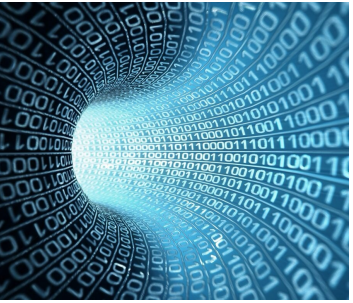

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Hadoop hands-on: using MapReduce / Spark on the cloud

Cécile Cavet

cecile.cavet at apc.univ-paris7.fr

Centre François Arago (FACe), Laboratoire APC, Université Paris Diderot
LabEx UnivEarthS

December 10, 2015

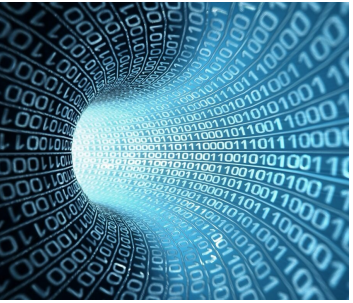

Plan

Big Data
- VVVV...
- in Astronomy
- Technologies

Hadoop
- Core
- Ecosystem
- Cluster

Cloud computing
- Why ?
- StratusLab
- SlipStream

Tutorial
- StratusLab
- SlipStream
- Hadoop
- HDFS
- MapReduce
- Spark

Links

1 Big Data

2 Hadoop

3 Cloud computing

4 Tutorial

5 Links

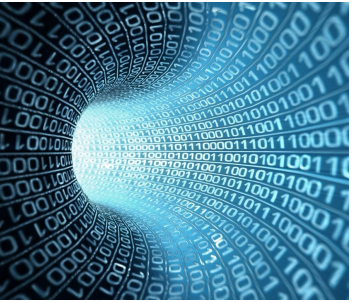

Big Data

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Management of huge volume of data:

- Collecte.
- Storage.
- Treatment.
- Visualisation.
- Analyse.

Big Data
vvvv...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Big Data bottleneck

4V of numerical data:

- Volume.
- Variety: diversity, heterogeneity.
- Velocity: producing, transmitting, and analyzing data.
- Veracity: noise, precision.

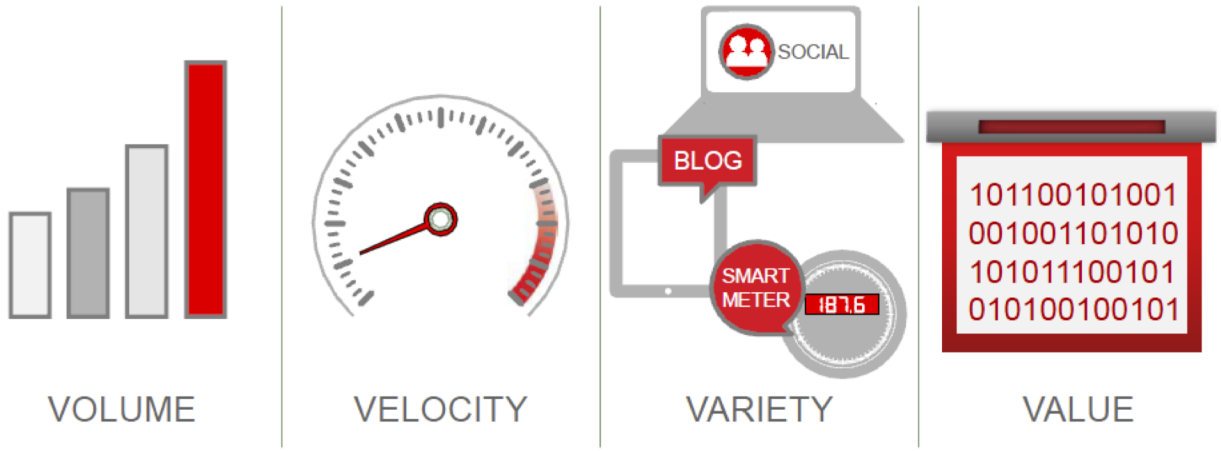

Big Data in Astronomy

Big Data
- VVVV...
- in Astronomy
- Technologies

Hadoop
- Core
- Ecosystem
- Cluster

Cloud computing
- Why ?
- StratusLab
- SlipStream

Tutorial
- StratusLab
- SlipStream
- Hadoop
- HDFS
- MapReduce
- Spark

Links

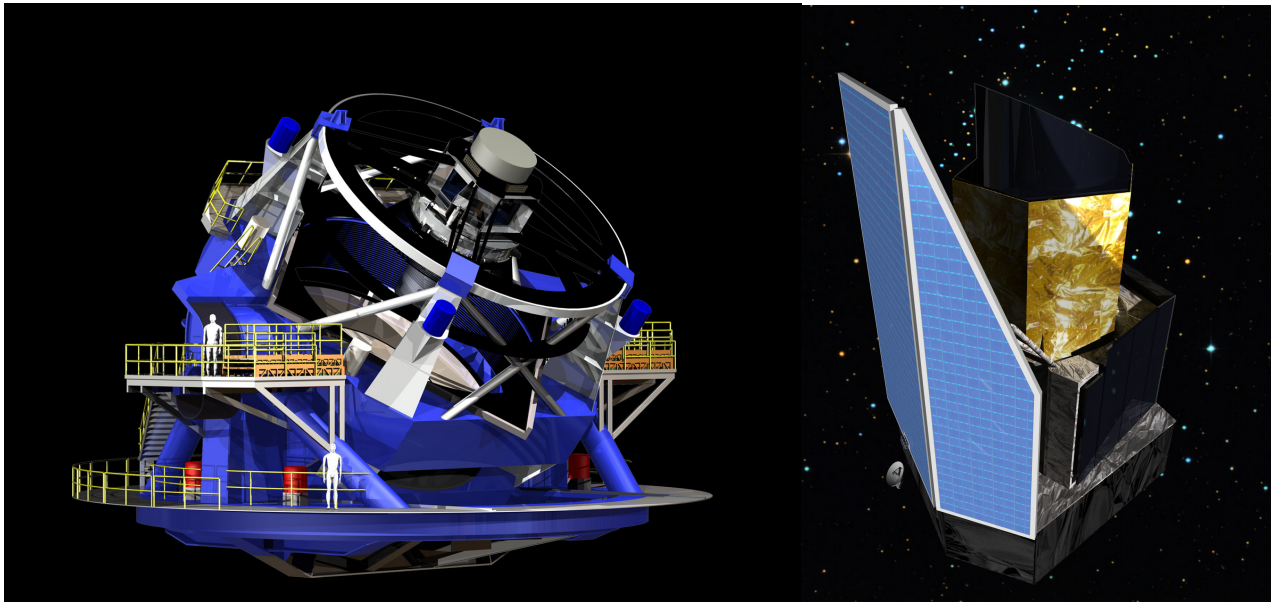

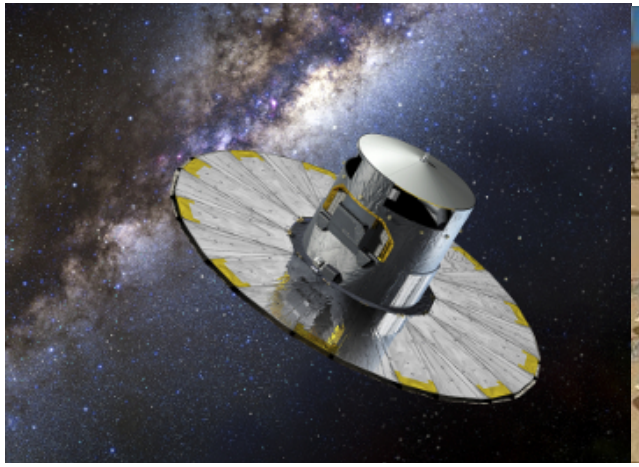

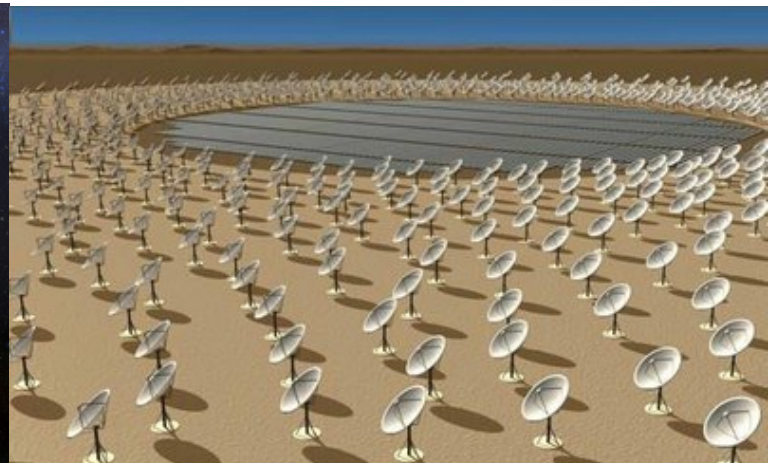

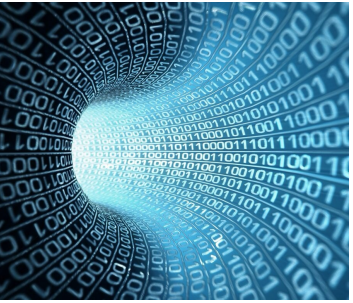

Big Data in Astronomy

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Sky Survey Projects	Data Volume
DPOSS (The Palomar Digital Sky Survey)	3 TB
2MASS (The Two Micron All-Sky Survey)	10 TB
GBT (Green Bank Telescope)	20 PB
GALEX (The Galaxy Evolution Explorer)	30 TB
SDSS (The Sloan Digital Sky Survey)	40 TB
SkyMapper Southern Sky Survey	500 TB
PanSTARRS (The Panoramic Survey Telescope and Rapid Response System)	~ 40 PB expected
LSST (The Large Synoptic Survey Telescope)	~ 200 PB expected
SKA (The Square Kilometer Array)	~ 4.6 EB expected

Figure: Data volumes of different sky survey projects (http://datascience.codata.org/articles/10.5334/dsj-2015-011/).

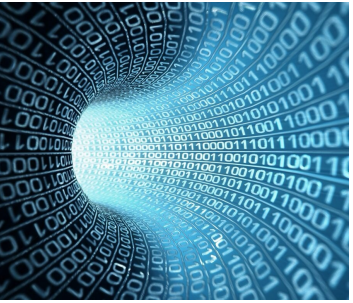

Big Data technologies

Big Data
- VVVV...
- in Astronomy
- Technologies

Hadoop
- Core
- Ecosystem
- Cluster

Cloud computing
- Why ?
- StratusLab
- SlipStream

Tutorial
- StratusLab
- SlipStream
- Hadoop
- HDFS
- MapReduce
- Spark

Links

- Hadoop ecosystem.
- NoSQL (non relational Data Base).
- Cloud computing.

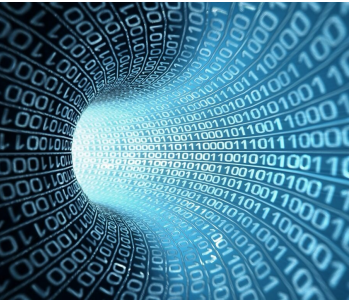

Hadoop core

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

- Google in 2004: proprietary solution.
- Apache since 2008: open-source solution[1].

- Distributed file system: **HDFS** (Hadoop DFS).
- Distributed computing: **Hadoop MapReduce**
  - parallel processing of large data sets.
  - execution framework ➡ Hadoop 1.0.
- Resource manager: **YARN** (Yet Another Resource Negotiator) ➡ Hadoop 2.0 (2013).

[1] http://hadoop.apache.org/

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

HDFS

- Large file storage (up to PB of data).
- Block storage ➡ block-sized chunks of default = 128 MB.
- Replication ➡ factor of default = 3.
- I/O optimized: write once, read many times.

HDFS

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

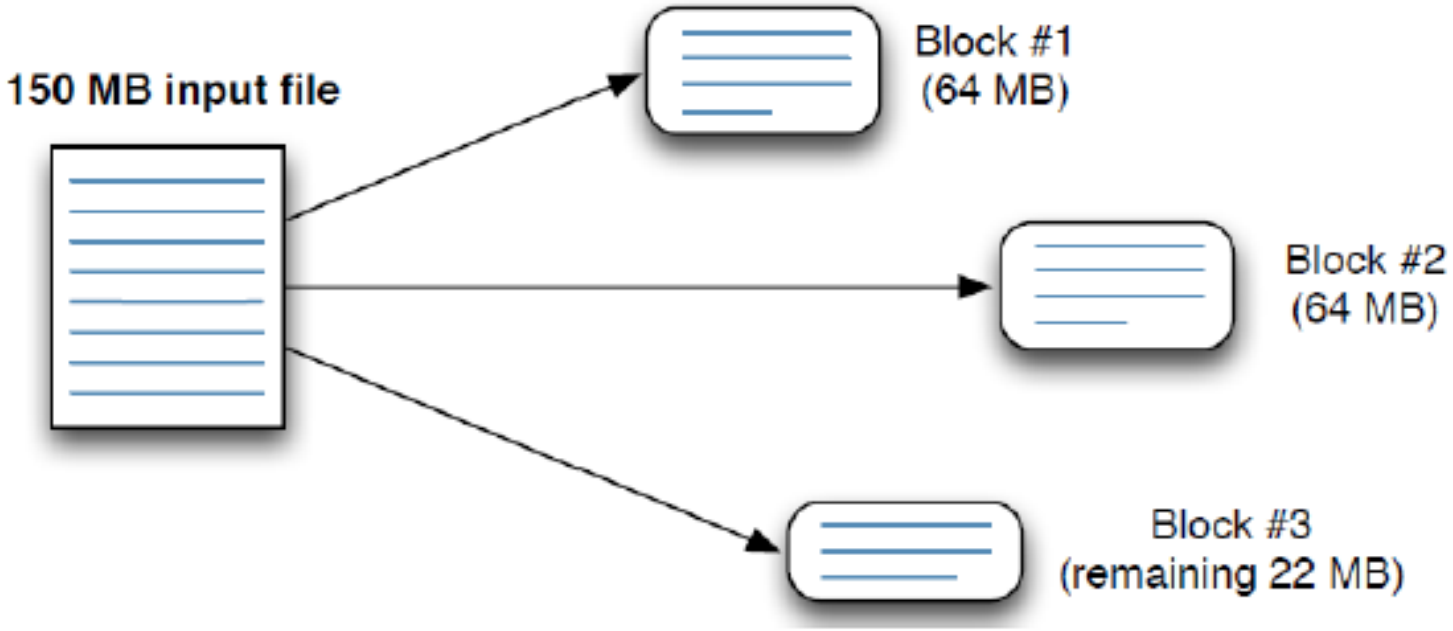

Figure: @INGENSI.

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

MapReduce

- Programming model: mapping, shuffling, sorting and reduction ➙ (key, value) functions.
- Algorithm for data analysis.
- Inherently parallel and implicit data flow.
- Data processing: batch, streaming.

MapReduce

Big Data
VVVV...
in Astronomy
Technologies

Hadoop
Core
Ecosystem
Cluster

Cloud computing
Why ?
StratusLab
SlipStream

Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark

Links

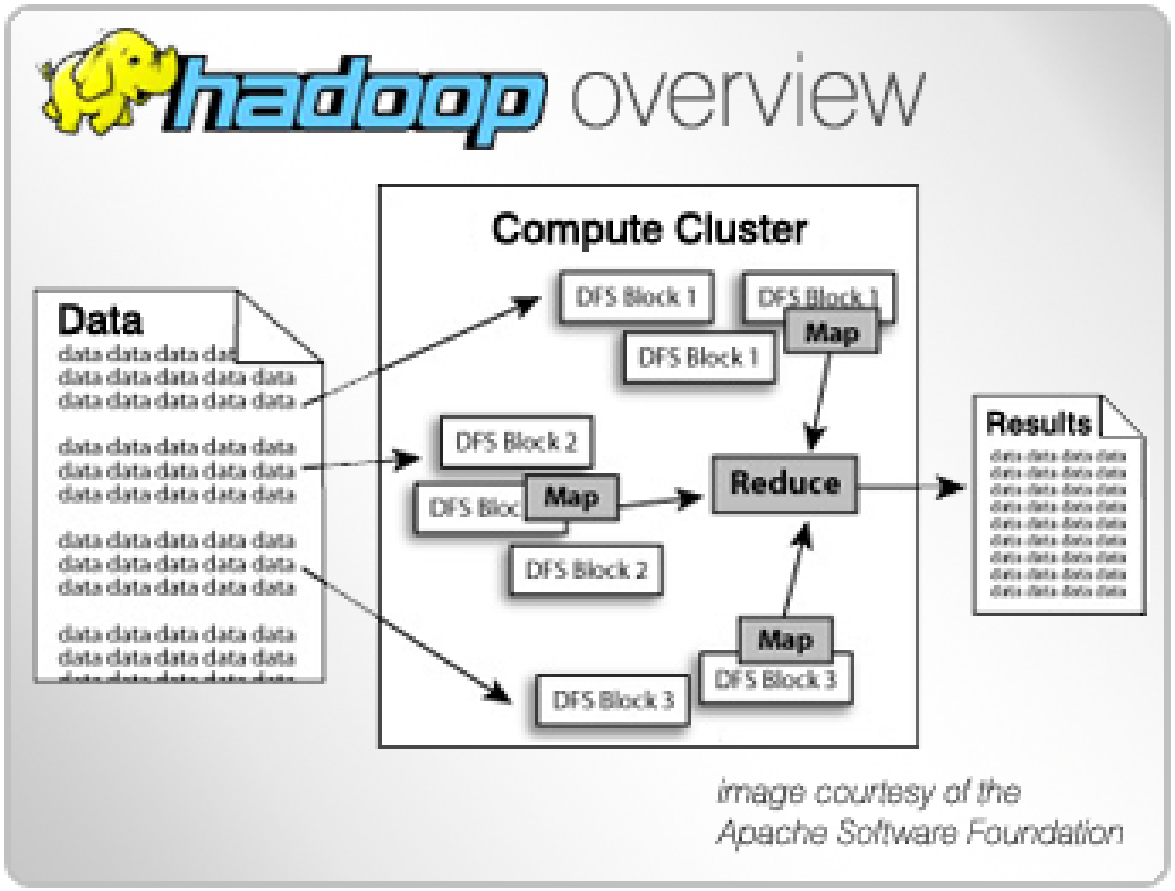

Figure: @Apache Hadoop.

YARN

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

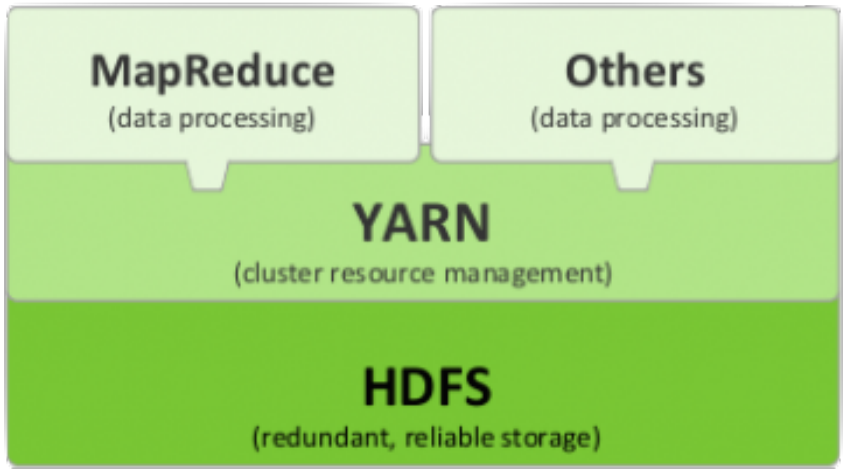

- Cluster resource manager: computing resources and job scheduling.
- General scheduler: based on containers of ressources (cpu and memory), several types (FIFO, CapacityScheduler...), services (MR, Spark...).
- High scalability.

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Hadoop API

- Java.
- Streaming: Python, Ruby...
- C++ (Hadoop Pipes).

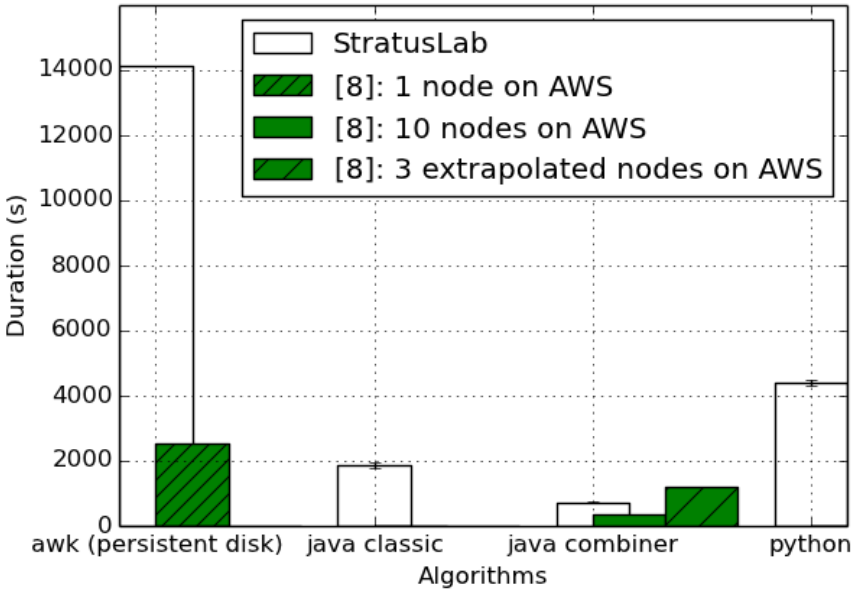

Figure: API performance comparison: analysis of meteorological data.

Hadoop ecosystem

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

- Distributed NoSQL DBMS: **HBase**.
- In-memory processing: **Spark** (2014).
- ...

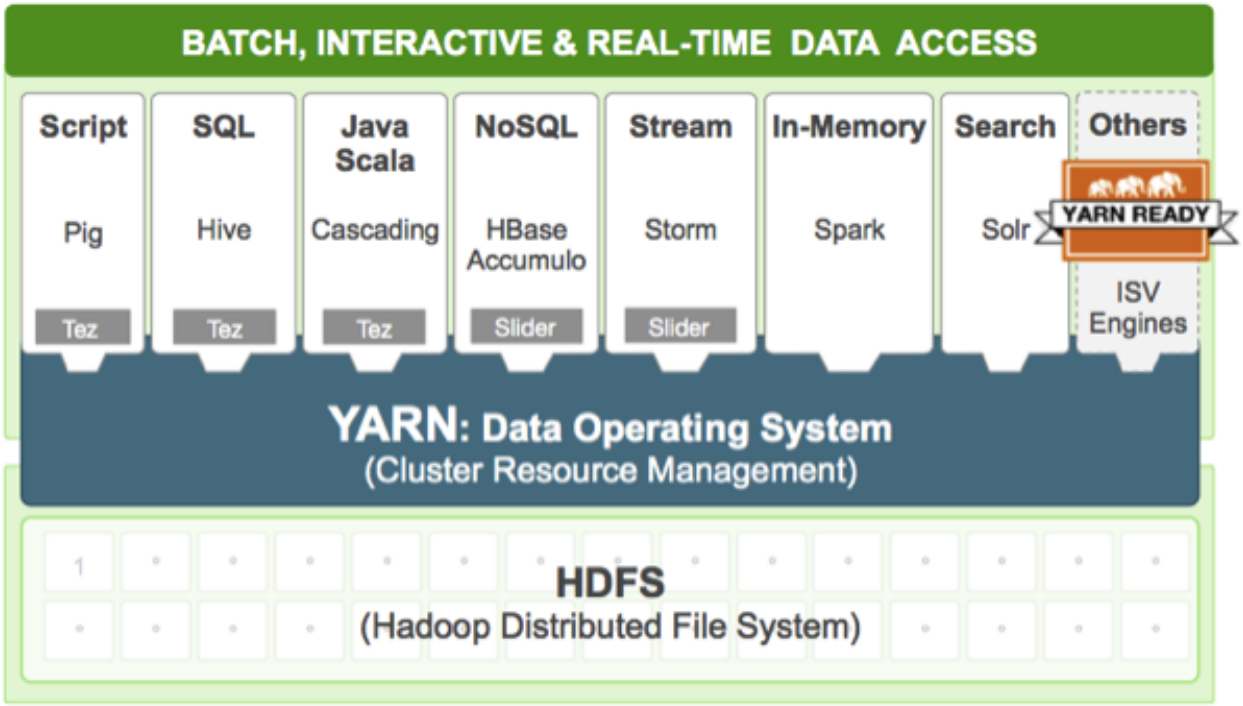

Figure: Cluster structure @Hortonworks.

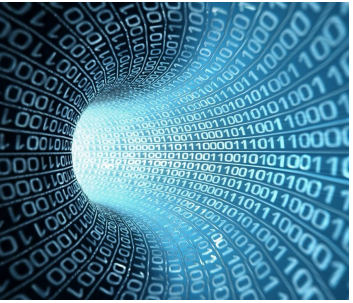

Spark

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Lightning-Fast Cluster Computing

- In-memory caching processing.
- Based on RDD (Resilient Distributed Datasets).
- **Spark** vs **MapReduce**: jobs faster (100 $\times$), easier to write and beyond (80 high level operators).

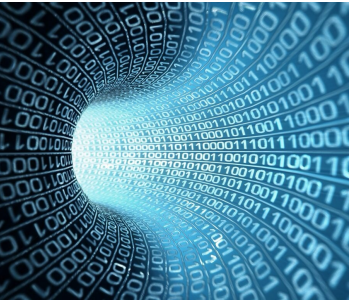

Spark

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

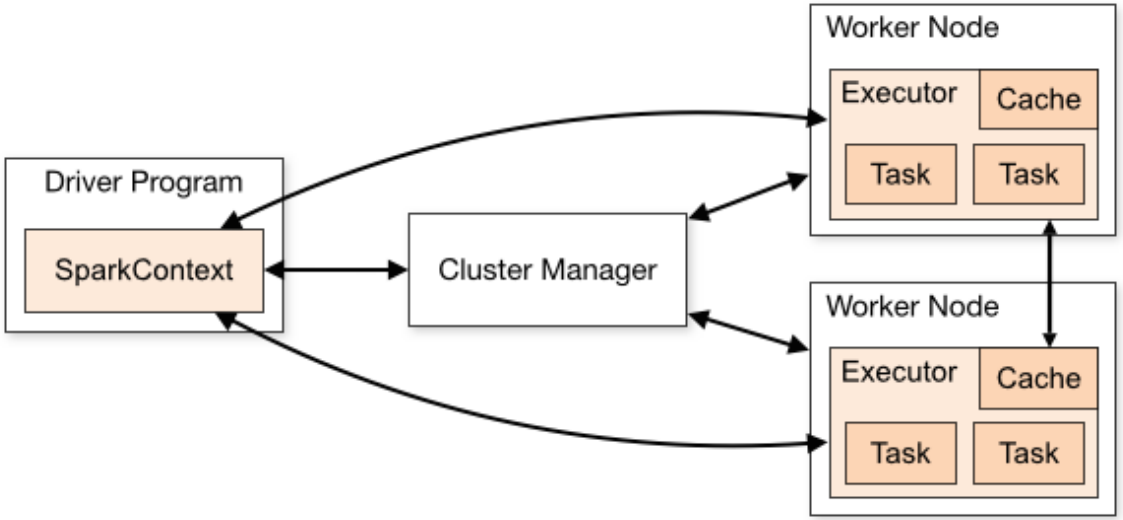

Figure: @Spark.

- API: Java, Scala, Python, R.
- Data processing: batch, interactive (developing Spark applications), streaming.
- Use: standalone, on Mesos or on YARN.

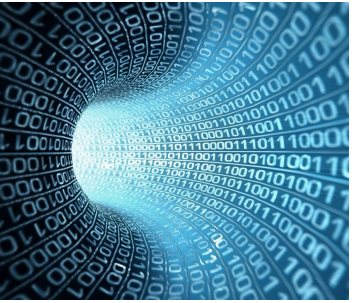

Hadoop distribution

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

- MapR.
- Cloudera.
- **Hortonworks**:

Hortonworks

  - Packaged distribution (Hadoop standalone distribution is instable).
  - Easy to install with scripts.
  - Optimised Hadoop parameters.

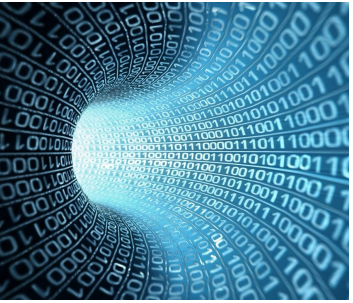

Hadoop cluster

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Properties:

- High-performance ➔ distributed computing and storage.
- Scalability ➔ horizontal: on-demand nodes; vertical: node ressources.
- Data locality ➔ brings the compute to the data.
- Fault-tolerant ➔ redundant storage, data replication.
- Portability: computing in JVM.
- High level operators.
- Monitoring.

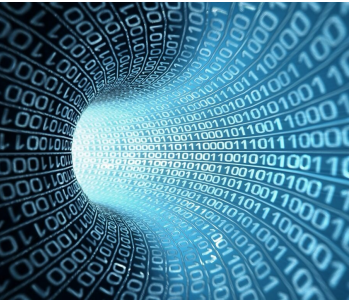

Hadoop cluster

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Services: HDFS

- **NameNode (NN)**:
  - catalog of block location.
  - block distribution on DN.
- **DataNode (DN)**: storage of data blocks send by NN.

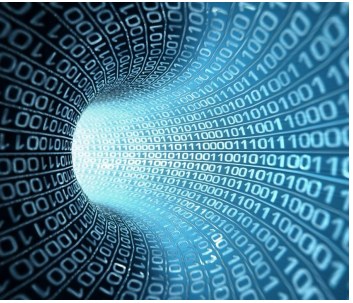

Hadoop cluster

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Services: YARN

- **ResourceManager (RM)**:
  - catalog of global resources.
  - resource distribution on MR jobs (application manager).
  - scheduler.
- **NodeManager (NM)**: catalog of node resources.
- **ApplicationMaster**:
  - resource negotiation with RM (1 per job).
  - containers starting.
  - task management with NM.
- **JobHistoryServer**: MR job history.

Hadoop cluster: overview

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

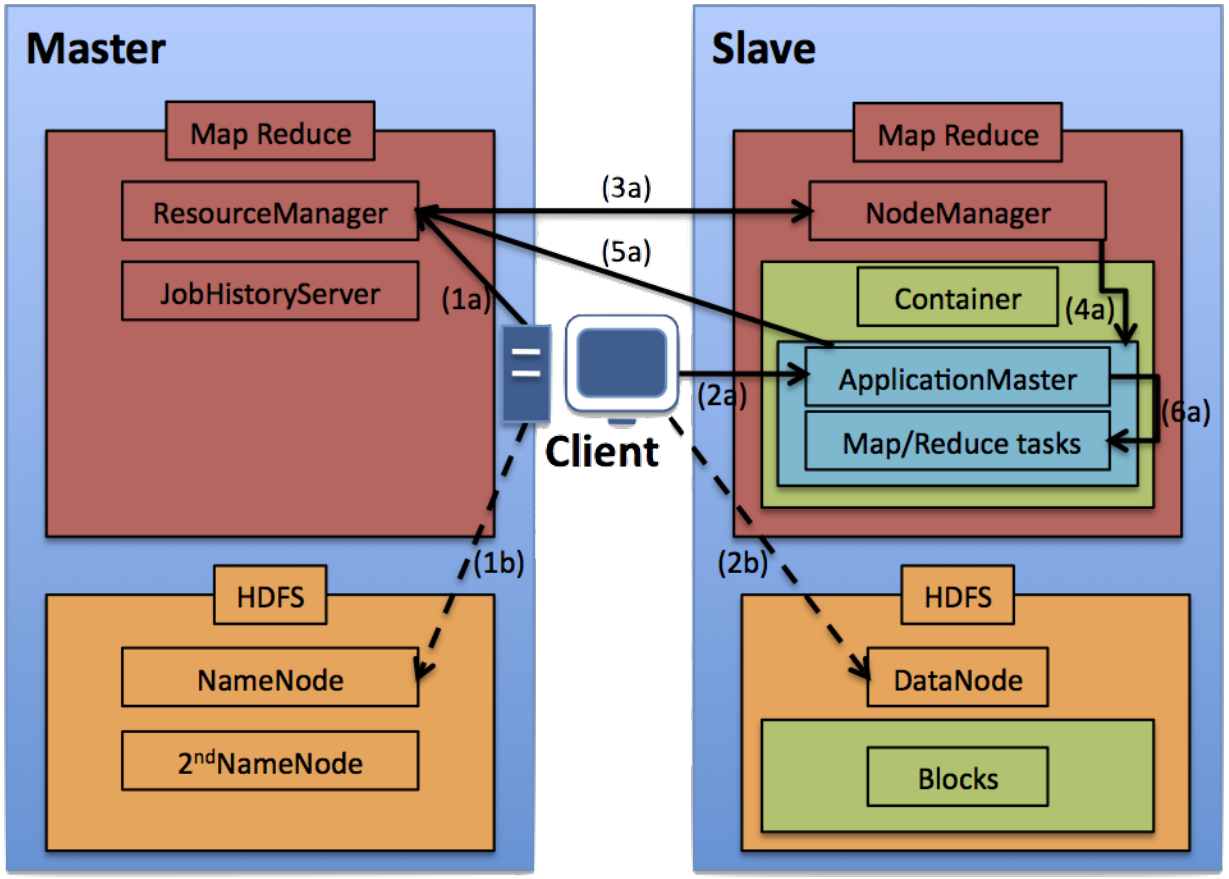

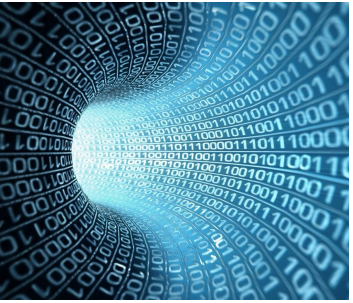

Cloud computing

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Why ?

- Few (no) Hadoop clusters in the academic community (Beowulf cluster with MPI in general) ➟ Hadoop technology too recent and too specific.
- **Hadoop-as-a-Service** already on the cloud:
  - On commercial clouds: AWS EC2 et S3 ➟ Amazon Elastic MapReduce[a] (PaaS type).
  - On academic cloud: StratusLab and OpenStack (IaaS type) ➟ Hadoop on-demand with SlipStream tool (PaaS type).

[a] http://aws.amazon.com/fr/elasticmapreduce/

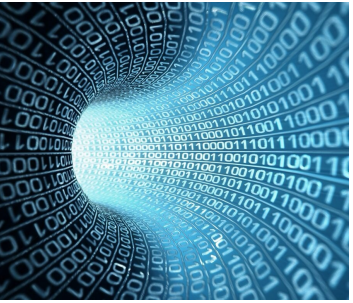

StratusLab Cloud

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

- Cloud manager: **StratusLab** with OpenNebula.
- Ressources**@LAL** (Orsay): 20 nodes, 630 cores, 1,1 TB of memory, 15 TB of space disk, 1 GbE/s of inter-nodes connexion.
- Services:
  - End-user client (http://stratuslab.eu/index.html).
  - MarketPlace (http://marketplace.stratuslab.eu).
  - Persistent disk (https://pdisk.lal.stratuslab.eu/svc-pdisk.html).
  - Resource monitoring (https://cloud.lal.stratuslab.eu/load.txt).

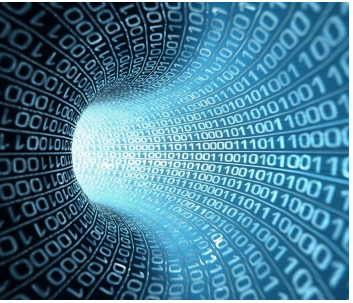

SlipStream

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

- PaaS.
- Sixsq private company (http://sixsq.com/products/slipstream/).
- Commercial connector for commercial clouds (i.e. AWS...).

➙ **Automatic deployment + Cloud infrastructure interoperability**.

And now: tutorial !

Big Data
VVVV...
in Astronomy
Technologies

Hadoop
Core
Ecosystem
Cluster

Cloud computing
Why ?
StratusLab
SlipStream

Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark

Links

Figure: @Software Carpentry.

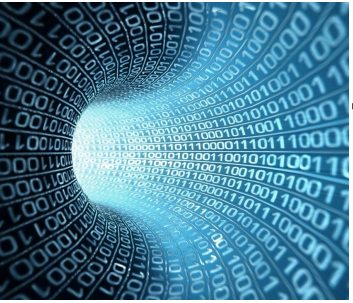

Tutorial

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Goal:

- Instanciate a cluster:
  - Create a Hadoop cluster.
  - On the StratusLab cloud.
  - Deployed by SlipStream.
- Use the cluster:
  - Discover HDFS.
  - Run analysis jobs:
    - MapReduce.
    - Spark.

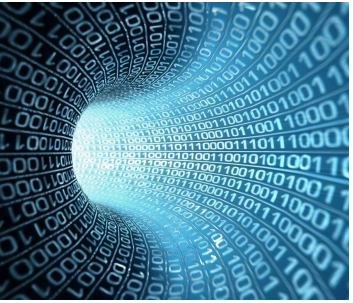

StratusLab

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

- **The cloud infrastructure.**
- TODO:
  - Check your account (if registered[2]).
  - Create SSH key (if not already done):

    ```
    $ ssh-keygen
    ```

  - Use SSH agent (not mandatory but else -i option):

    ```
    $ ssh-add ~/.ssh/id_rsa
    $ ssh-agent
    ```

[2] https://cloud.lal.stratuslab.eu/register/register/

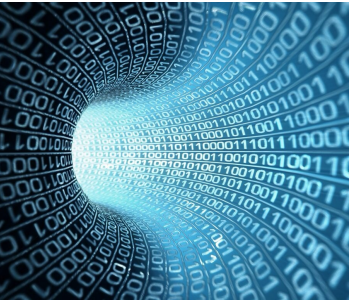

SlipStream

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

- **The environment manager.**
- API: Web interface.
- TODO:
  - Check your account (if registered[3]).
  - Register StratusLab's credentials and SSH public key.
    - General field: option ➡
      Keep running after deployment = **Always**.

[3] `https://slipstream.stratuslab.eu`

SlipStream: user credentials

Big Data
- VVVV...
- in Astronomy
- Technologies

Hadoop
- Core
- Ecosystem
- Cluster

Cloud computing
- Why ?
- StratusLab
- SlipStream

Tutorial
- StratusLab
- **SlipStream**
- Hadoop
- HDFS
- MapReduce
- Spark

Links

SlipStream. Dashboard Help cavet

Your profile

Regular user

/ Your profile Edit Delete

Summary

General

Description	Value
Select which cloud you want to use.	LAL
Keep running after deployment	Always
SSH Public Key(s) (keys must be separated by new line) Warning: Some clouds may take into account only the first key.	ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQCinoje6wMrzaeZw7I6gCLzE3yL3nEZlCvU RiKSMfiGIy/V99x0c8T0qyQN8crt9O1zFDrflkTuOoZZZasSHNiBV6FR2E3DU1X5qHZv

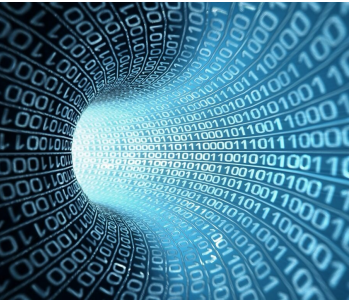

SlipStream: user credentials

Big Data
- VVVV...
- in Astronomy
- Technologies

Hadoop
- Core
- Ecosystem
- Cluster

Cloud computing
- Why ?
- StratusLab
- SlipStream

Tutorial
- StratusLab
- SlipStream
- Hadoop
- HDFS
- MapReduce
- Spark

Links

SlipStream. Dashboard Help cavet

Your profile

Regular user

/ Your profile Edit Delete

Summary

General

CC-IN2P3

IPHC

LAL

Description	Value
Number of VMs the user can start for this cloud	20
StratusLab account username	ccavet
StratusLab account password	•••••

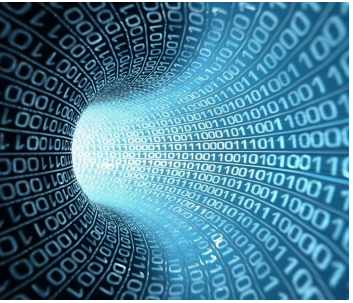

SlipStream

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

- TODO:
  - Launch the (small) Hadoop cluster:
    - 1 master + 1 slave: 2 CPU, 6 GB of memory, 30 GB of local disk, 10 GB of volatil disk.

SlipStream: Hadoop recipes

Big Data
- VVVV...
- in Astronomy
- Technologies

Hadoop
- Core
- Ecosystem
- Cluster

Cloud computing
- Why ?
- StratusLab
- SlipStream

Tutorial
- StratusLab
- **SlipStream**
- Hadoop
- HDFS
- MapReduce
- Spark

Links

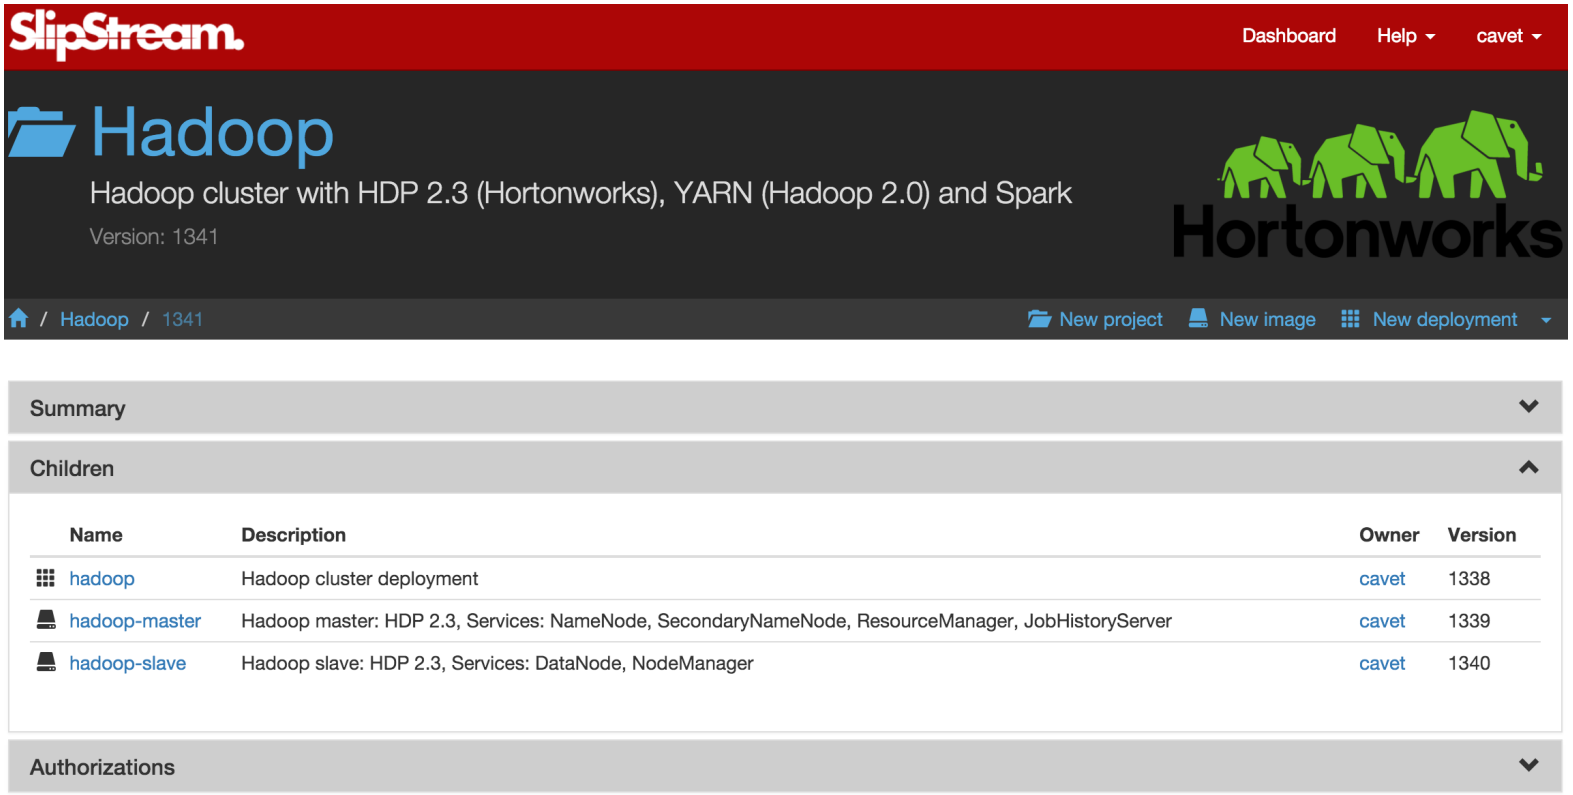

SlipStream: Hadoop deployment

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

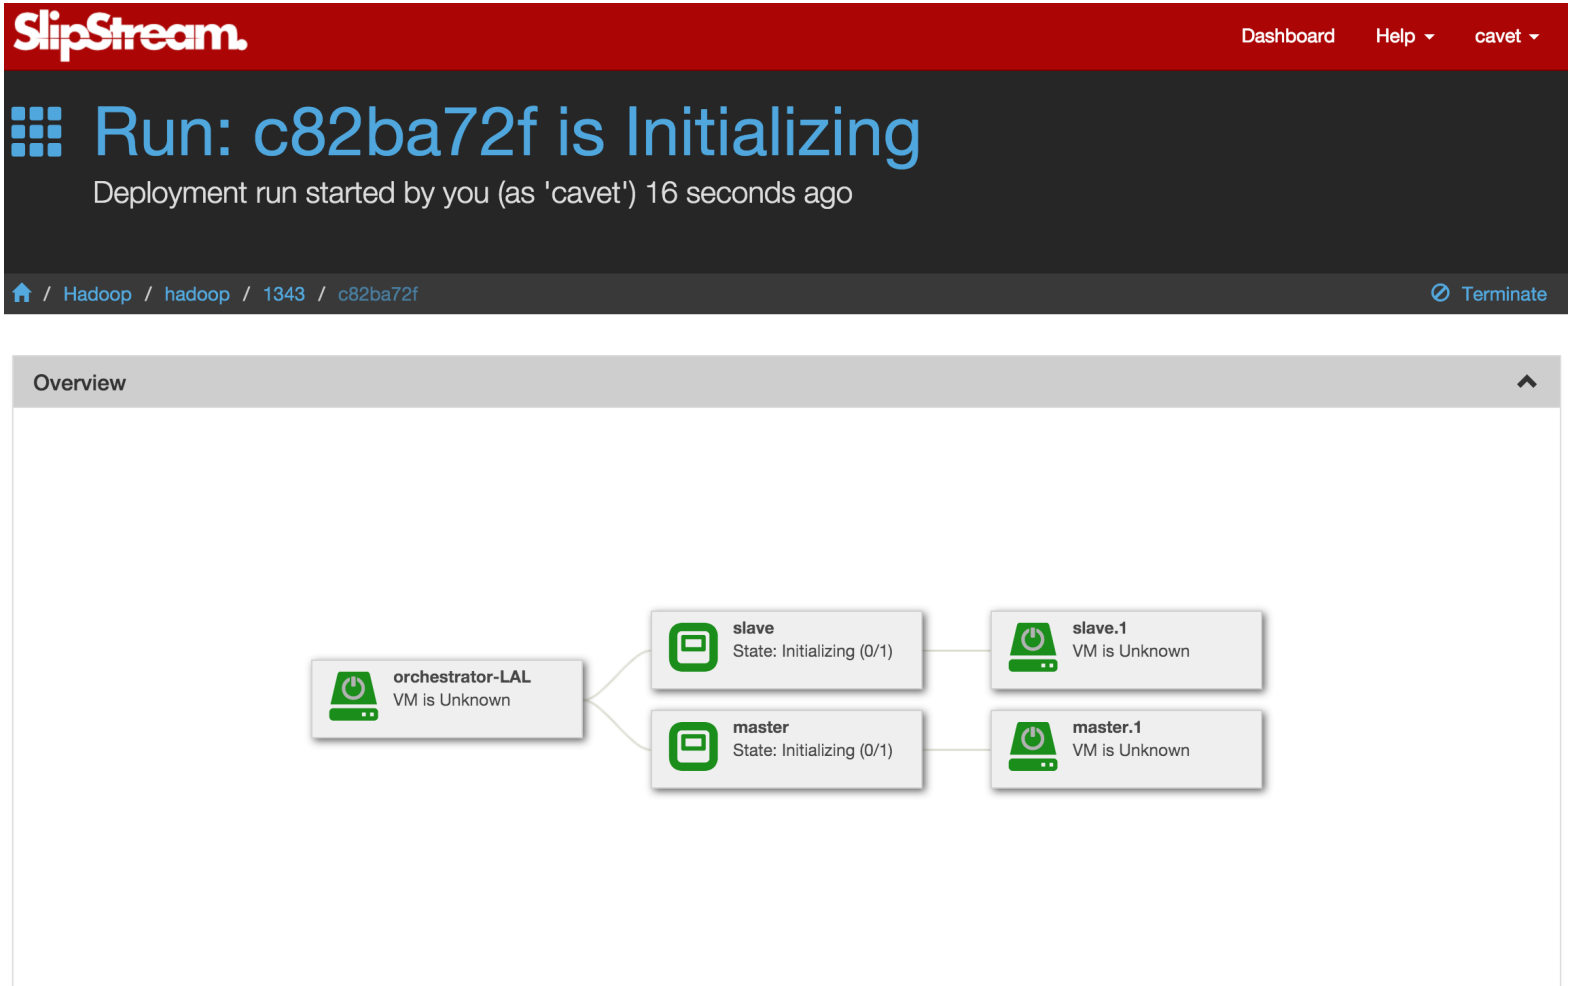

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Hadoop cluster

- **The working environment.**
- API: Command line + Web interface.
- TODO:
  - Wait the **Ready** state of SlipStream Run (∼10 min).
  - SSH connection as root user to the **master** node:

    ```
    $ ssh -Y root@onevm-XXX.lal.in2p3.fr
    ```
  - Check the cluster configuration on the Web interfaces.

Hadoop cluster overview

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Check cluster services:

```
root@master$ jps
  NameNode
  SecondaryNameNode
  ResourceManager
  JobHistoryServer
root@master$ ssh slave1
root@slave1$ jps
  DataNode
  NodeManager
root@slave1$ exit
```

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Hadoop cluster overview

Web interface access:

```
root$ firefox &
```

- http://master:50070 ➔ NameNode: HDFS manager.
- http://master:8088 ➔ ResourceManager: job manager.
- http://master:19888 ➔ JobHistoryServer: job history (accessible from ResourceManager).

➔ master and slave information.

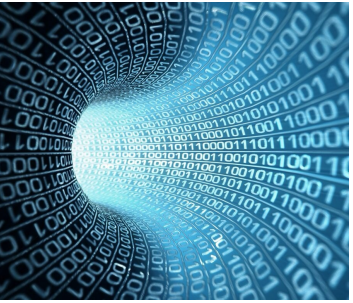

Hadoop Web interface: HDFS

Big Data
- VVVV...
- in Astronomy
- Technologies

Hadoop
- Core
- Ecosystem
- Cluster

Cloud computing
- Why ?
- StratusLab
- SlipStream

Tutorial
- StratusLab
- SlipStream
- Hadoop
- HDFS
- MapReduce
- Spark

Links

Namenode information | All Applications | JobHistory

master:50070/dfshealth.html#tab-overview

Hadoop | Overview | Datanodes | Datanode Volume Failures | Snapshot | Startup Progress | Utilities

Overview 'master:8020' (active)

Started:	Thu Dec 03 17:33:30 CET 2015
Version:	2.7.1.2.3.2.0-2950, r5cc60e0003e33aa98205f18bccaeaf36cb193c1c
Compiled:	2015-09-30T18:08Z by jenkins from (HEAD detached at 5cc60e0)
Cluster ID:	CID-a8fce6cc-d932-4099-94a8-2d084b1bc420
Block Pool ID:	BP-1918217548-134.158.75.180-1449160407971

Summary

Security is off.

Safemode is off.

13 files and directories, 2 blocks = 15 total filesystem object(s).

Heap Memory used 404.35 MB of 960 MB Heap Memory. Max Heap Memory is 960 MB.

Non Heap Memory used 35.03 MB of 36.13 MB Commited Non Heap Memory. Max Non Heap Memory is 214 MB.

Configured Capacity:	9.72 GB
DFS Used:	200.79 MB (2.02%)
Non DFS Used:	22.56 MB

Figure: NameNode.

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Hadoop Web interface: HDFS

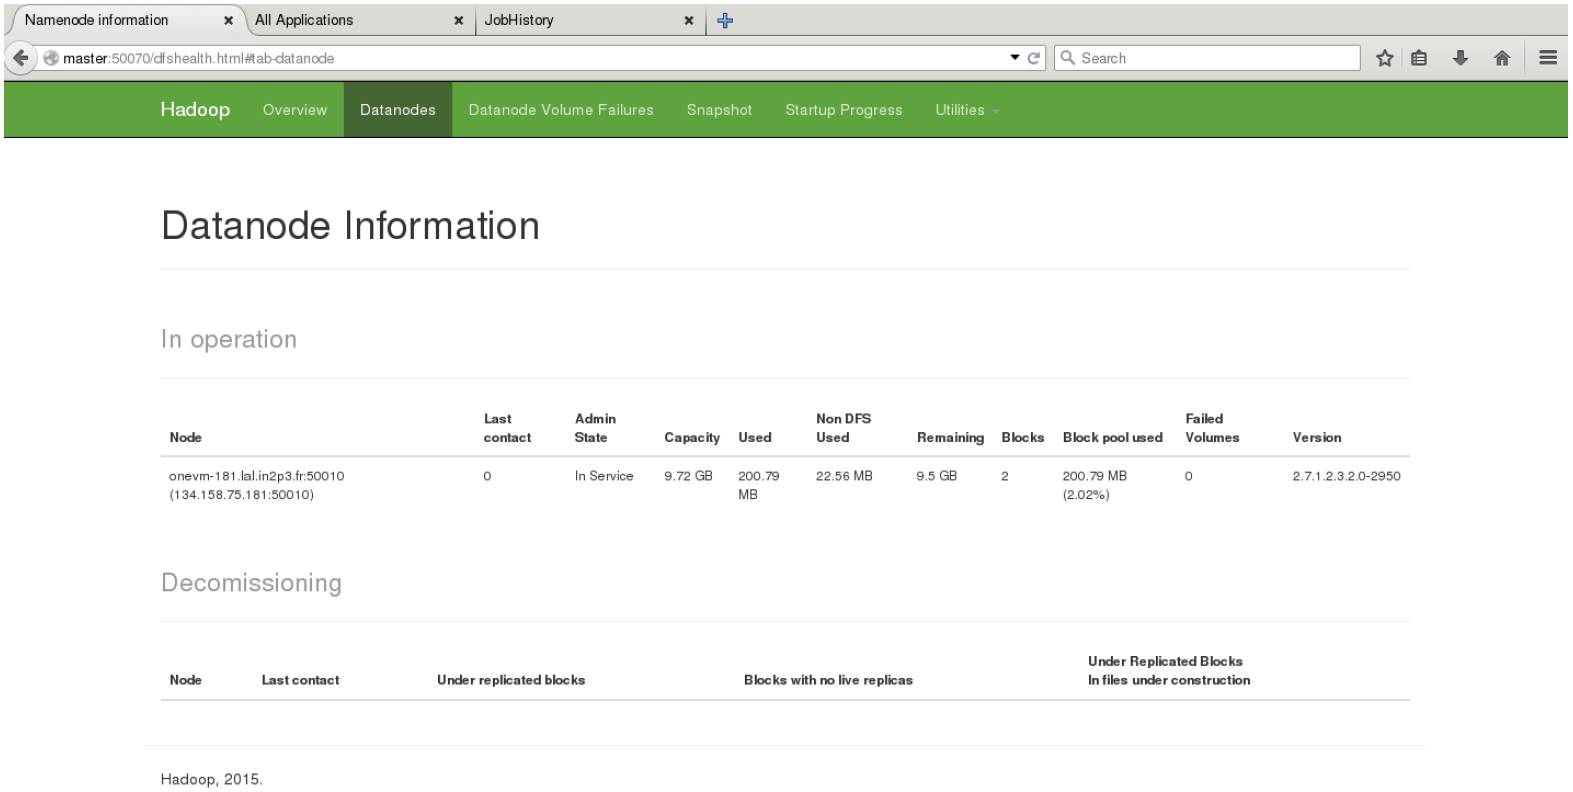

Figure: DataNode.

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Hadoop Web interface: job

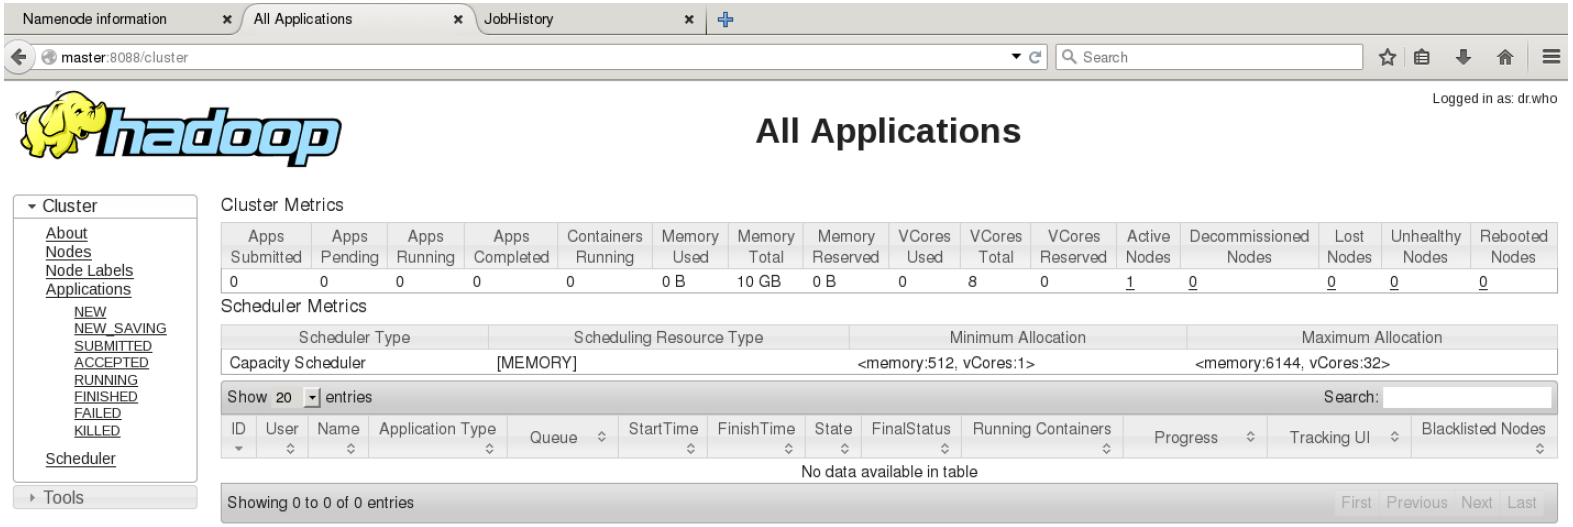

Figure: ResourceManager.

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Hadoop cluster overview

Users on Hadoop cluster:

- Cluster management: **hdfs, yarn, mapred**
  ➙ also command line API.
- Job running: **client**.

```
root$ su - client
client$ hdfs dfs
client$ mapred job
client$ yarn logs
```

Cluster use: HDFS

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

- TODO:
  - Create a new directory on HDFS and copy a local file inside.
  - Browse the directory with the Web interface.

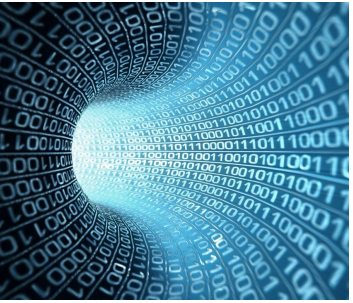

Hadoop command line

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

HDFS:

```
client$ hdfs dfs -mkdir data
client$ hdfs dfs -copyFromLocal \
  /usr/hdp/current/spark-client/README.md \
  data

client$ hdfs dfs -du -s -h data
3.5 K  data
client$ hdfs dfs -df -h
Filesystem              Size    Used Avail. Use%
hdfs://master:8020 9.7 G 200.8 M  9.5 G   2%
```

➡ the directory data is created on HDFS.

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Hadoop command line

HDFS:

```
client$ hdfs dfs -ls /
drwxrwxrwt   - yarn   hadoop 0 /app-logs
drwxr-xr-x   - hdfs   hadoop 0 /hdp
drwxr-xr-x   - mapred hadoop 0 /mr-history
drwxr-xr-x   - hdfs   hdfs   0 /user

client$ hdfs dfs -ls data
-rw-r--r-- 1 client client 3624 data/README.md
```

➟ the file README.md is replicated one time.

Hadoop Web interface for HDFS

Big Data
- VVVV...
- in Astronomy
- Technologies

Hadoop
- Core
- Ecosystem
- Cluster

Cloud computing
- Why ?
- StratusLab
- SlipStream

Tutorial
- StratusLab
- SlipStream
- Hadoop
- HDFS
- MapReduce
- Spark

Links

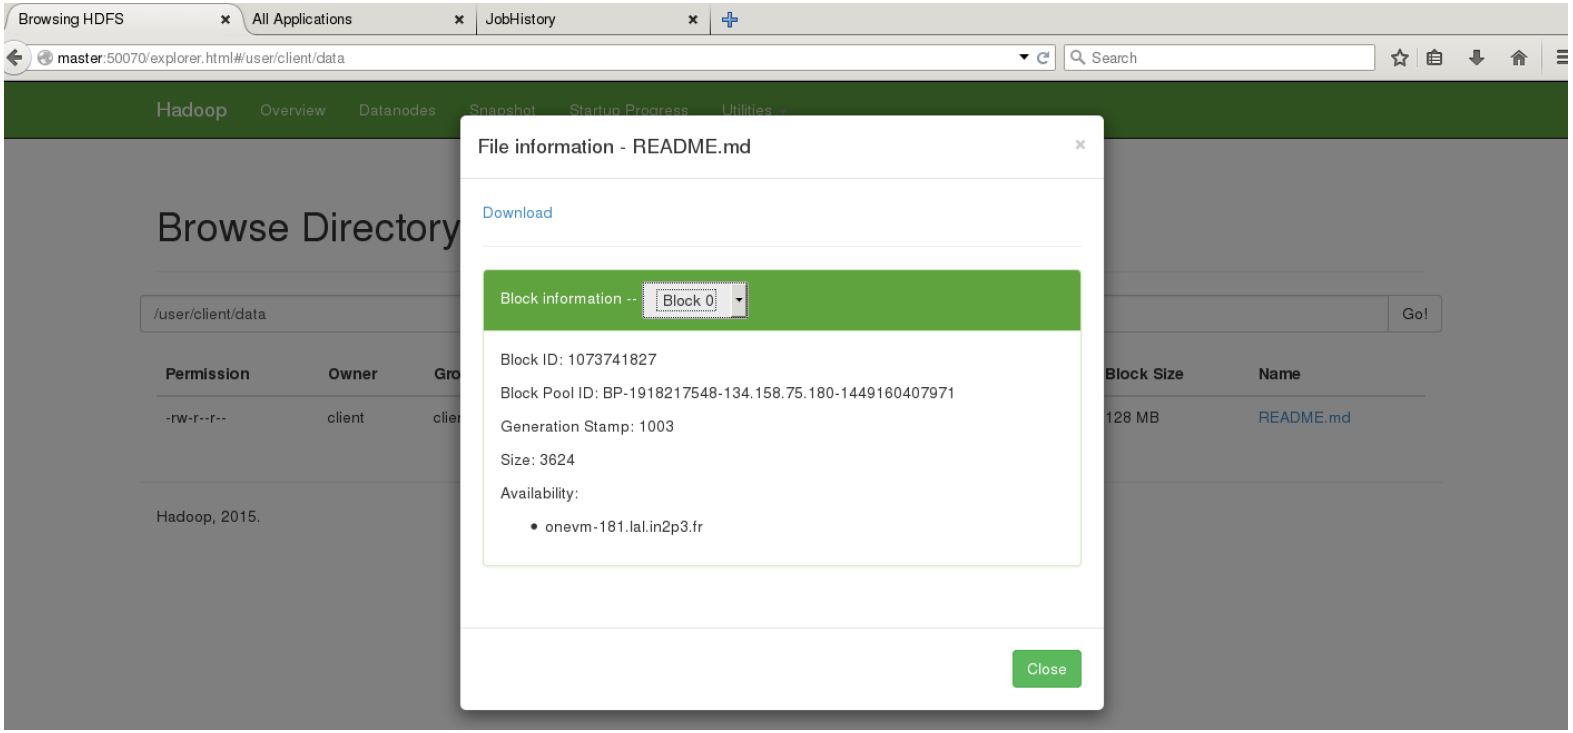

Figure: Web interface: file system.

Cluster use: MapReduce (MR)

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

- TODO:
  - Run a basic MR job:
    - Pi computation.
    - WordCount.
  - Run an analysis:
    - Meteo data.

MR job: Pi

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Execution:

```
client$ cd /usr/hdp/current && \
  cd hadoop-mapreduce-client
client$ yarn jar \
  hadoop-mapreduce-examples-*.jar \
  pi 10 1000000
Number of Maps  = 10
Samples per Map (precision) = 1000000
Wrote input for Map #0
...
INFO mapreduce.Job:  map 0% reduce 0%
...
Job Finished in 38.468 seconds
```

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Hadoop Web interface: job

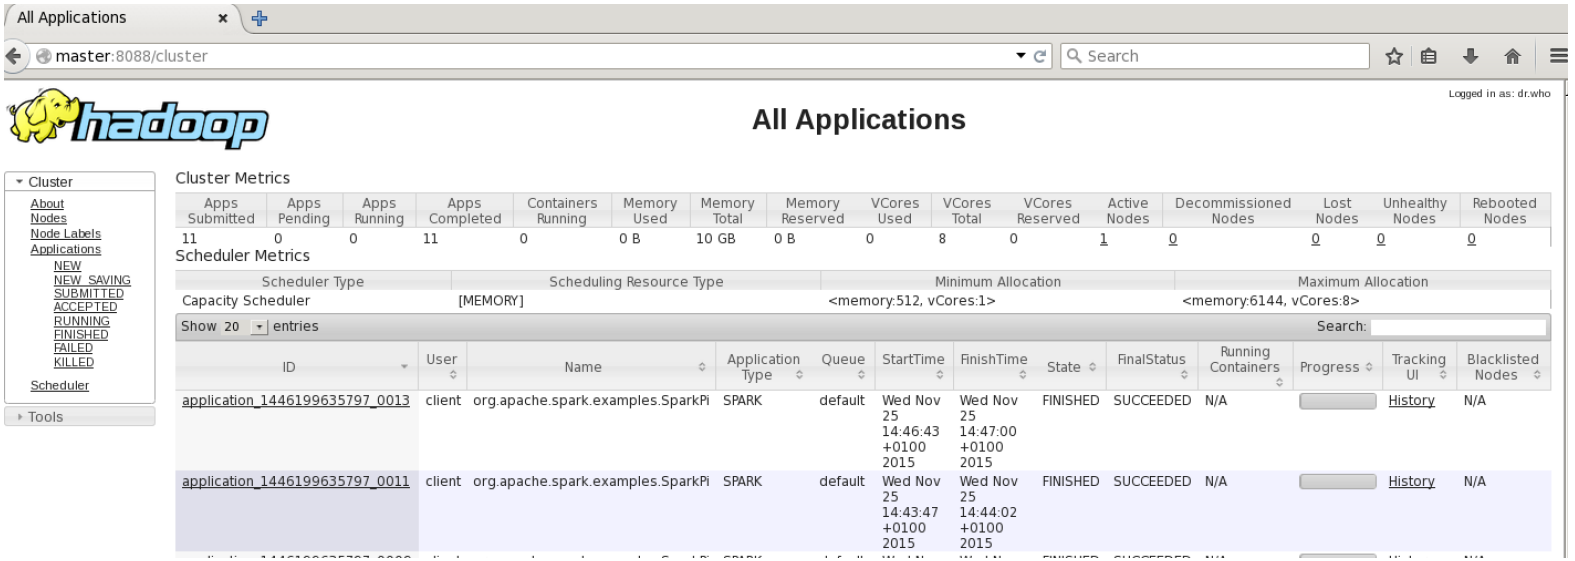

Figure: Web interface: ResourceManager.

MR job: WordCount with Python streaming

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Data:

- text file: **README.md**.

```
# Apache Spark
Spark is a fast and general cluster...
```

- already on HDFS:

```
$ hdfs dfs -ls data
-rw-r--r-- 1 client client 3624 data/README.md
```

MR job: WordCount with Python streaming

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Algorithm:

- HelloWorld of MR.
- count words in a text file.
- functions:
  - Map: (key, value) = (word, 1).
    `~/script/mapper.py`
  - Shuffle + sort: by word.
  - Reduce: by value.
    `~/script/reducer.py`

MR job: WordCount with Python streaming

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Test locally the algorithm:

```
$ cat /usr/hdp/current/spark-client/README.md \
  | ~/script/mapper.py | sort -k1,1 \
  | ~/script/reducer.py
>>>1
# 1
## 8
1000: 2
...
```

MR job: WordCount with Python streaming

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Execution:

```
client$ yarn jar hadoop-streaming.jar \
  -file ~/script/mapper.py \
  -mapper ~/script/mapper.py \
  -file ~/script/reducer.py \
  -reducer ~/script/reducer.py \
  -input data/README.md \
  -output output
```

➙ Options:

-input/ouput: HDFS path of I/O data for MR.

-mapper/reducer: scripts to execute as a m/r function.

-file: add a file to the distributed cache (for streaming).

MR job: WordCount with Python streaming

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Output:

➙ Job finished in 23.261 seconds.

- results:
  - dynamic resource allocation: 2 map tasks, 1 reduce task, 1 output file per reduce task.

```
client$ hdfs dfs -cat output/part-00000
"local" 1
"local[N]" 1
...
client$ hdfs dfs -rm -r -f output
```

MR job: WordCount with Python streaming

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Execution: adapt to job

```
client$ yarn jar hadoop-streaming.jar \
  -D dfs.replication=3 \
  -file ~/script/mapper.py ...

client$ hdfs dfs -ls output
-rw-r--r-- 3 client ...    0 output/_SUCCESS
-rw-r--r-- 3 client ... 3130 output/part-00000
```

➡ Options:
-D: change the configuration variables (see /etc/hadoop/conf/ directory).

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

MR job: analysis of meteo data

Data:

- weather data set (NCDC[a]) per year.
- semi-structured, record oriented data: **1901.gz**. 0029029070999991901010106004+64333+023450FM-12+000599999V0202701N015919999999N0000001N9-00781+99999...
  - date = 1901, air temperature (° Celsius x 10) = -0078, quality code = 1.
- from 1901 to 1949: < 1GB of data, small number of large files.

[a] http://www.ncdc.noaa.gov

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

MR job: analysis of meteo data

Data:

- copy to HDFS:

```
client$ hdfs dfs -copyFromLocal ~/data/*.gz \
  data
client$ hdfs dfs -df -h data
Filesystem              Size  Used Available Use%
hdfs://master:8020 9.7 G 1.1 G     8.6 G  11%
```

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

MR job: analysis of meteo data

Algorithm:

- Maximum temperature per year.
- Map: (key, value) = (year, temperature) for each row.
- Reduce: (key, value) = (year, max(temperature)).

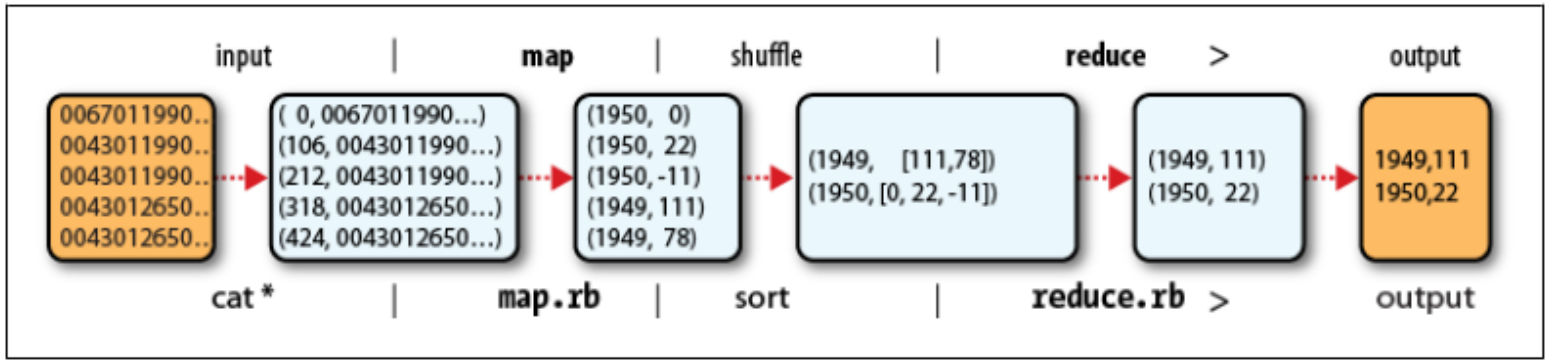

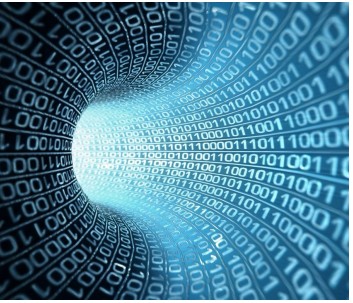

MR job: analysis of meteo data

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Execution:

```
client$ yarn jar hadoop-streaming.jar \
  -file ~/script/max_temperature_map.py \
  -mapper ~/script/max_temperature_map.py \
  -file ~/script/max_temperature_reduce.py \
  -reducer ~/script/max_temperature_reduce.py \
  -input data/*.gz -output output
```

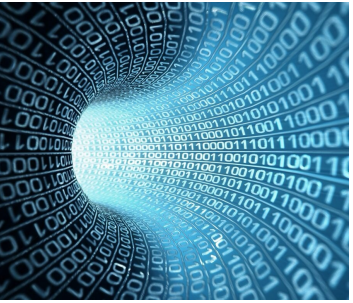

MR job: analysis of meteo data

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Output:

➙ Job finished in 6 min and 34 seconds.

- results:
  - 49 map tasks and 1 reduce task, 1 output file.

```
client$ hdfs dfs -cat output/part-00000
1901 317
1902 244
...
```

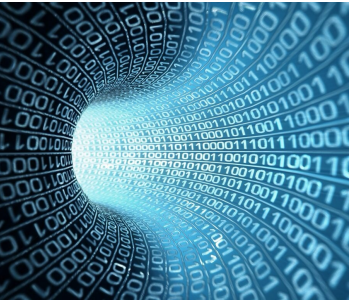

Cluster use: Spark

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Spark
Lightning-Fast Cluster Computing

- TODO:
  - Run a basic Spark job.
  - Run a Spark shell.
  - Run a Spark interactive job.

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Spark job: Pi

Execution:

```
client$ cd /usr/hdp/current/spark-client
client$ spark-submit \
  --class org.apache.spark.examples.SparkPi \
  --master yarn-cluster \
  lib/spark-examples-*.jar 100
```

➙ Options:
–class: the entry point for the application.
–master yarn-cluster: connect to a YARN cluster in cluster mode.

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Spark job: Pi

Output:

➙ Job finished in 34.39 seconds.

- results:

```
$ yarn logs -applicationId application_XXX | \
  grep Pi
Pi is roughly 3.1422368
```

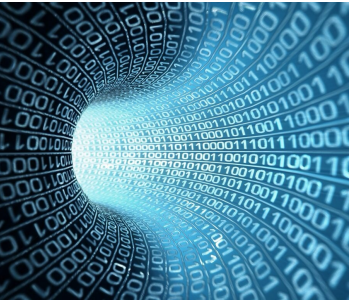

Hadoop command line: interactive job

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Spark shell:

```
client$ pyspark --master yarn-client
Welcome to
      ____              __
     / __/__  ___ _____/ /__
    _\ \/ _ \/ _ `/ __/  '_/
   /__ / .__/\_,_/_/ /_/\_\   version 1.4.1
      /_/
```

➡ Option:

–master yarn-client: connect to a YARN cluster in client mode.

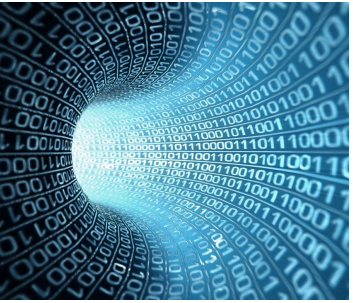

Spark job: WordCount

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Execution:

```
>>> file  = sc.textFile ("data/README.md")
>>> file.cache()
>>> file.count()
98
>>> file.first()
u'# Apache Spark'
>>> file.collect()
[u'# Apache Spark', u'', u'Spark is a fast...
```

- SparkContext: driver program as sc.
- RDD in cache with lazy evaluation.
- Actions (return a value): count, first, collect.

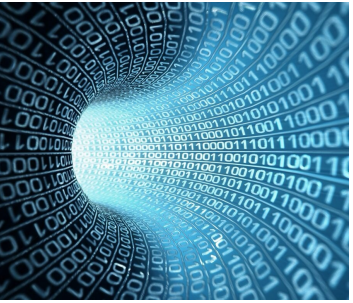

Spark job: WordCount

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Execution:

```
>>> counts = file \
    .flatMap(lambda line: line.split()) \
    .map(lambda  word: (word, 1)) \
    .reduceByKey(lambda  a, b: a + b)
```

➙ Transformations (create a new RDD):
- map/flatMap(func): pass each element of the RDD through a function func returning a single item/sequence.
- reduceByKey(func, [numTasks]): returns a dataset of (K, V) pairs where the values for each key are aggregated using the given reduce function func.

Spark job: WordCount

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Execution:

```
>>> counts.saveAsTextFile("output")
>>> counts.take(2)
[(u'all', 1), (u'when', 1)]
>>> counts.takeOrdered(2)
[(u'"local"', 1), (u'"local[N]"', 1)]
>>> for x in counts.collect(): \
print x
(u'all', 1)
(u'when', 1)
...
>>> exit()
```

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Spark job: WordCount

Output:

- results:
  - dynamic resource allocation: 2 executors (1 driver + 1 code executor), 2 output files.

```
client$ hdfs dfs -ls output
-rw-r--r--   1 client client output/_SUCCESS
-rw-r--r--   1 client client output/part-00000
-rw-r--r--   1 client client output/part-00001
```

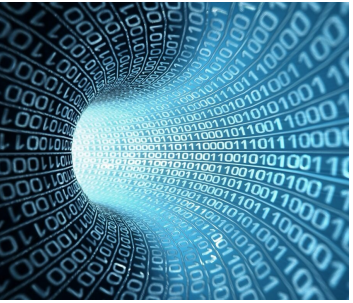

Cluster use: Spark

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Spark
Lightning-Fast Cluster Computing

- TODO:
  - Run an analysis.
  - Access to the YARN and Spark Web interface.

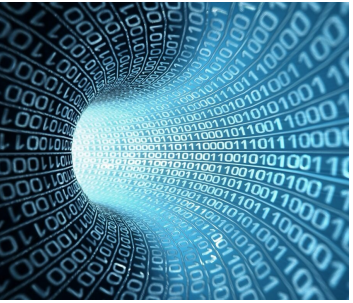

Spark job: analysis of meteo data

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Execution:

```
client$ spark-submit --master yarn-cluster \
  ~/script/max_temperature_spark.py
```

Output:

➙ Job Finished in 4 minutes and 57 seconds.

- results:
  - 3 executors and 49 output files.

```
client$ hdfs dfs -cat output/part-*
(u'1935', u'+0478')
(u'1944', u'+0507')
...
```

Hadoop command line: job

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Web interface access:

- http://master:8088/proxy/application_id ➙ Spark UI.

Hadoop Web interface: Spark

Big Data
- VVVV...
- in Astronomy
- Technologies

Hadoop
- Core
- Ecosystem
- Cluster

Cloud computing
- Why ?
- StratusLab
- SlipStream

Tutorial
- StratusLab
- SlipStream
- Hadoop
- HDFS
- MapReduce
- Spark

Links

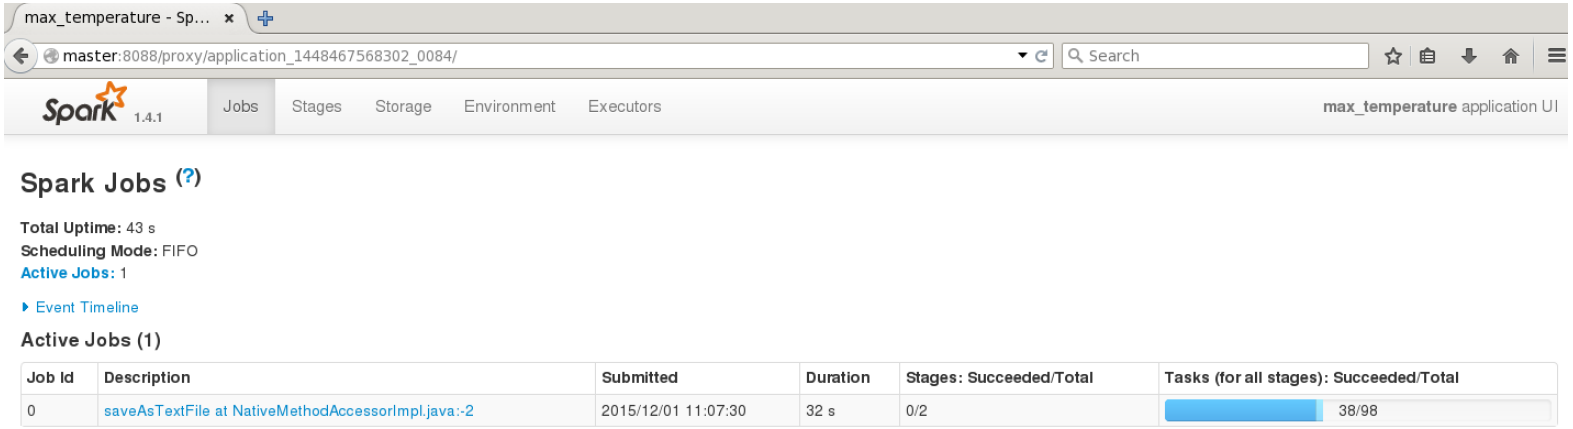

Figure: Web interface: jobs.

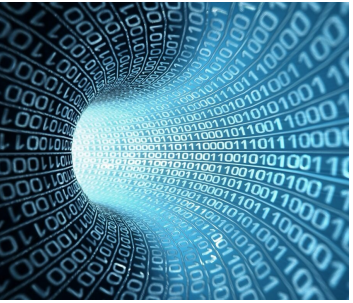

Hadoop Web interface: Spark

Big Data
- VVVV...
- in Astronomy
- Technologies

Hadoop
- Core
- Ecosystem
- Cluster

Cloud computing
- Why ?
- StratusLab
- SlipStream

Tutorial
- StratusLab
- SlipStream
- Hadoop
- HDFS
- MapReduce
- Spark

Links

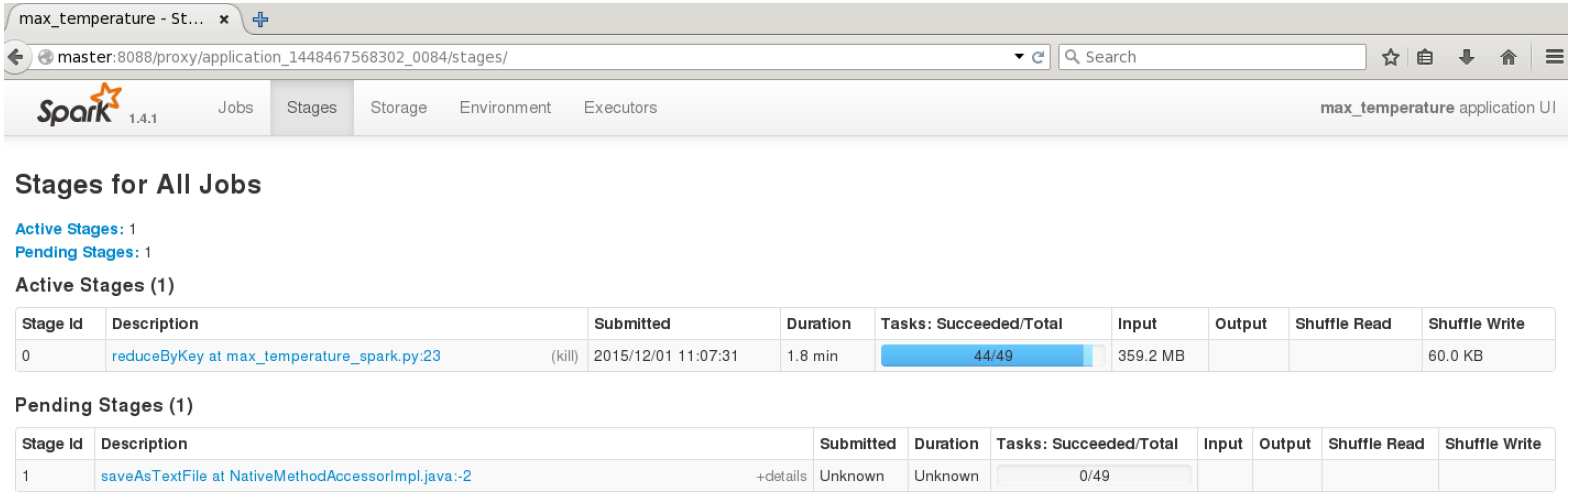

Figure: Web interface: stages.

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

Hadoop Web interface: Spark

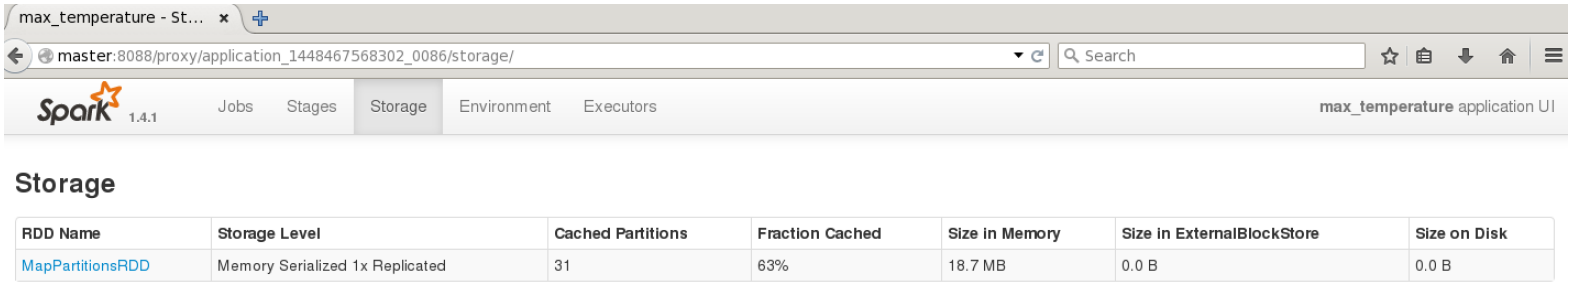

Figure: Web interface: storage.

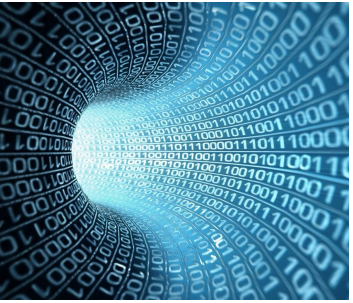

Hadoop Web interface: Spark

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

max_temperature - Ex...

master:8088/proxy/application_1448467568302_0086/executors/

Spark 1.4.1 | Jobs | Stages | Storage | Environment | Executors | max_temperature application UI

Executors (3)

Memory: 122.1 MB Used (1306.3 MB Total)
Disk: 0.0 B Used

Executor ID	Address	RDD Blocks	Storage Memory	Disk Used	Active Tasks	Failed Tasks	Complete Tasks	Total Tasks	Task Time	Input	Shuffle Read	Shuffle Write	Logs	Thread Dump
1	onevm-213.lal.in2p3.fr:43088	18	69.1 MB / 530.3 MB	0.0 B	1	0	18	19	27.0 s	34.0 MB	0.0 B	24.5 KB	stdout stderr	Thread Dump
2	onevm-213.lal.in2p3.fr:42634	18	53.0 MB / 530.3 MB	0.0 B	1	0	18	19	21.3 s	26.1 MB	0.0 B	24.5 KB	stdout stderr	Thread Dump
driver	134.158.75.213:52643	0	0.0 B / 245.7 MB	0.0 B	0	0	0	0	0 ms	0.0 B	0.0 B	0.0 B		Thread Dump

Figure: Web interface: executor.

Terminate the cluster

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

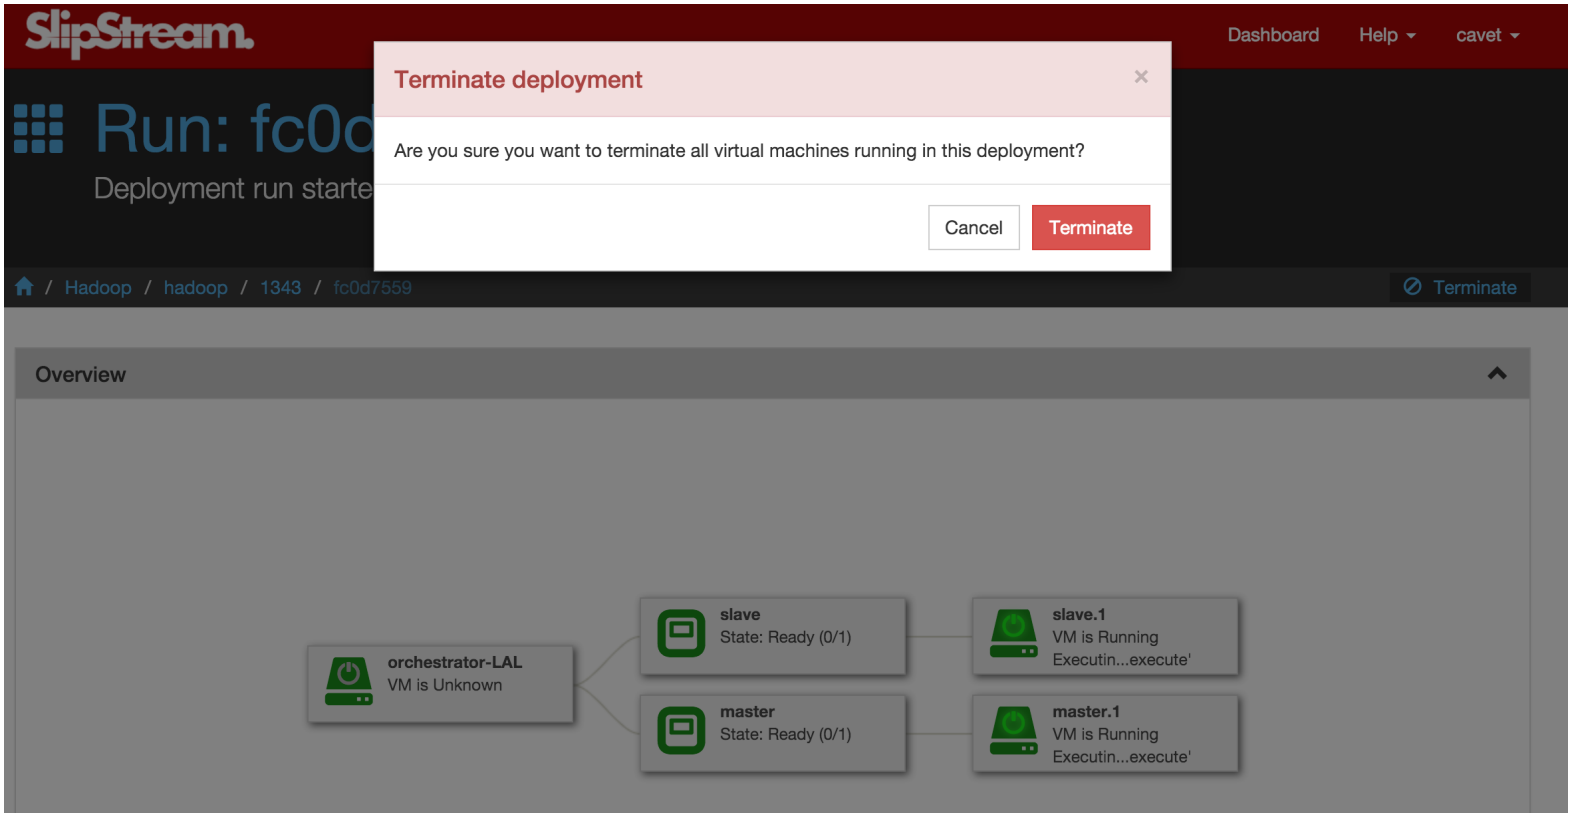

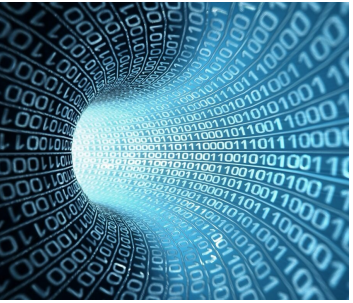

Several documents

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

- HPC Today n°13: `http://www.hpctoday.fr/files/docs/mag/HPCReview-current.pdf`
- T. White, Hadoop: The Definitive Guide, Sebastopol: O'Reilly Media, Inc. (2010)
- J.R. Owens, J. Lentz, B. Femiano, Hadoop par la pratique, Pearson (2014)

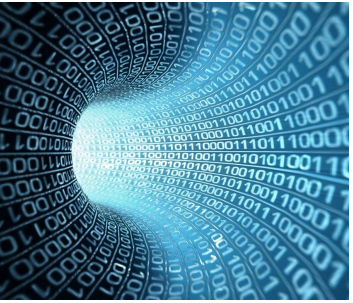

Several links

Big Data
VVVV...
in Astronomy
Technologies
Hadoop
Core
Ecosystem
Cluster
Cloud computing
Why ?
StratusLab
SlipStream
Tutorial
StratusLab
SlipStream
Hadoop
HDFS
MapReduce
Spark
Links

- StratusLab: http://stratuslab.readthedocs.org/en/latest/
- SlipStream: http://ssdocs.sixsq.com/en/latest/
- Spark: http://spark.apache.org/, https://jaceklaskowski.gitbooks.io/mastering-apache-spark/
- MPI vs Spark: cons MPI: http://www.dursi.ca/hpc-is-dying-and-mpi-is-killing-it/, pro MPI: http://www.sciencedirect.com/science/article/pii/S1877050915017895

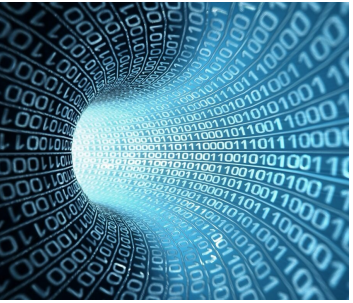

Big Data
- VVVV...
- in Astronomy
- Technologies

Hadoop
- Core
- Ecosystem
- Cluster

Cloud computing
- Why ?
- StratusLab
- SlipStream

Tutorial
- StratusLab
- SlipStream
- Hadoop
- HDFS
- MapReduce
- Spark

Links

Hadoop command line: HDFS

```
root$ su - hdfs
```

- **Ressource report**:

```
$ hdfs dfsadmin -report
Configured Capacity: 4.79 GB
Present Capacity:    4.78 GB
DFS Remaining:       4.78 GB
...
```