Camel Kafka Connector

Metamorphosis: when Kafka meets Camel

About

- **Jakub Scholz (twitter: @scholzj)**
  - Principal Software Engineer in the Red Hat Messaging and IoT team with a long-term experience in messaging, mainly on Apache Kafka. One of the core maintainers of the Strimzi project, which delivers several operators and tools for running Apache Kafka on Kubernetes and OpenShift.
- **Otavio Piske (twitter: @otavio021)**
  - Senior Software Engineer at RedHat's Fuse Team. Involved with messaging and integration technologies for the last few years. Currently he is contributing to the Camel Kafka Connector project.

Agenda

- Introduction
- The basics: what is Camel, Kafka and Kafka Connect
- The Camel Kafka Connector Project
- Demos
- Closing comments

Introduction

Integrating

Systems

What is Apache Camel?

- The Swiss knife of integration
    - One of the largest and most active Apache projects
- Open source **integration framework**
- **Flexible**: data formats, products, protocols and tools

Camel Concepts

- Routes and Enterprise Integration Patterns (EIP)
  - Models for designing and developing integration solutions
    - Content-based router, Splitter, Aggregator, etc.
- Camel components:
  - Primary extension point: allow talking to external systems
  - Products and tools: Salesforce, AWS S3 Storage Service,, etc
  - Protocols: AMQP, HTTP, SCP, etc

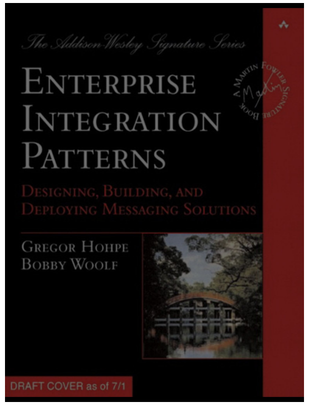
The Addison-Wesley Signature Series
Enterprise Integration Patterns
Designing, Building, and Deploying Messaging Solutions
Gregor Hohpe
Bobby Woolf
DRAFT COVER as of 7/1

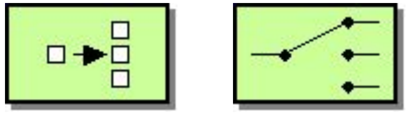

Demo: Apache Camel

JMS-to-Kafka

Broker

What is Apache Kafka?

- Most people know Kafka as
  - a **distributed streaming platform**
  - or a **pub/sub messaging broker**
- But it is more than that, it is an **ecosystem**
  - Multiple components part of Apache Kafka
  - Many 3rd party integrations
  - One of them is **Kafka Connect**

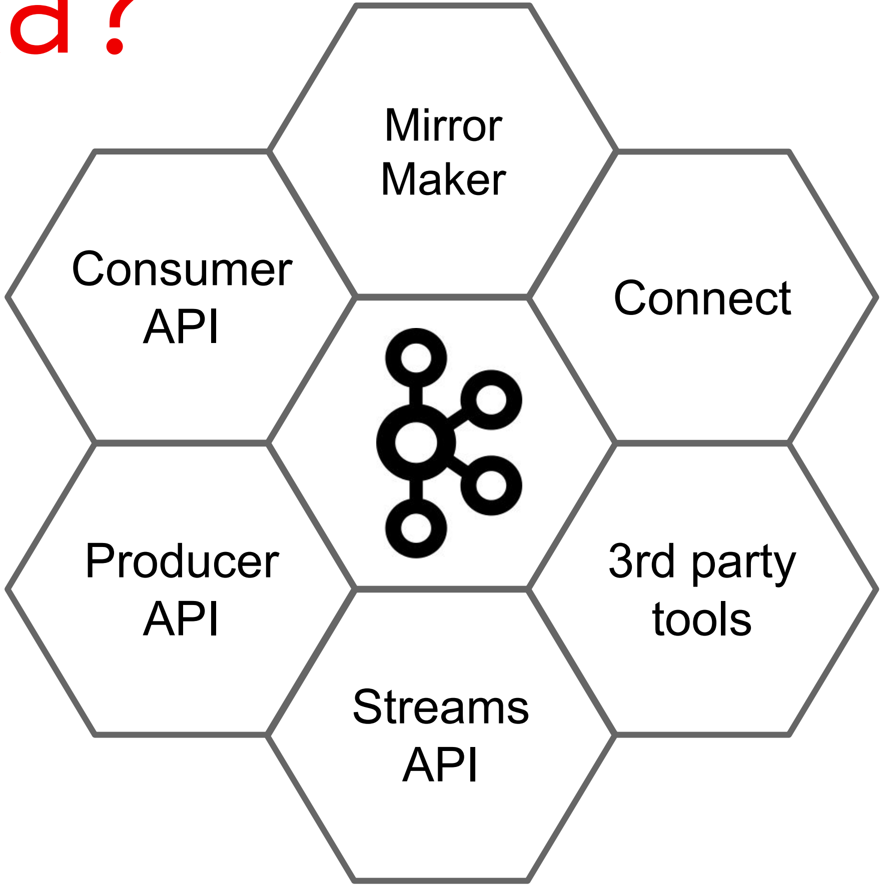

Wait, what is Kafka Connect?

- It is a **framework** that helps you to integrate Kafka with other systems
  - A user can define **source** and **sink** connectors to stream data in, respectively out, of Kafka brokers
  - Connectors are **plugable** - you can use one of the many connectors available or write your own
- **Distributed** and **scalable** by default
- Automatic **offset** management
- Simple transformations
- Streaming / batch integration

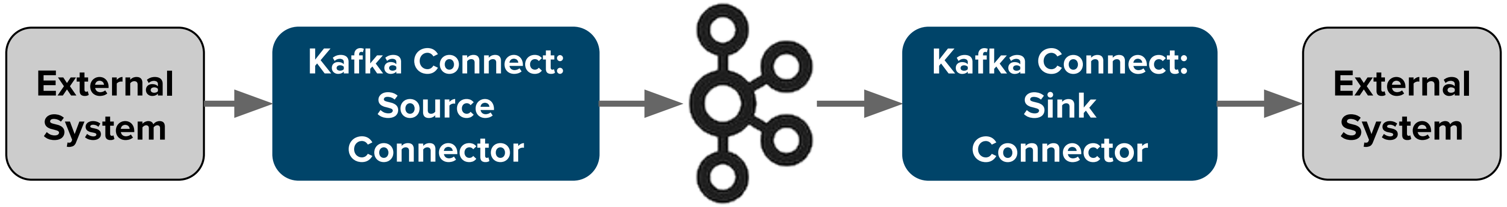
External System
Kafka Connect: Source Connector
Kafka Connect: Sink Connector
External System

Key Kafka Connect concepts

- **Connector:** generally refers to a source/sink implemented using Kafka Connect API
- **SinkConnector / SourceConnector:** responsible for setting up the connector and partitioning the work by creating SinkTask / SourceTask.
- **SinkTask / SourceTask:** responsible to handle the actual work.
- **Key / value converter:** a component able to convert the key or value of a kafka message from a format to another.
- **Transformer:** a component able to manipulate a message.

Standalone mode

- Only a single Connect worker
- Configuration and offsets stored locally
- Suitable as a local adapter

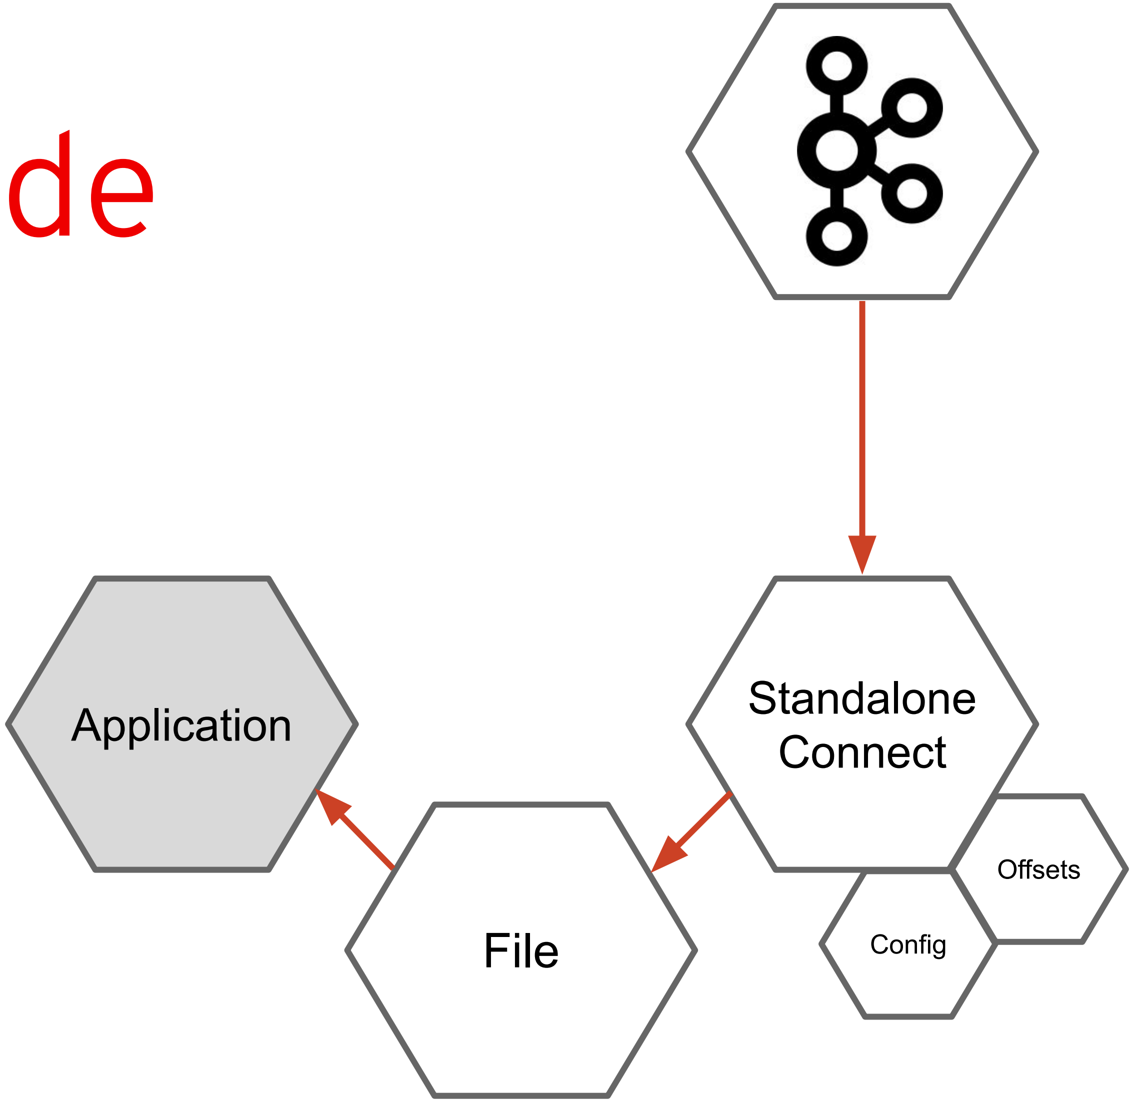

Distributed mode

- Multiple workers to distribute the work
- Configuration and offsets stored in Kafka
- One of the workers is a leader
- Balancing tasks between workers

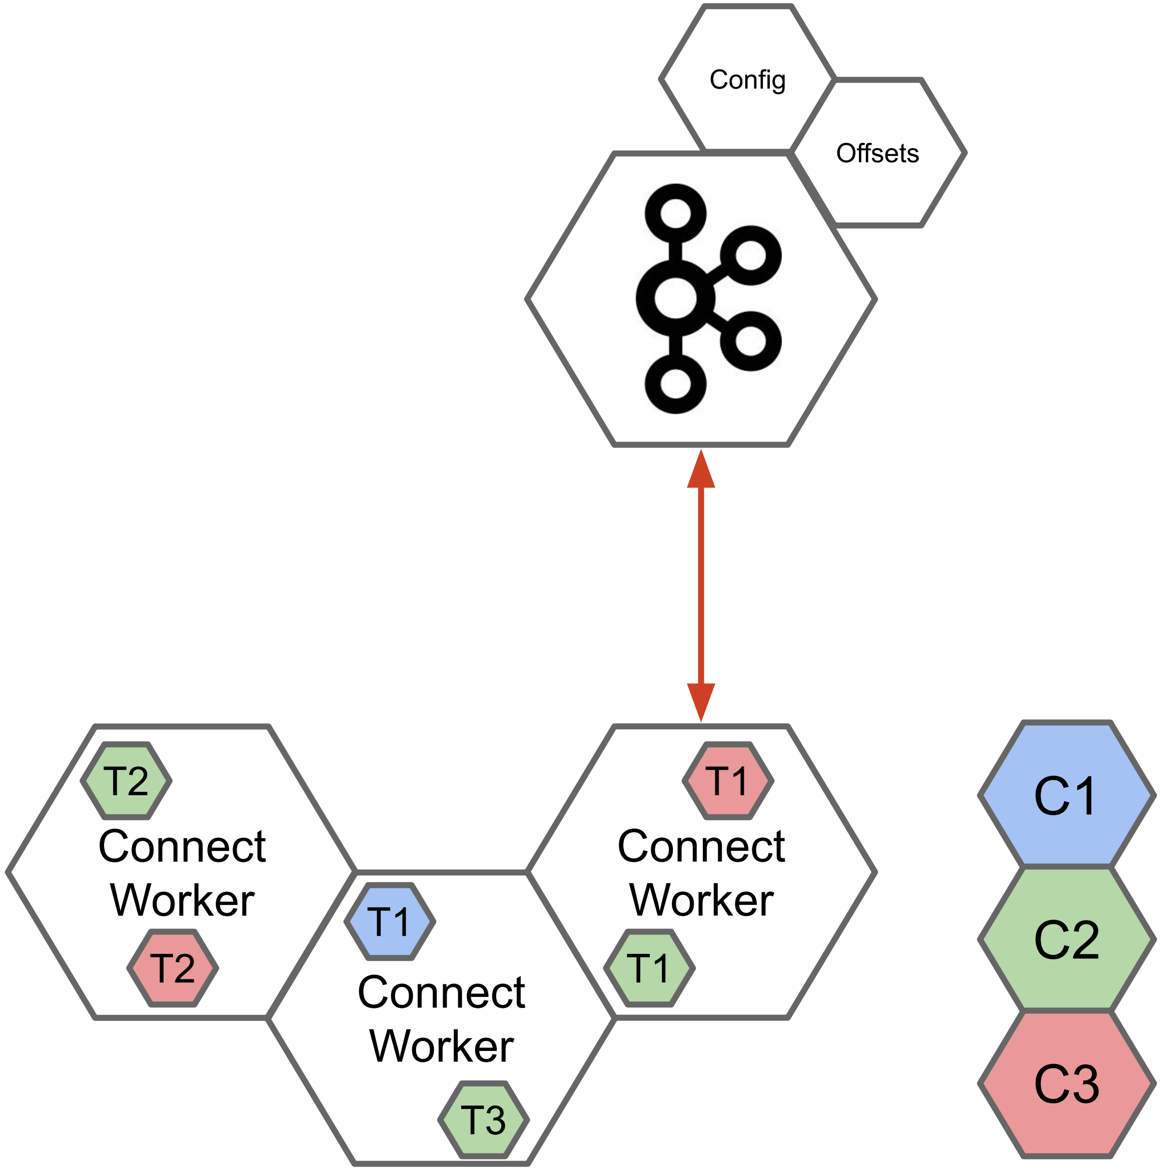

Metamorphosis

When Kafka

meets the

Camel

What is Camel Kafka Connector?

- A Kafka Connector built on top of Apache Camel
- Started as an internal proof-of-concept
  - A sub-project of the Apache Camel
  - Donated by Red Hat to the ASF on December 2019
- Reuses in a **simple way** most of the Camel components as Kafka sink and sources

Available connectors

- Initial focus on **selected** set of components:

Amazon AWS Kinesis	Cassandra CQL	JMS
Amazon AWS S3	ElasticSearch	Telegram
Amazon AWS SNS	File	Syslog
Amazon AWS SQS	HTTP	

* Sink/Source availability may depend on the source/sink system and type of the connector.

Project status

- Well received by the community
- Initial support for a limited set of components
    - Extend to support all/most of/a selected list of 250+ camel components.
- Working towards a stable release
- Open Issues and TODO

Demo: Camel Kafka Connector

Standalone

integration

demo

Red Hat

Demo: Camel Kafka Connector

Integrating systems with Strimzi on OpenShift

Conclusion

- Combines the features of two great Apache projects
    - Opens the doors for many new connectors
    - Experience and maturity of the Apache Camel project with enterprise integration
    - Simplicity and distributed nature of Kafka Connect
- Existing Kafka Connect users get a lot of new options and integrations
- Existing Camel users get jump-start into the Kafka world

Slides

http://jsch.cz/metamorphosis

Demo repositories

- https://github.com/orpiske/ckc-demos
- https://github.com/scholzj/devconf-2020-supersonic-subatomic-apache-kafka

Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.

in linkedin.com/company/red-hat

youtube.com/user/RedHatVideos

f facebook.com/redhatinc

twitter.com/RedHat

Red Hat