Kafka Kloner

Michael Natkovich

Kafka Kloner

A Dynamic High-Speed Inter-Cluster Kafka Replicator

Propane replicates around 150 billion events per day cross colo across its clusters with an average latency of less than 2 seconds.

YAHOO!

Motivation

- Demand for Yahoo Streaming in multiple data centers
- Kafka MirrorMaker operability issues

Finally, we borrowed the MirrorMaker approach and implemented replication using Storm.

MirrorMaker vs Kloner

	MirrorMaker	Kloner
Performance	High	High
Operability	Low	High
Fault Tolerance	Low	High
Topic Management	Low	High
Monitoring	Low	High

Overview

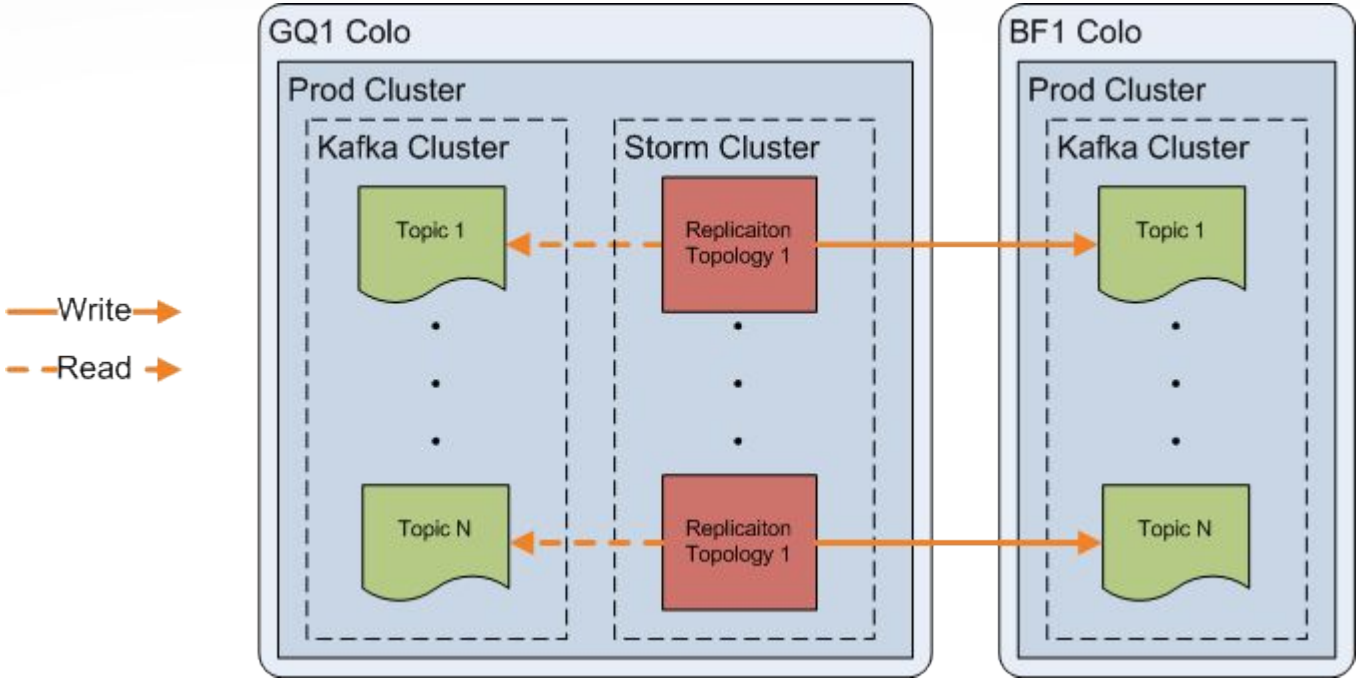

YAHOO!

Features

Design

- Storm based approach
- One topology per topic being cloned
- One worker per topic partition
- One kafka consumer + producer per worker

Core Logic

- Same as MirrorMaker
  - Read; Commit Offsets; Send Message;

Features

Manager

- Runs based on a DSL input with list of topics and their destinations
- Polls Storm for list of topics already being replicated and number of workers
- Polls Kafka clusters for topic metadata such as partitions per topic

  Then,

  1. Starts replication by starting new topologies
  2. Scales replication by rebalancing topologies (Restarts with updated worker count)
  3. Stops replication by stopping topologies

- The Manager can be a cron or simply invoked after DSL update

Features

Security

- Cross colo data transfers are encrypted
- Kafka SSL + Yahoo Certificate Authority.

Monitoring & Alerting

- Yahoo Monitoring Framework
- Kloner msg rates
- Kafka msg rates
- Kloner consumer lags

Numbers

Clusters	2
Topics	15
Msgs/sec	1 million
Bytes/sec	1.5 GB
Replication Lag	< 2 seconds

Limitations

- No exact replication
  - Offsets differ between clusters
- No replay or spooling
  - No message level failure handling
  - But we are tracking number of failed events
  - There is rate limiting

QUESTIONS