

How to setup or install Apache Flink

Setup or install Apache Flink

Setup or Install Apache Flink – Flink could be set up on the following Operating Systems:

- Linux
- Mac OS X
- Windows

Prerequisites to run Apache Flink

For Flink to run, the environment/machine should have the following:

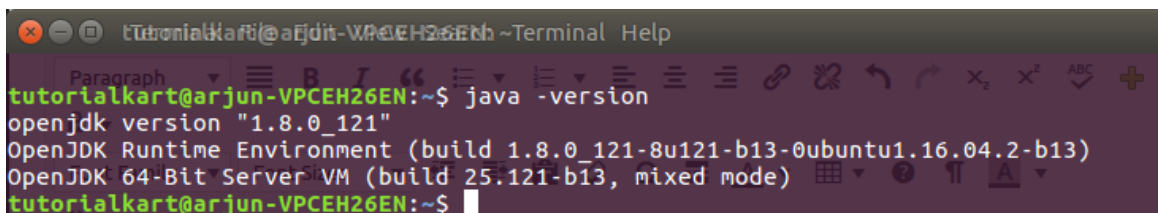
1. Java 7 (or higher)

How to check if Java is available

To check if Java is available, open the command prompt or terminal and enter the following command :

```
java -version
```

The output would look like:



```
tutorialkart@arjun-VPCEH26EN:~$ java -version
openjdk version "1.8.0_121"
OpenJDK Runtime Environment (build 1.8.0_121-8u121-b13-0ubuntu1.16.04.2-b13)
OpenJDK 64-Bit Server VM (build 25.121-b13, mixed mode)
tutorialkart@arjun-VPCEH26EN:~$
```

Checking Java availability in PC/Machine using "java -version"

In this case, the machine has Java with the version of 1.8.

If the output for the command "java -version" is not as shown in the above image, Java is not installed or the Path to Java executable is not set properly. To set up Java in your machine, please refer [How to setup java](#).

Download a binary of Apache Flink

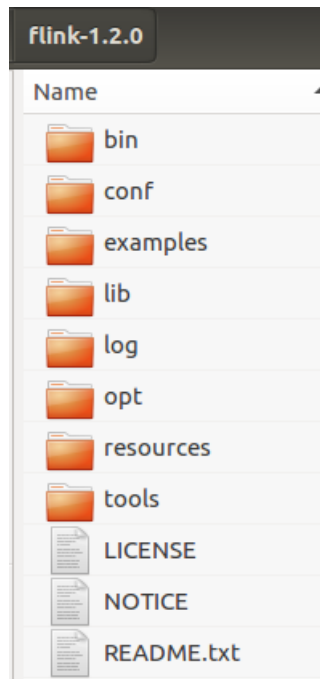
Once your PC or machine is ready with Java, we have to download a binary of Apache Flink. Prior to download, we have to choose a binary of Apache Flink, based on the requirements. If you have data stored in Hadoop, pick one based on the Hadoop version the data is available in. Binaries are available for download at the official download page[<http://flink.apache.org/downloads.html>].

If you don't require Hadoop, any of the binary at the download page should do the job.

For this tutorial, we chose the Hadoop version 2.7 and Scala 2.11 at the link [http://www.apache.org/dyn/closer.lua/flink/flink-1.2.0/flink-1.2.0-bin-hadoop27-scala_2.11.tgz].

The download file is over 100MB, so we shall wait for the file to download. Once the binary is download in zip format, unzip and extract the contents.

The contents of the folder should look like the following



Apache Flink folder contents

How to Start a Local Apache Flink Cluster

In Windows: Inside the bin folder, start-local.bat has all the necessary scripts to start the local cluster. To start the local cluster, navigate to /flink-folder>/bin/ folder and open command prompt from the bin folder. Run the command: `start-local.bat` . Or just double clicking the start-local.bat file, should start the local cluster.

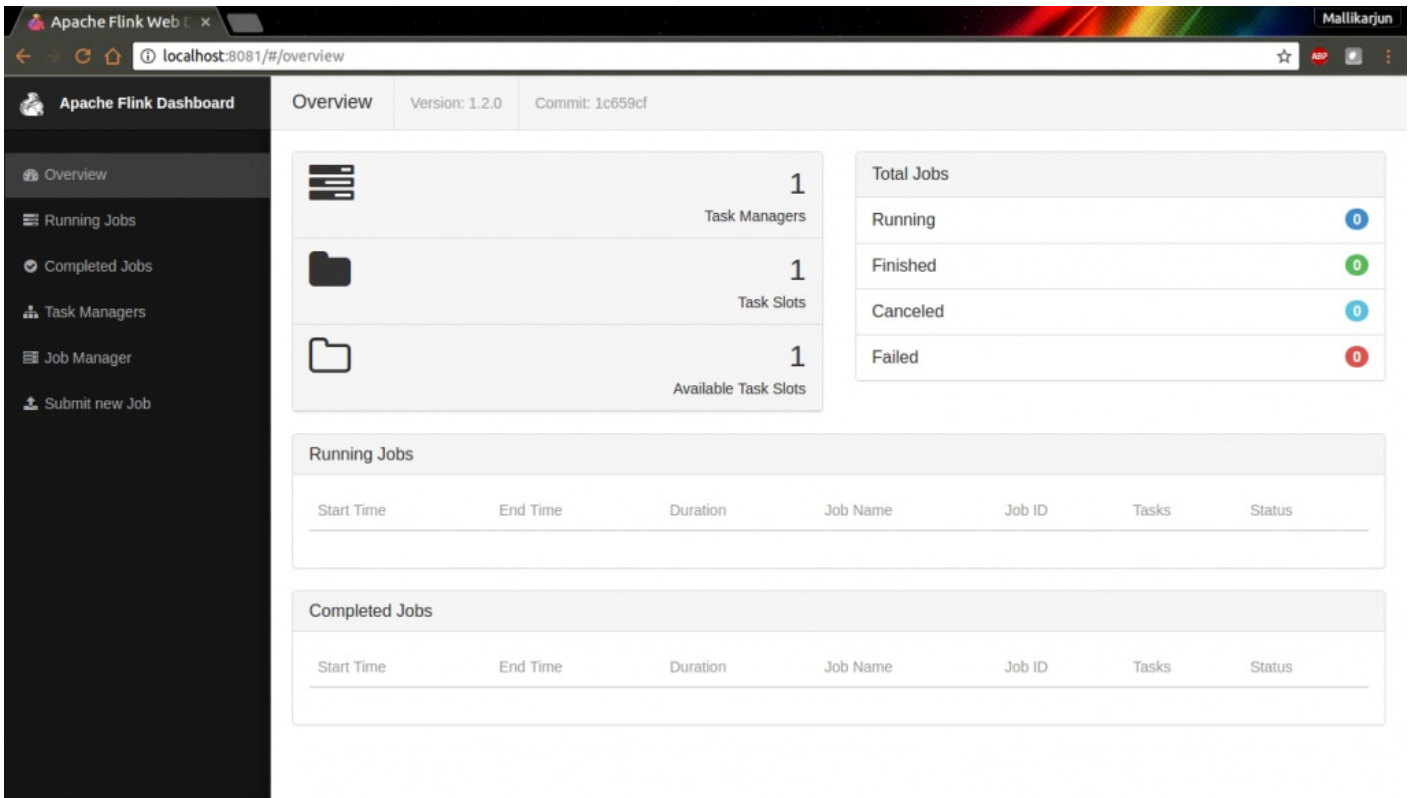
In Ubuntu: Inside the bin folder, start-local.sh has all the necessary scripts to start the local cluster.

Open the terminal from Flink folder and run the command `./bin/start-local.sh` as shown below:

```
Terminal File Edit View Search Terminal Help
tutorialkart@arjun-VPCEH26EN:~$ cd flink-1.2.0/
tutorialkart@arjun-VPCEH26EN:~/flink-1.2.0$ ./bin/start-local.sh
Starting jobmanager daemon on host arjun-VPCEH26EN.
tutorialkart@arjun-VPCEH26EN:~/flink-1.2.0$
```

Start local Flink cluster using command

jobmanger daemon should be started on your PC/machine with a message "Starting jobmanager daemon on host <hostname>". The jobmanager daemon which is running can be viewed through a browser at <http://localhost:8081/>



Apache Flink jobmanager overview

Apache Flink jobmanager overview could be seen in the browser as above.

How to stop Apache Flink local cluster

In Windows, running the command `stop-local.bat` in the command prompt from the `<flink-folder>/bin/` folder should stop the jobmanager daemon and thus stopping the cluster.

In Ubuntu, running the command `./bin/stop-local.sh` in the terminal from bin folder should stop the jobmanager daemon and thus stopping the cluster.

```
tutorialkart@arjun-VPCEH26EN:~/flink-1.2.0$ ./bin/stop-local.sh
Stopping jobmanager daemon (pid: 8944) on host arjun-VPCEH26EN.
tutorialkart@arjun-VPCEH26EN:~/flink-1.2.0$
```

Stop Apache Flink Local cluster

So, now we are able to start or stop a stop a Flink local cluster, and thus came to the end of the topic setup or install Apache Flink.

Conclusion

In this Flink Tutorial, we have seen how to set up or install the Apache Flink to run as a local cluster. In our next tutorial, we shall observe **how to submit a job to the Apache Flink local cluster**.

Apache Flink

↳ [Learn Apache Flink](#)

↳ [Apache Flink - Setup and Installation](#)

Apache Tutorials

↳ [Learn Apache Hadoop](#)

↳ [Learn Apache Spark](#)

↳ [Learn Apache Flink](#)

↳ [Learn Apache Beam](#)

↳ [Learn Apache MXNet](#)

↳ [Learn Apache PDFBox](#)

↳ [Learn Apache OpenNLP](#)