



Kafka, an Unconventional Messaging System

Jun Rao

jrao@linkedin.com




Different Use Cases

- All kinds of log data
 - user activity, software, etc
 - for both offline & online consumption
- Mostly care about
 - simplicity, throughput, scalability,

Throughput & Efficiency

- Batch API
- No caching
- Zero-copy transfer
- Stateless broker (rewindable consumption)

Performance: Producer





- <http://incubator.apache.org/kafka/>
- ~10K lines in Scala
- To do: intra-cluster replication
 - <https://issues.apache.org/jira/browse/KAFKA-50>