

Audit Trail Proposal

Motivation

As a system Kafka has a fairly simple description of correctness: all messages sent to a topic should be delivered to each consumer group consuming that topic. The goal of this facility is to continuously monitor that this is occurring. A facility similar to this exists at LinkedIn.

Audit Messages

Clients that support auditing do so by periodically reporting a summary of the messages they have sent or received. These messages are then aggregated to check that all messages sent were received. Auditing is an optional facility--clients that don't support it simply will not be able to use the audit monitoring tool.

Participants report by offset interval. A simple scheme would be to have all participants report every 10k offsets. The downside of this is that some topics may take a very long time to receive 10000 messages per partition

The format of this message is json in the form

```
{
  "guid": "51656274-a86a-4dff-b824-8e8e20a6348f", // a unique identifier for this audit message
  "time": 1348335709389, // time at which the audit began
  "host": "host-name", // the host name of the machine sending the audit
  "facility": "datacenter-name", // the data center of the machine sending the audit
  "client": "application-name", // the client application name
  "tier": "producer", // the name of the logical tier (producer, broker, or other consumer group)
  "offset_begin": 22157000, // the beginning offset this covers
  "offset_end": 22239000, // the ending offset this covers
  "summary": 1234234 // a hash summary of the messages seen.
}
```

General Cluster UI

There are several uses for a web UI related to Kafka, we may as well fold these into a single tool. We did a code dump of the current audit tool as part of KAFKA-260, but if we are going to expand the functionality it might make sense to clean up that code, and move it into Scala.

In addition to audit monitoring the following would be useful:

1. Cluster information: Which machines have which topics and partitions? Which are leaders?
2. Administrative commands: Add topic, move partitions, etc.

A facility I would hesitate to add would be general monitoring and timeseries graphs as most people have a way to do that and it is perhaps better to integrate with those.