

# FreeBSD Administration - Support #891

## Install Kafka on FreeBSD

04/14/2017 09:52 PM - Daniel Curtis

<b>Status:</b>	Closed	<b>Start date:</b>	04/16/2017
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Daniel Curtis	<b>% Done:</b>	100%
<b>Category:</b>	Distributed Computing	<b>Estimated time:</b>	1.00 hour
<b>Target version:</b>	FreeBSD 10	<b>Spent time:</b>	1.50 hour

### Description

This is a guide for setting up Apache Kafka on FreeBSD 10.

## Prepare the Environment

- Make sure the system is up to date:

```
pkg update && pkg upgrade
```

- Install a few dependencies:

```
pkg install bash git gradle openjdk7 sudo
```

- Create a symlink of bash:

```
ln -s /usr/local/bin/bash /bin/bash
```

- Install Zookeeper and ensure its running, refer to issue [#890](#)

```
service zookeeper start
```

- Create the kafka user:

```
pw useradd -n kafka -m -c "Kafka" -s /bin/sh
```

## Install Kafka

- Download Kafka from GitHub:

```
git clone https://github.com/apache/kafka.git /usr/local/kafka
```

- Bootstrap Kafka:

```
cd /usr/local/kafka  
gradle
```

- Build the Kafka jar:

```
./gradlew jar
```

- Create log file directory for Kafka:

```
mkdir /var/log/kafka
```

- Set the kafka user permissions:

```
chown -R kafka /usr/local/kafka  
chown -R kafka /var/log/kafka
```

- Edit the Kafka config:

```
vi /usr/local/kafka/config/server.properties
```

- And modify the following parameters:

```
# The id of the broker. This must be set to a unique integer for each broker.  
broker.id=1  
  
# Increase the message size limit  
message.max.bytes=20000000  
replica.fetch.max.bytes=30000000  
  
log.dirs=/var/log/kafka  
listeners=PLAINTEXT://192.168.1.104:9092  
  
zookeeper.connect=localhost:2181  
#zookeeper.connect=192.168.1.101:2181,192.168.1.102:2181,192.168.1.103:2181
```

- Test start the Kafka server:

```
/usr/local/kafka/bin/kafka-server-start.sh /usr/local/kafka/config/server.properties
```

- Once finished testing press CTRL+C to stop the kafka server.

## Init Script

- Create a kafka init script:

```
vi /usr/local/etc/rc.d/kafka
```

- And add the following:

```
#!/bin/sh  
#  
# PROVIDE: kafka  
# REQUIRE: zookeeper LOGIN  
#  
# Add the following lines to /etc/rc.conf to enable odoo server  
#  
#  
# kafka_enable (bool): Set to "NO" by default,  
#                       Set it to "YES" to enable kafka server  
#  
# kafka_config (str):  The path to the kafka server configuration file
```

```

# (defaults to /usr/local/kafka/config/server.properties)
#
# kafka_user (str): The user to run kafka server as. Set to "root" by default

. /etc/rc.subr

name=kafka
command=/usr/local/kafka/bin/kafka-server-start.sh
rcvar=kafka_enable

load_rc_config $name

kafka_enable="${kafka_enable-"NO"}"
kafka_config="${kafka_config-"/usr/local/kafka/config/server.properties"}"
kafka_user="${kafka_user-"root"}"
kafka_flags="${kafka-"${kafka_config}"}"

# /etc/rc.subr use $pidfile (not ${name}_pidfile)
pidfile="${kafka_pidfile}"

required_files="${kafka_config}"

start_cmd="su - ${kafka_user} -c '${command} ${kafka_flags}' &"
stop_cmd="${name}_stop"
status_cmd="${name}_status"
getval_cmd="${name}_getval"

kafka_stop()
{
    echo "Stopping ${name}."
    pkill -TERM -f ${name}
}

kafka_status()
{
    kafka_pid=$(pgrep -f ${name})

    if [ -n "${kafka_pid}" ]
    then
        echo "${name} running with pid: $kafka_pid"
    else
        echo "${name} not running? (pid not found)"
    fi
}

command_args=" >/dev/null 2>&1 &"

load_rc_config $name
run_rc_command "$1"

```

- Make the script executable:

```
chmod +x /usr/local/etc/rc.d/kafka
```

- Start and enable kafka at boot:

```
echo 'kafka_enable="YES"' >> /etc/rc.conf
echo 'kafka_user="kafka"' >> /etc/rc.conf
service kafka start
```

## Resources

- <https://github.com/apache/kafka>
- <http://kafka.apache.org/documentation.html#quickstart>
- <https://github.com/btccom/btcpool/blob/master/docs/INSTALL-Kafka.md>

## History

---

### #1 - 04/15/2017 12:48 AM - Daniel Curtis

- Description updated

### #2 - 04/15/2017 01:06 AM - Daniel Curtis

- Description updated

- Category set to Distributed Computing

- Status changed from New to In Progress

- % Done changed from 0 to 50

### #3 - 04/17/2017 06:04 PM - Daniel Curtis

- Description updated

- Status changed from In Progress to Resolved

- % Done changed from 50 to 100

### #4 - 06/04/2017 08:59 PM - Daniel Curtis

- Status changed from Resolved to Closed