



Kafka On YARN (KOYA)

An Open Source Initiative to Integrate Kafka & YARN

Goals

- **Use capabilities of YARN for Kafka broker management**
- **Automate management tasks (alternative to Kafka command line utilities)**
- **Automate broker recovery**
- **Make it easy to deploy and configure Kafka clusters**

Features

- **YARN application master with full HA support**
 - Monitor brokers, server health status
 - Detect failed brokers.
 - Restart failed brokers.
 - Deploy replacement broker to new host on host failure.
- **Policy for dealing with repeated broker failure on same host.**
- **Sticky location – allocate containers on previous host when possible (based on broker information from ZooKeeper)**
- **Create on-demand Kafka cluster**
- **Ability to push Kafka code to new hosts (for new broker or broker recovery on new host)**
- **Admin web service to view broker status, topics, partitions, leader/replica information, consumer info etc.) and do configuration changes**

Architecture

