

Restful API Proposal

I am working on a proposal to satisfy the Restful Proxy project listed on the Projects page. (<http://incubator.apache.org/kafka/projects.html>)

My working knowledge of Kafka is about 3 hours old and I have actually written zero code yet, it is my job to do so for Box so I will update this as my knowledge increases. There may be some fundamental issues with this proposal at this time.

General Idea

Create a Restful API which will allow producers and consumers to communicate with a Kafka server via HTTP calls

Consumers

URL Format:

GET /messages/<topic>/<consumer>/<offset> - Retrieve the messages for a consumer, from the topic, starting with an offset

DELETE /messages/<topic>/<consumer>/<offset> - ACK consumption of message for a consumer, topic, offset

Producers

PUT /messages/<topic>

Request Body will become the message

Known Missing Items:

- I know that this does not take into account partitions.
- It would be nice to be able to implement the async producer and have an HTTP callback happen from within the RESTServer (producer specifies a URL to be requested as callback)
- If an HTTP call back wasn't available, have the service return a task ID (MD5 of the message or something), and then allow the client to query the server to see if the callback from Kafka has happened yet.