



Kafka

Kafka
A pub-sub messaging system
designed for web-scale
stream processing

by [John Kirklin](#)
Published on [Medium](#) by [John Kirklin](#)



Overview
Kafka is a distributed, high
throughput, persistent
pub-sub messaging system



Goals
• Persistent
• Highly performant
• Distributed model



Design principles
• Simple API
• Low overhead
• Distributed model



Current status
github
• open sourced, submitted to be Apache
incubator project
• deployed in production for a year
• Hosted at [github <https://github.com/kafka-dev/kafka>](https://github.com/kafka-dev/kafka)
• Outside contributions arriving

LinkedIn is hiring!
www.linkedin.com/company/kafka-dev



Kafka

Kafka
A pub-sub messaging system
designed for web-scale
stream processing

Kafka Github
Download Apache Kafka



Overview
Kafka is a distributed, high
throughput, persistent
pub-sub messaging system

Kafka Goals

- Persistent
- High performance
- Distributed model

What is Kafka's place in the universe?

Message Queue
(Messaging)

Kafka
goes
here.

Log aggregation
(Flume)

- Design principles**
- Simple API
 - Low overhead
 - Distributed model



ETL and Apache Hadoop



Current status
github
2014, CC BY

- open sourced, submitted to be Apache incubator project
- deployed in production for a year
- Hosted at [github](https://github.com/apache/kafka) <https://github.com/apache/kafka>
- Outside contributions arriving

LinkedIn is hiring!

www.linkedin.com/company/apache-kafka

Kafka

A pub-sub messaging system
designed for web-scale
stream processing

Neha Narkhede
Distributed Systems at LinkedIn



Neha Narkhede YOU

Senior Software Engineer at LinkedIn

San Francisco Bay Area | Information Technology and Services

Share profile

Ask for recommendations

Create your profile in another language

Neha Narkhede



Why Facebook Needs Sheryl Sandberg [businessweek.com](#) · via Baq Haidri

Mark Zuckerberg's No. 2 provides "adult supervision" at the company, trying to keep growth at an optimum level.

8 days ago · Like (1) · Comment (1) · Share · See all activity

Work With Us at LinkedIn



Web Developer

LinkedIn [Apply now](#)



Product Manager

LinkedIn [Apply now](#)



Software Engineer

LinkedIn [Apply now](#)

[Learn more about careers](#)

Current	<ul style="list-style-type: none">Senior Software Engineer, Distributed Systems at LinkedIn
Past	<ul style="list-style-type: none">Member of Technical Staff at Oracle Corporation Summer Intern at Oracle Corporation
Education	<ul style="list-style-type: none">Georgia Institute of TechnologyUniversity of PuneSt. Anne's School
Recommendations	3 people have recommended Neha
Connections	353 connections
Twitter	nehanarkhede
Public Profile	http://www.linkedin.com/pub/neha-narkhede/16/35a/9a2

Share

PDF

Print

Neha's Activity

Neha Narkhede has added the G
14 minutes ago · Add GitHub to your profile

Neha Narkhede likes this update

Chris Saccheri Newspaper ma
big, bold headline. <http://flic.kr/p/>

3 hours ago · Unlike · Comment (5)

Neha Narkhede likes this update

Hari Krishnan All set to close
[@LinkedInIndia](#) with our monthl
headed to Wankhede stadium to
via Twitter

6 hours ago · Unlike · Comment (3)

Neha Narkhede likes this update

Hirji Delhonte LinkedIn on NYS



Timelapse of L

GitHub Projects

nehanarkhede/kafka

forked from [kafka-dev/kafka](#)

A distributed messaging service · May 17, 2011

nehanarkhede/voldemort

forked from [voidemort/voldemort](#)

An open source clone of Amazon's Dynamo. · Oct 04, 2010

One sentence:

Kafka is a distributed, high throughput, persistent pub-sub messaging system

Kafka at LinkedIn

Tracking - User activity data

Page views Network updates
Profile Views Account logins
Search Invitations
Ad impressions

More than 60 types...

Service metrics

CPU/disk usage statistics
Healthcheck on services
Reporting/graphs

For many consumers

News feed
Ops
Security
Hadoop

Volume

- billions of events
- couple terabytes of data

Katka at LinkedIn

Tracking - User activity data

Page views

Network updates

Profile Views

Account logins

Search

Invitations

Ad impressions

More than 60 types...

Service metrics

CDU I/disk usage statistics

Profile Views

Account logins

Search

Invitations

Ad impressions

More than 60 types...

Service metrics

CPU/disk usage statistics

Healthcheck on services

Reporting/graphs

For many consumers

News feed

Ops

Healthcheck on services

Reporting/graphs

For many consumers

News feed

Ops

Security

Hadoop

Volume

- billions of events

IPO linked in

Search within checked filters

Try these searches

Information Technology and Services, San Francisco Bay Area

Information Technology and Services

Search for your company, competitors or favorite product.

Network

- By Me (0)
1st Connections (0)
2nd Connections (2)
3rd + Everyone (15)

Company

- JP Work Solutions MRI Network (62)
JR Teetern Enterprises (62)
Xeesm (62)
LinkedIn (44)
Adwyce (13)
Google (10)
Twitter (5)
Facebook (2)

Industry

- Internet (9)
Marketing and Advertising (3)
Telecommunications (2)
Computer Software (1)
Accounting (1)

Time

- Last Minute (0)
Last Hour (0)
Last Day (17)
Last Week (65)
Last 2 Weeks (8)

Location

- San Francisco Bay Area (7)
Greece (2)
Greater Detroit Area (2)
United Arab Emirates (1)
Zürich Area, Switzerland (1)

School

17 Updates Save this search Share this search



Michal Lorenc Congratulations LinkedIn on great IPO. Share some love, would you? - for less than 1.5 of your new shares you can help me reach my \$1K goal for Race For The Cure.



ActiveGiving.com Donations active.com

8 hours ago Like (1) Comment Share

Kyle Jusick likes this

Add a comment...



Joichi Ito Congratulations to @quixotic and the LinkedIn team on a nice IPO via Twitter

22 hours ago Like (7) Comment (2)

Scott Schlegel, Takashi Ishihara and 5 others like this



James Miao inspiring 22 hours ago



Gérard van Schip Love LinkedIn! 22 hours ago

Add a comment...



Borja Díaz de Terán RT @petrussola : Masterclass of IPO by @alleyinsider http://goo.gl/KjOX9 You may need extra help from #wikipedia to understand evth... http://goo.gl/6aDAg :) via Twitter



LINKEDIN IPO: Stock Priced Too Low businessinsider.com

A huge gift from Wall Street to its professional clients -- at the expense of the company.

22 hours ago Like Comment Share



Elliott Selzer Good read from @businessinsider on what LinkedIn's surge in price really means for everyones wallets. http://goo.gl/4dxBF via Twitter




LINKEDIN IPO: Stock Priced Too Low businessinsider.com

A huge gift from Wall Street to its professional clients -- at the expense of the company.

23 hours ago Like Comment Share

Show more...

Trending Links



LinkedIn goes ballistic following IPO - theregister.co.uk

Industry Top Headlines

- Internet
Marketing and Advertising
Telecommunications
Computer Software

News feed

Ops

Security

Hadoop

Volume

- billions of events
- couple terabytes of data

Goals:

- Persistent
- Highly performant
- Distributed model

What is Kafka's place in the universe?

Kafka
goes
here.

Message Queues
(ActiveMQ)

Log aggregators
(Flume)

Message Queues

(ActiveMQ)

Log aggregators (Flume)

What is Kafka's place in the universe?


Kafka
goes
here.

Message Queues
(ActiveMQ)

Log aggregators
(Flume)

Design principles:

- Simple API
- Low overhead
- Distributed model



File - /Users/nmarkhed/network/Projects/KafkaProject/kafka/core/src/main/scala/kafka/utils/ProducerExample.scala

```
import java.util.Properties
import kafka.message.Message
import kafka.producer.{ProducerData, ProducerConfig, Producer}

object ProducerExample {
  def main(args: Array[String]): Unit = {
    // set the right config options
    val props = new Properties
    props.put("zk.connect", "127.0.0.1:2121")

    val producer = new Producer[Message, Message](new ProducerConfig(props))


    // Kafka is not in the serialization business
    val data = "The Metamorphosis".getBytes
    val topic = "Novels"

    val sendData = new ProducerData[Message, Message](topic, new Message(data))

    // send the data to Kafka
    producer.send(sendData)
  }
}
```

Easy Kafka integration hint:

```
Logger logger = Logger.getLogger(KafkaAppender.class)
logger.info("Straight from your app to Kafka!")
```



```
import java.util.Properties
import kafka.consumer.{ConsumerConfig, Consumer}

object ConsumerExample {
  def main(args: Array[String]): Unit = {
    // set the right config options
    val props = new Properties
    props.put("zk.connect", "127.0.0.1:2121")
    props.put("groupid", "short-stories")

    val consumer = Consumer.create(new ConsumerConfig(props))

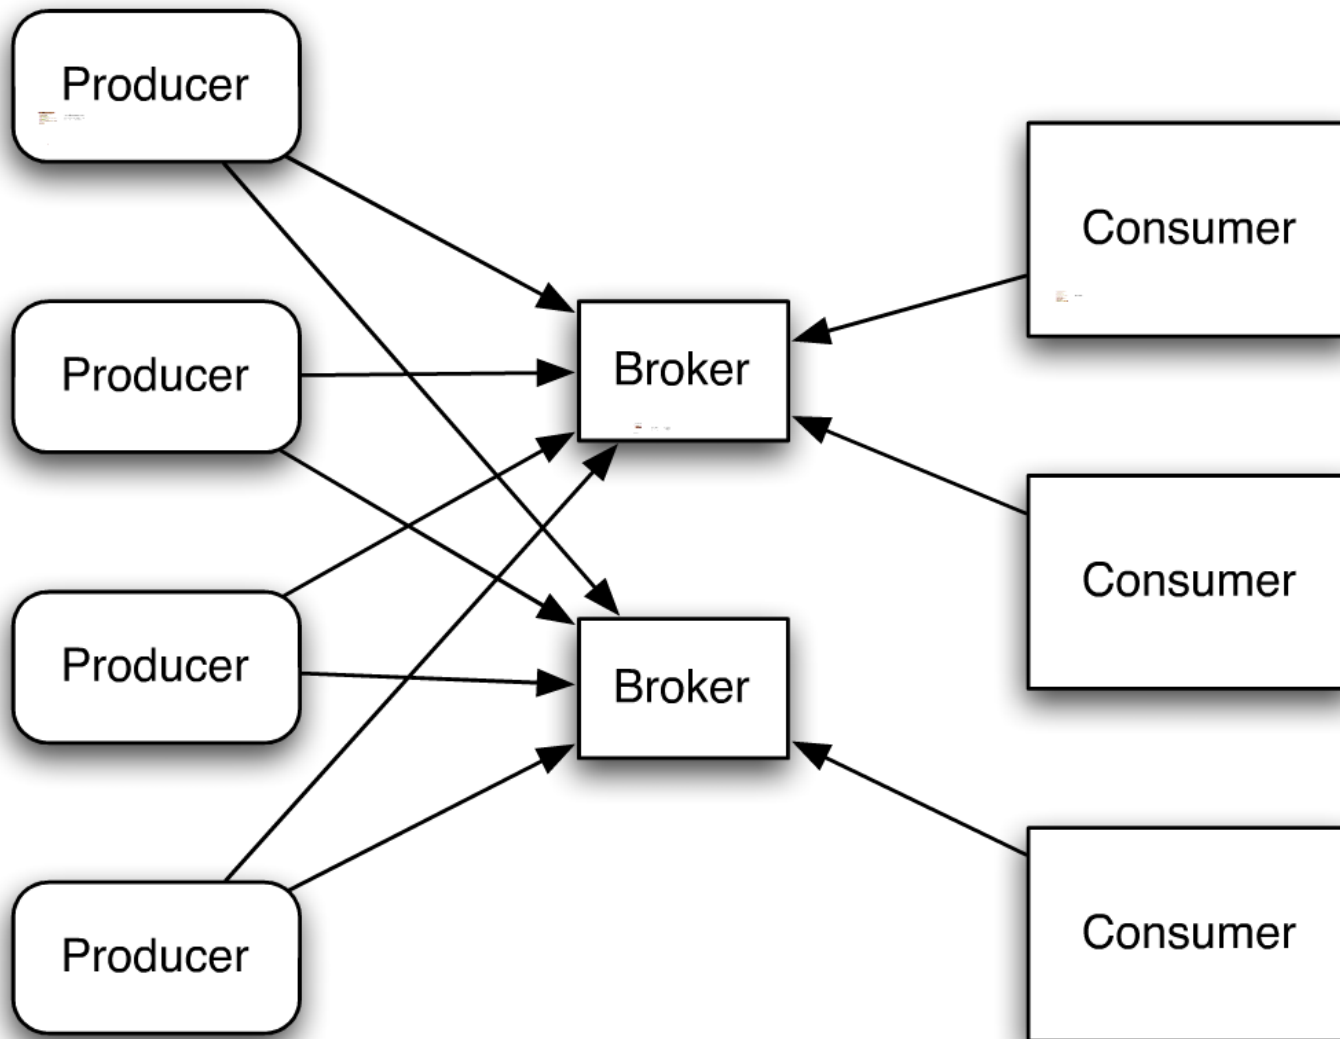
    val topic = "Novels"
    val topicMap = Predef.Map(topic -> 1)

    // create the message streams
    val streams = consumer.createMessageStreams(topicMap)
    val novelStreams = streams.get(topic).get

    // we just have one stream
    val novelStream = novelStreams.head

    for(novel <- novelStream)
      println("Whoohoo ! Time to read " + prettyPrintMessage(novel))
  }
}
```

Whoohoo ! Time to read A common confusion
Whoohoo ! Time to read A dream
Whoohoo ! Time to read A visit to a mine



Storage format

Brokers maintain efficient log files of all the messages they receive

```
[nnarkhed@nnarkhed-md kafka-logs]$ tree -s
.
|-- [      4096] novels-0
|   |-- [119878441] 00000000000000000000.kafka
|   |-- [      4096] short stories-0
|       |-- [119817774] 00000000000000000000.kafka
2 directories, 2 files
```

Messages indexed by 64-byte offset in the message/partition stream

```
00000000000000000000.kafka
0000000000092342342235.kafka
00000195329587682492345.kafka
```

No object level caching

- Relies on OS pagecache
- Access pattern provides $O(1)$ access to most data

Storage format

Brokers maintain efficient log files of all the messages they receive

```
[nnarkhed@nnarkhed-md kafka-logs]$ tree -s
.
|-- [      4096]  novels-0
|   |-- [119878441]  000000000000000000000000.kafka
|-- [      4096]  short stories-0
|   |-- [119817774]  000000000000000000000000.kafka

2 directories, 2 files
```

Messages indexed by 64-byte offset in the message/partition stream

```
000000000000000000000000.kafka
00000000000092342342235.kafka
00000195329587682492345.kafka
```

```
-- [ 4090] SHORT STORIES-V  
^-- [119817774] 000000000000000000000000.kafka
```

2 directories, 2 files

Messages indexed by 64-byte offset in the message/partition stream

```
000000000000000000000000.kafka  
00000000000092342342235.kafka  
00000195329587682492345.kafka
```

No object level caching

- Relies on OS pagecache

00000I95329587682492345.kafka

No object level caching

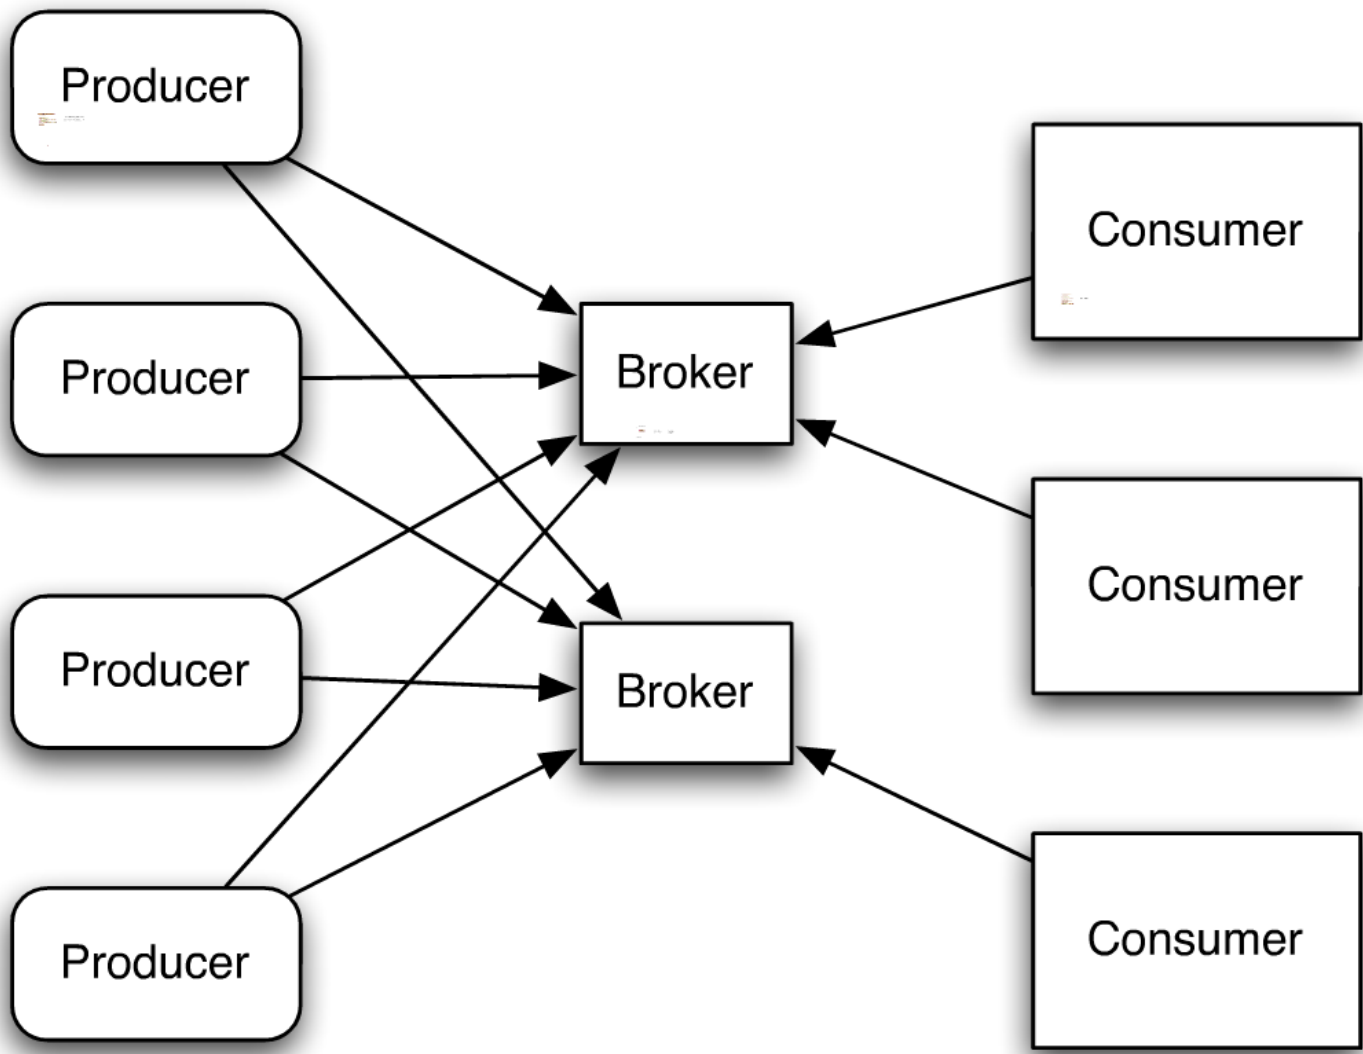
- Relies on OS pagecache
- Access pattern provides $O(1)$ access to most data

Network layer

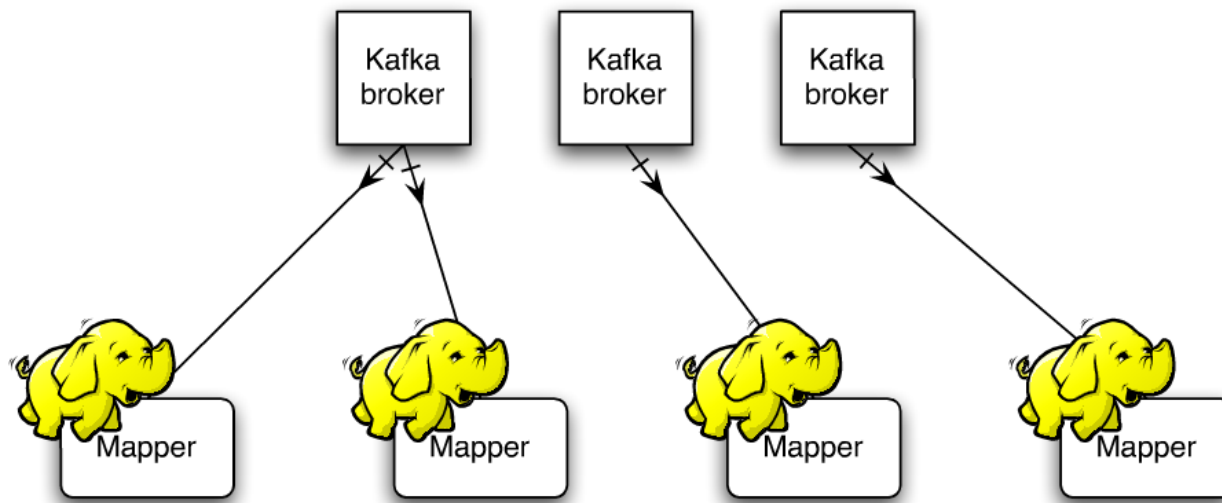
- batch API
- zero copy transfer

Log retention

- No state at the broker
- Time based retention policy
- Consumer rewind



ETL and Apache Hadoop



Current status



- open sourced, submitted to be Apache incubator project
- deployed in production for a year
- Hosted at [github:https://github.com/kafka-dev/kafka](https://github.com/kafka-dev/kafka)
- Outside contributions arriving

LinkedIn is hiring!

nnarkhede@linkedin.com

<http://www.linkedin.com/in/nehanarkhede>

@nehanarkhede