



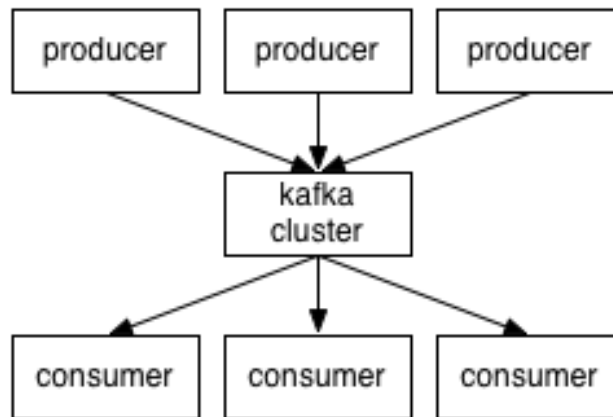
# Apache Kafka Installation and Basic Usage



Emre Övünç  
Intern – Innovera  
[info@emreovunc.com](mailto:info@emreovunc.com)

## What is Kafka ?

Kafka is a distributed, partitioned, replicated commit log service. It provides the functionality of a messaging system, but with a unique design.



### **Producers:**

Producers publish data to the topics of their choice. The producer is responsible for choosing which message to assign to which partition within the topic.

### **Consumers:**

Messaging traditionally has two models: queuing and publish-subscribe. In a queue, a pool of consumers may read from a server and each message goes to one of them; in publish-subscribe the message is broadcast to all consumers. Kafka offers a single consumer abstraction that generalizes both of these—the *consumer group*.

### **Clusters:**

Servers which are running Kafka are called “broker” and also many brokers creates “Cluster” together.

## Installation

First of all, I download Kafka from <http://kafka.apache.org/downloads.html>. Then, by using shell, I go to the file that I downloaded Kafka and extract it.

### 0.10.0.0

- **Release Notes**
- Source download: `kafka-0.10.0.0-src.tgz` (asc, md5)
- Binary downloads:
  - Scala 2.10 - `kafka_2.10-0.10.0.0.tgz` (asc, md5)
  - Scala 2.11 - `kafka_2.11-0.10.0.0.tgz` (asc, md5)

We build for multiple versions of Scala. This only matters if you are using Scala and you want a version for the same Scala version you use. Otherwise any version should work (2.11 is recommended).

```
root@Emre0vunc:~# cd Downloads/
root@Emre0vunc:~/Downloads# tar -xzf kafka_2.11-0.10.0.0.tgz
root@Emre0vunc:~/Downloads# cd kafka_2.11-0.10.0.0/
root@Emre0vunc:~/Downloads/kafka_2.11-0.10.0.0# ls
bin  config  libs  LICENSE  NOTICE  site-docs
```

Kafka is running through Zookeeper. So that, I should run zookeeper server, then I can run Kafka server.

```
root@Emre0vunc:~/Downloads/kafka_2.11-0.10.0.0# bin/zookeeper-server-start.sh
config/zookeeper.properties
[2016-07-14 22:55:41,630] INFO Reading configuration from: config/zookeeper.pr
operties (org.apache.zookeeper.server.quorum.QuorumPeerConfig)
[2016-07-14 22:55:41,632] INFO autopurge.snapRetainCount set to 3 (org.apache.
zookeeper.server.DatadirCleanupManager)
[2016-07-14 22:55:41,632] INFO autopurge.purgeInterval set to 0 (org.apache.zo
ookeeper.server.DatadirCleanupManager)
[2016-07-14 22:55:41,632] INFO Purge task is not scheduled. (org.apache.zookee
per.server.DatadirCleanupManager)
[2016-07-14 22:55:41,633] WARN Either no config or no quorum defined in config
, running in standalone mode (org.apache.zookeeper.server.quorum.QuorumPeerMa
in)
```

```
root@Emre0vunc:~/Downloads/kafka_2.11-0.10.0.0# bin/kafka-server-start.sh config/server.properties
[2016-07-14 22:56:51,009] INFO KafkaConfig values:
  advertised.host.name = null
  metric.reporters = []
  quota.producer.default = 9223372036854775807
  offsets.topic.num.partitions = 50
  log.flush.interval.messages = 9223372036854775807
  auto.create.topics.enable = true
  controller.socket.timeout.ms = 30000
  log.flush.interval.ms = null
  principal.builder.class = class org.apache.kafka.common.security.auth.DefaultPrincipalBuilder
```

## Usage

### Creating a topic:

Let's create a topic named "deneme" and list all topics.

```
root@Emre0vunc:~/Downloads/kafka_2.11-0.10.0.0# bin/kafka-topics.sh --create --zookeeper localhost:2181 --topic deneme --partitions 1 --replication-factor 1
Created topic "deneme".
root@Emre0vunc:~/Downloads/kafka_2.11-0.10.0.0# bin/kafka-topics.sh --list --zookeeper localhost:2181
deneme
test
```

### Sending a message:

I use the command line client to send messages to the Kafka cluster.

```
root@Emre0vunc:~/Downloads/kafka_2.11-0.10.0.0# bin/kafka-console-producer.sh --broker-list localhost:9092 --topic deneme
Hello World
My second message ...
█
```

### Reading messages:

Now I can switch consumer way and read messages.

```
root@Emre0vunc:~/Downloads/kafka_2.11-0.10.0.0# bin/kafka-console-consumer.sh --zookeeper localhost:2181 --topic deneme --from-beginning
Hello World
My second message ...
```

Note that you can see the list of your topics. As you can from the picture, there are two topics.

```
root@Emre0vunc:~/Downloads/kafka_2.11-0.10.0.0# bin/kafka-topics.sh --zookeeper localhost:2181 --describe
Topic:deneme    PartitionCount:1      ReplicationFactor:1   Configs:
   Topic: deneme  Partition: 0         Leader: 0              Replicas: 0           Isr: 0
Topic:test      PartitionCount:1      ReplicationFactor:1   Configs:
   Topic: test   Partition: 0         Leader: 0              Replicas: 0           Isr: 0
```

## Server Properties:

You can change the server properties by using “server.properties” configuration file that is in the ‘config’ folder.

```
root@EmreOvunc: ~/Downloads/kafka_2.11-0.10.0.0/config
GNU nano 2.5.3 File: server.properties
# The interval at which log segments are checked to see if they can be deleted according
# to the retention policies
log.retention.check.interval.ms=300000

##### Zookeeper #####

# Zookeeper connection string (see zookeeper docs for details).
# This is a comma separated host:port pairs, each corresponding to a zk
# server. e.g. "127.0.0.1:3000,127.0.0.1:3001,127.0.0.1:3002".
# You can also append an optional chroot string to the urls to specify the
# root directory for all kafka znodes.
zookeeper.connect=localhost:2181

# Timeout in ms for connecting to zookeeper
zookeeper.connection.timeout.ms=6000
```

## References:

<https://kafka.apache.org/documentation.html>