

American University Kyiv

**Teaching Strategies for
Solution Architecture:
Utilizing Open-Source
Software Projects to
Enhance Professional
Learning**

Ivan Manzhula

MSE | 2024

AUK Experience



Practical Cases on Architectural Katas

Group work and collaboration

Industry experience and knowledge sharing

Ivan Manzhula

MSE | 2024



Problem



Gap between industry and academia

Lack of good teaching materials. Problems are usually explored in isolation, from the beginning ignoring pre-existing solution cases, redesign issues, assuming stable requirements. Reality differs significantly


Uncertainty

Complex domain research goes with non-obvious problems. Students usually face real tradeoffs, risks, demands from wide group of stakeholders after years of working experiences with high probability to fail first designs

Ivan Manzhula

MSE | 2024

Expectation



Software Architecture Course must combine modern theory and practice.

Students can gain experience working with systems larger than a few hundred lines of code.

Course in Software Architecture must emphasize soft skills such as brainstorming, facilitation, negotiation, leadership, decision making, risk analysis etc.

Literary Review



Huge theoretical base

- Software Architecture in Practice
- Patterns of Software Architecture
- Software Architecture: Foundations, Theory, and Practice
- Just Enough Software Architecture: A Risk-Driven Approach
- Software Systems Architecture: Working with Stakeholders Using Viewpoints and Perspectives
- Design It! From Programmer to Software Architect
- Software Architecture: Visual Lecture Notes



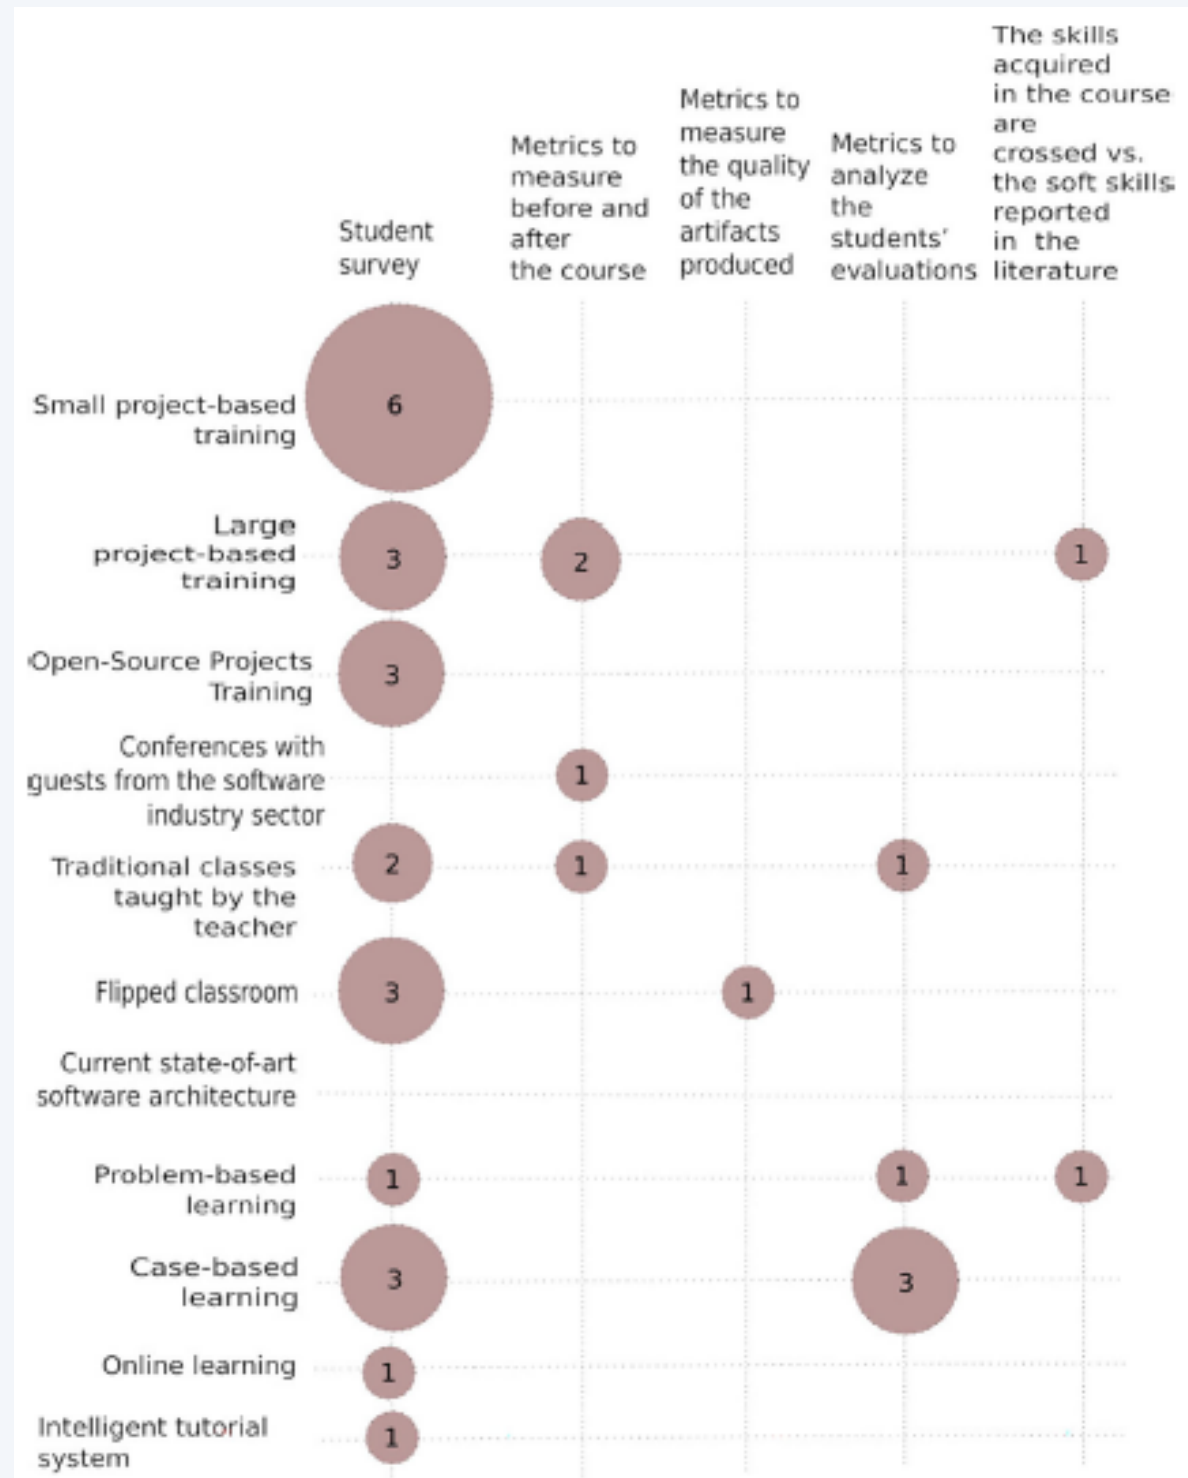
Research Specific

- A Better Way to Teach Software Architecture
- Teaching Reform and Practice of Software Architecture for Postgraduates in Universities
- Training software architects suiting software industry needs

Ivan Manzhula

MSE | 2024

Teaching Strategies



EXAMPLE:

1. Divide students into groups / teams
2. Assign roles: senior / cognitive / junior architects
3. Senior architects - make decisions
4. Cognitive architects - challenge these decisions
5. Junior architects - model these decisions

Use existing mature open-source projects

Pros 1

Real-world settings, long-term project allows study of changing requirements impact and architecture refactoring

Pros 2


Learning, recovering architecture can potentially help students find bugs or implement new features, thus contributing to the community and bring real industrial impact

Pros 3

Opportunity to communicate with open-source and industrial practitioners

Pros 4

Numerous design decisions within one project, observability of real tradeoffs



Challenges



Complexity: can be resolved by proactive extraction of design decisions from source code and their usage for teaching examples



Course structure changes: new content organization and additional materials. Design decisions can be collected based on quality attributes, viewpoints, software development lifecycles



Event Driven: Apache Kafka

Source Code: Java vs Scala

Complex Structure

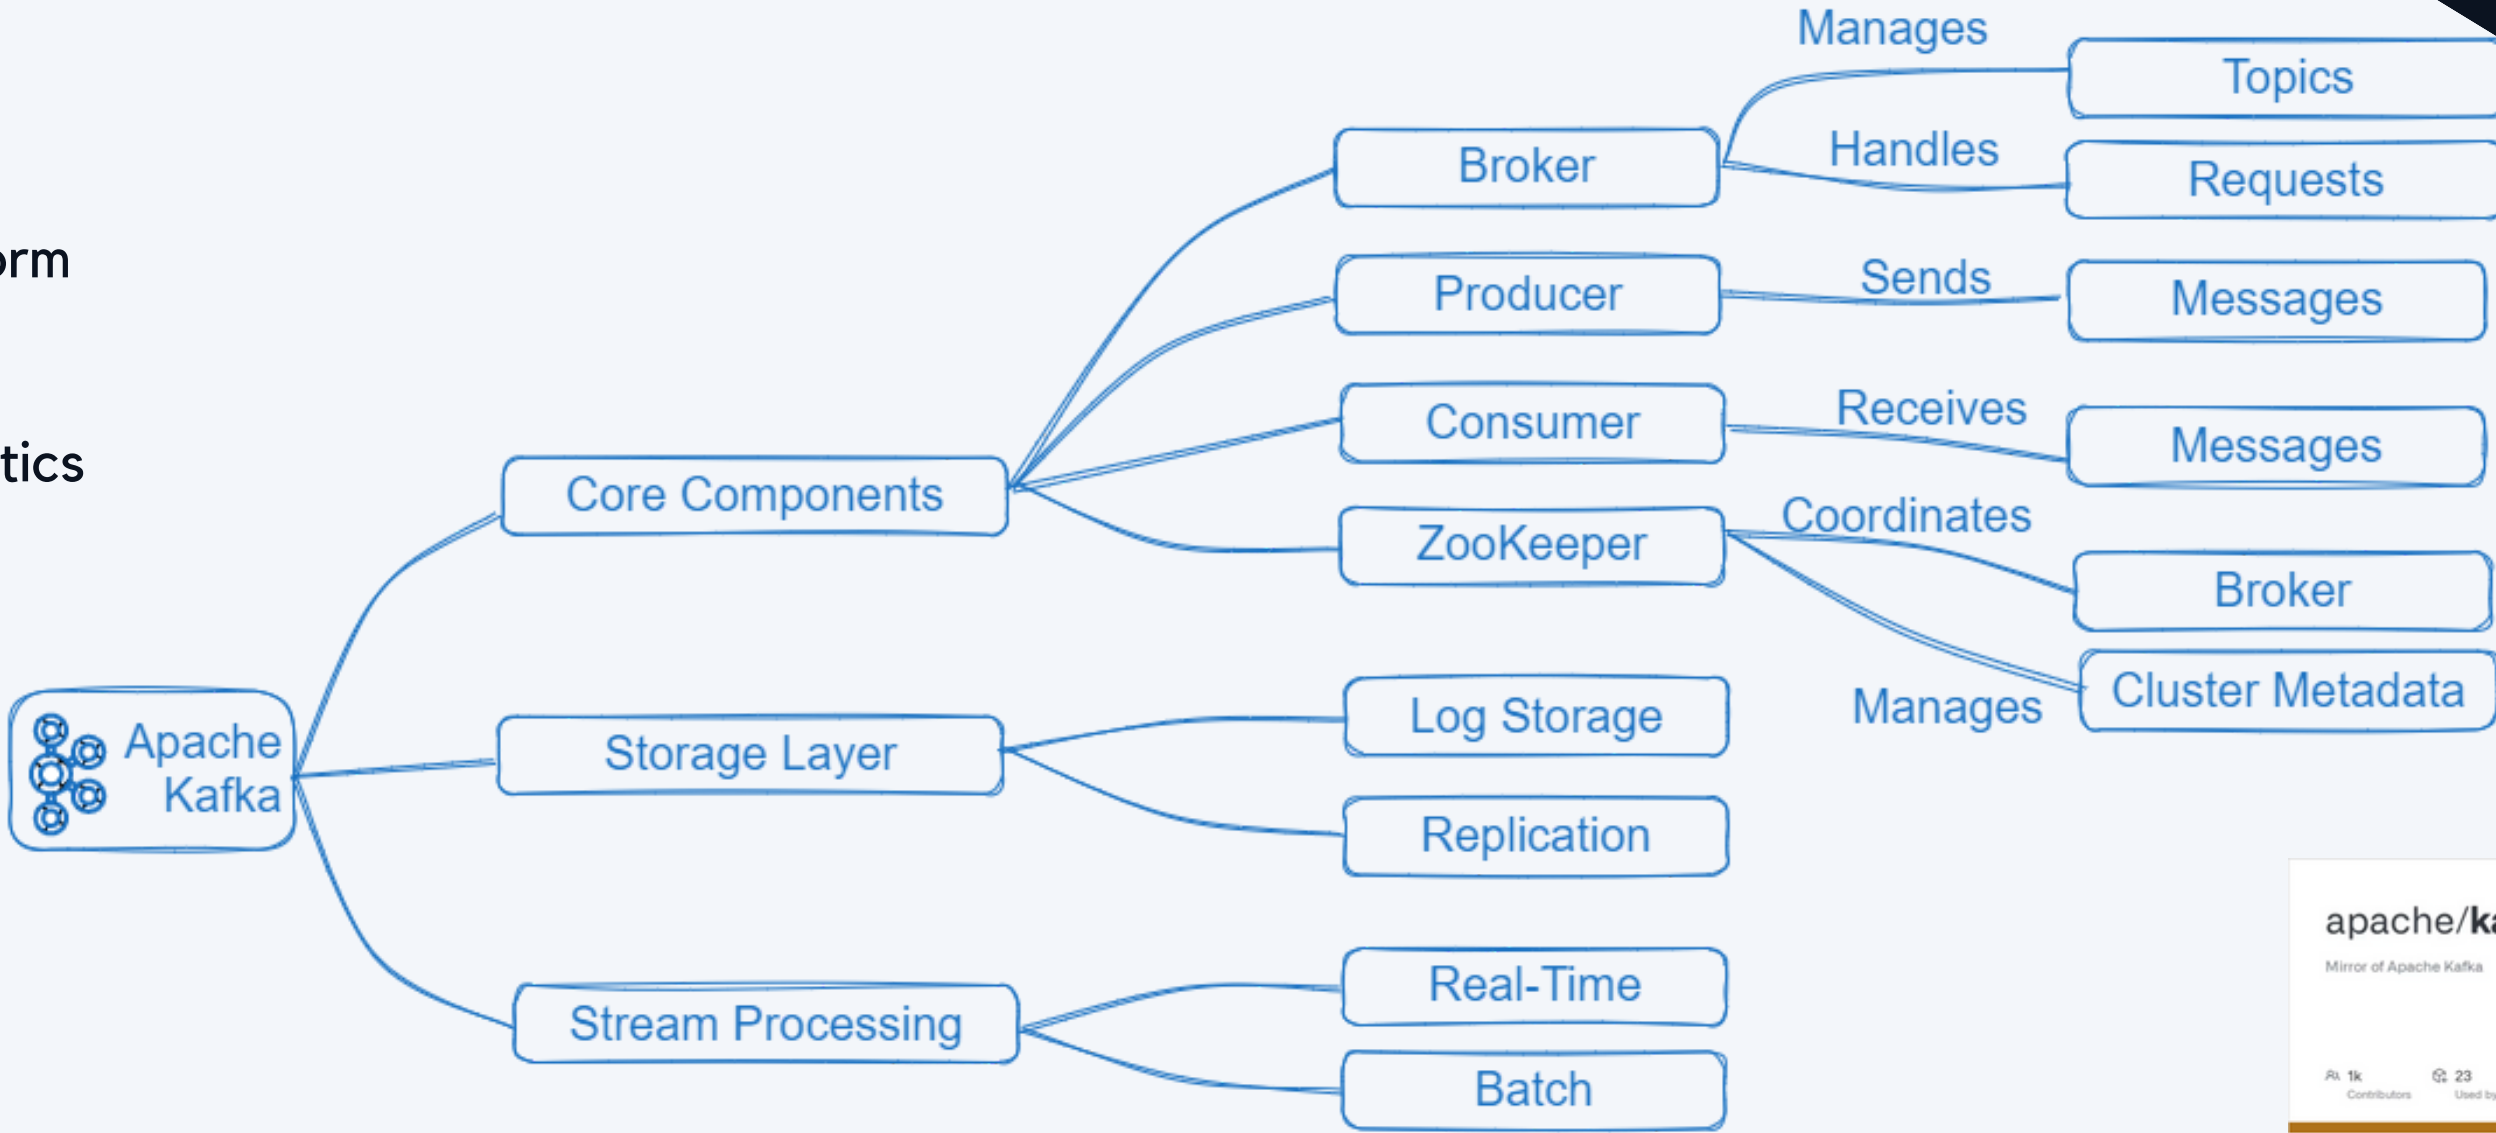
Multiple Design Patterns Combined

Popular and Widely Used

Distributed streaming platform

Topic-based data model

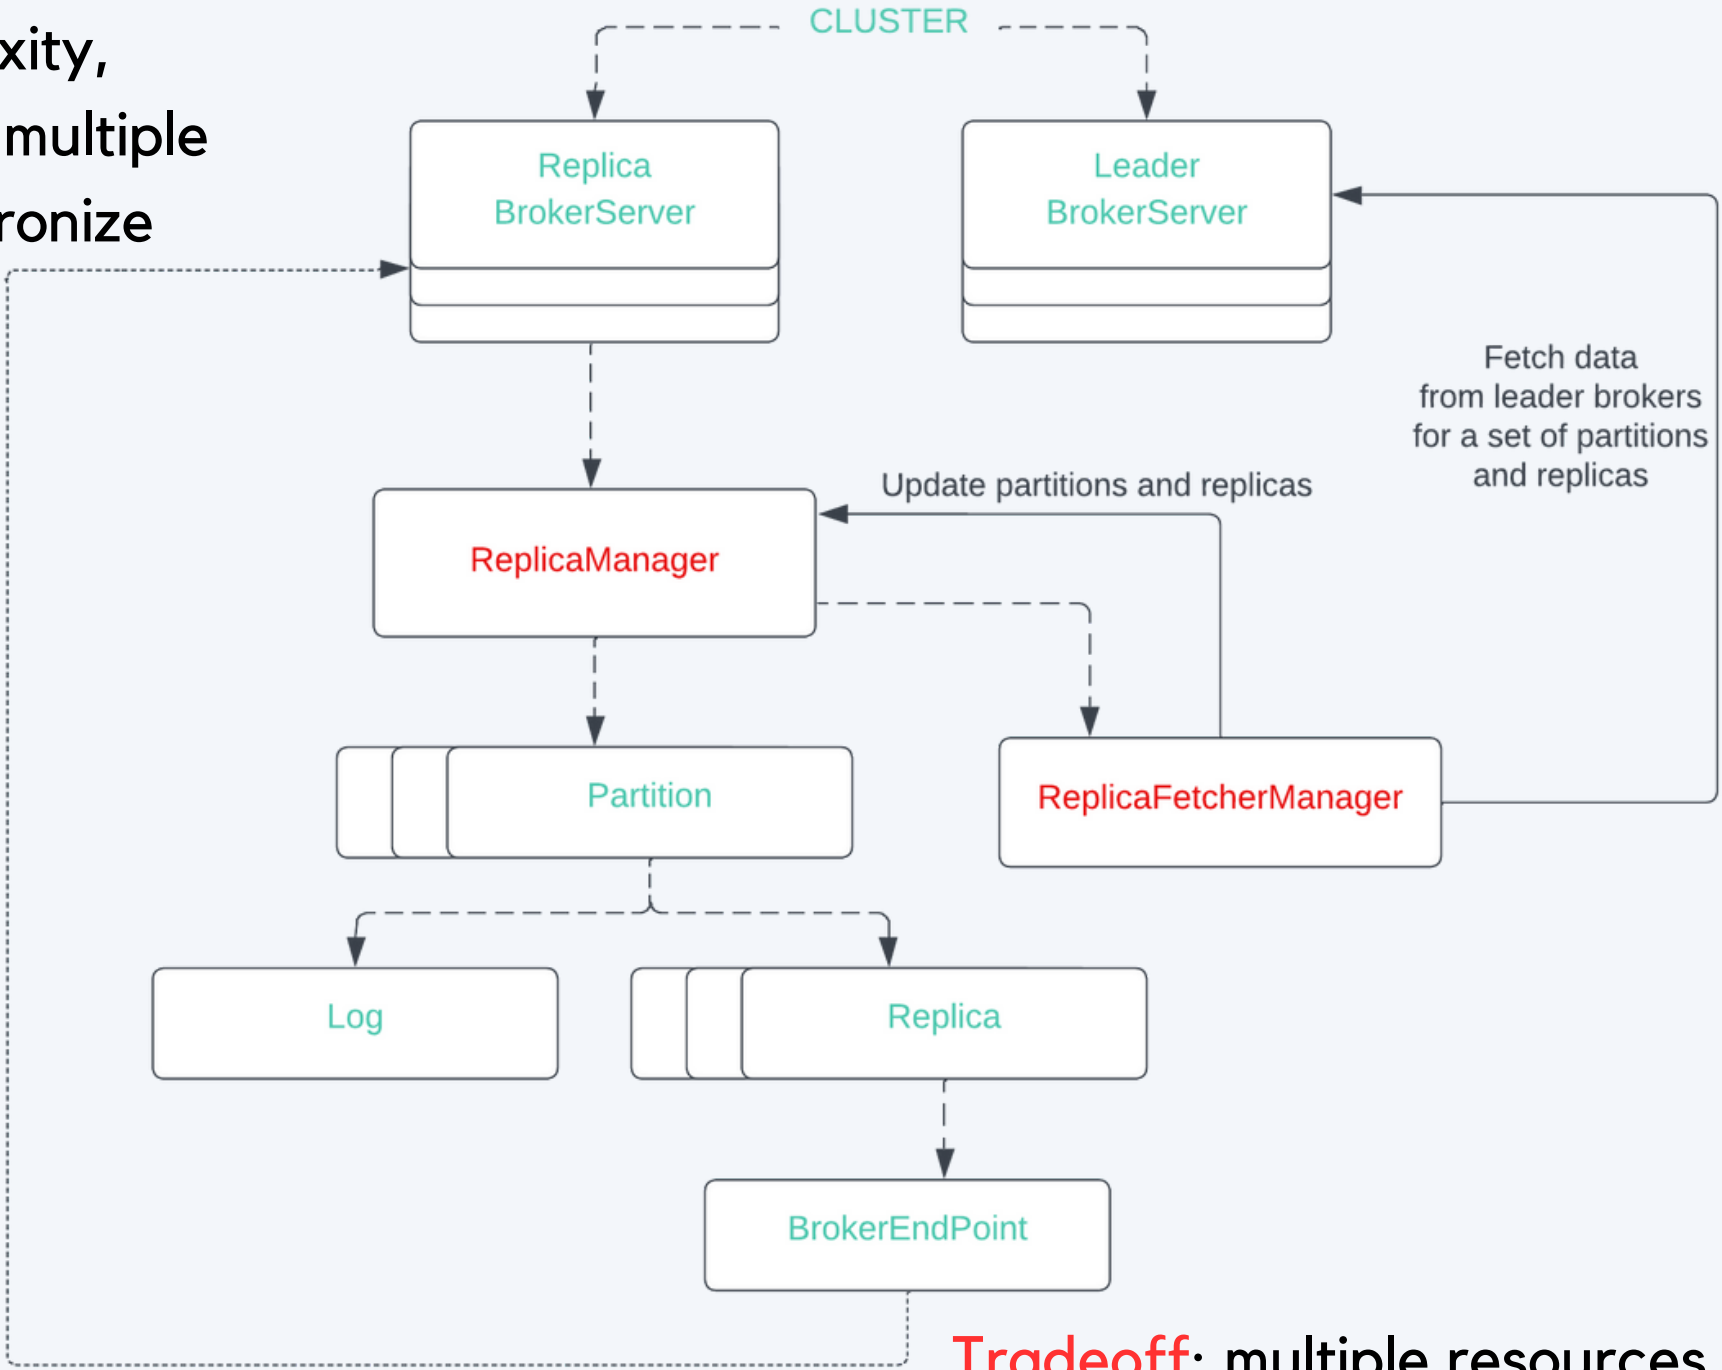
At least one delivery semantics



High Availability: Replication

Data is replicated across multiple brokers within the Kafka cluster to ensure data availability and prevent data loss in case of hardware failures

Tradeoff: complexity, necessity to manage multiple replicas and synchronize them



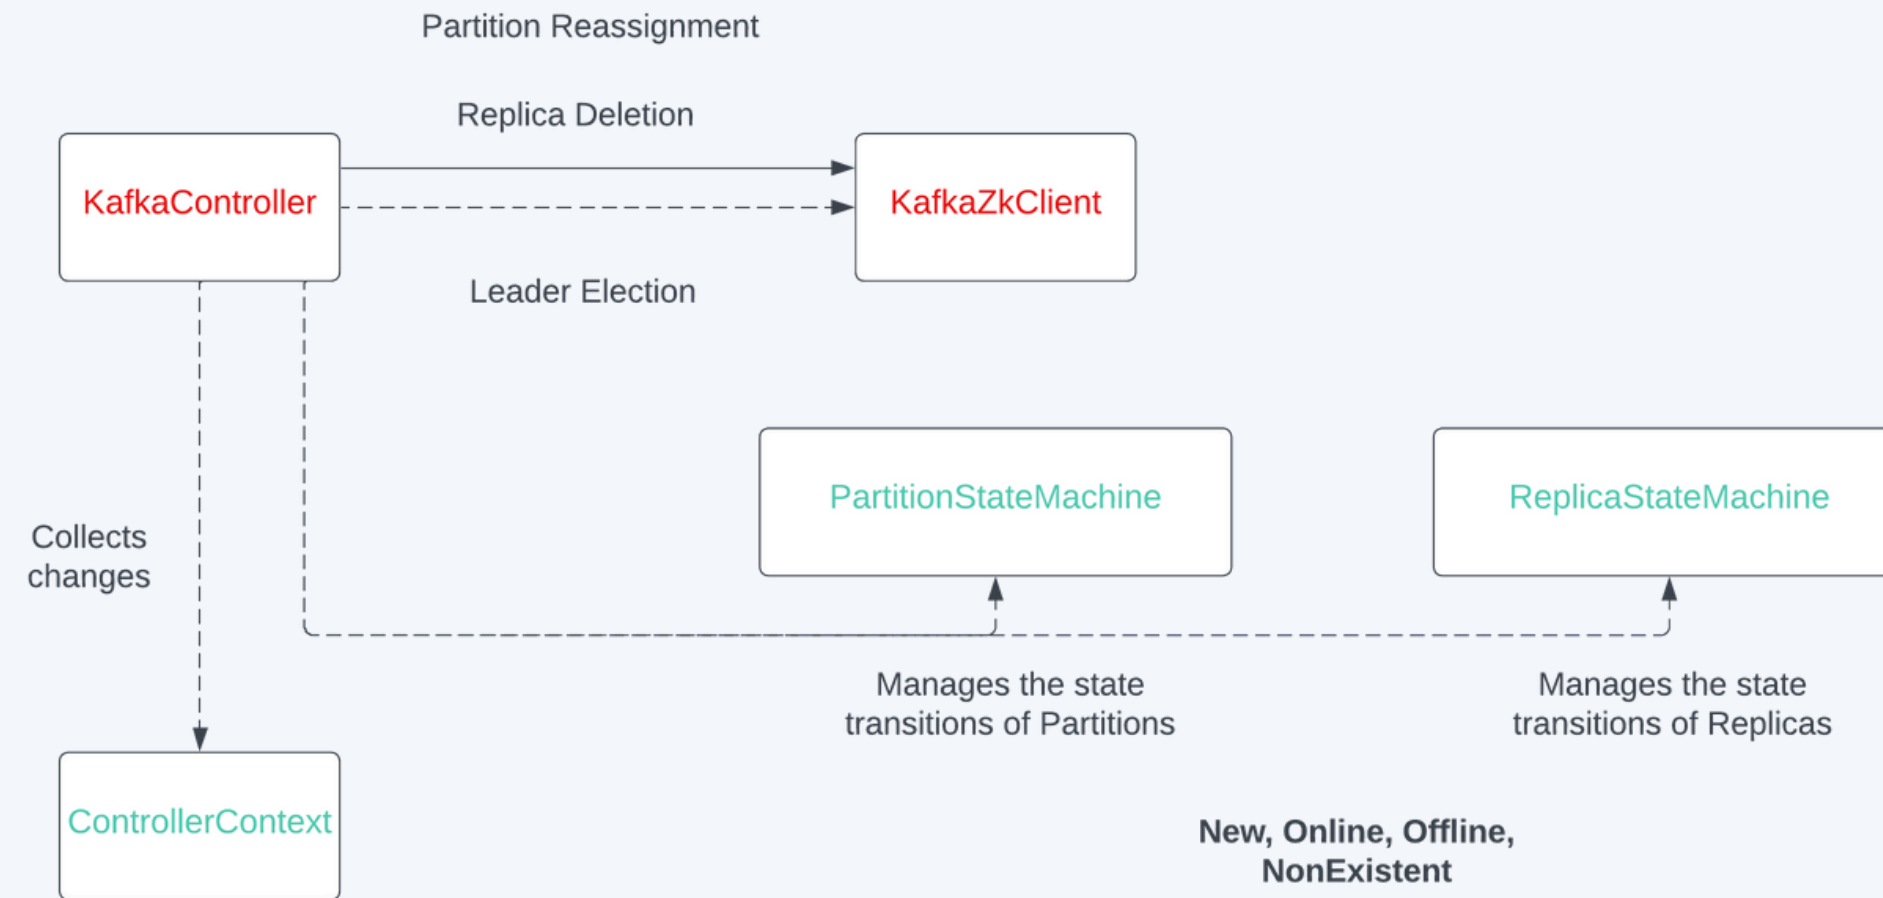
LEGEND

- Entity** -- Object, contains data
- Worker** -- Service, performs actions
- -- Action
- - -> -- Includes
- ⋯ - - -> -- Reference / Association
- BrokerServer** -- Kafka Broker Class (Leader(s) or Replica(s))
- ReplicaManager** -- Class, handles data replication logic for the broker
- ReplicaFetcherManager** -- Class, handles update of replicas from the leader brokers
- Partition** -- Class, partition of a topic
- Log** -- Object, stores the actual data for the partition
- Replica** -- Class, replica of the partition
- BrokerEndPoint** -- Object, stores the host and port of the broker replica

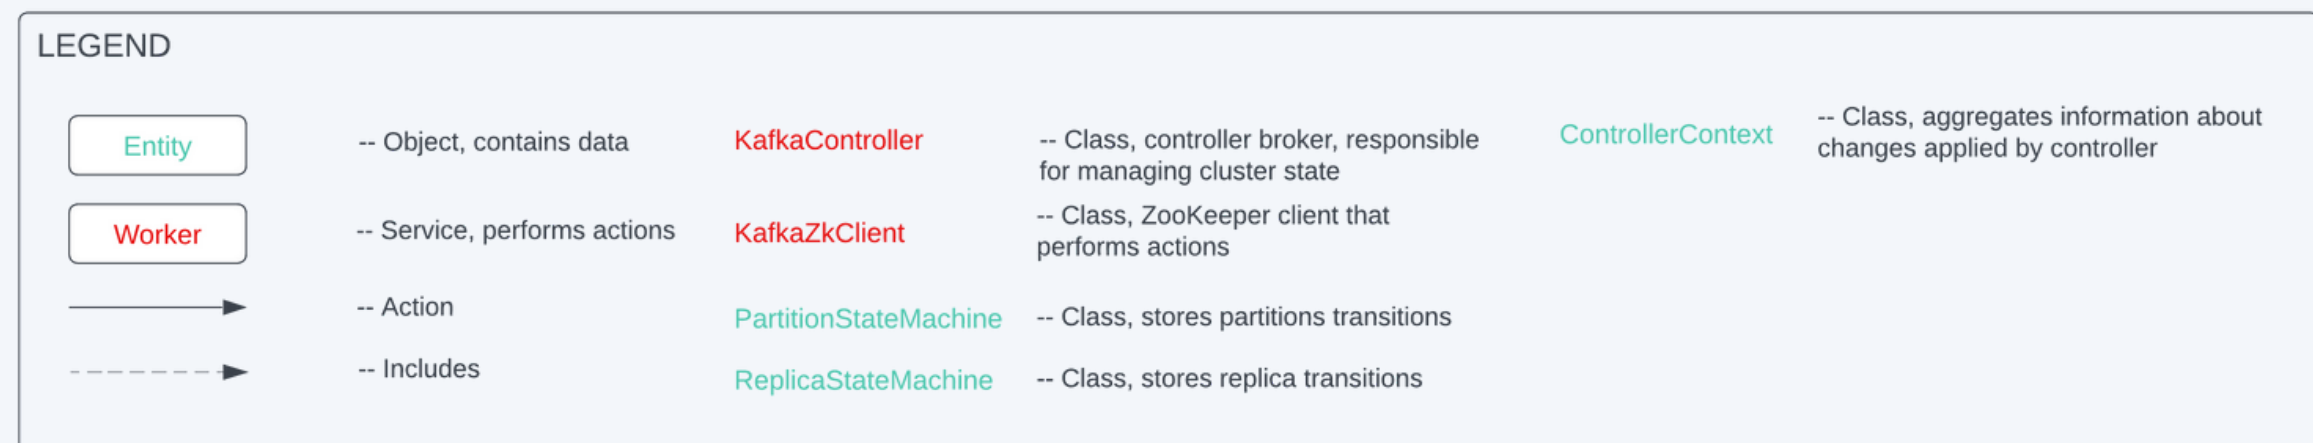
Tradeoff: multiple resources reflecting same information in many objects

High Availability: Leader election

For each topic partition, a leader broker is responsible for replicating data to other followers and handling read/write requests. This ensures data availability and fault tolerance even if the leader fails.

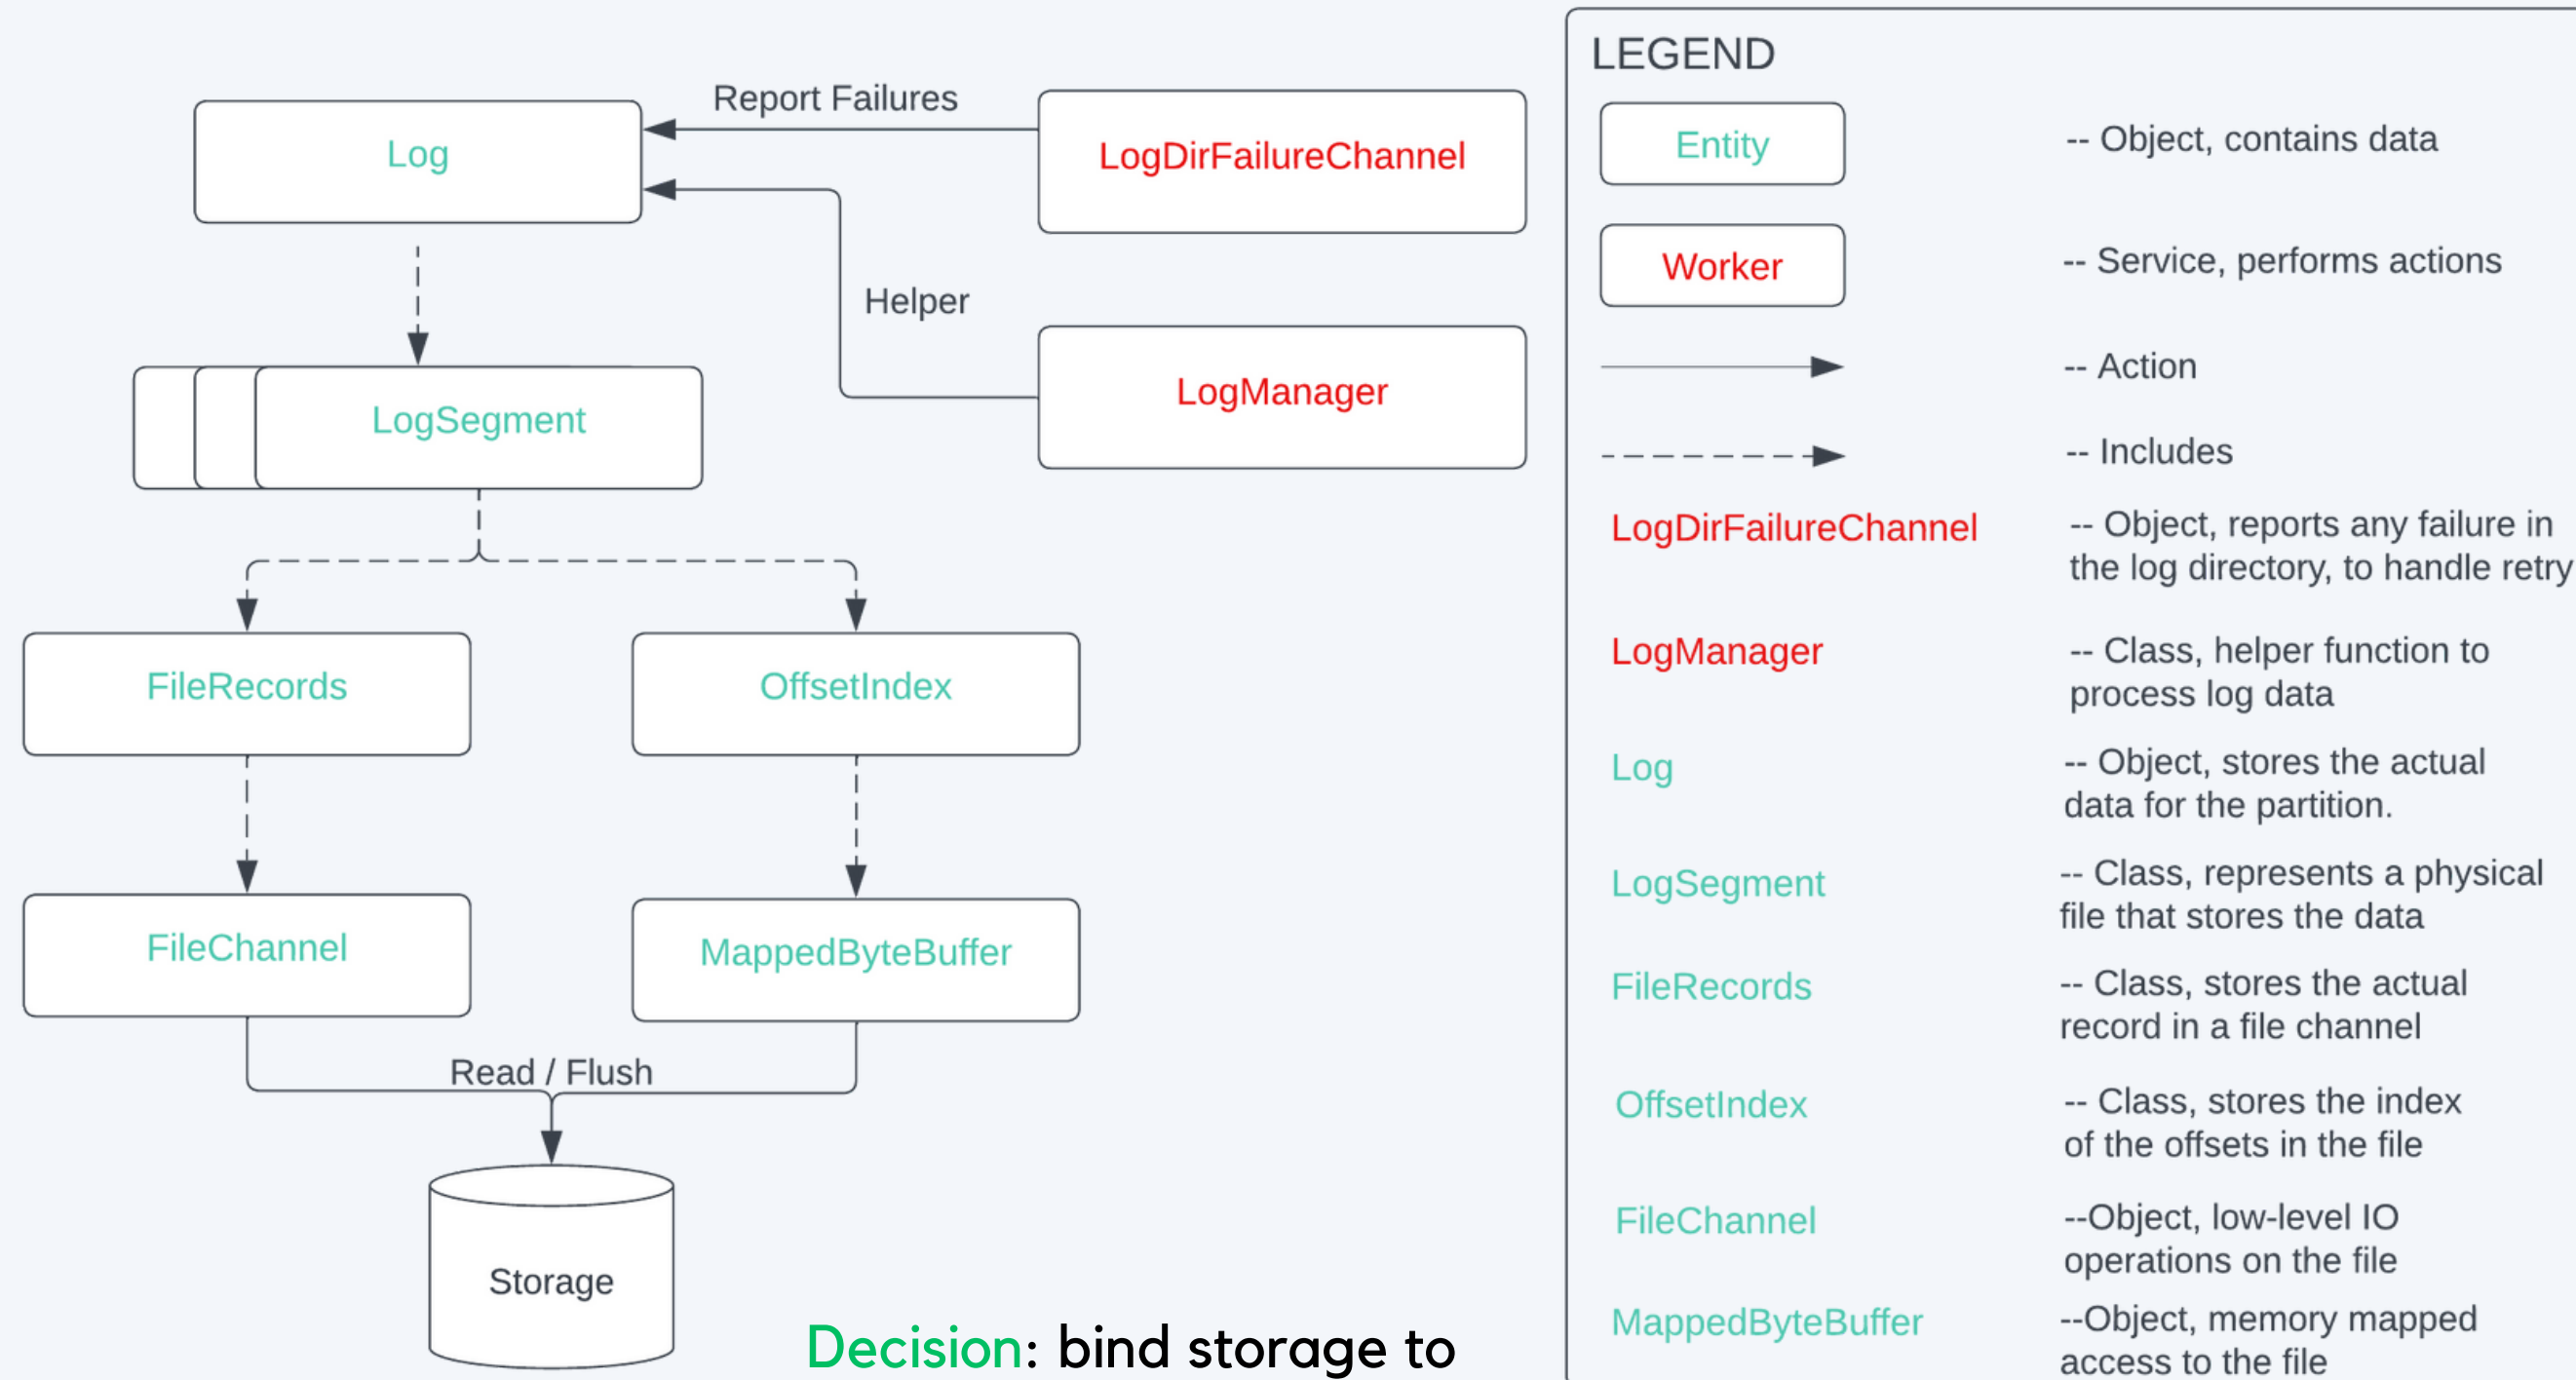


Tradeoff: despite states and caches leader election can impact performance



High Availability: Data persistence

Data is written to disk and periodically flushed to ensure durability even if the system crashes

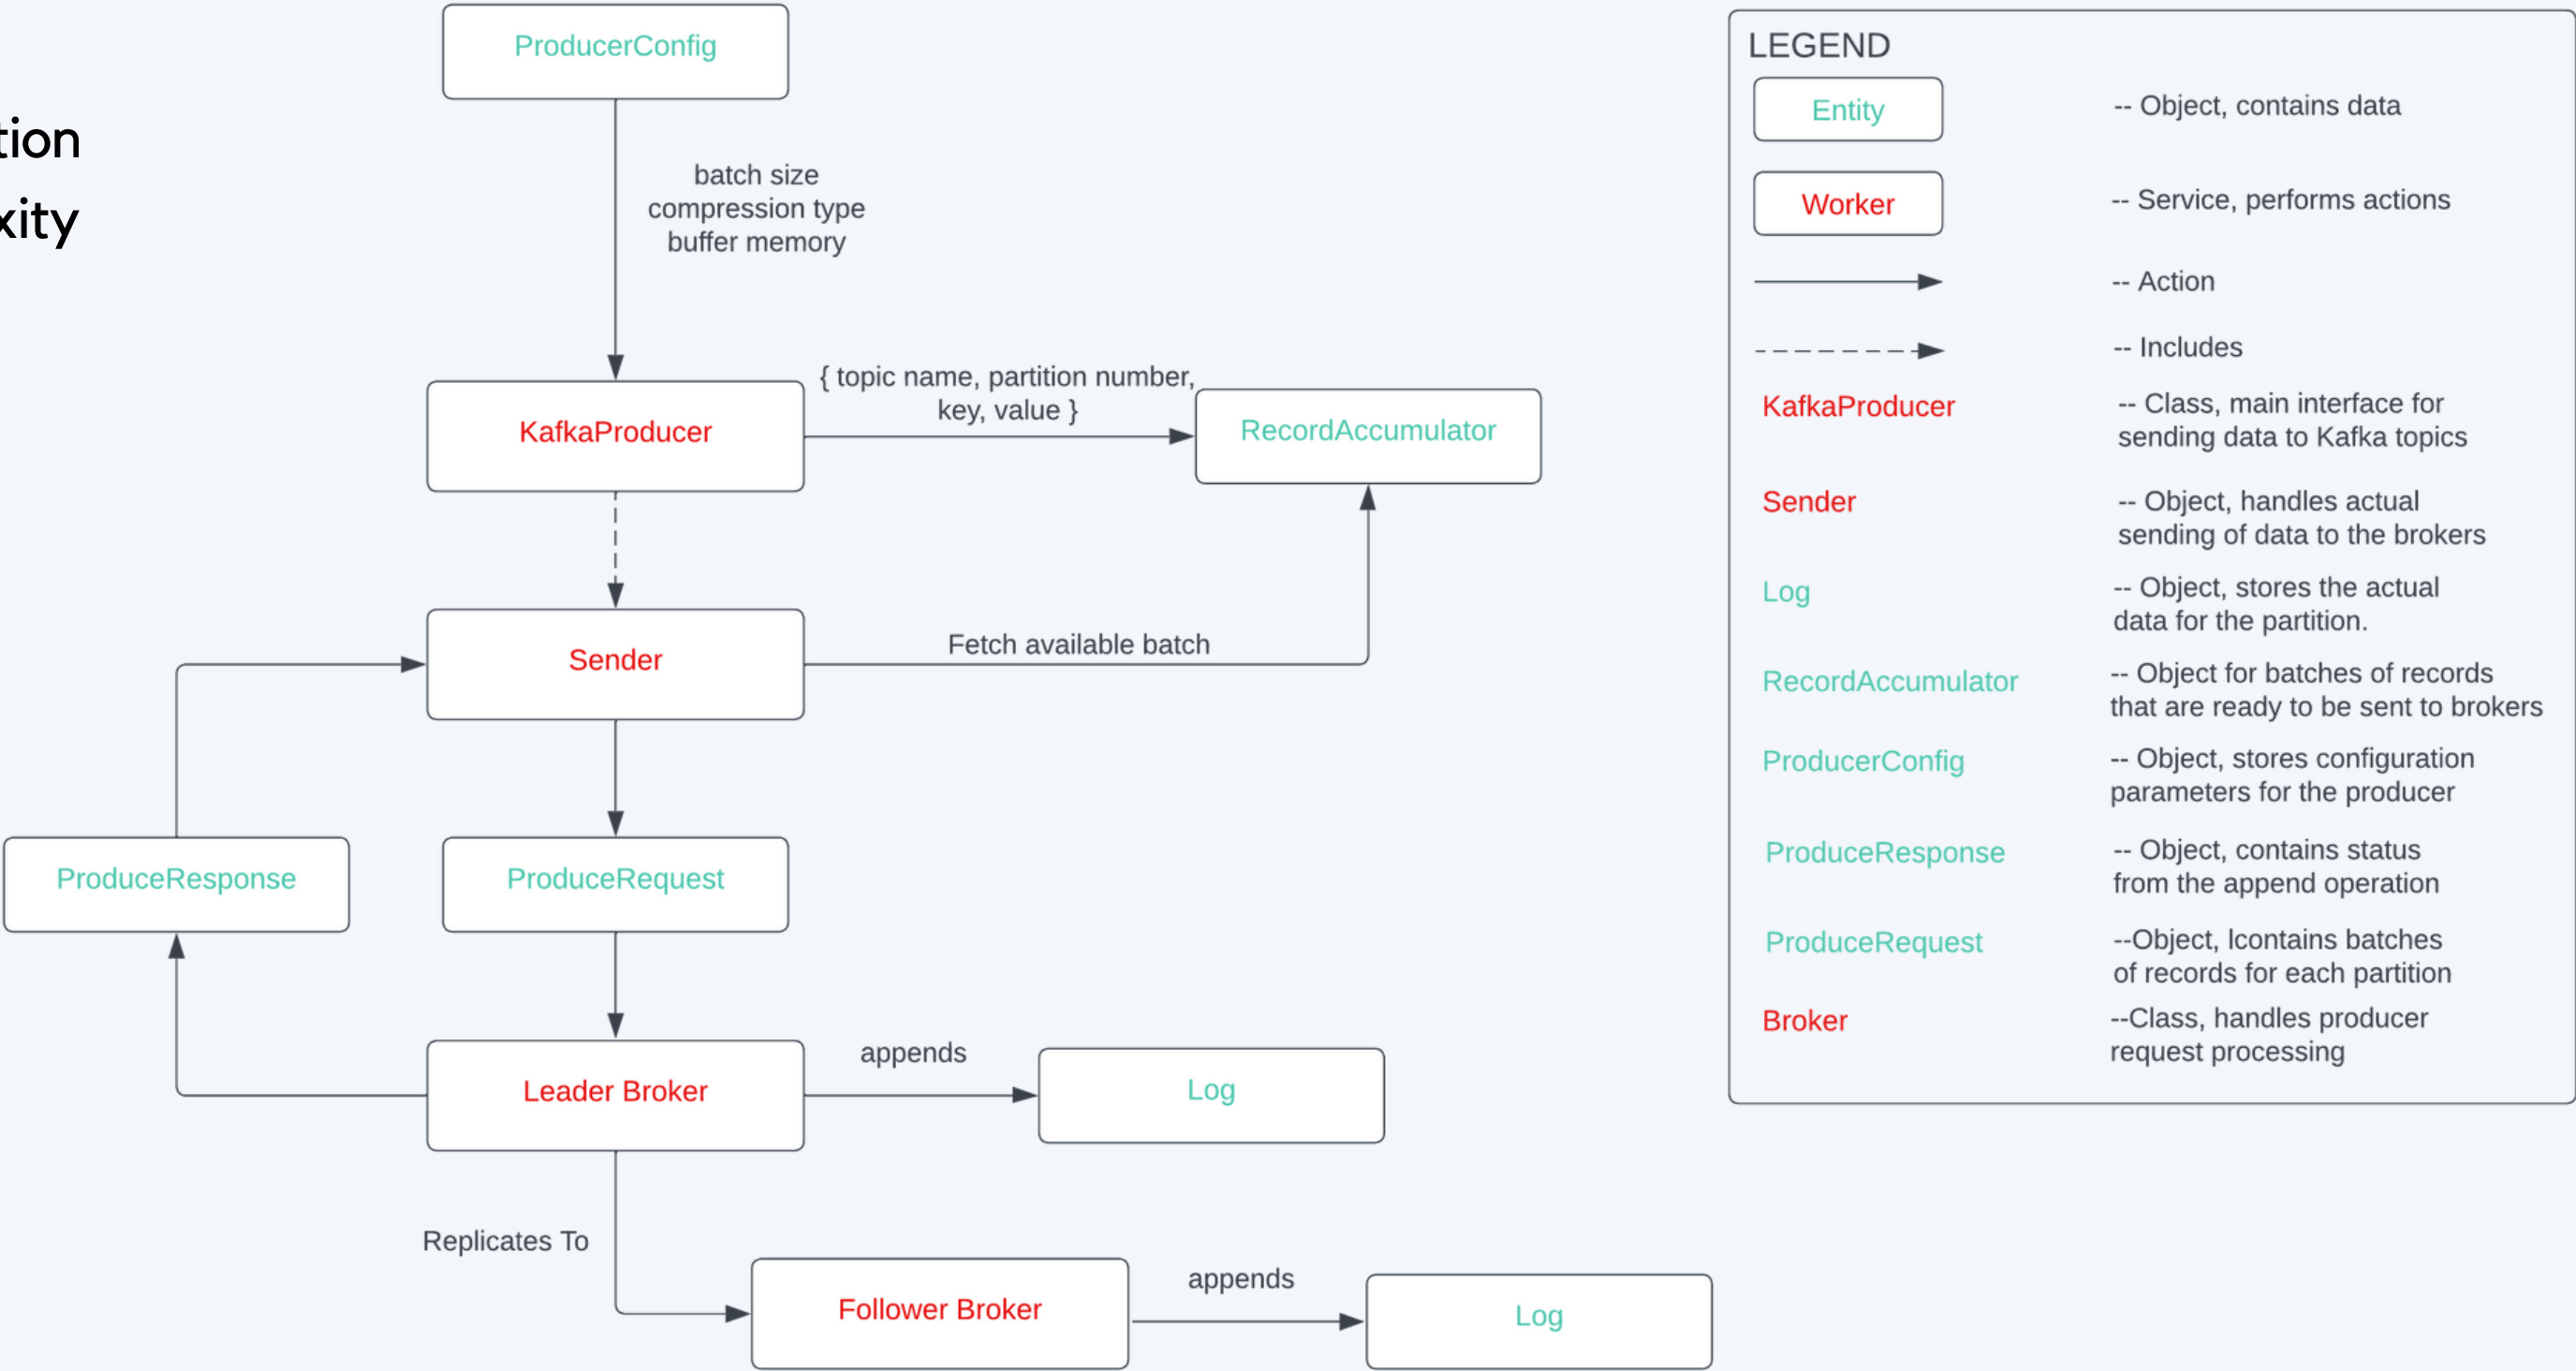


Decision: bind storage to data object instead of using central resources

High Performance: Producer batching

Producers can batch data into larger messages for more efficient network transfer and processing

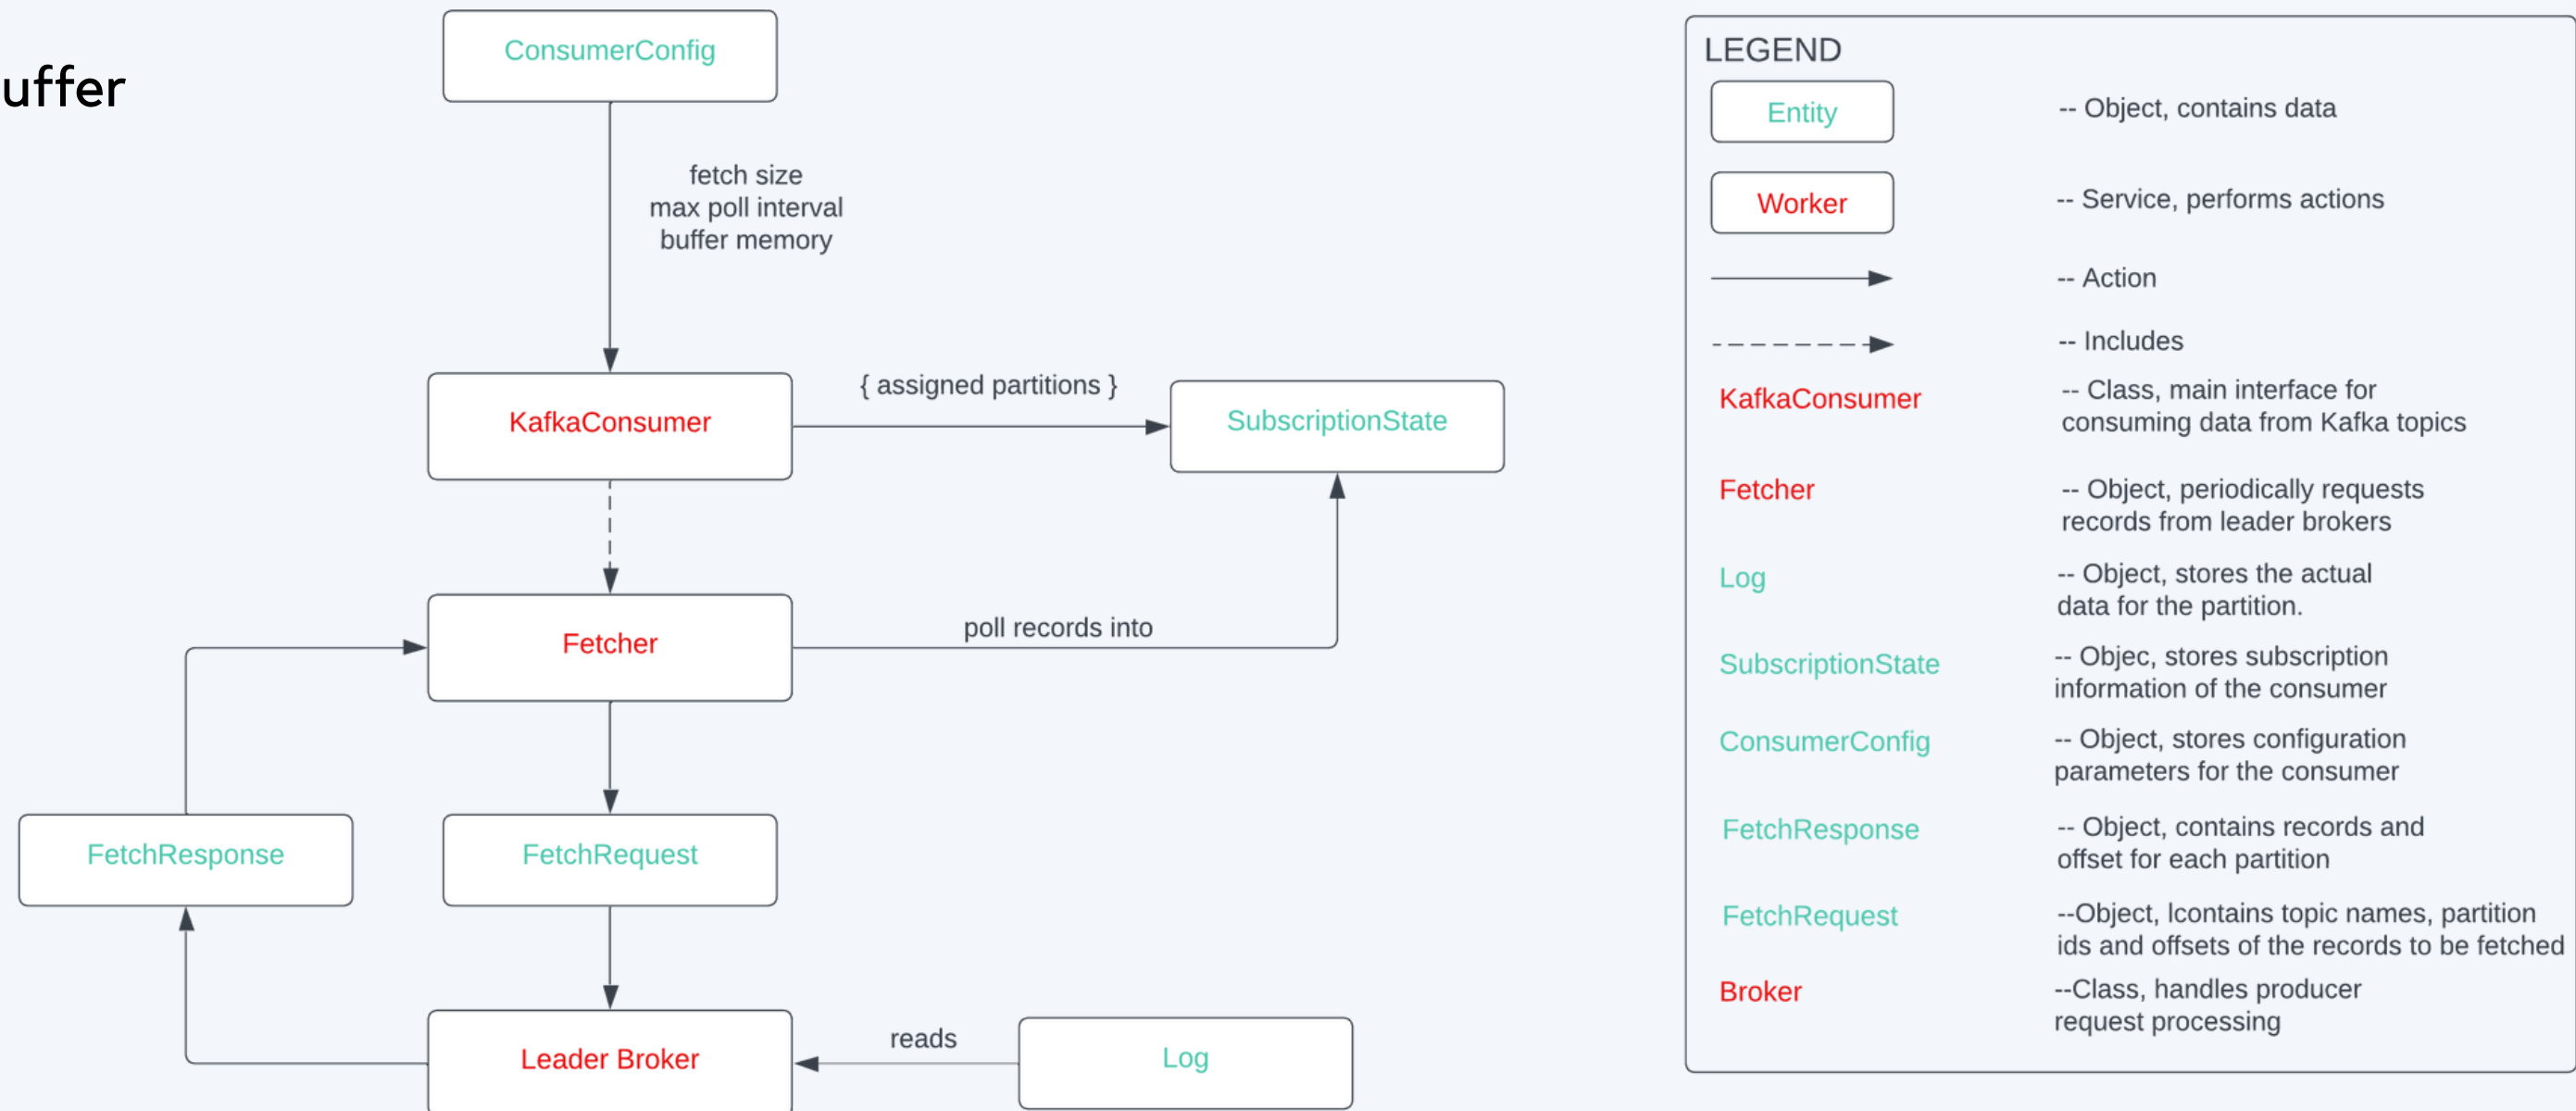
Tradeoff: communication and recovery complexity



High Performance: Consumer buffering

Consumers can buffer data to reduce the number of disk reads and improve consumer performance

Decision: use buffer




Distributed Database: Apache Cassandra

Source Code: Java

NoSQL Database

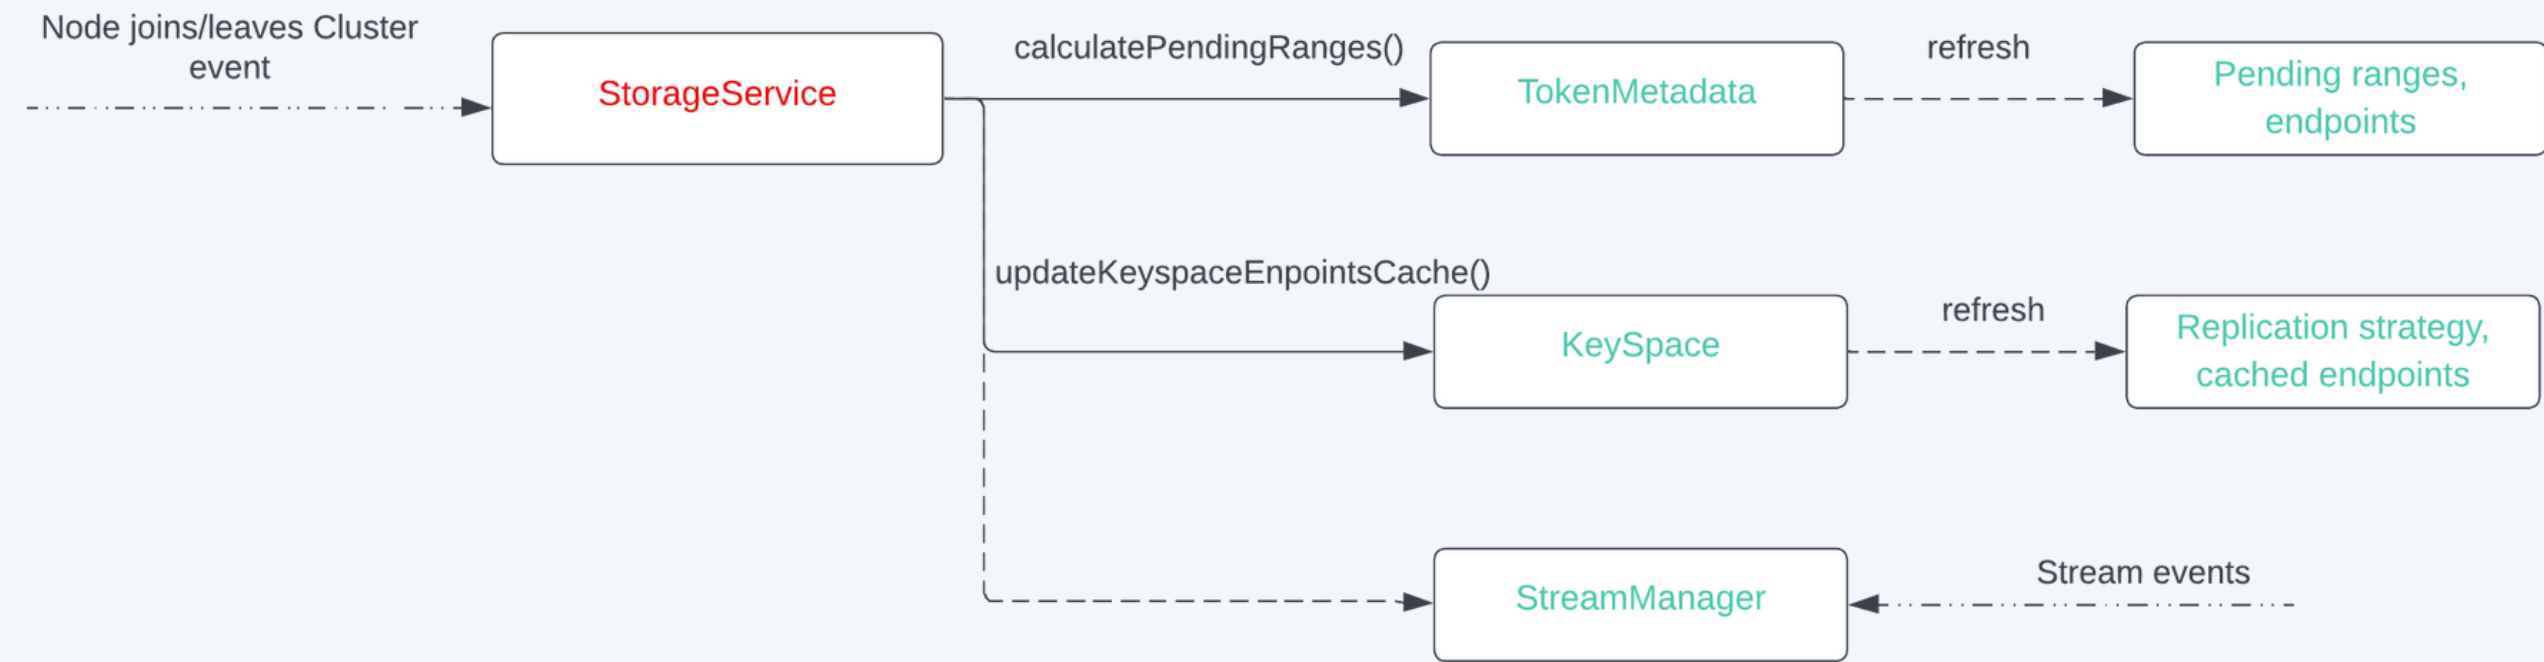
Row-based data model

Tunable consistency: eventual consistency \leftrightarrow strong data integrity

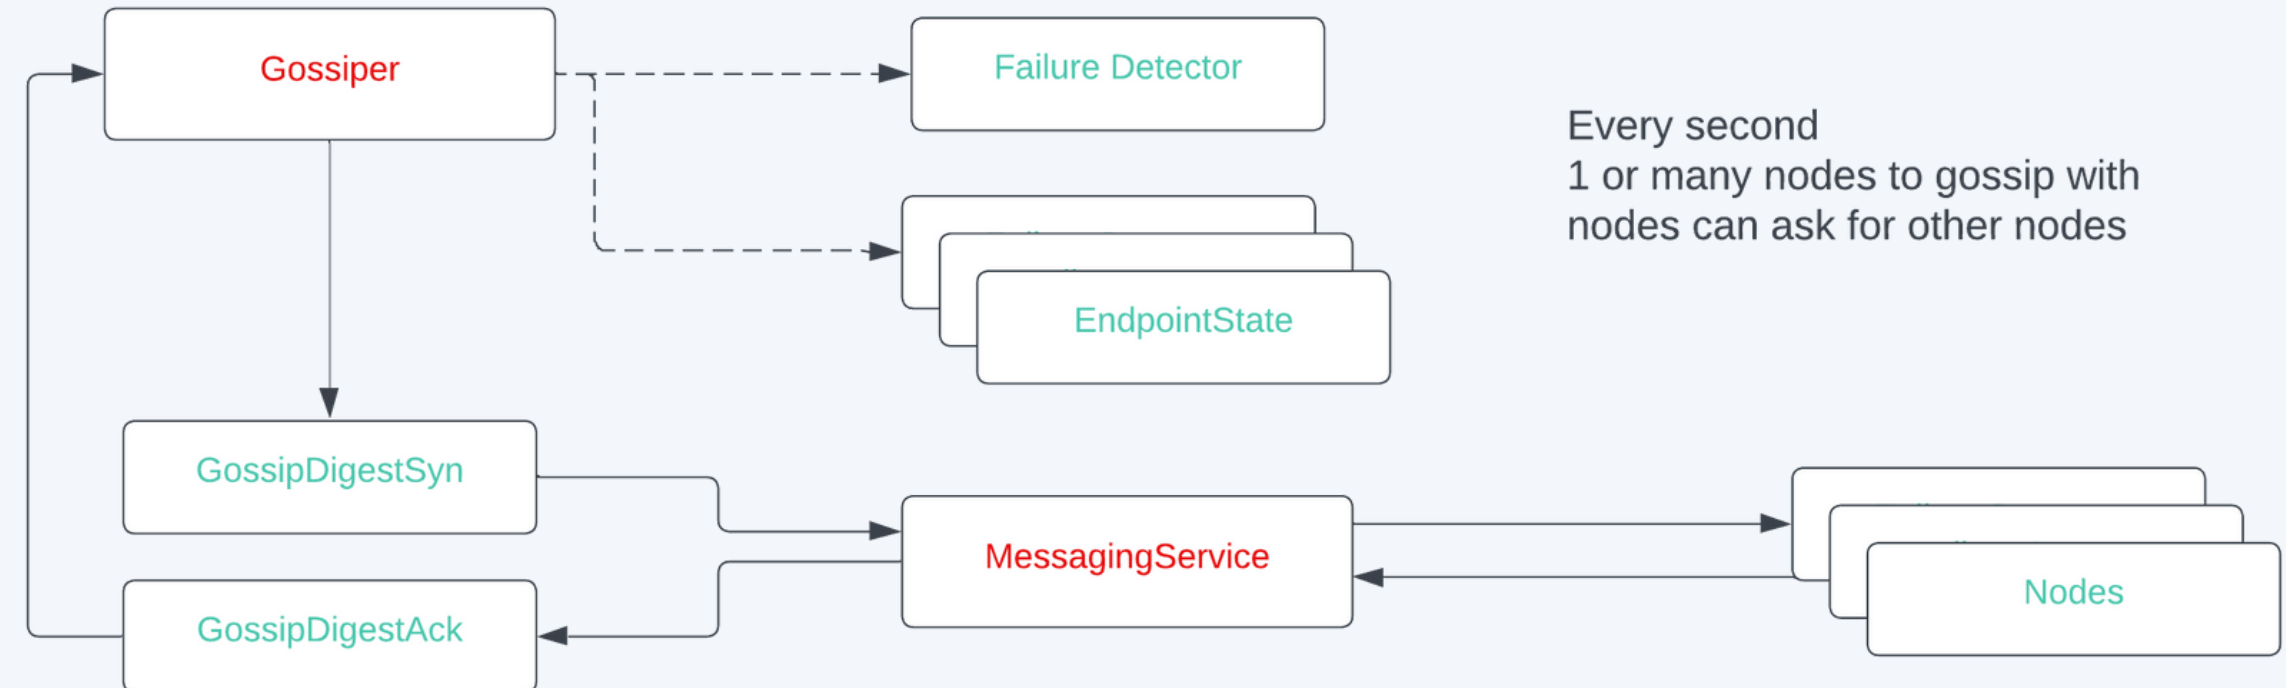


High Availability: Automatic failover and Gossip protocol

When a node fails, Cassandra automatically rebalances the data and workload across remaining nodes, minimizing downtime and data loss. Nodes continuously communicate with each other through a gossip protocol, sharing information about their state and ensuring data consistency.



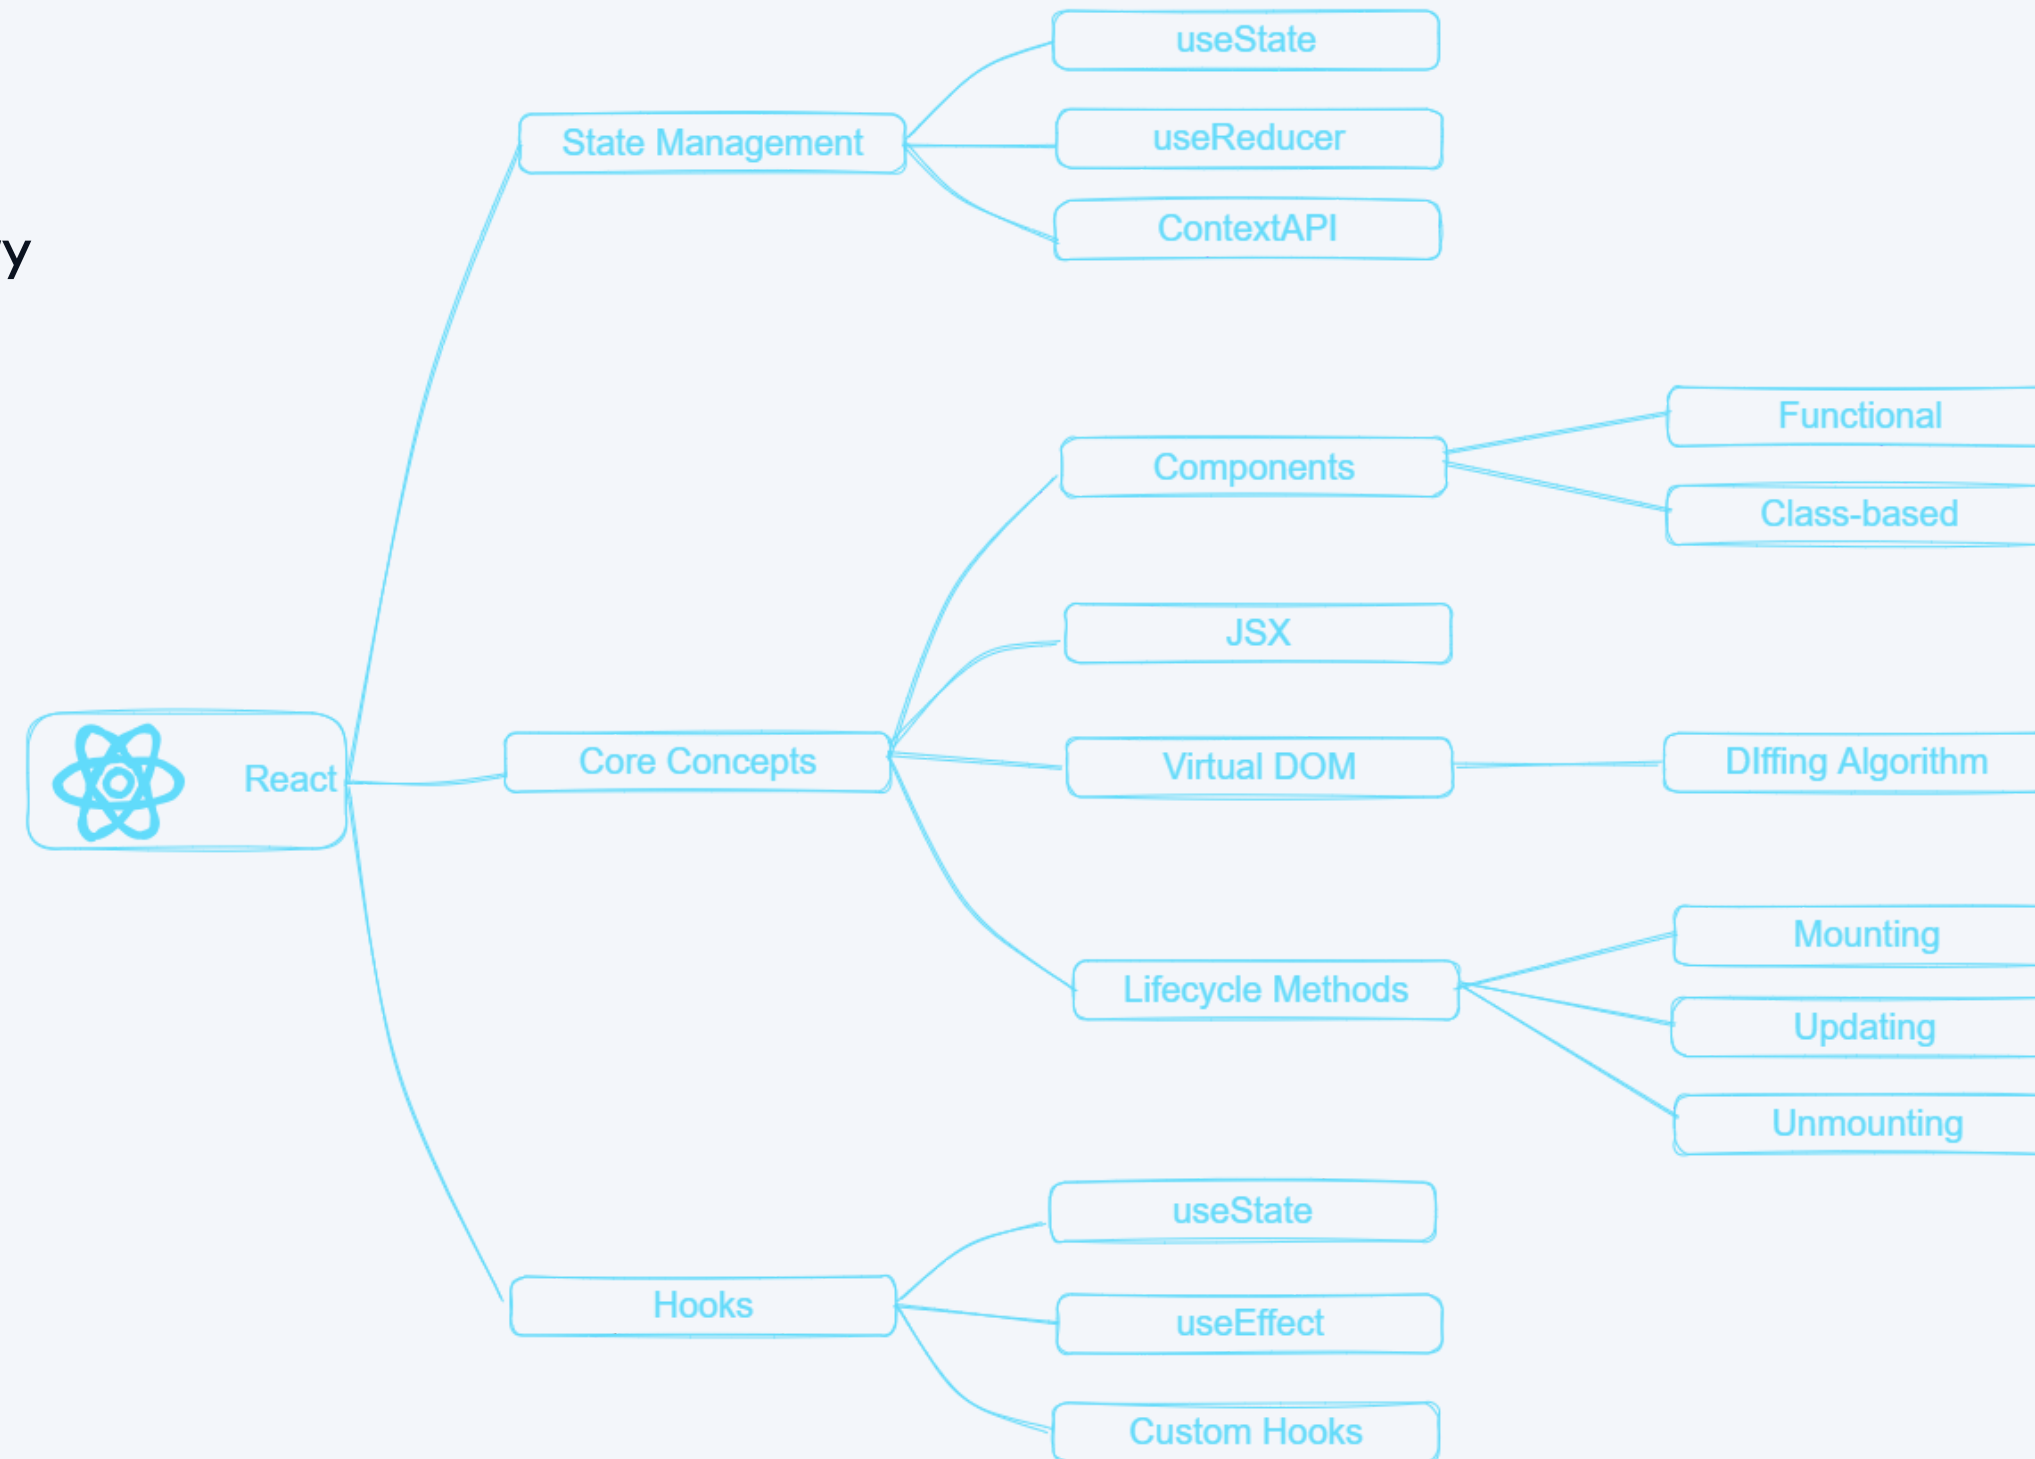
Tradeoff: temporary data inconsistency during failover; failover spikes when nodes handle increased workload



UI Library: React

Source Code: JavaScript

Most Popular Front-End Library



Ivan Manzhula

MSE | 2024

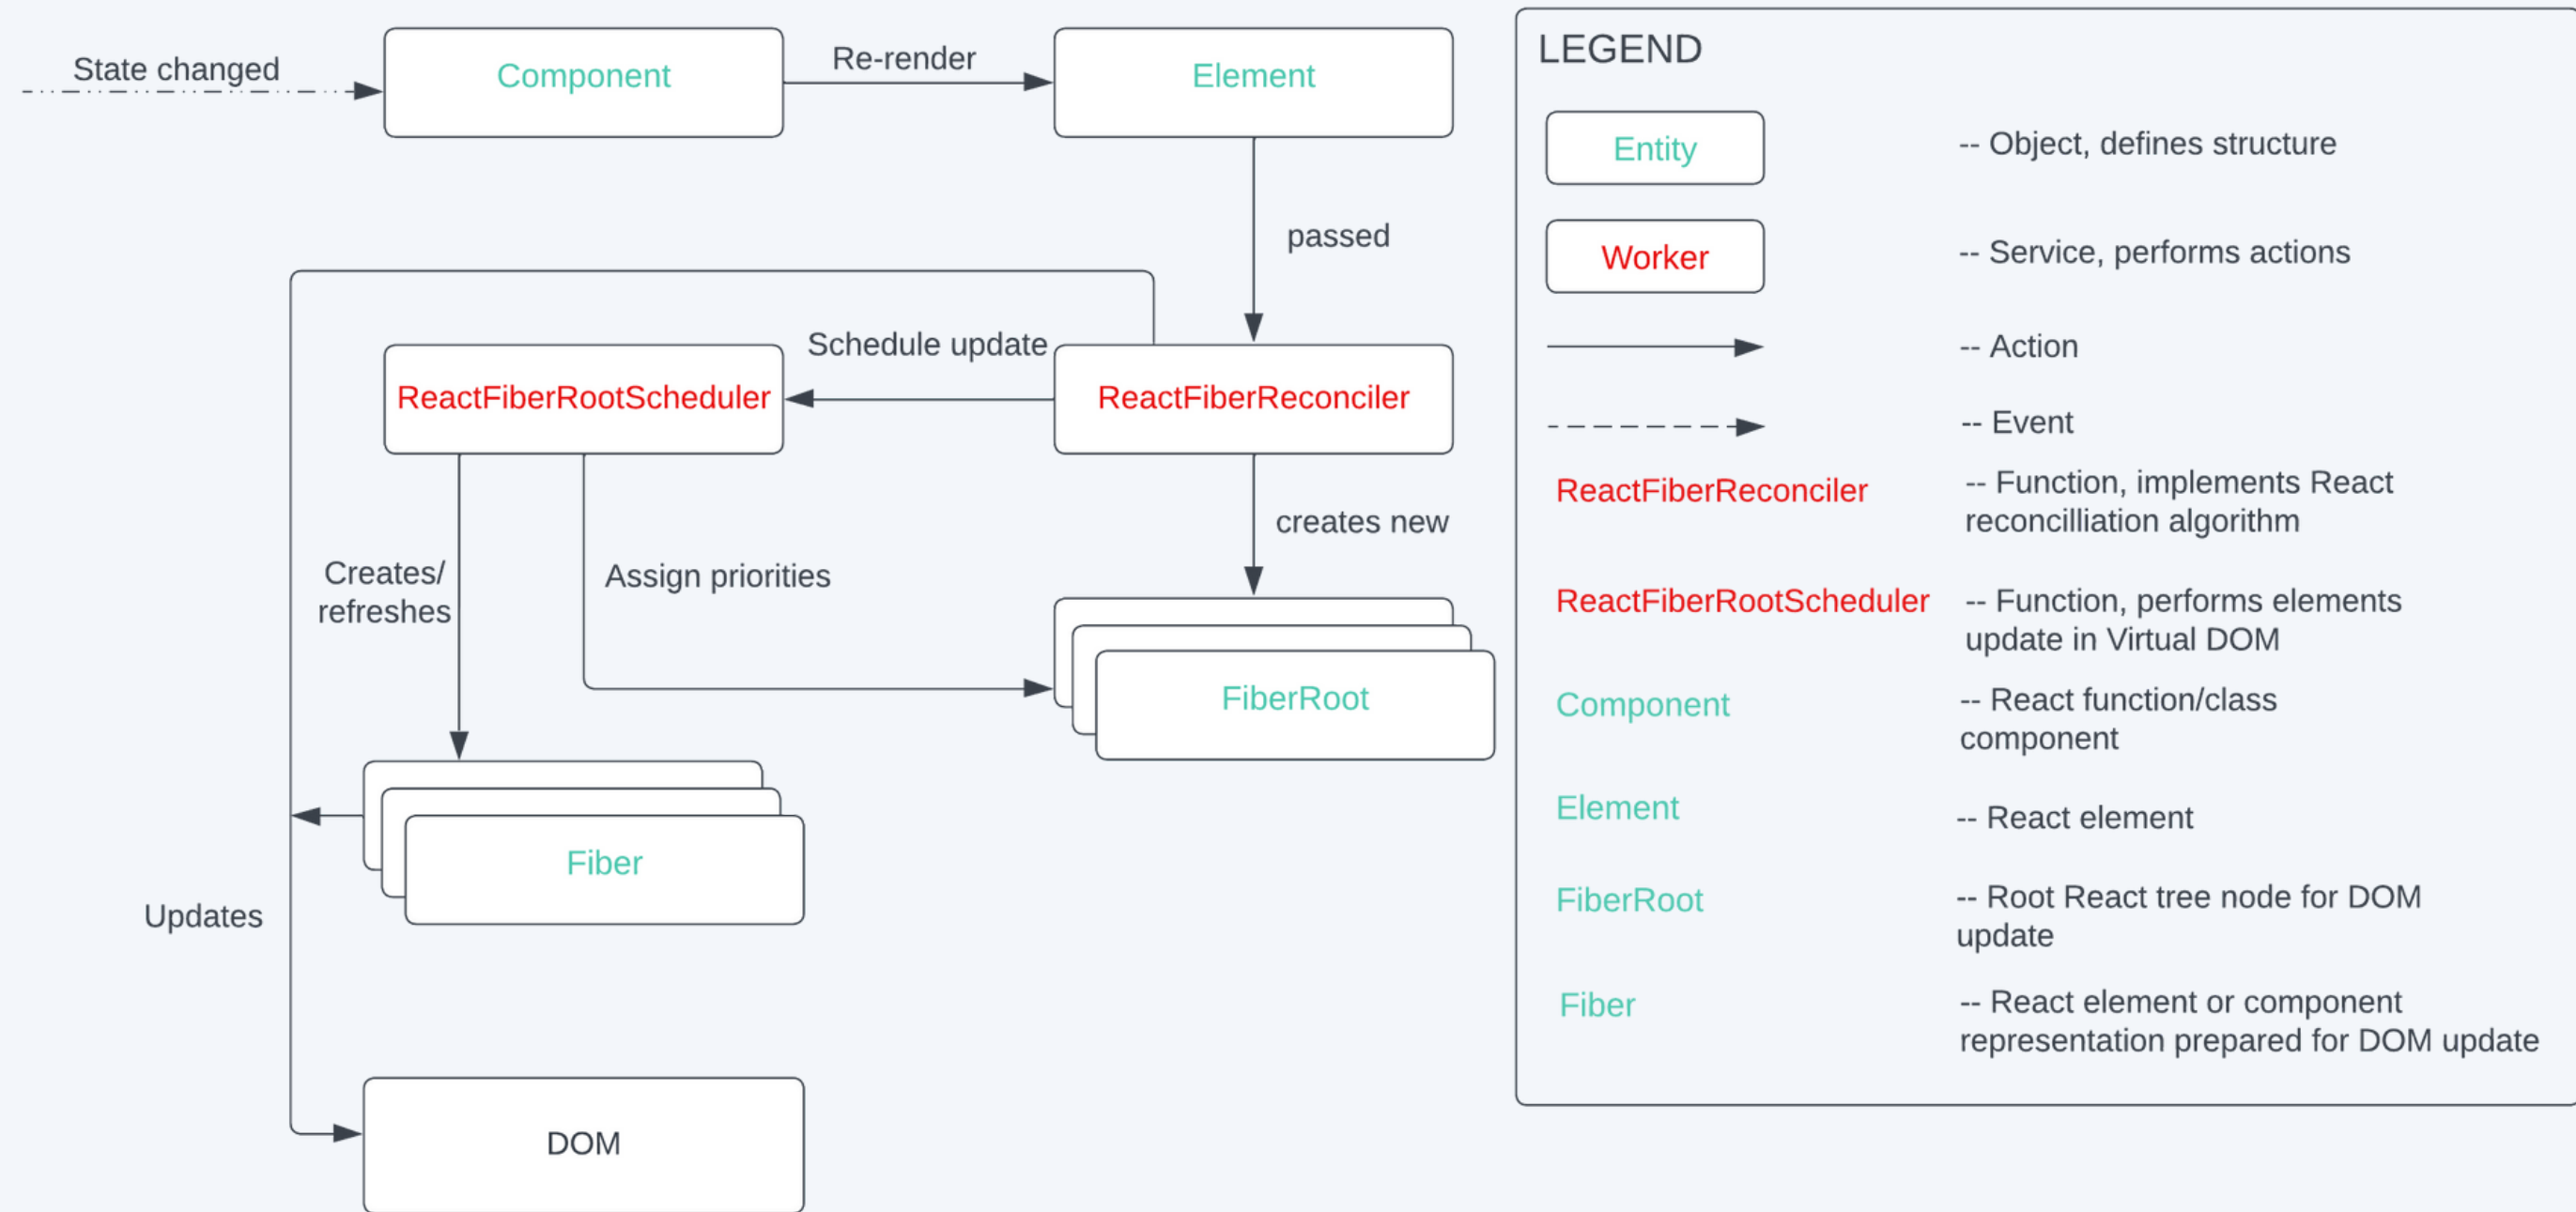
Performance: Virtual DOM

React uses a virtual DOM, which is an in-memory representation of the UI

When the state changes, React efficiently updates the virtual DOM and then re-renders only the parts of the UI that have changed

Decision: scheduling and updates based on Unit of Work

Tradeoff: Virtual DOM maintenance - requires more resources



American University Kyiv

Thank You

For Your Attention

Ivan Manzhula

MSE | 2024