

Hortonworks Data Platform

Spark QuickStart Guide

(October 13, 2015)

Hortonworks Data Platform: Spark QuickStart Guide

Copyright © 2012-2015 Hortonworks, Inc. Some rights reserved.

The Hortonworks Data Platform, powered by Apache Hadoop, is a massively scalable and 100% open source platform for storing, processing and analyzing large volumes of data. It is designed to deal with data from many sources and formats in a very quick, easy and cost-effective manner. The Hortonworks Data Platform consists of the essential set of Apache Hadoop projects including MapReduce, Hadoop Distributed File System (HDFS), HCatalog, Pig, Hive, HBase, ZooKeeper and Ambari. Hortonworks is the major contributor of code and patches to many of these projects. These projects have been integrated and tested as part of the Hortonworks Data Platform release process and installation and configuration tools have also been included.

Unlike other providers of platforms built using Apache Hadoop, Hortonworks contributes 100% of our code back to the Apache Software Foundation. The Hortonworks Data Platform is Apache-licensed and completely open source. We sell only expert technical support, [training](#) and partner-enablement services. All of our technology is, and will remain, free and open source.

Please visit the [Hortonworks Data Platform](#) page for more information on Hortonworks technology. For more information on Hortonworks services, please visit either the [Support](#) or [Training](#) page. Feel free to [contact us](#) directly to discuss your specific needs.



Except where otherwise noted, this document is licensed under
Creative Commons Attribution ShareAlike 4.0 License.
<http://creativecommons.org/licenses/by-sa/4.0/legalcode>

Table of Contents

1. Introduction	1
2. Prerequisites	3
3. Installing Spark	4
4. Validating Spark	8
4.1. Run the Spark Pi example	8
4.2. Run the WordCount Example	9
5. Installing Spark with Kerberos	12
5.1. Accessing the Hive Metastore in Secure Mode	13
6. Using Spark with HDFS	14
7. Troubleshooting Spark	15

List of Tables

- 1.1. Spark Support in HDP, Ambari 2
- 1.2. Spark Feature Support by Version 2
- 2.1. Spark Prerequisites 3

1. Introduction

Hortonworks Data Platform supports Apache Spark 1.3.1, a fast, large-scale data processing engine.

Deep integration of Spark with YARN allows Spark to operate as a cluster tenant alongside other engines such as Hive, Storm, and HBase, all running simultaneously on a single data platform. YARN allows flexibility: you can choose the right processing tool for the job. Instead of creating and managing a set of dedicated clusters for Spark applications, you can store data in a single location, access and analyze it with multiple processing engines, and leverage your resources. In a modern data architecture with multiple processing engines using YARN and accessing data in HDFS, Spark on YARN is the leading Spark deployment mode.

Spark Features

Spark on HDP supports the following features:

- Spark Core
- Spark on YARN
- Spark on YARN on Kerberos-enabled clusters
- Spark History Server
- Spark MLLib
- Support for Hive 0.13, including the `collect_list` UDF

The following features are available as technical previews:

- Spark DataFrame API
- ORC file support
- Spark SQL
- Spark Streaming
- Spark SQL Thrift Server
- Dynamic Executor Allocation

The following features and tools are not officially supported in this release:

- ML Pipeline API
- SparkR
- Spark Standalone

- GraphX
- iPython
- Zeppelin

Spark on YARN uses YARN services for resource allocation, running Spark Executors in YARN containers. Spark on YARN supports workload management and Kerberos security features. It has two modes:

- YARN-Cluster mode, optimized for long-running production jobs.
- YARN-Client mode, best for interactive use such as prototyping, testing, and debugging. Spark Shell runs in YARN-Client mode only.

The following tables summarize Spark versions and feature support across HDP and Ambari versions.

Table 1.1. Spark Support in HDP, Ambari

HDP	Ambari	Spark
2.2.4	2.0.1	1.2.1
2.2.6	2.1.1	1.2.1
2.2.8	2.1.1	1.3.1
2.2.9	2.1.1	1.3.1

Table 1.2. Spark Feature Support by Version

Feature	1.2.1	1.3.1
Spark Core	Yes	Yes
Spark on YARN	Yes	Yes
Spark on YARN, Kerberos-enabled clusters	Yes	Yes
Spark History Server	Yes	Yes
Spark MLLib	Yes	Yes
Hive 0.1.3, including collect_list UDF		Yes
ML Pipeline API (PySpark)		
DataFrame API		TP
ORC Files		TP
Spark SQL	TP	TP
Spark Streaming	TP	TP
Spark SQL Thrift Server		TP
Dynamic Executor Allocation		TP
SparkR		
Spark Standalone		
GraphX		

TP: Tech Preview

2. Prerequisites

Before installing Spark, make sure your cluster meets the following prerequisites.

Table 2.1. Spark Prerequisites

Prerequisite	Description
Cluster Stack Version	HDP 2.2.4 or later stack
(Optional) Ambari	Version 2.0.0 or later
Components	Spark requires HDFS and YARN.



Note

If you used the tech preview, save any configuration changes you made to the tech preview environment. Install Spark, and then update the configuration with your changes.

3. Installing Spark

To install Spark manually, see [Installing and Configuring Apache Spark](#) in the Manual Installation Guide.

To install Spark on a Kerberized cluster, first read [Installing Spark with Kerberos](#) (the next topic in this Quick Start Guide).

The remainder of this section describes how to install Spark using Ambari. (For general information about installing HDP components using Ambari, see [Adding a Service](#) in the Ambari Documentation Suite.)

The following diagram shows the Spark installation process using Ambari.



To install Spark using Ambari, complete the following steps:

1. Choose the Ambari "Services" tab.

In the Ambari "Actions" pulldown menu, choose "Add Service." This will start the Add Service Wizard. You'll see the Choose Services screen.

Select "Spark", and click "Next" to continue.

Add Service Wizard

ADD SERVICE WIZARD

Choose Services

Assign Masters

Assign Slaves and Clients

Customize Services

Configure Identities

Review

Install, Start and Test

Summary

Choose Services

Choose which services you want to install on your cluster.

<input type="checkbox"/> Service	Version	Description
<input checked="" type="checkbox"/> HDFS	2.6.0.2.2	Apache Hadoop Distributed File System
<input checked="" type="checkbox"/> YARN + MapReduce2	2.6.0.2.2	Apache Hadoop NextGen MapReduce (YARN)
<input checked="" type="checkbox"/> Tez	0.5.2.2.2	Tez is the next generation Hadoop Query Processing framework written on top of YARN.
<input checked="" type="checkbox"/> Hive	0.14.0.2.2	Data warehouse system for ad-hoc queries & analysis of large datasets and table & storage management service
<input type="checkbox"/> HBase	0.98.4.2.2	Non-relational distributed database and centralized service for configuration management & synchronization
<input checked="" type="checkbox"/> Pig	0.14.0.2.2	Scripting platform for analyzing large datasets
<input checked="" type="checkbox"/> Sqoop	1.4.5.2.2	Tool for transferring bulk data between Apache Hadoop and structured data stores such as relational databases
<input checked="" type="checkbox"/> Oozie	4.1.0.2.2	System for workflow coordination and execution of Apache Hadoop jobs. This also includes the installation of the optional Oozie Web Console which relies on and will install the ExtJS Library .
<input checked="" type="checkbox"/> ZooKeeper	3.4.6.2.2	Centralized service which provides highly reliable distributed coordination
<input type="checkbox"/> Falcon	0.6.0.2.2	Data management and processing platform
<input type="checkbox"/> Storm	0.9.3.2.2	Apache Hadoop Stream processing framework
<input type="checkbox"/> Flume	1.5.2.2.2	A distributed service for collecting, aggregating, and moving large amounts of streaming data into HDFS
<input type="checkbox"/> Ambari Metrics	0.1.0	A system for metrics collection that provides storage and retrieval capability for metrics collected from the cluster
<input type="checkbox"/> Kafka	0.8.1.2.2	A high-throughput distributed messaging system
<input type="checkbox"/> Knox	0.5.0.2.2	Provides a single point of authentication and access for Apache Hadoop services in a cluster
<input type="checkbox"/> Ranger	0.4.0	Comprehensive security for Hadoop
<input checked="" type="checkbox"/> Slider	0.60.0.2.2	A framework for deploying, managing and monitoring existing distributed applications on YARN.
<input checked="" type="checkbox"/> Spark	1.2.0.2.2	Apache Spark is a fast and general engine for large-scale data processing.

[Next →](#)

(Starting with HDP 2.2.4, Ambari will install Spark version 1.2.1, not 1.2.0.2.2.)

- Ambari will display a warning message. Confirm that your cluster is running HDP 2.2.4 or later, and then click "Proceed".

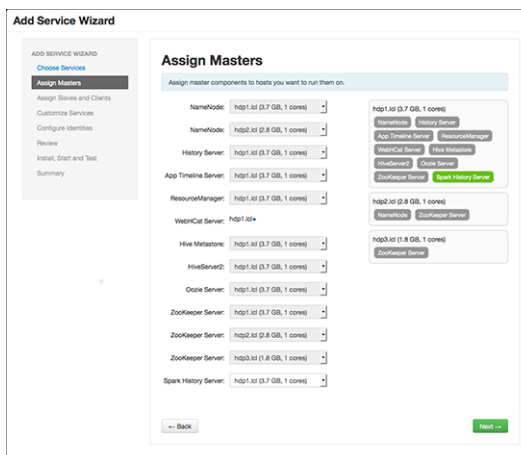


Note

You can reconfirm component versions in Step 6 before finalizing the upgrade.

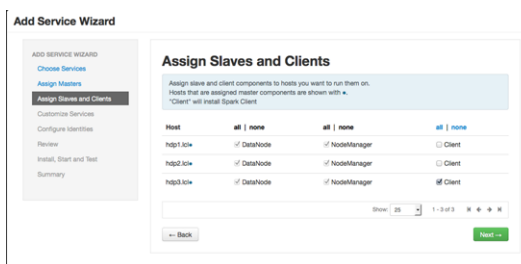
- On the Assign Masters screen, choose a node for the Spark History Server.

Click "Next" to continue.



4. On the Assign Slaves and Clients screen, specify the machine(s) that will run Spark clients.

Click "Next" to continue.



5. On the Customize Services screen there are no properties that must be specified. We recommend that you use default values for your initial configuration. Click "Next" to continue. (When you are ready to customize your Spark configuration, see [Apache Spark 1.2.1 properties.](#))

6. Ambari will display the Review screen.




Important

On the Review screen, make sure all HDP components are version 2.2.4 or later.

Click "Deploy" to continue.

7. Ambari will display the Install, Start and Test screen. The status bar and messages will indicate progress.



8. When finished, Ambari will present a summary of results. Click "Complete" to finish installing Spark.

**Caution**

Ambari will create and edit several configuration files. Do not edit these files directly if you configure and manage your cluster using Ambari.

4. Validating Spark

To validate the Spark installation, run the following Spark jobs:

- [Spark Pi example](#)
- [WordCount example](#)

4.1. Run the Spark Pi example

The Pi program tests compute-intensive tasks by calculating pi using an approximation method. The program “throws darts” at a circle – it generates points in the unit square ((0,0) to (1,1)) and sees how many fall within the unit circle. The result approximates pi.



To run Spark Pi:

1. Log on as a user with HDFS access—for example, your `spark` user (if you defined one) or `hdfs`. Navigate to a node with a Spark client and access the `spark-client` directory:

```
su hdfs  
  
cd /usr/hdp/current/spark-client
```

2. Submit the Spark Pi job:

```
./bin/spark-submit --class org.apache.spark.examples.SparkPi --  
master yarn-cluster --num-executors 3 --driver-memory 512m --  
executor-memory 512m --executor-cores 1 lib/spark-examples*.jar  
10
```

The job should complete without errors. It should produce output similar to the following:

```
15/04/10 17:29:35 INFO Client:  
  client token: N/A  
  diagnostics: N/A  
  ApplicationMaster host: N/A  
  ApplicationMaster RPC port: 0  
  queue: default  
  start time: 1428686924325  
  final status: SUCCEEDED  
  tracking URL: http://blue1:8088/proxy/  
application_1428670545834_0009/  
  user: hdfs
```

To view job status in a browser, copy the URL tracking from the job output and go to the associated URL.

3. Job output should list the estimated value of pi. In the following example, output was directed to stdout:

```
Log Type: stdout
Log Upload Time: 22-Mar-2015 17:13:33
Log Length: 23
Pi is roughly 3.142532
```

4.2. Run the WordCount Example

WordCount is a simple program that counts how often a word occurs in a text file.

1. Select an input file for the Spark WordCount example. You can use any text file as input.
2. Upload the input file to HDFS. The following example uses `log4j.properties` as the input file:

```
su hdfs

cd /usr/hdp/current/spark-client/

hadoop fs -copyFromLocal /etc/hadoop/conf/log4j.properties /tmp/
data
```

3. Run the Spark shell:

```
./bin/spark-shell --master yarn-client --driver-memory 512m --
executor-memory 512m
```

You should see output similar to the following:

```
Spark assembly has been built with Hive, including Datanucleus jars on
classpath
15/03/30 17:42:41 INFO SecurityManager: Changing view acls to: root
15/03/30 17:42:41 INFO SecurityManager: Changing modify acls to: root
15/03/30 17:42:41 INFO SecurityManager: SecurityManager: authentication
disabled; ui acls disabled; users with view permissions: Set(root); users
with modify permissions: Set(root)
15/03/30 17:42:41 INFO HttpServer: Starting HTTP Server
15/03/30 17:42:41 INFO Utils: Successfully started service 'HTTP class
server' on port 55958.
Welcome to

  ____
 /  __/  /  __/  /  __/  /  __/
_\\  \\/_  \\/_  \\/_  \\/_  \\/_  \\/_
/___/  ._/\\_./_/\\_/ /_/\\_\\_
  /_/

version 1.2.1

Using Scala version 2.10.4 (Java HotSpot(TM) 64-Bit Server VM, Java 1.7.
0_67)
Type in expressions to have them evaluated.
Type :help for more information.
15/03/30 17:42:47 INFO SecurityManager: Changing view acls to: root
15/03/30 17:42:47 INFO SecurityManager: Changing modify acls to: root
15/03/30 17:42:47 INFO SecurityManager: SecurityManager: authentication
disabled; ui acls disabled; users with view permissions: Set(root); users
with modify permissions: Set(root)
15/03/30 17:42:48 INFO Slf4jLogger: Slf4jLogger started
15/03/30 17:42:48 INFO Remoting: Starting remoting
```

```
15/03/30 17:42:48 INFO Remoting: Remoting started; listening on addresses :
[akka.tcp://sparkDriver@green4:33452]
15/03/30 17:42:48 INFO Utils: Successfully started service 'sparkDriver' on
port 33452.
15/03/30 17:42:48 INFO SparkEnv: Registering MapOutputTracker
15/03/30 17:42:48 INFO SparkEnv: Registering BlockManagerMaster
15/03/30 17:42:48 INFO DiskBlockManager: Created local directory at /
tmp/spark-a0fdb1ce-d395-497d-bf6f-1cf00ae253b7/spark-52dfe754-7f19-4b5b-
bd73-0745a1f6d158
15/03/30 17:42:48 INFO MemoryStore: MemoryStore started with capacity 265.4
MB
15/03/30 17:42:48 WARN NativeCodeLoader: Unable to load native-hadoop
library for your platform... using builtin-java classes where applicable
15/03/30 17:42:49 INFO HttpFileServer: HTTP File server directory
is /tmp/spark-817944df-07d2-4205-972c-elb877ca4869/spark-280ea9dd-
e40d-4ec0-8ecf-8c4b159dafaf
15/03/30 17:42:49 INFO HttpServer: Starting HTTP Server
15/03/30 17:42:49 INFO Utils: Successfully started service 'HTTP file
server' on port 56174.
15/03/30 17:42:49 INFO Utils: Successfully started service 'SparkUI' on port
4040.
15/03/30 17:42:49 INFO SparkUI: Started SparkUI at http://green4:4040
15/03/30 17:42:49 INFO Executor: Starting executor ID <driver> on host
localhost
15/03/30 17:42:49 INFO Executor: Using REPL class URI: http://172.23.160.
52:55958
15/03/30 17:42:49 INFO AkkaUtils: Connecting to HeartbeatReceiver: akka.
tcp://sparkDriver@green4:33452/user/HeartbeatReceiver
15/03/30 17:42:49 INFO NettyBlockTransferService: Server created on 47704
15/03/30 17:42:49 INFO BlockManagerMaster: Trying to register BlockManager
15/03/30 17:42:49 INFO BlockManagerMasterActor: Registering block manager
localhost:47704 with 265.4 MB RAM, BlockManagerId(<driver>, localhost,
47704)
15/03/30 17:42:49 INFO BlockManagerMaster: Registered BlockManager
15/03/30 17:42:49 INFO SparkILoop: Created spark context..
Spark context available as sc.

scala>
```

4. Submit the job. At the scala prompt, type the following commands, replacing node names, file name and file location with your own values:

```
val file = sc.textFile("/tmp/data")

val counts = file.flatMap(line => line.split(" ")).map(word =>
(word, 1)).reduceByKey(_ +_)

counts.saveAsTextFile("/tmp/wordcount")
```

5. To view the output from within the scala shell:

```
counts.toArray().foreach(println)
```

To view the output using HDFS:

- a. Exit the scala shell:

```
scala > exit
```

b. View WordCount job results:

```
hadoop fs -l /tmp/wordcount
```

You should see output similar to the following:

```
/tmp/wordcount/_SUCCESS  
/tmp/wordcount/part-00000  
/tmp/wordcount/part-00001
```

c. Use the HDFS cat command to list WordCount output. For example:

```
hadoop fs -cat /tmp/wordcount/part*
```

5. Installing Spark with Kerberos

Spark jobs are submitted to a Hadoop cluster as YARN jobs. The developer creates a Spark application in a local environment, and tests it in a single-node Spark Standalone cluster on their developer workstation.

When a job is ready to run in a production environment, there are a few additional steps if the cluster is Kerberized:

- The Spark History Server daemon needs a Kerberos account and keytab to run in a Kerberized cluster.
- When you enable Kerberos for a Hadoop cluster with Ambari, Ambari configures Kerberos for the Spark History Server and automatically creates a Kerberos account and keytab for it. For more information, see [Configuring Ambari and Hadoop for Kerberos](#).
- If you are not using Ambari, or if you plan to enable Kerberos manually for the Spark History Server, see [Creating Service Principals and Keytab Files for HDP](#) in the Manual Install Guide.
- To submit Spark jobs in a Kerberized cluster, the account (or person) submitting jobs needs a Kerberos account & keytab.
- When access is authenticated without human interaction – as happens for processes that submit job requests – the process would use a headless keytab. Security risk is mitigated by ensuring that only the service who should be using the headless keytab has the permissions to read it.
- An end user should use their own keytab when submitting a Spark job.

Setting Up Principals and Keytabs for End User Access to Spark

In the following example, user \$USERNAME runs the Spark Pi job in a Kerberos-enabled environment:

```
su $USERNAME
kinit USERNAME@YOUR-LOCAL-REALM.COM
cd /usr/hdp/current/spark-client/
./bin/spark-submit --class org.apache.spark.examples.SparkPi \
  --master yarn-cluster \
  --num-executors 3 \
  --driver-memory 512m \
  --executor-memory 512m \
  --executor-cores 1 \
  lib/spark-examples*.jar 10
```

Setting Up Service Principals and Keytabs for Processes Submitting Spark Jobs

The following example shows the creation and use of a headless keytab for a spark service user account that will submit Spark jobs on node blue1@example.com:

1. Create a Kerberos service principal for user spark:

```
kadmin.local -q "addprinc -randkey spark/blue1@EXAMPLE.COM"
```


2. Create the keytab:

```
kadmin.local -q "xst -k /etc/security/keytabs/spark.keytab
spark/blue1@EXAMPLE.COM"
```

3. Create a spark user and add it to the hadoop group. (Do this for every node of your cluster.)

```
useradd spark -g hadoop
```

4. Make spark the owner of the newly-created keytab:

```
chown spark:hadoop /etc/security/keytabs/spark.keytab
```

5. Limit access: make sure user spark is the only user with access to the keytab:

```
chmod 400 /etc/security/keytabs/spark.keytab
```

In the following steps, user `spark` runs the Spark Pi example in a Kerberos-enabled environment:

```
su spark
kinit -kt /etc/security/keytabs/spark.keytab spark/blue1@EXAMPLE.COM
cd /usr/hdp/current/spark-client/
./bin/spark-submit --class org.apache.spark.examples.SparkPi \
  --master yarn-cluster \
  --num-executors 1 \
  --driver-memory 512m \
  --executor-memory 512m \
  --executor-cores 1 \
  lib/spark-examples*.jar 10
```

5.1. Accessing the Hive Metastore in Secure Mode

Requirements for accessing the Hive Metastore in secure mode (with Kerberos):

- The Spark thrift server must be co-located with the Hive thrift server.
- The `spark` user must be able to access the Hive keytab.
- In yarn-client mode on a secure cluster you can use `HiveContext` to access the Hive Metastore. (`HiveContext` is not supported for yarn-cluster mode on a secure cluster.)

6. Using Spark with HDFS

Specifying Compression

To specify compression in Spark-shell when writing to HDFS, use code similar to:

```
rdd.saveAsHadoopFile("/tmp/spark_compressed",  
"org.apache.hadoop.mapred.TextOutputFormat",  
compressionCodecClass="org.apache.hadoop.io.compress.GzipCodec")
```

7. Troubleshooting Spark

When you run a Spark job, you will see a standard set of console messages. In addition, the following information is available:

- A list of running applications, where you can retrieve the application ID and check the application log:

```
yarn application -list
```

```
yarn logs -applicationId <app_id>
```

- Check the Spark environment for a specific job:

```
http://<host>:8088/proxy/<job_id>/environment/
```

Specific Issues

The following paragraphs describe specific issues and possible solutions:

Issue: Job stays in "accepted" state; it doesn't run. This can happen when a job requests more memory or cores than available.

Solution: Assess workload to see if any resources can be released. You might need to stop unresponsive jobs to make room for the job.

Issue: Insufficient HDFS access. This can lead to errors such as the following:

```
"Loading data to table default.testtable
Failed with exception
Unable to move sourcehdfs://blue1:8020/tmp/hive-spark/hive_2015-03-04_
12-45-42_404_3643812080461575333-1/-ext-10000/kv1.txt to destination
hdfs://blue1:8020/apps/hive/warehouse/testtable/kv1.txt"
```

Solution: Make sure the user or group running the job has sufficient HDFS privileges to the location.

Issue: Wrong host in Beeline, shows error as invalid URL:

```
Error: Invalid URL: jdbc:hive2://localhost:10001 (state=08S01,code=0)
```

Solution: Specify the correct Beeline host assignment.

Issue: Error: closed SQLContext.

Solution: Restart the Thrift server.