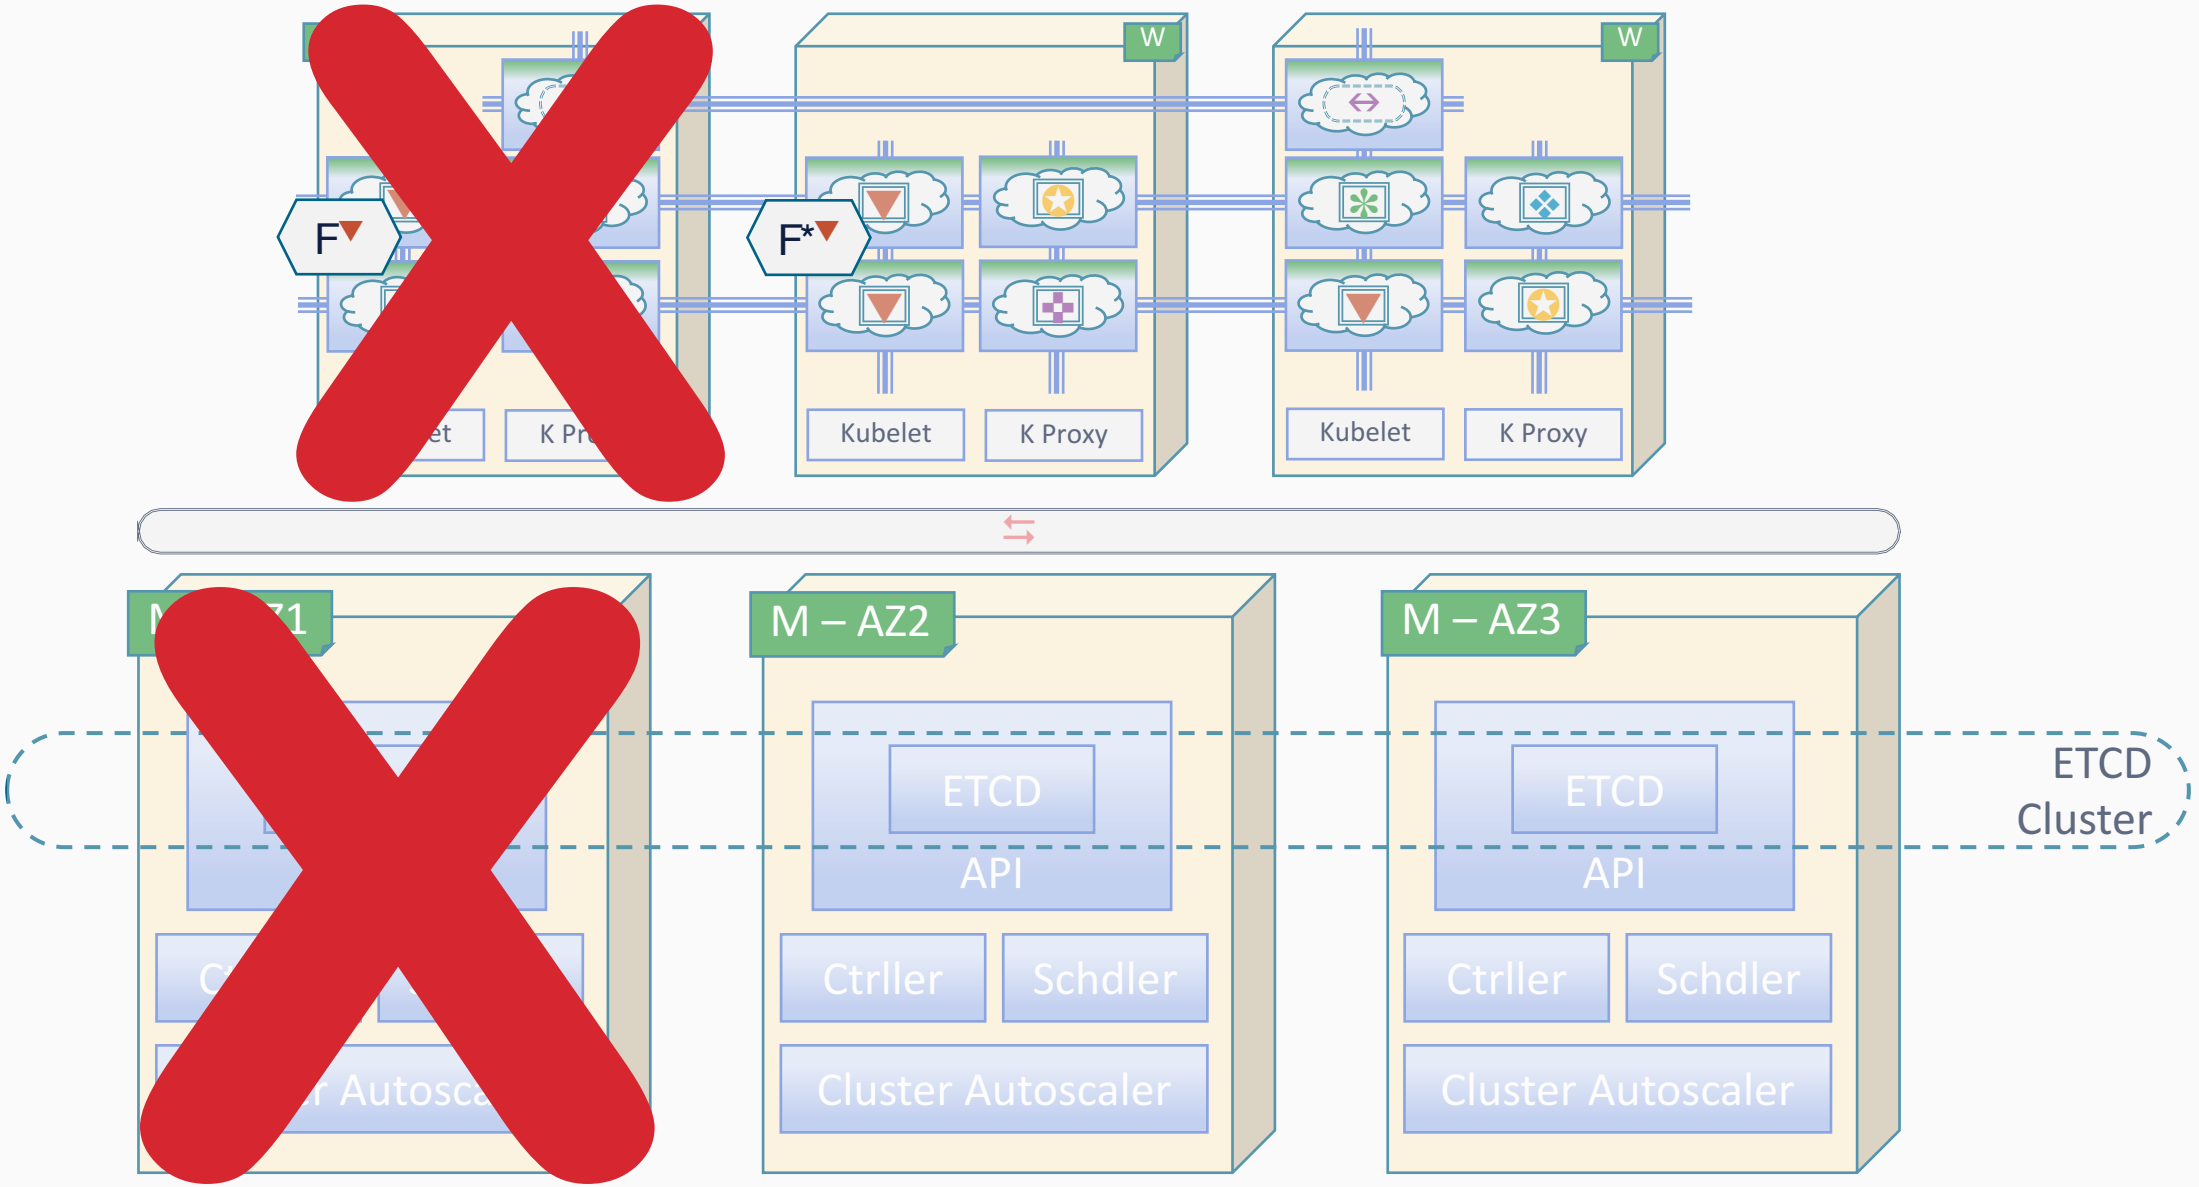

Exactly Once Processing

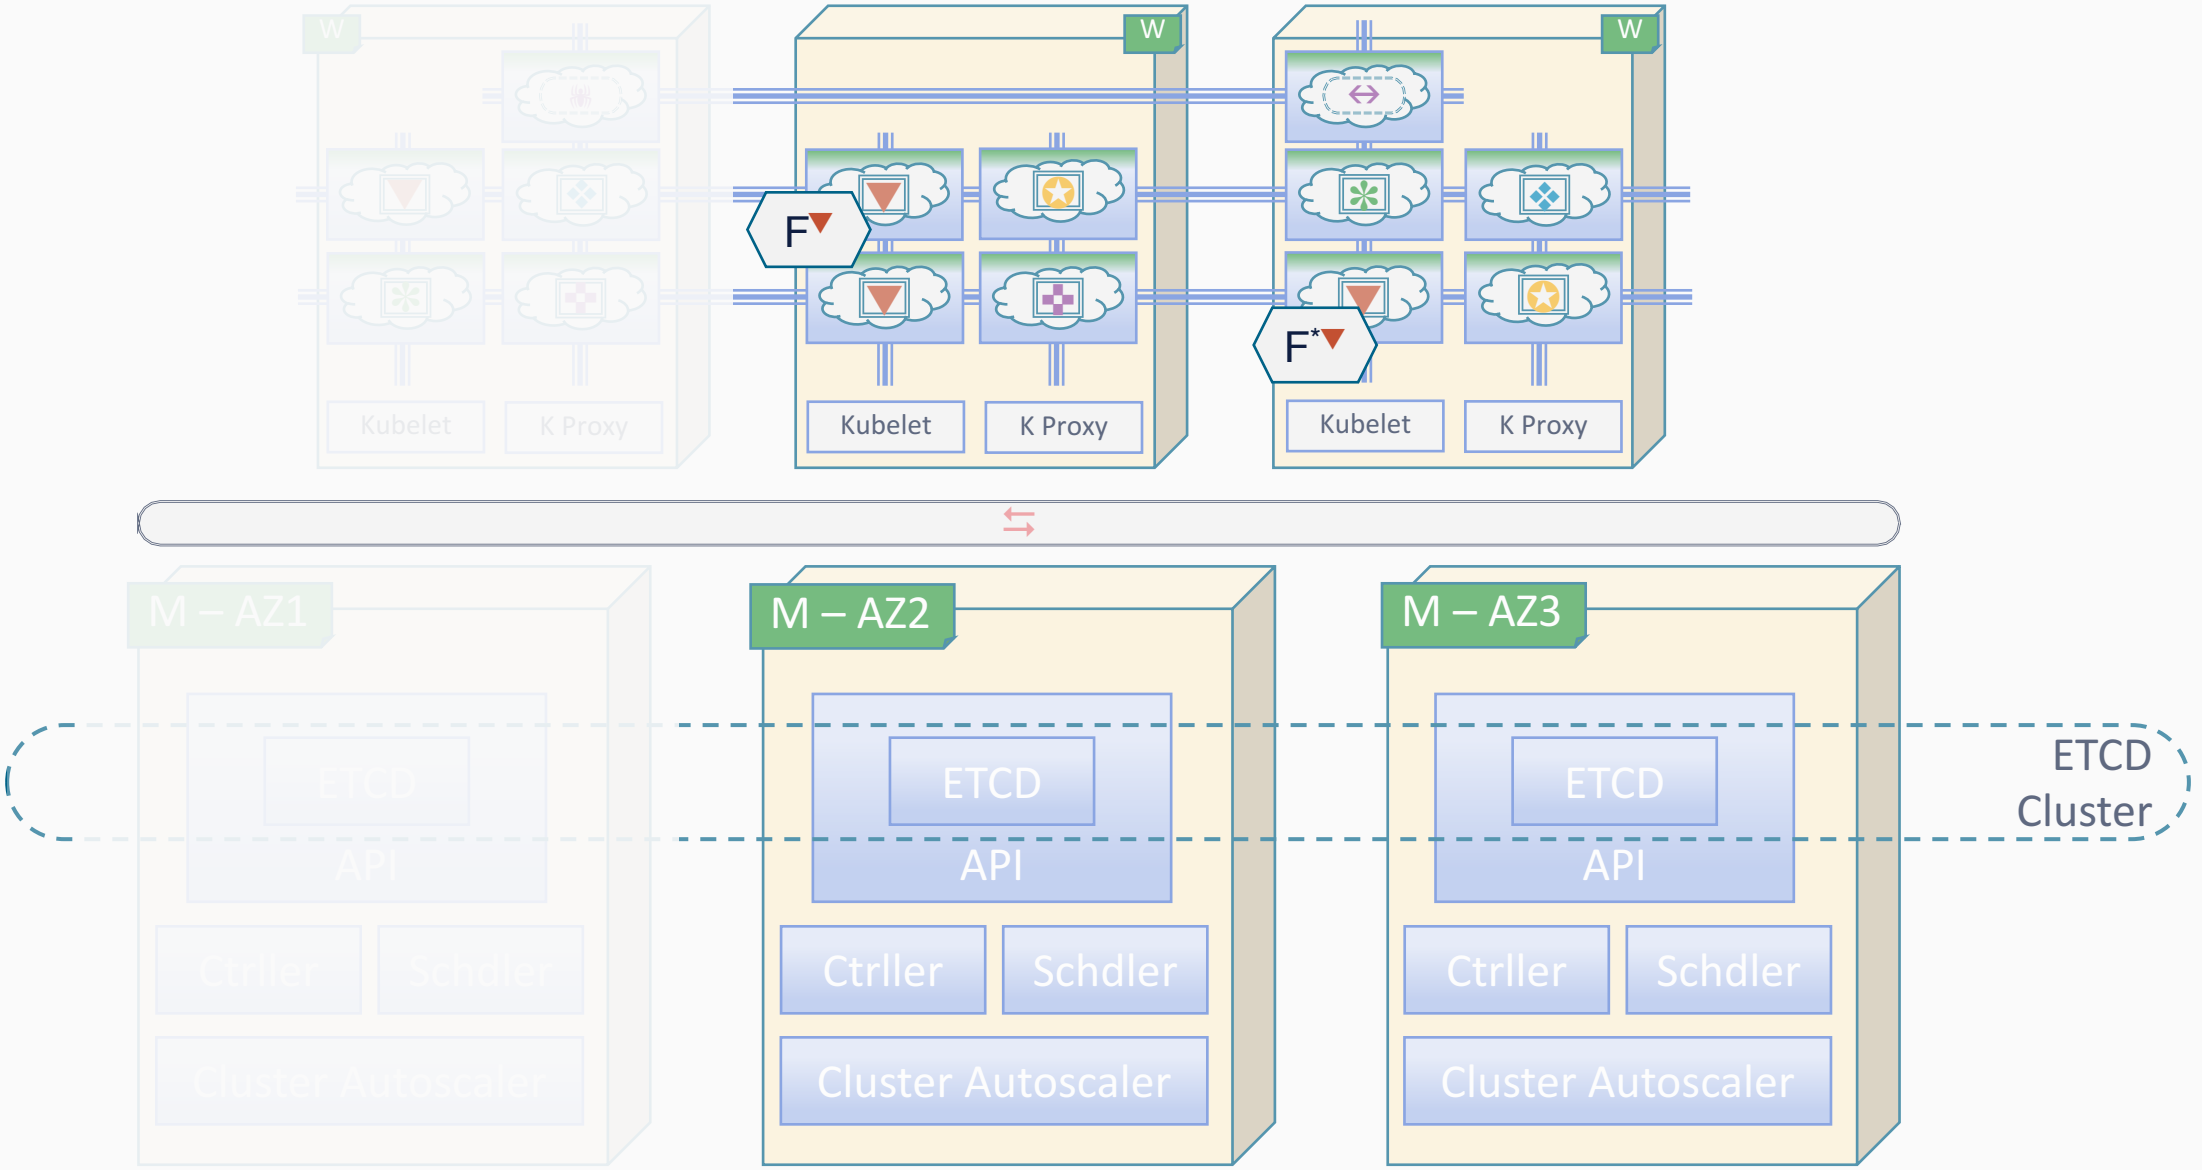

Exactly Once Processing

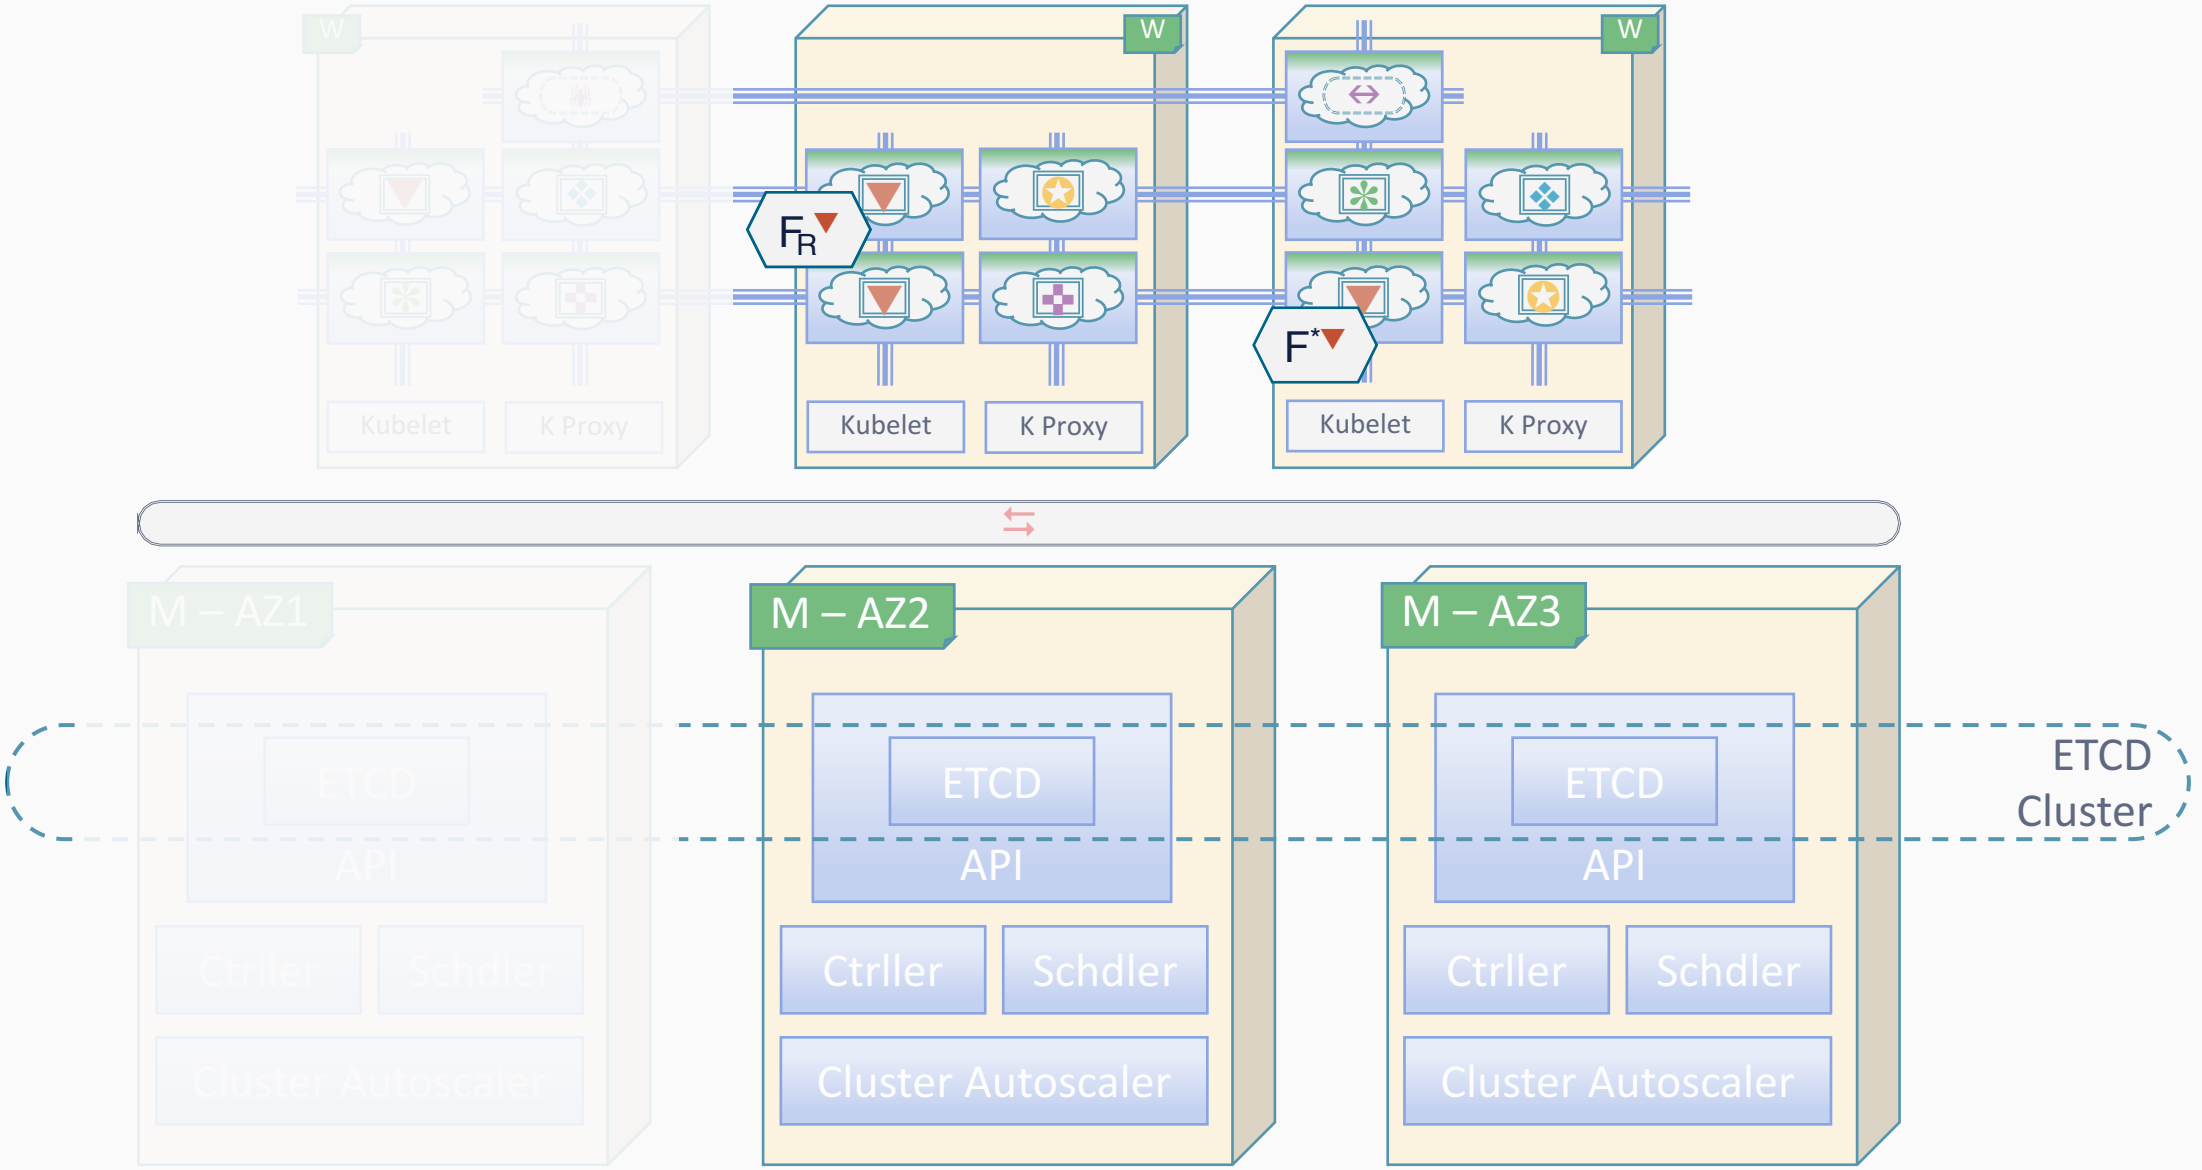

Exactly Once Processing

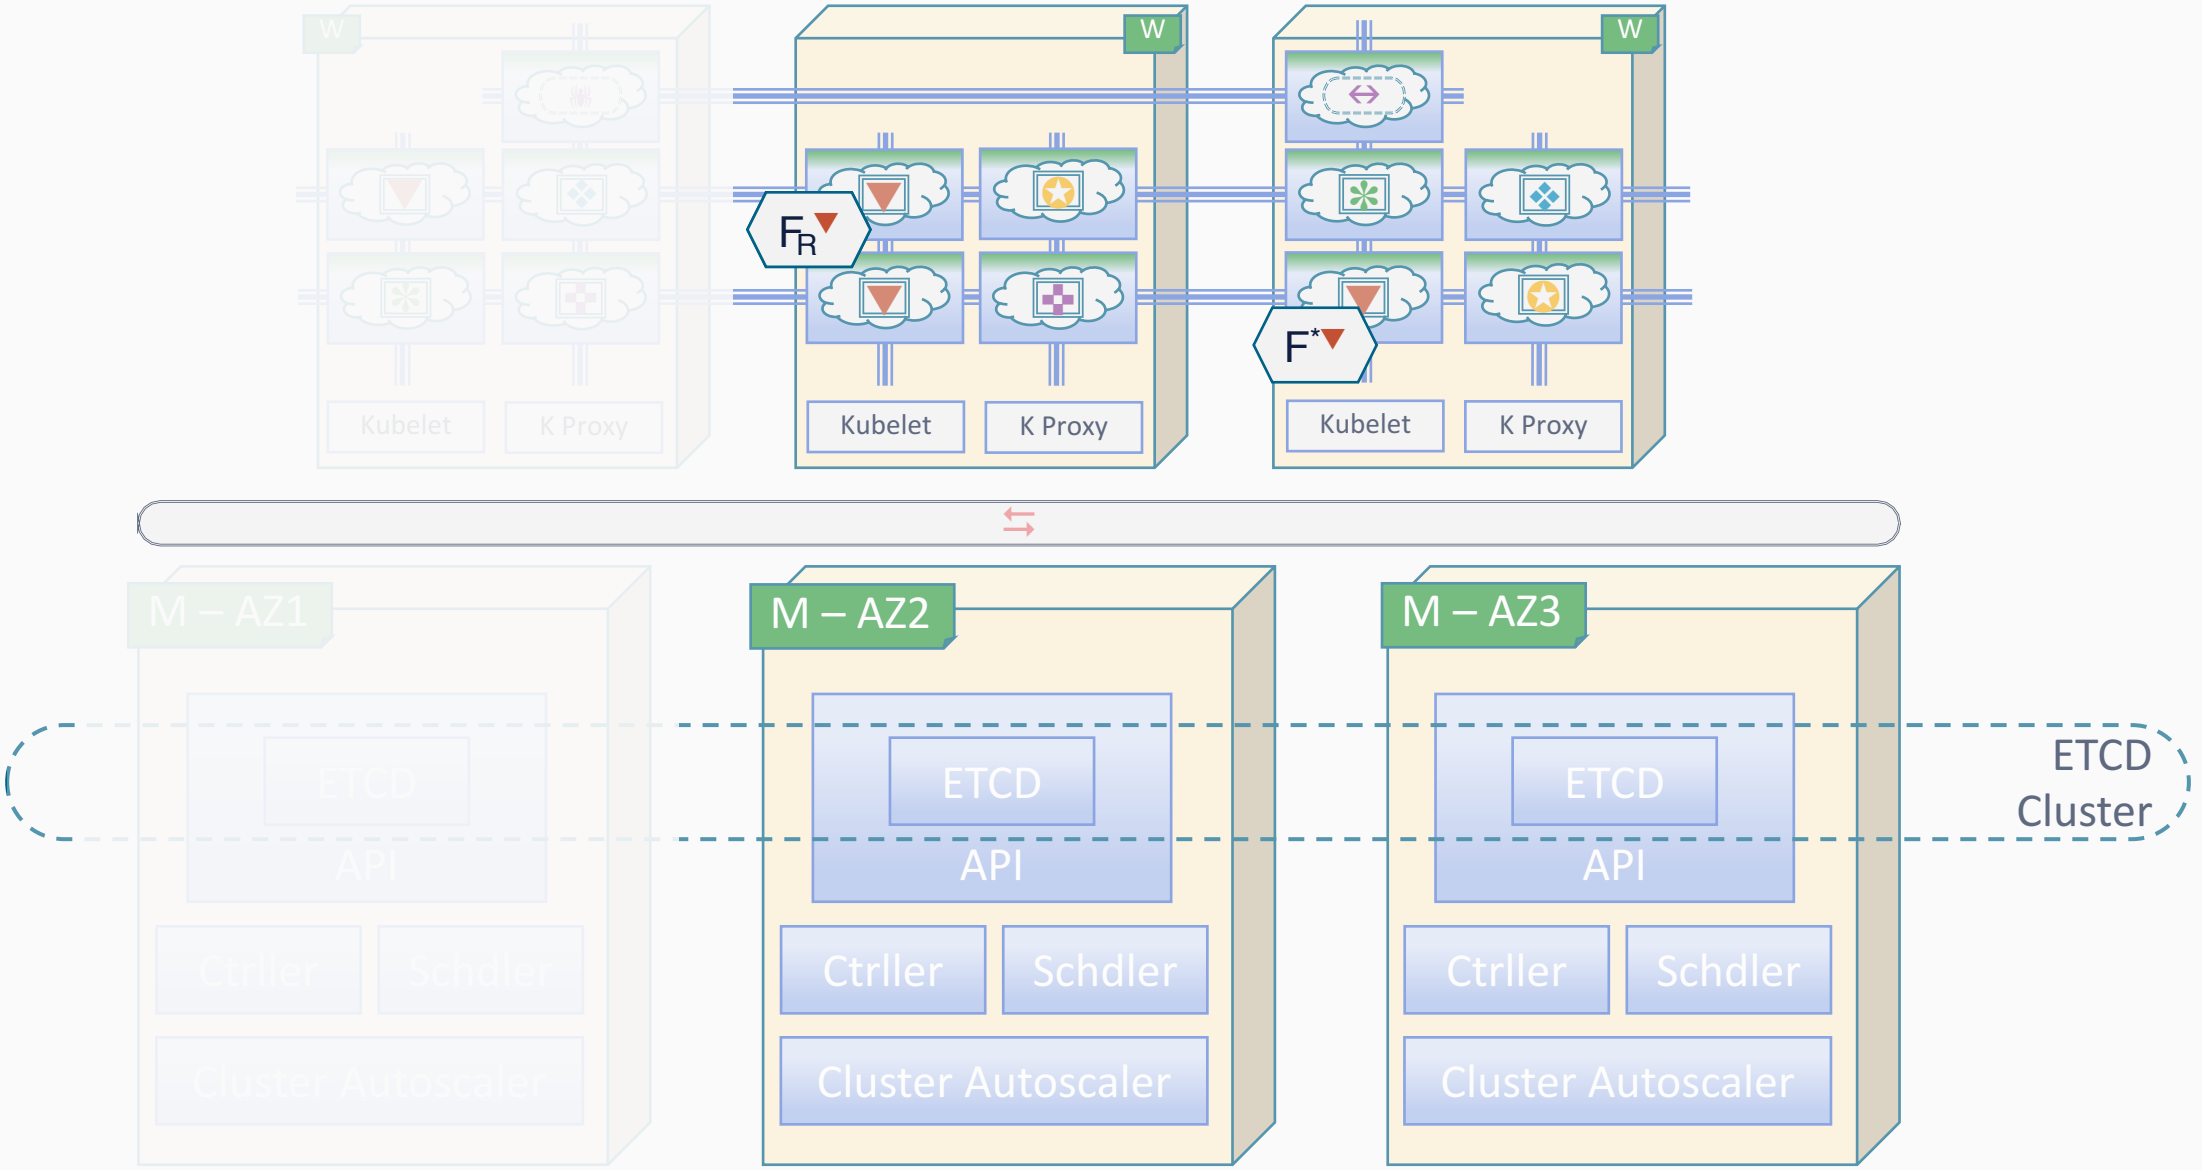

Exactly Once Processing

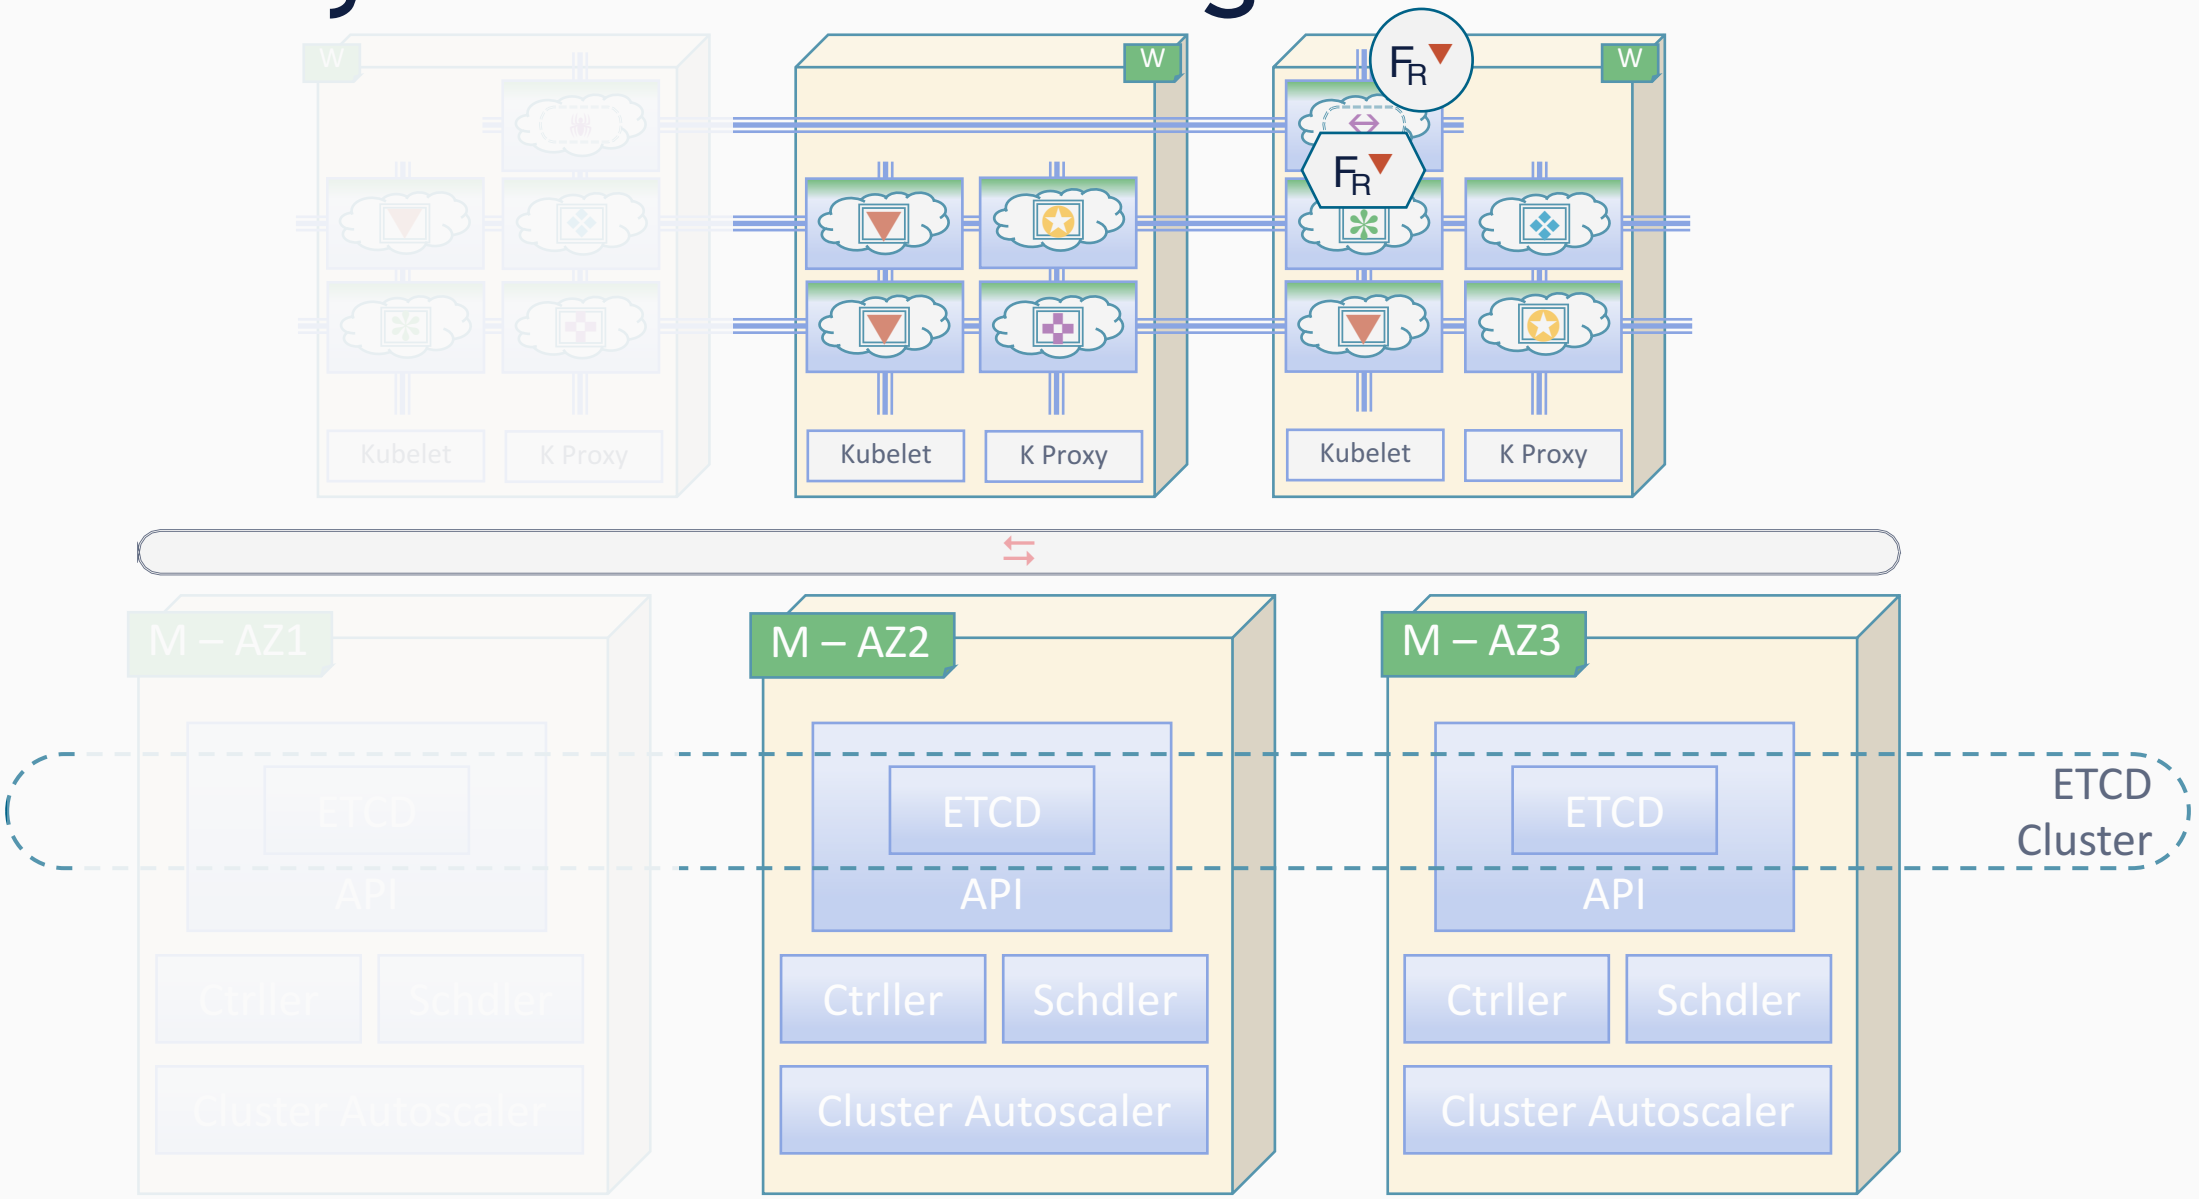

Dynamic Re-Configuration

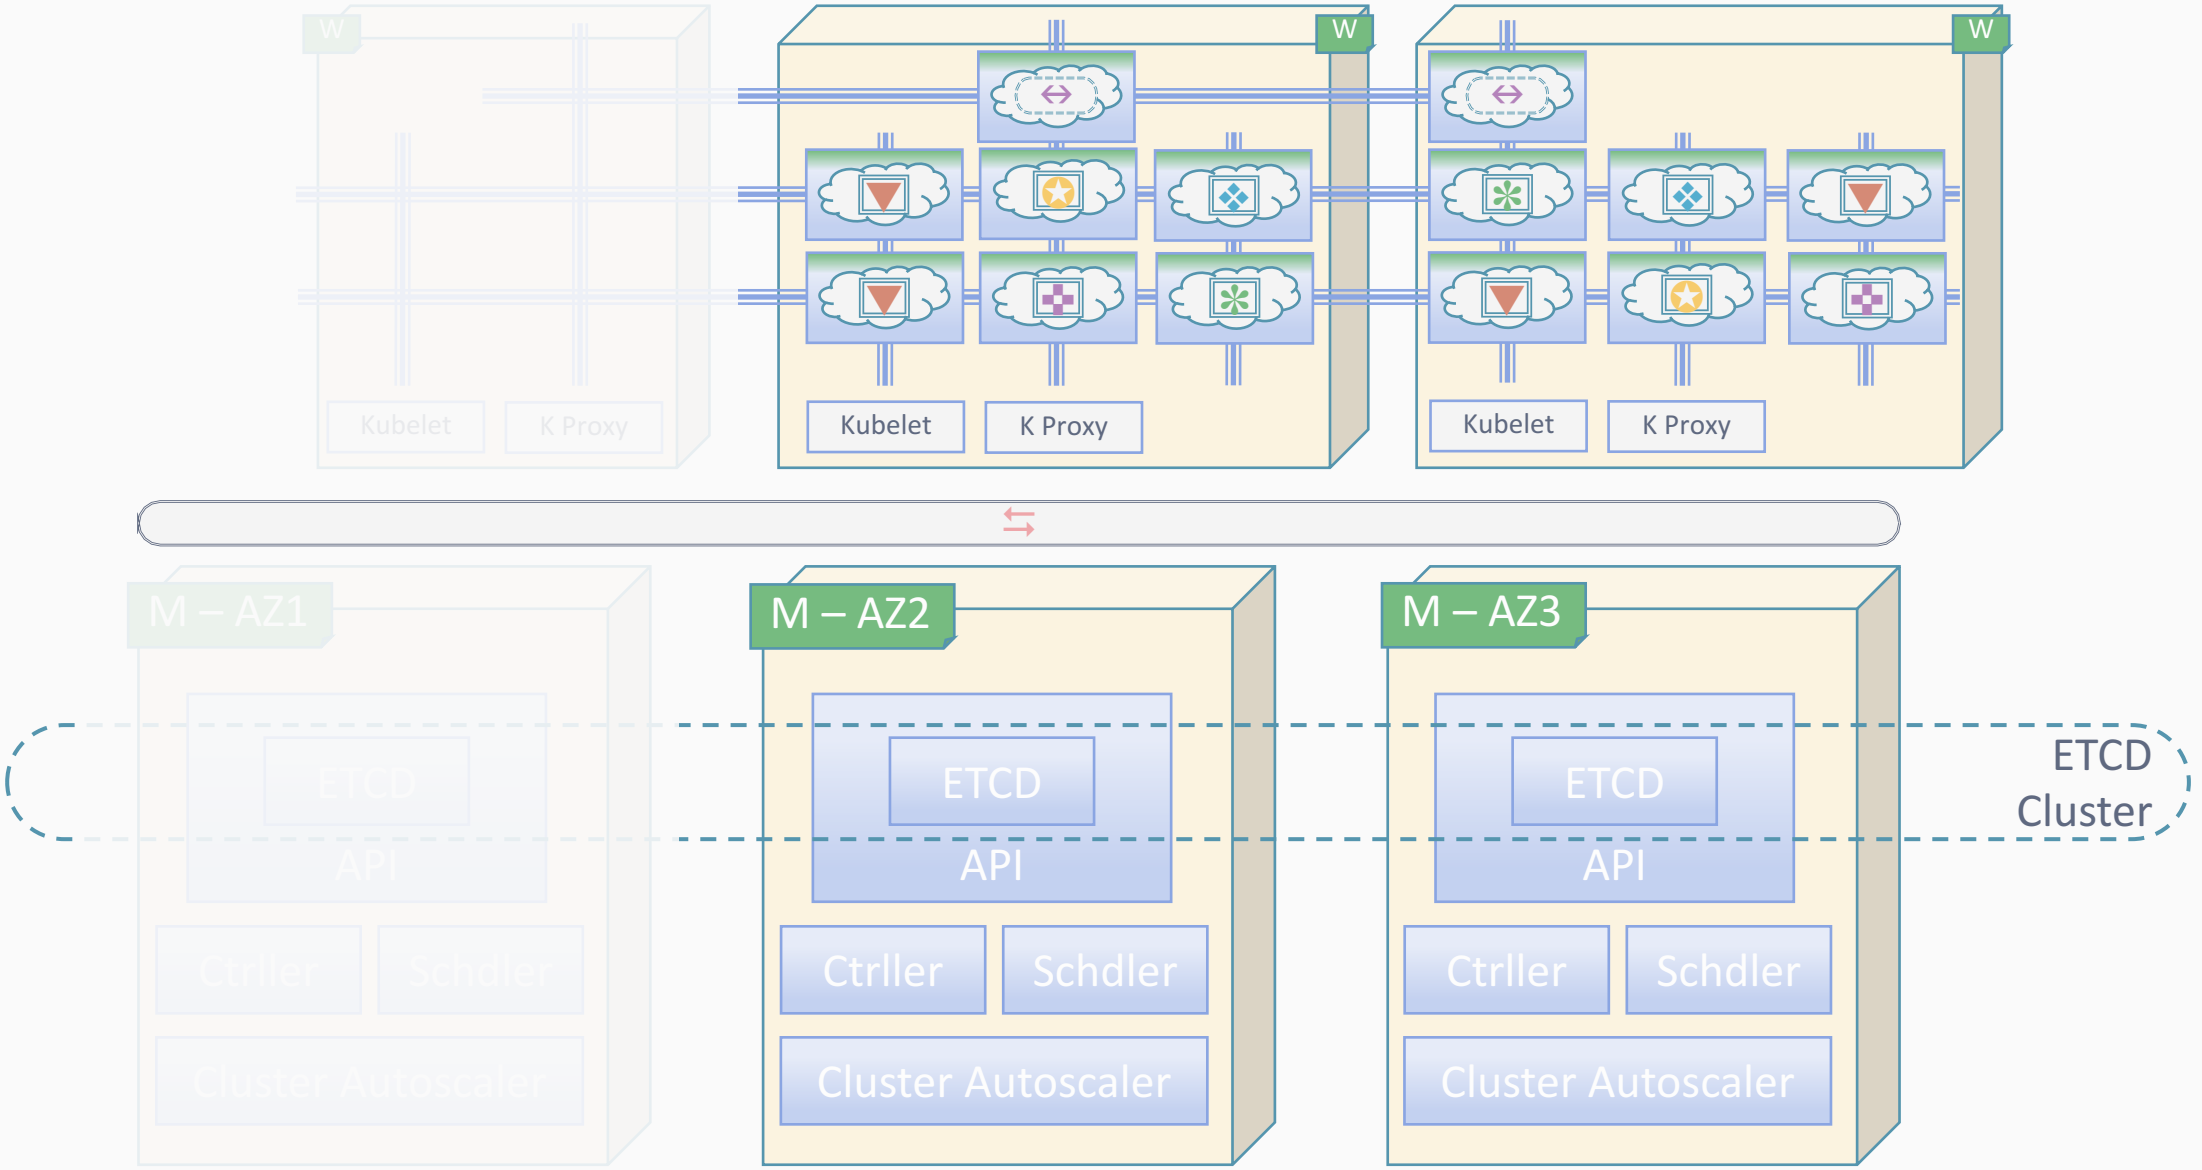

Multi-Datacenter Setup

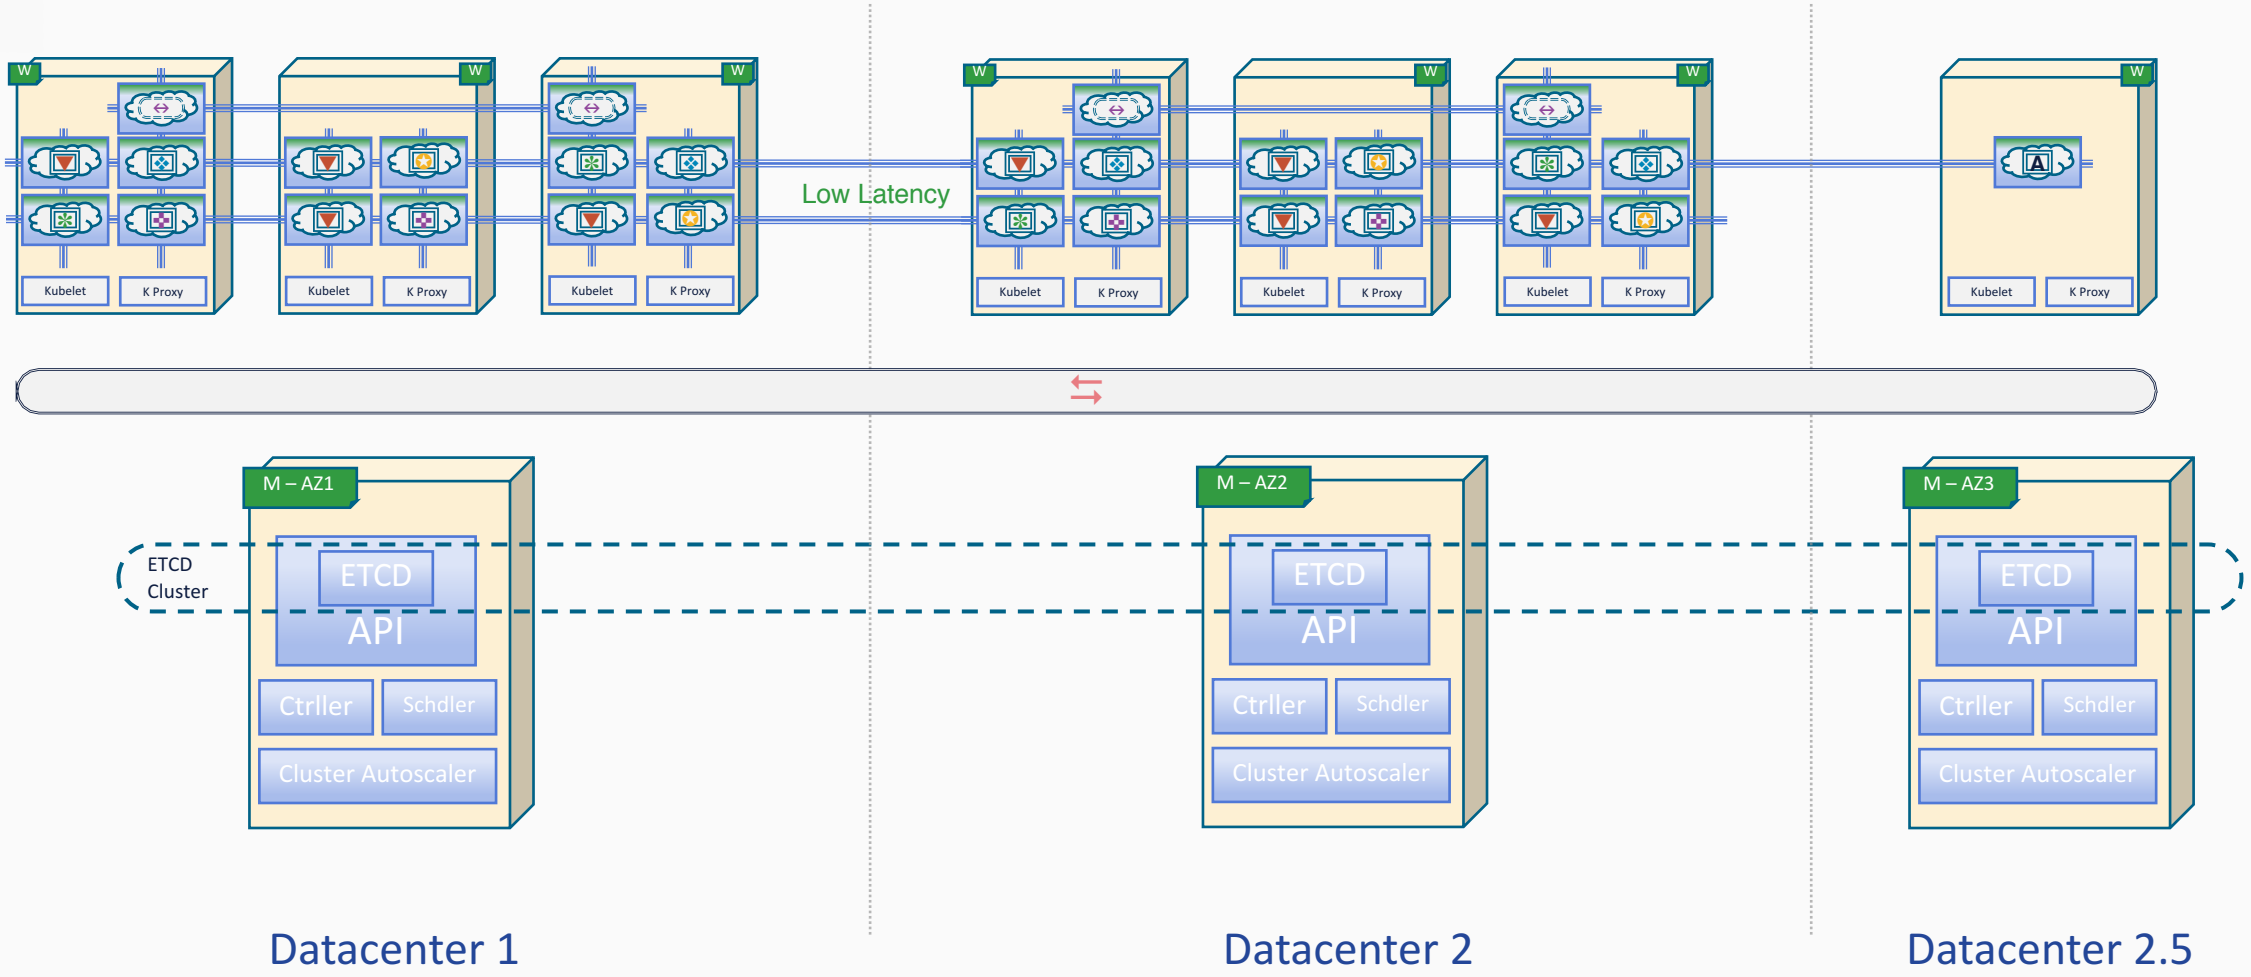

The Glue in a larger µServices architecture

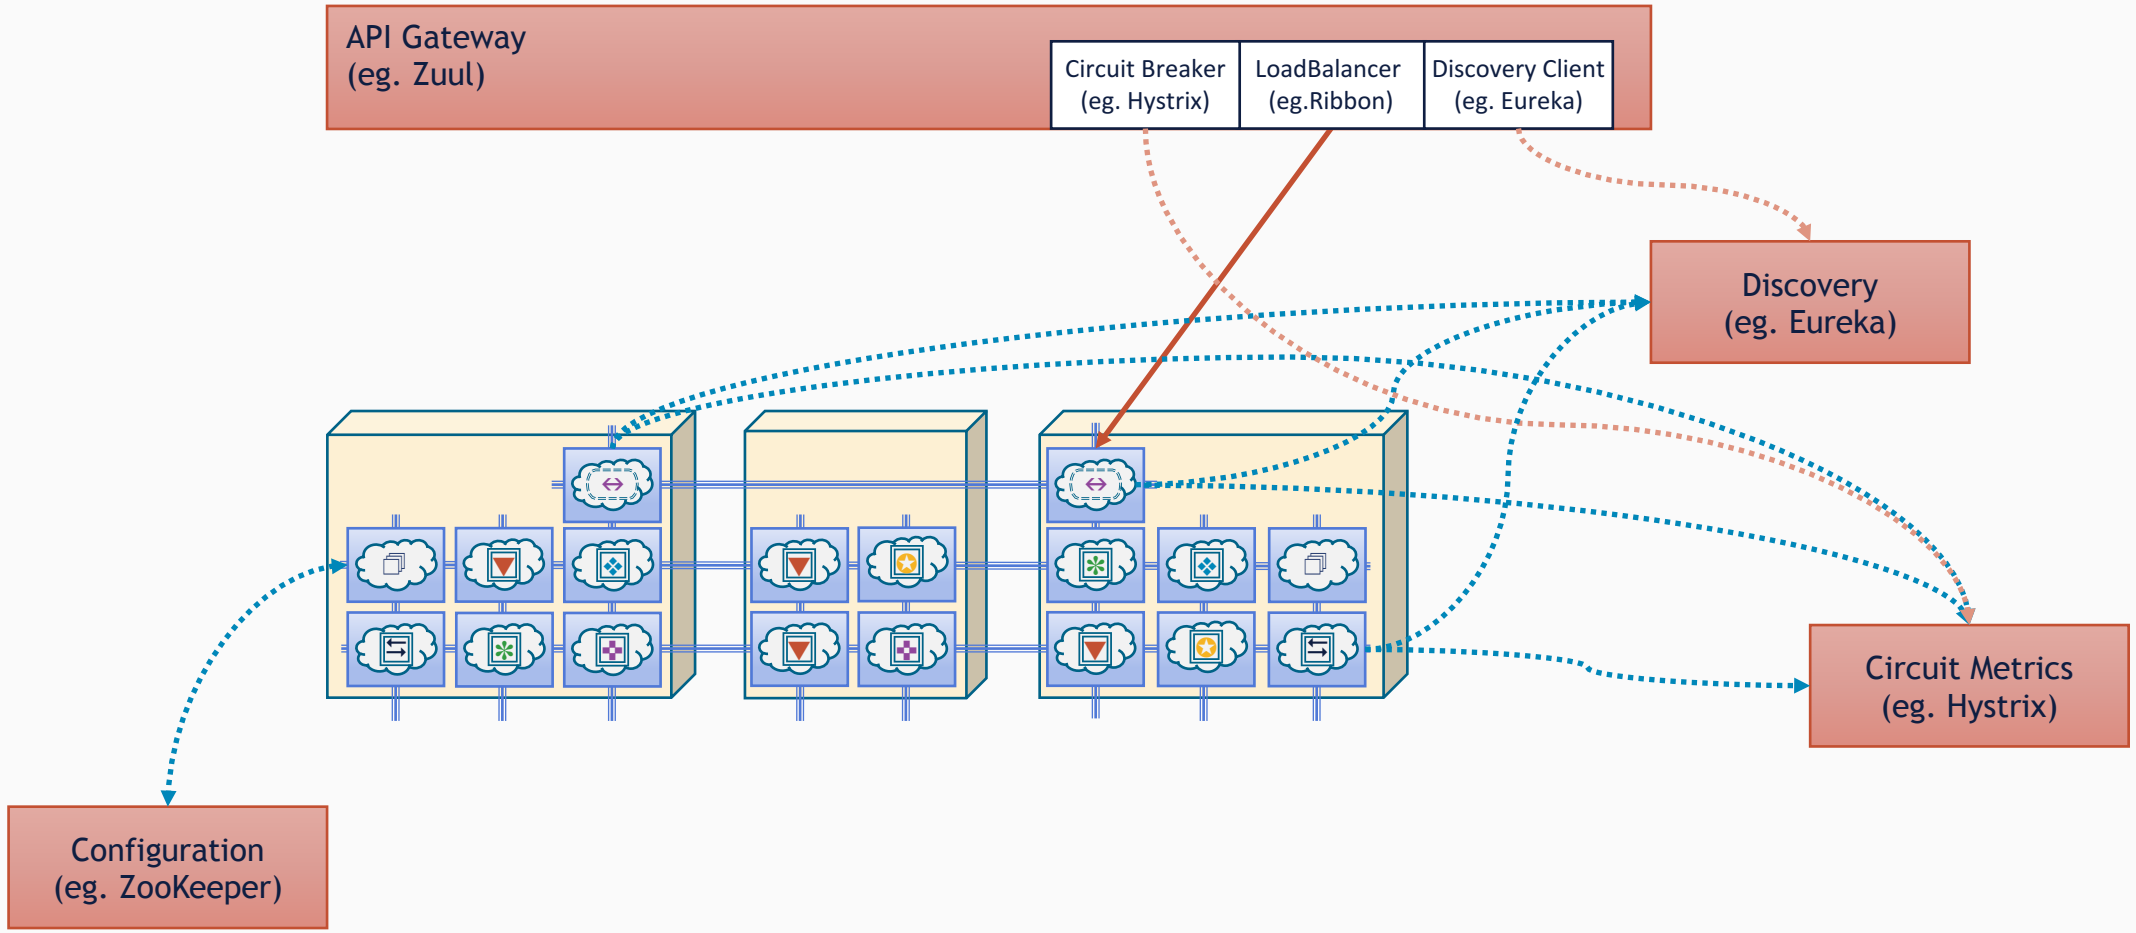

The Glue in a larger µServices architecture

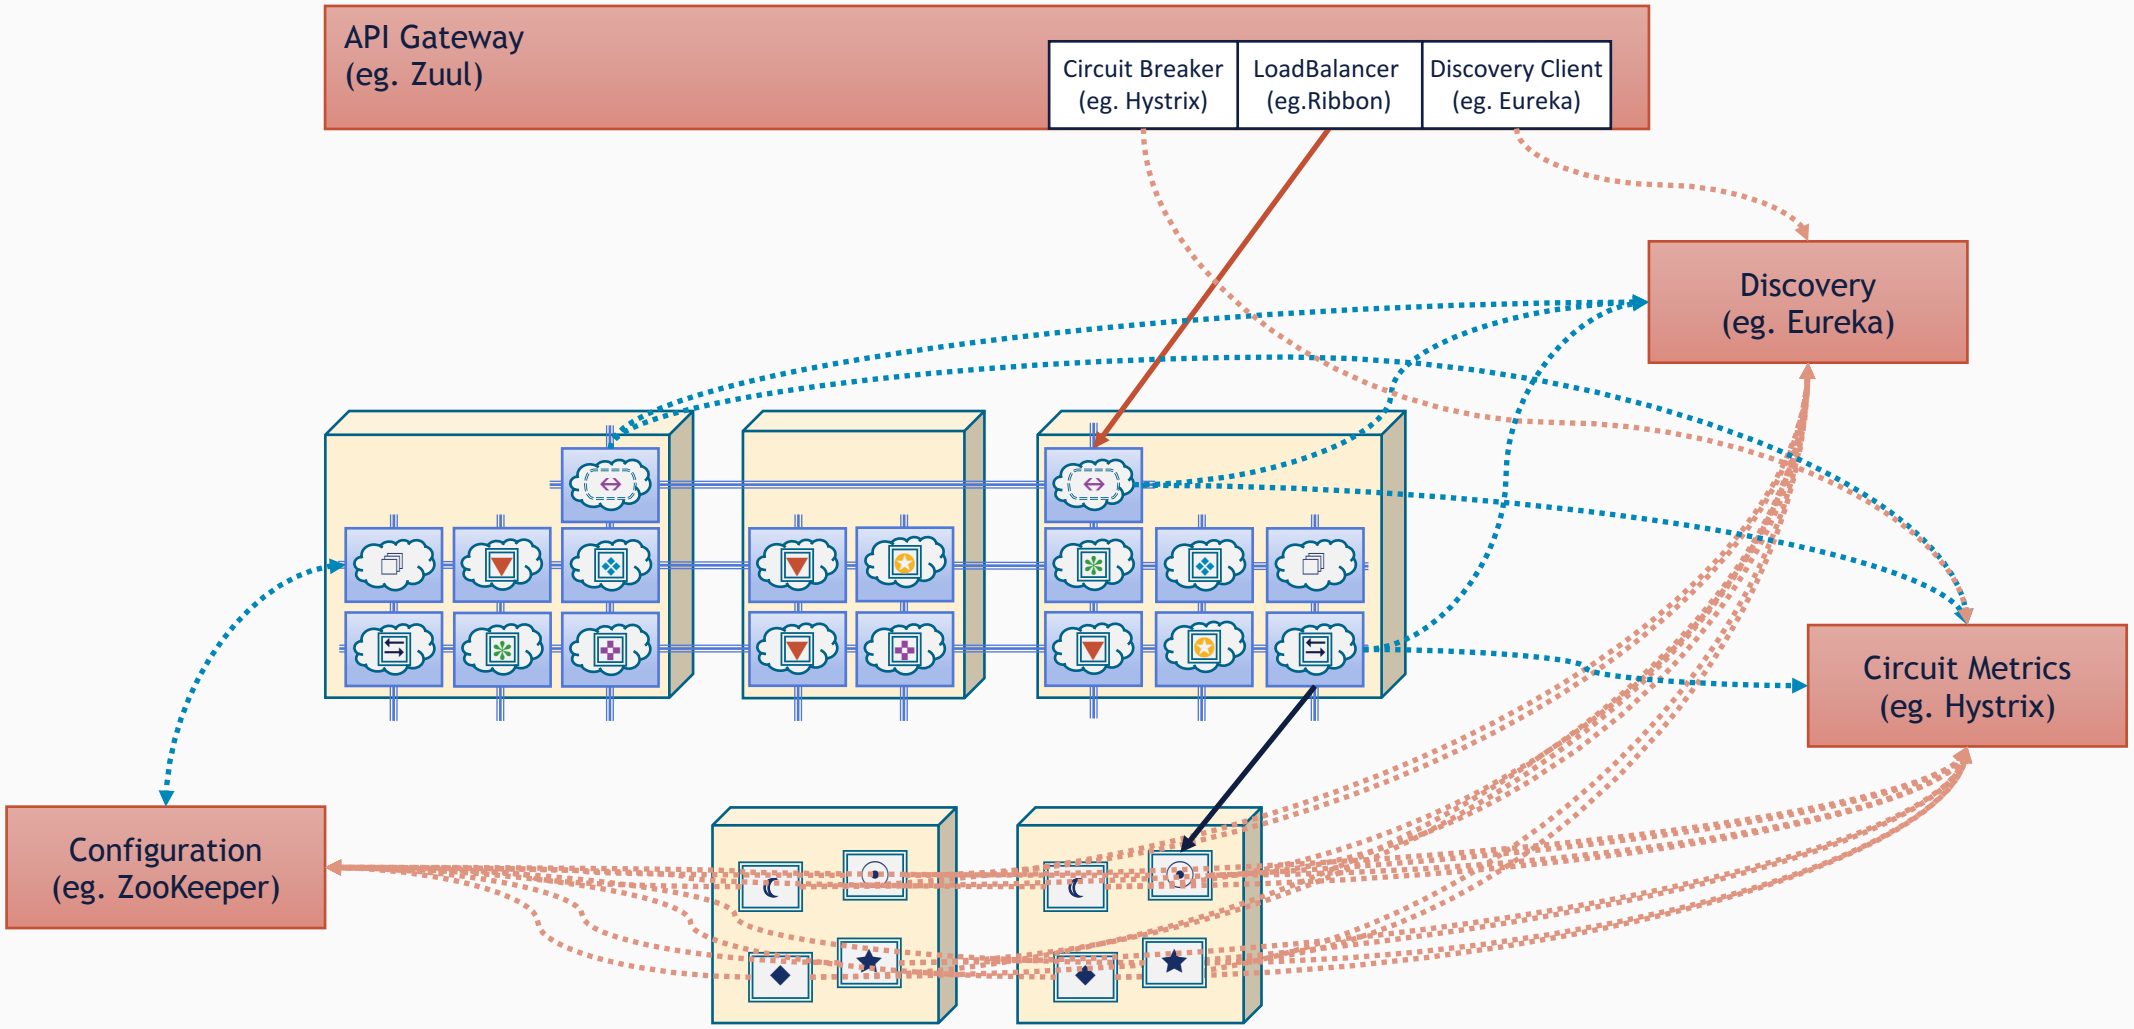


EVENT DRIVEN JOURNEYS
CONTAINER BASED
THE GLUE
THE GLUE
THE GLUE
IN MEMORY DATA GRID

www.theglue.com
www.linkedin.com/in/svenbeauprez

THE GLUE