Where Apache Kafka meets OpenShift

oc whoami

```
$ oc whoami
rmarting, jromanmartin
$ oc describe user rmarting jromanmartin
Name: Jose Roman Martin Gil
Created: 43 years ago
Labels: father, husband, friend, runner, curious, red hatter,
developer (in any order)
Annotations: Principal Middleware Architect @ Red Hat
Identities:
     GitHub:   https://github.com/rmarting
     LinkedIn: https://www.linkedin.com/in/jromanmartin/
```

@jromanmartin

Spoiler!!

It is a true love story!

What is Apache Kafka?

Apache Kafka is a distributed system designed for streams. It is built to be an horizontally-scalable, fault-tolerant, commit log, and allows distributed data streams and stream processing applications.

What is Apache Kafka?

- Developed at LinkedIn back in 2010, open sourced in 2011
- Designed to be fast, scalable, durable and available
- Distributed by nature
- Data partitioning (sharding)
- High throughput / low latency
- Ability to handle huge number of consumers

Main Use Cases

Messaging

Replacement of traditional message broker, has better throughput, built-in partitioning, replication, and fault-tolerance. Provides strong durability.

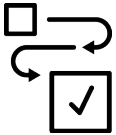

Website Activity Tracker

Rebuild user activity tracking pipeline as a set of real-time publish-subscribe feeds. Activity is published to central topics with one topic per activity type.

Metrics

Aggregation of statistics from distributed applications to produce centralized feeds of operational data.

Main Use Cases

Stream Processing

Enables continuous, real-time applications built to react to, process, or transform streams.

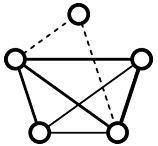

Data Integration

Captures streams of events or data changes and feeds these to other data systems.

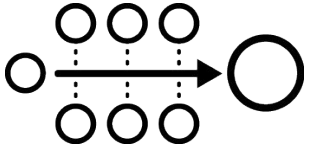

Log Aggregation

Abstracts details of files an gives event data as stream of messages. Offers good performance, stronger durability guarantees due to replication.

Apache Kafka Core Components

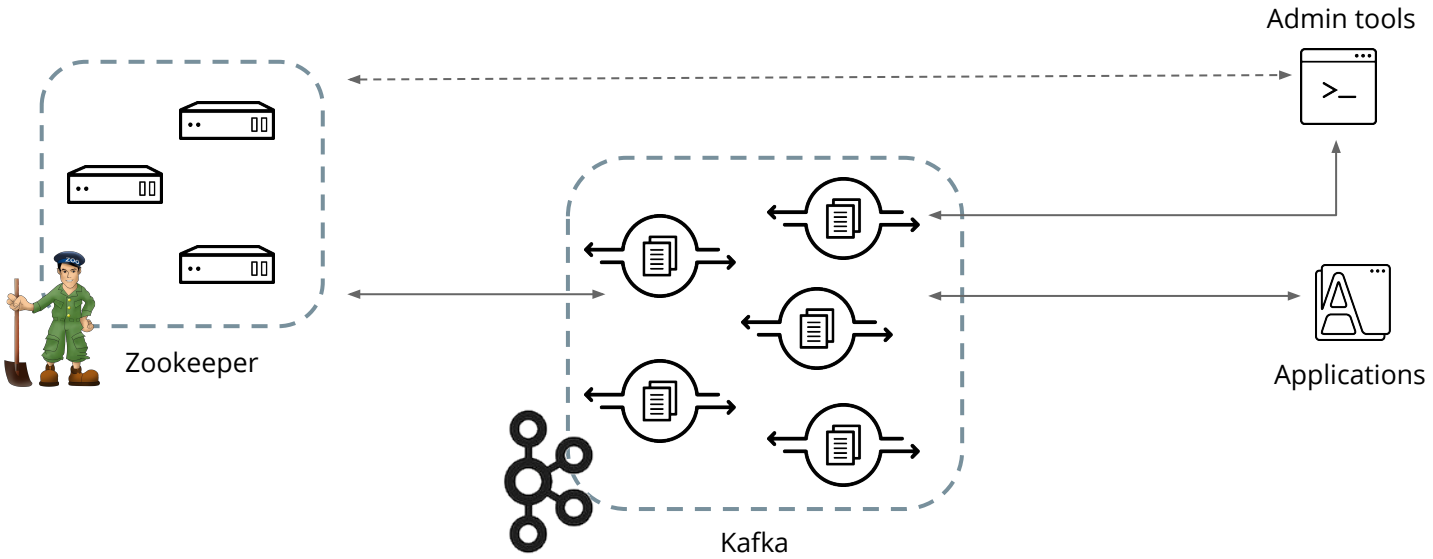

STRIMZI

Topics and Partitions

- Messages/records are sent to/received from topic
  - Topics are split into one or more partitions
  - All actual work is done on partition level, topic is just a virtual object
- Each message is writing only into one selected partition
  - Partitioning is usually done based on the message key
  - Messaging ordering within the partition is fixed
- Retention policies
  - Based on size/message age
  - Compacted based on message key
- Replication
  - Each partition can exist in one or more backup copies to achieve HA

Topics and Partitions - Producers

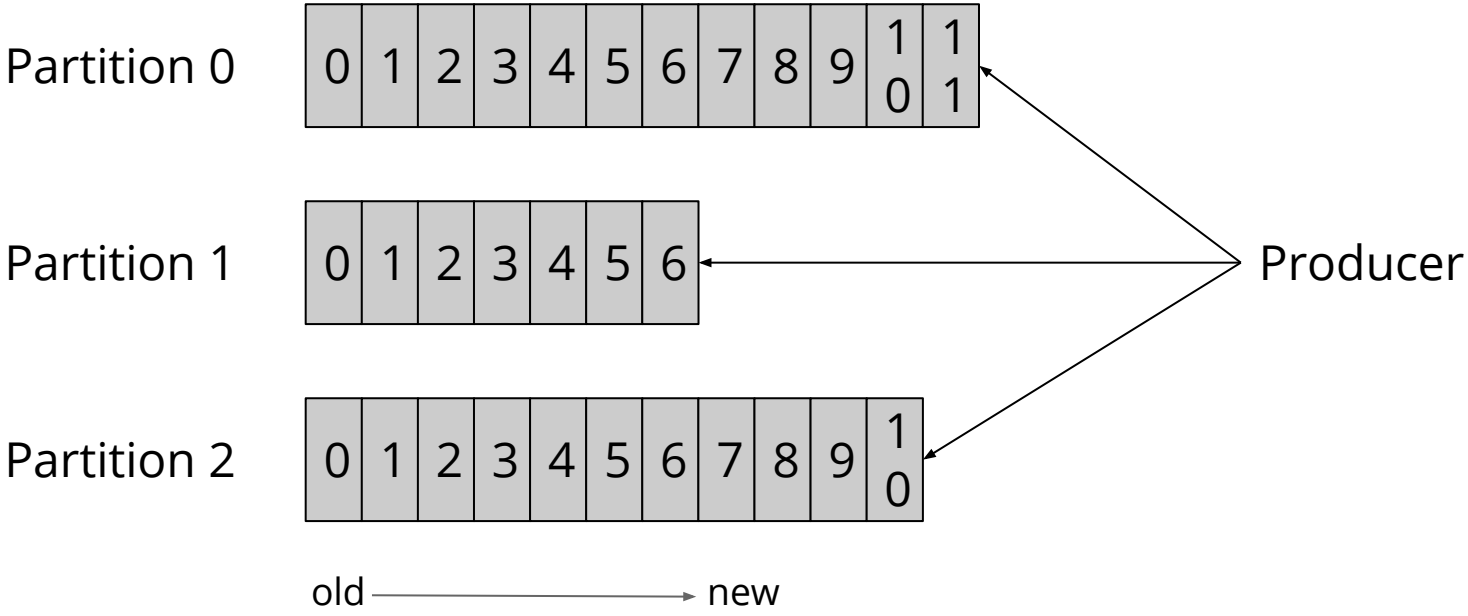

Topics and Partitions - Consumers

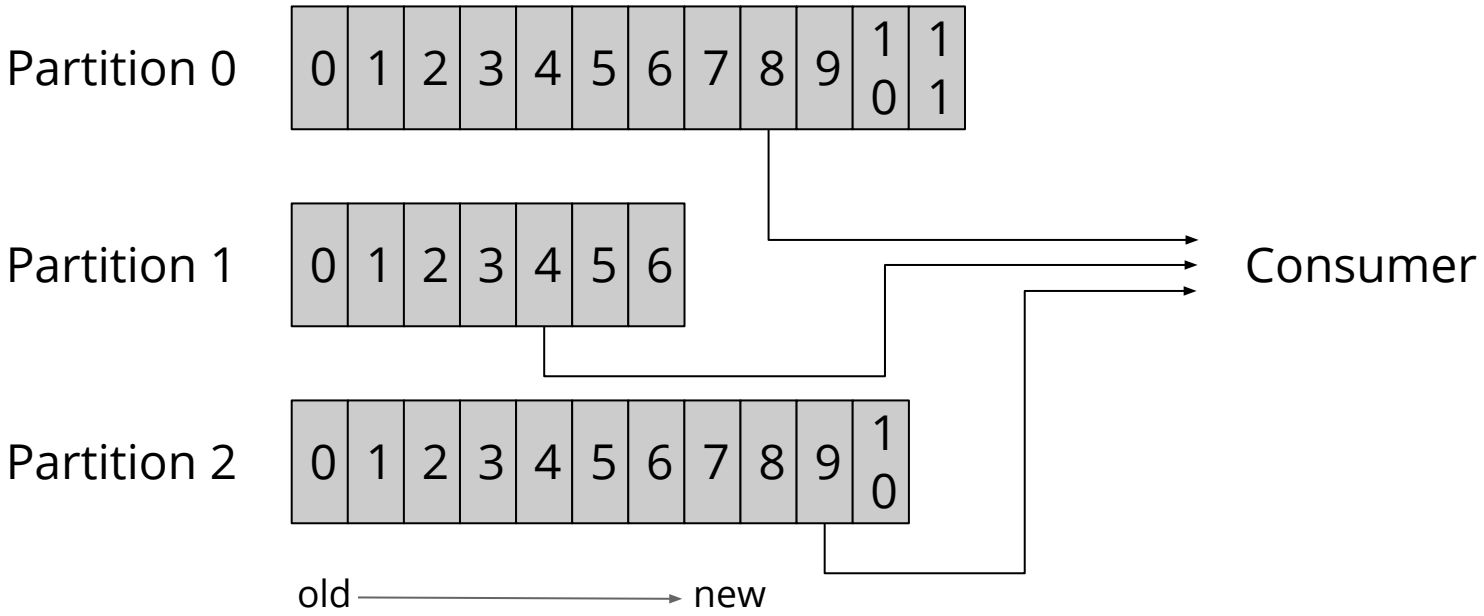

High Availability

Replication from leaders to followers

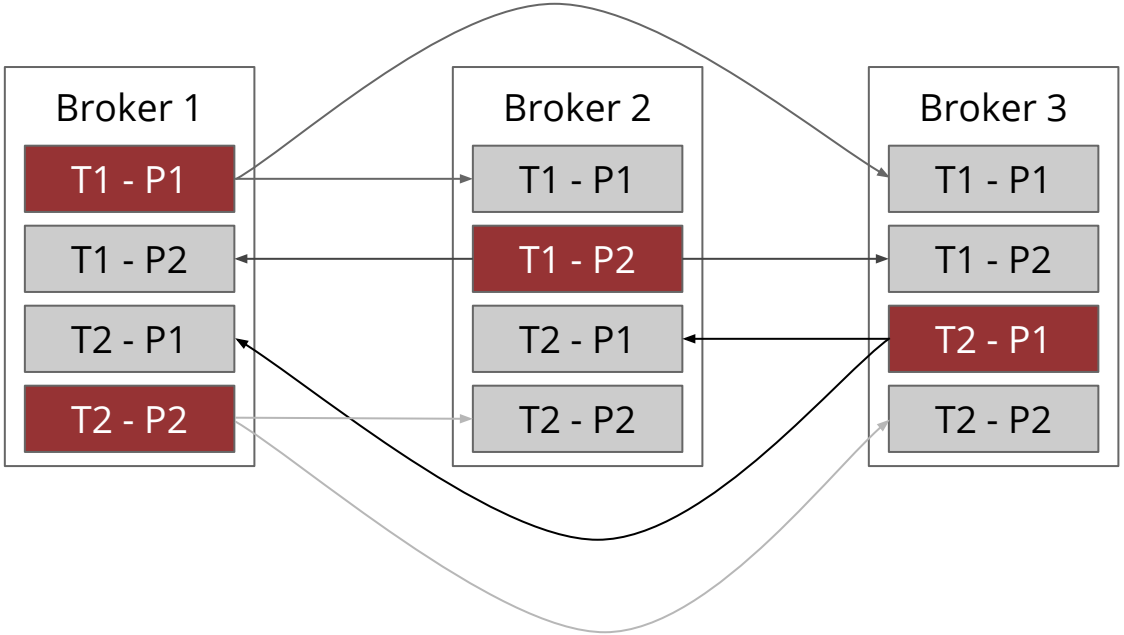

STRIMZI

High Availability

New leaders on different brokers

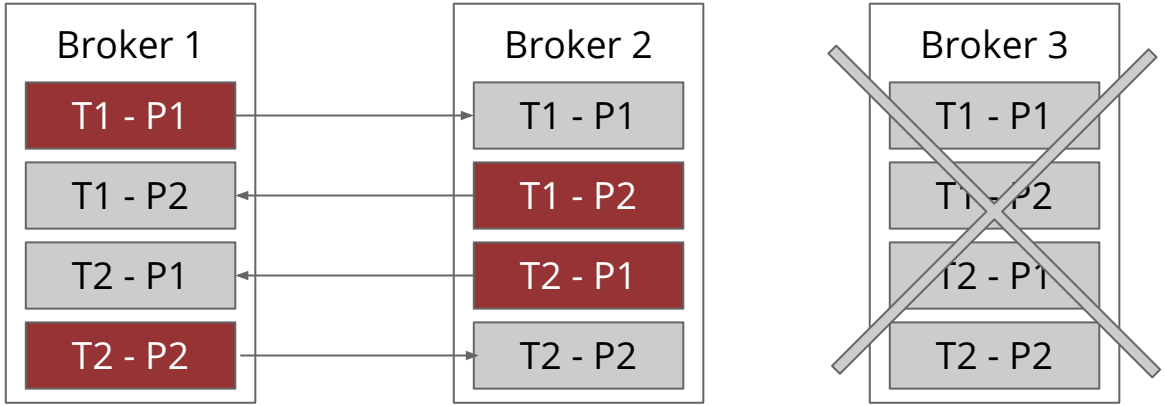

High Availability

Reading and writing to leaders

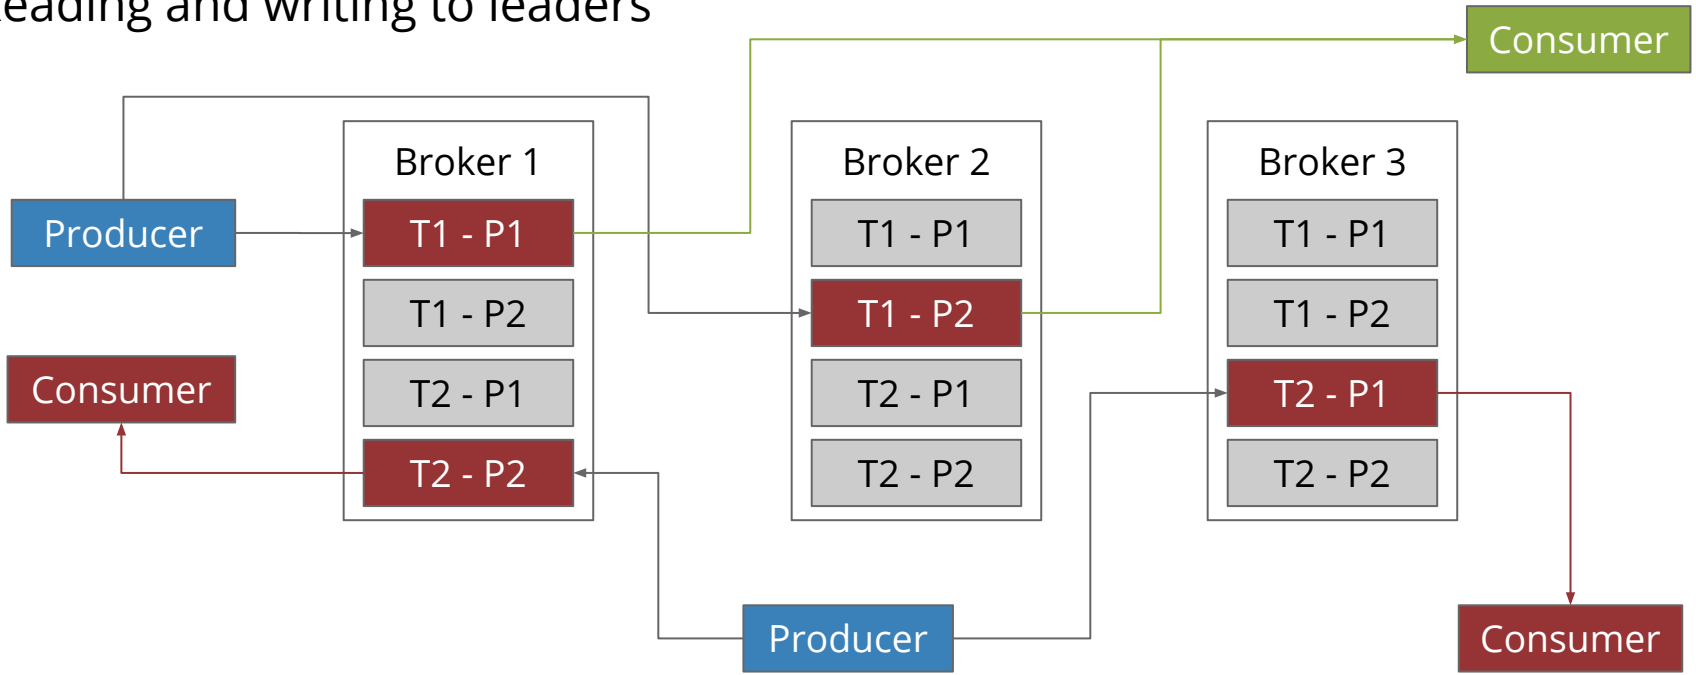

STRIMZI

Consumer Groups

Partitions Assignment

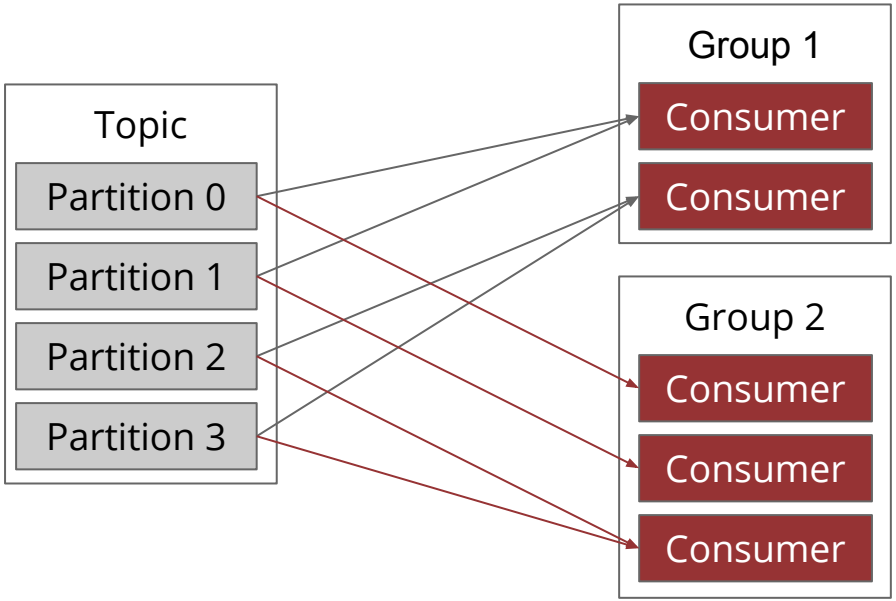

Consumer Groups

Rebalancing

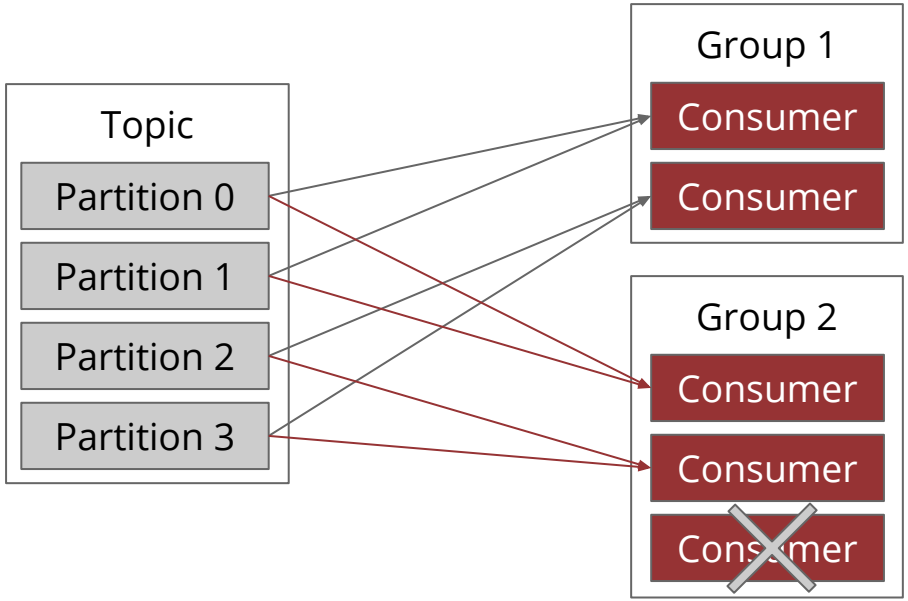

Consumer Groups

Max parallelism & Idle Consumer

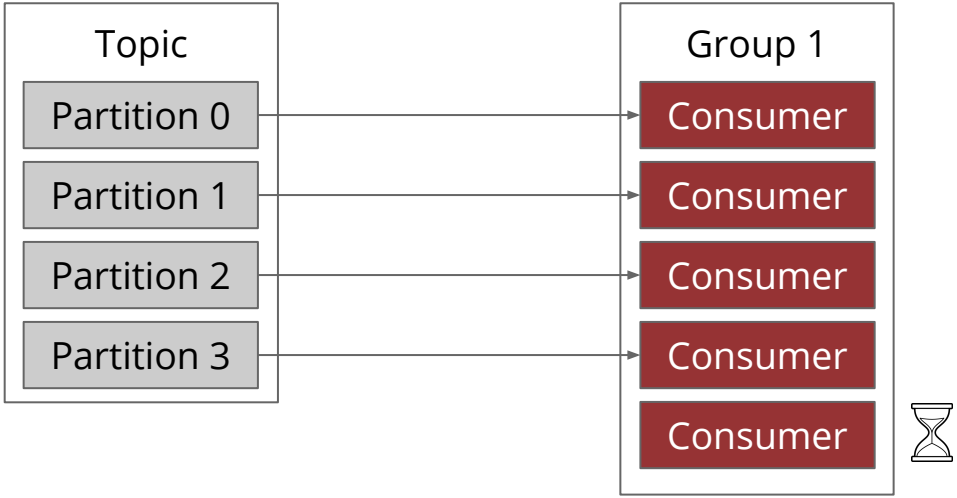

Applications

- They are really “smart” (unlike “traditional” messaging)
- Configured with a “bootstrap servers” list for fetching first metadata
  - Where are interested topics? Connect to broker which holds partition leaders
  - Producer specifies destination partition
  - Consumer handles messages offsets to read
  - If error happens, refresh metadata (something is changed in the cluster)
- Batching on producing/consuming

What is OpenShift?

A distribution of Kubernetes optimized for continuous application development and multi-tenant deployment.

OpenShift enables Developer Productivity

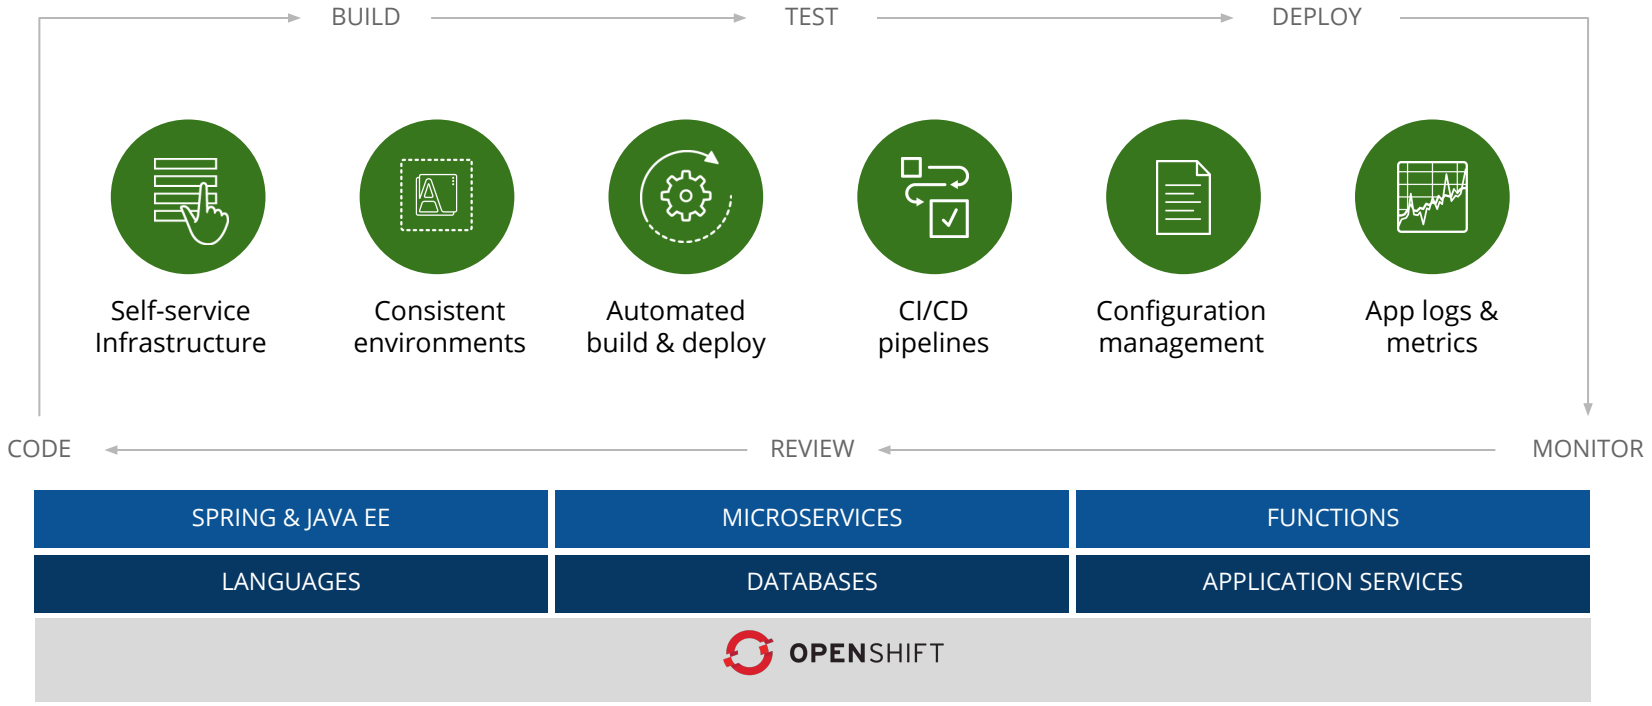

STRIMZI

But ...

- Apache Kafka is **stateful** which means we require:
  - A stable broker identity
  - A way for the brokers to discover each other on the network
  - Durable broker state
  - The ability to recover broker state after a failure
- All the above are true for Zookeeper as well and ...
  - Each node has the configuration of the others
  - To have notes able to communicate each others
- Accessing Kafka is not so simple
- OpenShift provides several services ...

Who is going to help me?

Strimzi is our hero!

Strimzi is a set of enabling services that allow Apache Kafka to work in OpenShift as a first class citizen, be installed easily and configured and managed simply.

What is Strimzi?

- Simplifies Apache Kafka deployment on OpenShift
- Uses the OpenShift native mechanisms for:
  - Provisioning Kafka clusters
  - Managing topics and users
- Provides:
  - Linux Container images for running Apache Kafka and Zookeeper
  - Tooling for managing and configuring Apache Kafka clusters, topics and users
- Follows the Kubernetes Operator Model
- Available from ***OperatorHub.io***

OperatorHub.io

- Home for the Kubernetes community to share Operators
- Many different Operators available ...
  - Big Data, Cloud Provider, Database, Logging & Tracing, Monitoring, ...
- ... And growing up !!!

Kubernetes Operators

- Application-specific controller is used to create, configure and manage other complex application:
  - The controller contains specific domain/application knowledge
  - Usually used for stateful applications (databases, ...) which are non-trivial to operate on Kubernetes/OpenShift
- Controller operates based on input from Config Maps or Custom Resource Definitions
  - User describes the desired state
  - Controller applies this state to the application

Operator Pattern

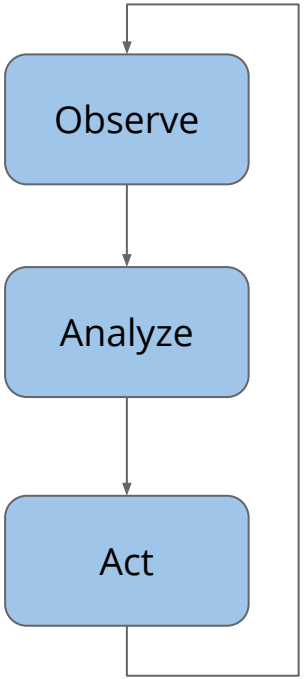

- Monitor the current state of the application
- Compare the actual state to the desire state
- Resolve any differences between actual and desired state

Strimzi Operators

Kafka Operator

Topic Operator

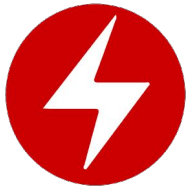

User Operator

Strimzi Operators in a glance

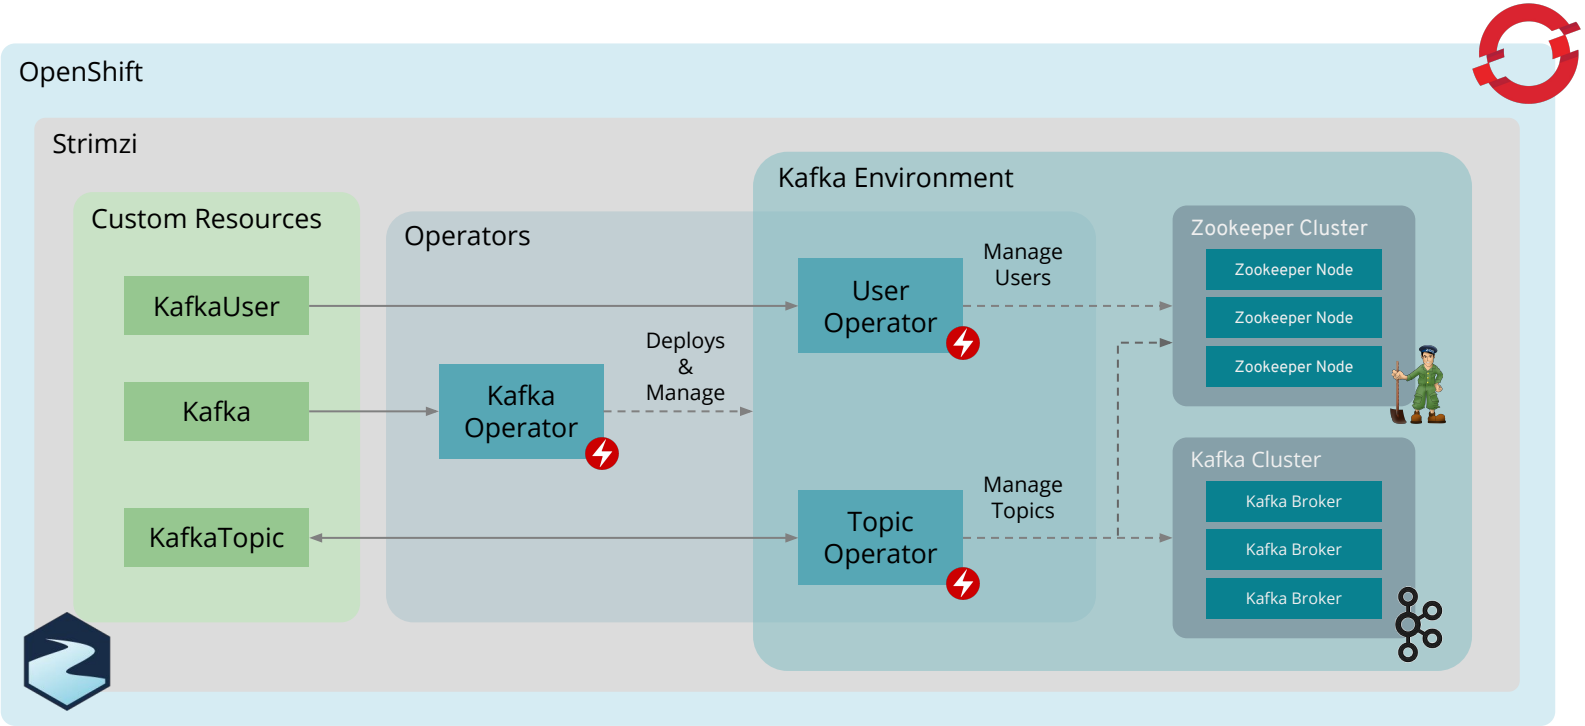

Kafka Operator

- Creating and managing Apache Kafka clusters
  - Number of Zookeeper, Kafka Brokers and Kafka Connect nodes
  - Configuration of Kafka and Kafka Connect
- Responsible for
  - Deployment
  - Storage
  - Scale up / Scale Down
  - Metrics
  - Healthchecks
- One Cluster Operator can manage several clusters in parallel
  - Can cover one or more projects

Topic Operator

- Creating and managing Kafka topics
- Some Kafka components (Streams, Connect) often create their own topics
  - Bi-directional synchronization
  - Changes done directly in Kafka/Zookeeper are applied to Custom Resource Definitions
  - Changes done in Custom Resource Definitions are applied to Kafka topics
- Topic Operator solves this by using 3-way diff
  - Our own Zookeeper-based store
  - Apache Kafka / Zookeeper
  - Custom Resource Definitions

User Operator

- Creating and managing Kafka users
- Unlike Topic Operator, does not sync any changes from the Kafka cluster with OCP
- It is not expected that the users will be managed directly in Kafka cluster in parallel with the User operator
- User Credentials managed in a Secret
- Manages authorization rules

Strimzi Custom Resource Definitions

- Extensions of the Kubernetes API to store a collection of API objects
- Allows to introduce an own API into a cluster
- Strimzi CRD
  - Kafka: Definition of a Kafka Cluster
  - KafkaTopic: Definition of a Kafka Topic
  - KafkaUser: Definition of a User

Kafka CRD

```yaml
apiVersion: kafka.strimzi.io/v1alpha1
kind: Kafka
metadata:
  name: meetup-cluster
spec:
  kafka:
    replicas: 2
    # listeners, config, metrics, jvmOptions, resources
    storage:
      type: persistent-claim
      size: 12Gi
  zookeeper:
    replicas: 3
    # jvmOptions, resources
    storage:
      type: persistent-claim
      size: 5Gi
  entityOperator:
    # topicOperator, userOperator, tlsSidecar
```

STRIMZI

KafkaTopic & KafkaUser CRD

```yaml
apiVersion: kafka.strimzi.io/v1alpha1
kind: KafkaTopic
metadata:
  name: meetings
  labels:
    strimzi.io/cluster: meetup-cluster
spec:
  partitions: 10
  replicas: 2
```

```yaml
apiVersion: kafka.strimzi.io/v1alpha1
kind: KafkaUser
metadata:
  name: roman
  labels:
    strimzi.io/cluster: meetup-cluster
spec:
  authentication:
    type: tls
  authorization:
    type: simple
    acls:
      - resource:
          type: topic
          name: meetings
          patternType: literal
        operation: Read
```

STRIMZI

Strimzi is much more ...

STRIMZI

Demo Time !!!

Get Started !!!

- Strimzi:
  - https://strimzi.io
  - https://github.com/strimzi
- OpenShift:
  - https://www.openshift.com
  - http://learn.openshift.com
- Apache Kafka:
  - https://kafka.apache.org
- OperatorHub.io:
  - https://www.operatorhub.io

Questions?

Thank you!