

---

# **sparkhpc Documentation**

***Release 0.1***

**Rok Roskar**

**Oct 01, 2018**



---

## Contents

---

<b>1</b>	<b>Installation</b>	<b>3</b>
1.1	From pypi . . . . .	3
1.2	From source . . . . .	3
<b>2</b>	<b>Usage</b>	<b>5</b>
2.1	Command line . . . . .	5
2.2	Python code . . . . .	6
2.3	Jupyter notebook . . . . .	6
<b>3</b>	<b>Dependencies</b>	<b>7</b>
3.1	Python . . . . .	7
3.2	System configuration . . . . .	7
3.3	Job templates . . . . .	7
<b>4</b>	<b>Using other schedulers</b>	<b>9</b>
<b>5</b>	<b>Jupyter notebook</b>	<b>11</b>
5.1	Setup . . . . .	11
5.2	Launching . . . . .	11
<b>6</b>	<b>Contributing</b>	<b>13</b>
<b>7</b>	<b>API</b>	<b>15</b>
	<b>Python Module Index</b>	<b>19</b>



This package tries to greatly simplify deploying and managing [Apache Spark](#) clusters on HPC resources.



# CHAPTER 1

---

## Installation

---

### 1.1 From pypi

```
$ pip install sparkhpc
```

### 1.2 From source

```
$ python setup.py install
```

This will install the python package to your default package directory as well as the `sparkcluster` and `hpcnotebook` command-line scripts.





There are two options for using this library: from the command line or directly from python code.

## 2.1 Command line

### 2.1.1 Get usage info

```
Usage: sparkcluster [OPTIONS] COMMAND [ARGS]...

Options:
  --scheduler [lsf|slurm]  Which scheduler to use
  --help                  Show this message and exit.

Commands:
  info    Get info about currently running clusters
  launch  Launch the Spark master and workers within a...
  start   Start the spark cluster as a batch job
  stop    Kill a currently running cluster ('all' to...

$ sparkcluster start --help
Usage: sparkcluster start [OPTIONS] NCORES

    Start the spark cluster as a batch job

Options:
  --walltime TEXT          Walltime in HH:MM format
  --jobname TEXT           Name to use for the job
  --template TEXT          Job template path
  --memory-per-executor INTEGER Memory to reserve for each executor (i.e. the
                           JVM) in MB
  --memory-per-core INTEGER Memory per core to request from scheduler in
                           MB
```

(continues on next page)

(continued from previous page)

<code>--cores-per-executor</code>	<code>INTEGER</code>	Cores per executor
<code>--spark-home</code>	<code>TEXT</code>	Location of the Spark distribution
<code>--wait</code>		Wait until the job starts
<code>--help</code>		Show this message and exit.

## 2.1.2 Start a cluster

```
$ sparkcluster start 10
```

## 2.1.3 Get information about currently running clusters

```
$ sparkcluster info
----- Cluster 0 -----
Job 31454252 not yet started

$ sparkcluster info
----- Cluster 0 -----
Number of cores: 10
master URL: spark://10.11.12.13:7077
Spark UI: http://10.11.12.13:8080
```

## 2.1.4 Stop running clusters

```
$ sparkcluster stop 0
Job <31463649> is being terminated
```

## 2.2 Python code

```
from sparkhpc import sparkjob
import findspark
findspark.init() # this sets up the paths required to find spark libraries
import pyspark

sj = sparkjob.sparkjob(ncores=10)

sj.wait_to_start()

sc = sj.start_spark()

sc.parallelize(...)
```

## 2.3 Jupyter notebook

sparkhpc gives you nicely formatted info about your jobs and clusters in the jupyter notebook - see the [example notebook](#).

### 3.1 Python

- `click`
- `findspark`

These are installable via `pip install`.

### 3.2 System configuration

- Spark installation in `~/spark` OR wherever `SPARK_HOME` points to
- java distribution (set `JAVA_HOME`)
- `mpirun` in your path

### 3.3 Job templates

Simple job templates for the currently supported schedulers are included in the distribution. If you want to use your own template, you can specify the path using the `--template` flag to `start`. See the [included templates](#) for an example. Note that the variable names in curly braces, e.g. `{jobname}` will be used to inject runtime parameters. Currently you must specify `walltime`, `ncores`, `memory`, `jobname`, and `spark_home`. If you want to significantly alter the job submission, the best would be to subclass the relevant scheduler class (e.g. `LSFSparkCluster`) and override the `submit` method.



---

## Using other schedulers

---

The LSF and SLURM schedulers are currently supported. However, adding support for other schedulers is rather straightforward (see the `LSFSparkJob` and `SLURMSparkJob` implementations as examples). Please submit a pull request if you implement a new scheduler or get in touch if you need help!

To implement support for a new scheduler you should subclass `SparkCluster`. You must define the following *class* variables:

- `_peek()` (function to get stdout of the current job)
- `_submit_command` (command to submit a job to the scheduler)
- `_job_regex` (regex to get the job ID from return string of submit command)
- `_kill_command` (scheduler command to kill a job)
- `_get_current_jobs` (scheduler command to return jobid, status, jobname one job per line)

Note that `_get_current_jobs` should return a custom formatted string where the output looks like this:

```
JOB_NAME STAT JOBID
sparkcluster PEND 31610738
sparkcluster PEND 31610739
sparkcluster PEND 31610740
```

Depending on the scheduler's behavior, you may need to override some of the other methods as well.



## CHAPTER 5

---

### Jupyter notebook

---

Running Spark applications, especially with python, is really nice from the comforts of a [Jupyter notebook](#). This package includes the `hpcnotebook` script, which will setup and launch a secure, password-protected notebook for you.

```
$ hpcnotebook
Usage: hpcnotebook [OPTIONS] COMMAND [ARGS]...

Options:
  --port INTEGER  Port for the notebook server
  --help          Show this message and exit.

Commands:
  launch  Launch the notebook
  setup   Setup the notebook
```

### 5.1 Setup

Before launching the notebook, it needs to be configured. The script will first ask for a password for the notebook and generate a self-signed ssh certificate - this is done to prevent other users of your cluster to stumble into your notebook by chance.

### 5.2 Launching

On a computer cluster, you would normally either obtain an interactive job and issue the command below, or use this as a part of a batch submission script.

```
$ hpcnotebook launch
To access the notebook, inspect the output below for the port number, then point your
➞ browser to https://1.2.3.4:<port_number>
```

(continues on next page)

(continued from previous page)

```
[TerminalIPythonApp] WARNING | Subcommand `ipython notebook` is deprecated and will
↳be removed in future versions.
[TerminalIPythonApp] WARNING | You likely want to use `jupyter notebook` in the future
[I 15:43:12.022 NotebookApp] Serving notebooks from local directory: /cluster/home/
↳roskarr
[I 15:43:12.022 NotebookApp] 0 active kernels
[I 15:43:12.022 NotebookApp] The Jupyter Notebook is running at: https://[all ip
↳addresses on your system]:8889/
[I 15:43:12.022 NotebookApp] Use Control-C to stop this server and shut down all
↳kernels (twice to skip confirmation).
```

In this case, you could set up a port forward to host 1.2.3.4 and instruct your browser to connect to `https://1.2.3.4:8889`.

Inside the notebook, it is straightforward to set up the `SparkContext` using the `sparkhpc` package (see above).



## CHAPTER 6

---

### Contributing

---

Please submit an issue if you discover a bug or have a feature request! Pull requests also very welcome.



---

```
class sparkhpc.lsfsparkjob.LSFSparkJob (clusterid=None,      jobid=None,      ncores=4,
                                         cores_per_executor=1,  walltime='00:30',  mem-
                                         ory_per_core=2000,  memory_per_executor=None,
                                         jobname='sparkcluster',  template=None,  ex-
                                         tra_scheduler_options="",      config_dir=None,
                                         spark_home=None, scheduler=None)
```

Bases: *sparkhpc.sparkjob.SparkJob*

Class for submitting spark jobs with the LSF scheduler

```
class sparkhpc.slurmsparkjob.SLURMSparkJob (walltime='00:30', **kwargs)
```

Bases: *sparkhpc.sparkjob.SparkJob*

Class for submitting spark jobs with the SLURM scheduler

See the *SparkJob* class for keyword descriptions.

```
class sparkhpc.sparkjob.SparkJob (clusterid=None,      jobid=None,      ncores=4,
                                   cores_per_executor=1,  walltime='00:30',  mem-
                                   ory_per_core=2000,  memory_per_executor=None,
                                   jobname='sparkcluster',  template=None,  ex-
                                   tra_scheduler_options="",      config_dir=None,
                                   spark_home=None, scheduler=None)
```

Bases: object

Generic SparkJob class

To implement other schedulers, you must simply extend this class and define some class variables:

- *\_peek\_command* (command to get stdout of current job)
- *\_submit\_command* (command to submit a job to the scheduler)
- *\_job\_regex* (regex to get the job ID from return string of submit command)
- *\_kill\_command* (scheduler command to kill a job)
- *\_get\_current\_jobs* (scheduler command to return jobid, status, jobname one job per line)

See the LSFSparkJob class for an example.

```
__init__(clusterid=None, jobid=None, ncores=4, cores_per_executor=1, walltime='00:30',
         memory_per_core=2000, memory_per_executor=None, jobname='sparkcluster', template=None,
         extra_scheduler_options="", config_dir=None, spark_home=None, scheduler=None)
```

Creates a SparkJob

Parameters:

**clusterid: int** if a spark cluster is already running, initialize this SparkJob with its metadata

**jobid: int** same as *clusterid* but using directly the scheduler job ID

**ncores: int** number of cores to request

**walltime: string** walltime in *HH:MM* format as a string

**memory\_per\_core: int** memory to request per core from the scheduler in MB

**memory\_per\_executor: int** memory to give to each spark executor (i.e. the jvm part) in MB If using pyspark and python workers need a lot of memory, this should be less than *memory\_per\_core \* ncores*.

**jobname: string** name for the job - only used for the scheduler

**template: file path** custom template to use for job submission

**extra\_scheduler\_options: string** A string with custom options for the scheduler

**config\_dir: directory path** path to spark configuration directory

**spark\_home:** path to spark directory; default is the *SPARK\_HOME* environment variable, and if it is not set it defaults to *~/spark*

**scheduler: string** specify manually which scheduler you want to use; usually the automatic determination will work fine so this should not be used

Example usage:

```
from sparkhpc.sparkjob import sparkjob import findspark findspark.init() # this sets up the paths
required to find spark libraries import pyspark

sj = sparkjob(ncores=10)

sj.wait_to_start()

sc = pyspark.SparkContext(master=sj.master_url())

sc.parallelize(...)
```

**classmethod current\_clusters()**

Determine which Spark clusters are currently running or in the queue

**job\_started()**

Check whether the job is running already or not

**master\_ui()**

Get the UI address of the Spark master

**master\_url()**

Get the URL of the Spark master

```
start_spark(spark_conf=None, executor_memory=None, profiling=False,
            graphframes_package='graphframes:graphframes:0.3.0-spark2.0-s_2.11',
            extra_conf=None)
```

Launch a SparkContext

#### Parameters

**spark\_conf: path** path to a spark configuration directory

**executor\_memory: string** executor memory in java memory string format, e.g. '4G' If *None*, *memory\_per\_executor* is used.

**profiling: boolean** whether to turn on python profiling or not

**graphframes\_package: string** which graphframes to load - if it isn't found, spark will attempt to download it

**extra\_conf: dict** additional configuration options

**stop ()**

Stop the current job

**submit ()**

Write job file to current working directory and submit to the scheduler

**wait\_to\_start (timeout=60)**

Wait for the job to start or until timeout, whichever comes first

`sparkhpc.sparkjob.start_cluster (memory, cores_per_executor=1, timeout=30, spark_home=None)`

Start the spark cluster

This is the script used to launch spark on the compute resources assigned by the scheduler.

#### Parameters

**memory: string** memory specified using java memory format

**timeout: int** time in seconds to wait for the master to respond

**spark\_home: directory path** path to base spark installation



### S

`sparkhpc`, [17](#)  
`sparkhpc.lsfsparkjob`, [15](#)  
`sparkhpc.slurmsparkjob`, [15](#)  
`sparkhpc.sparkjob`, [15](#)





## Symbols

`__init__()` (sparkhpc.sparkjob.SparkJob method), 16

## C

`current_clusters()` (sparkhpc.sparkjob.SparkJob class method), 16

## J

`job_started()` (sparkhpc.sparkjob.SparkJob method), 16

## L

LSFSparkJob (class in sparkhpc.lsfsparkjob), 15

## M

`master_ui()` (sparkhpc.sparkjob.SparkJob method), 16

`master_url()` (sparkhpc.sparkjob.SparkJob method), 16

## S

SLURMSparkJob (class in sparkhpc.slurmsparkjob), 15

sparkhpc (module), 17

sparkhpc.lsfsparkjob (module), 15

sparkhpc.slurmsparkjob (module), 15

sparkhpc.sparkjob (module), 15

SparkJob (class in sparkhpc.sparkjob), 15

`start_cluster()` (in module sparkhpc.sparkjob), 17

`start_spark()` (sparkhpc.sparkjob.SparkJob method), 16

`stop()` (sparkhpc.sparkjob.SparkJob method), 17

`submit()` (sparkhpc.sparkjob.SparkJob method), 17

## W

`wait_to_start()` (sparkhpc.sparkjob.SparkJob method), 17