Hadoop on the Cloud: the SlipStream deployment tool

Cécile Cavet

François Arago Centre, APC, Univ. Paris Diderot, CNRS/IN2P3, CEA/Irfu, Obs. de Paris, Sorbonne Paris Cité,

13 rue Watt, 75013, Paris, France

cecile.cavet at apc.univ-paris7.fr

Abstract

Cloud computing offers IT resources on-demand. The IaaS (Infrastructure-as-a-Service) layer provides virtual machines (VMs), storage and network. In order to overcome the difficulties carried by complex workflow managements, we have used the SlipStream solution. This service allows to automatically deploy environments on interoperable Cloud platforms such as StratusLab. We have deployed an Hadoop cluster in order to demonstrate the Cloud suitability for Big Data applications. Benchmarks were realised to check the cluster performance.

SlipStream

The **SlipStream**™ PaaS (Platform-as-a-Service) solution [6] allows to automatically deploy environnements on VMs. **SlipStream** supports:

- Recipes for specific nodes (master and slave):
  - Disk image customisation.
  - Shared parameters.
- Deployment of complex workflows:
  - Node number.
  - Multi-Cloud.

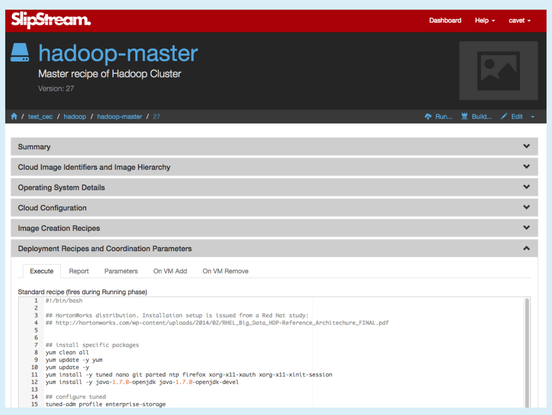

SlipStream Web interface: Hadoop master node recipe.

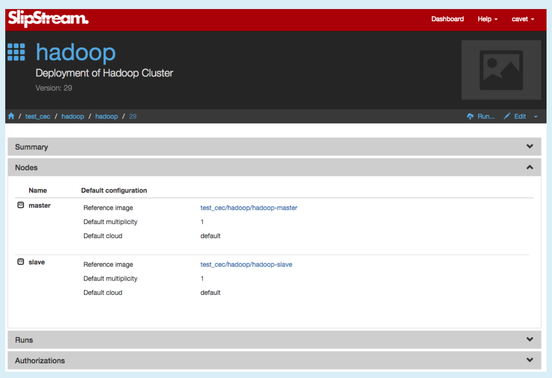

SlipStream Web interface: Hadoop deployment.

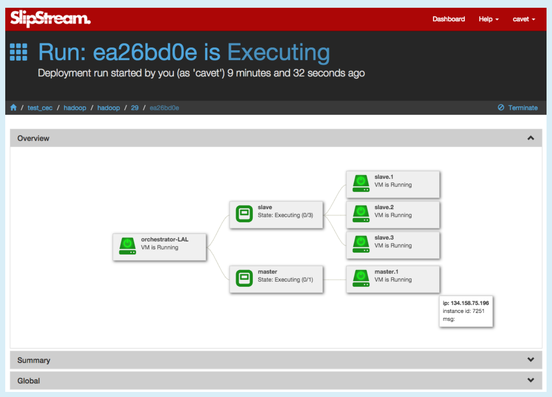

SlipStream Web interface: Hadoop virtual cluster in running phase.

► **Hadoop cluster is ready in 10 minutes !**

StratusLab Cloud

StratusLab [7] offers an academic public IaaS Cloud since 2010. Virtual computing resources such as VMs are:

- generated by **KVM** hypervisor and managed by **OpenNebula** virtual infrastructure manager.
- supported by physical machines @LAL: **16 nodes**, **440 CPUs**, **772 GB** of memory, **15 TB** of storage, **1 GbE/s** network.
- provided in OS disk images by the **MarketPlace** [4] and manageable by the **StratusLab client**.

Big Data

Techniques to treat huge volume of data are encompassed in **Big Data** terminology. **Big Data** technology is efficient for **unstructured** or **semi-structured data**.

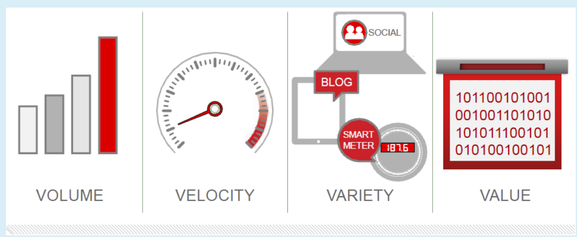

4 V of Big Data.

Hadoop history

The **Hadoop** solution has been developed since 2008. **Hadoop** cluster nodes hosted both computing and storage services:

- Distributed File System **HDFS**.
- Distributed algorithm **MapReduce**.
- Huge ecosystem (HBase, Pig...).

Hadoop cluster

Hadoop 2.0 is a recent version using **Yet Another Resource Negotiator** (**YARN**).

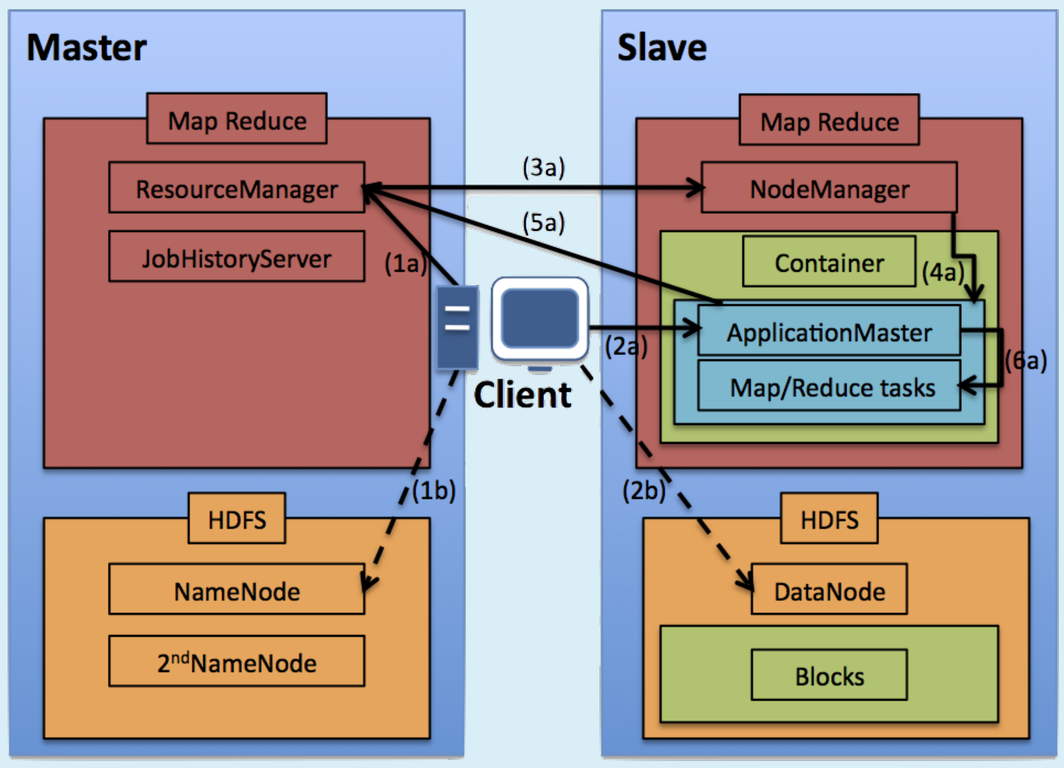

Hadoop cluster with YARN services.

Hortonworks distribution

Advantages [2]:

- Packaged distribution.
- Easy to install with shell scripts.
- Optimised **Hadoop** parameters.

Parameter	Type	Value
I/O Buffer size	**Core**	128 KB
Block size	**HDFS**	128 MB
Replication factor	–	2
Sort factor	**MapRed**	100
Sort memory	–	200 MB
Codec	–	Snappy

Main Hadoop parameters.

Virtual cluster setup

- **OS tuning** [3]: files, FS, network, huge pages.
- **Master node**:
  - Services: **NameNode**, **SecondaryNameNode**, **ResourceManager**, **JobHistoryServer**.
  - Resources: **1 node**, **4 CPUs**, **8 GB** of memory, **5 GB** of local disk, **10 GB** of ephemeral disk, **50 GB** of persistent disk.
- **Slave node**:
  - Services: **DataNode**, **NodeManager**.
  - Resources: **3 nodes**, **4 CPUs**, **16 GB** of memory, **5 GB** of local disk, **150 GB** of persistent disk (**2** disks).
- **Hadoop parameters tuning**.

Hadoop Job: scientific example

Big data application with **MapReduce** algorithm on meteorological data (NCDC) [8].

- **Input data**: **36 GB** of semi-structured, record oriented data.
- **Algorithm**: max(temperature).

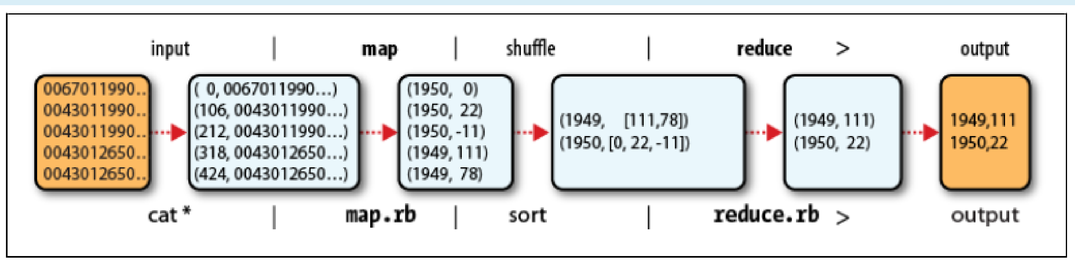

MapReduce job workflow [8].

Performance Benchmark

Benchmarks to check the cluster setup.
Methodology:

- Running 5 times the benchmark.
- Memory cleaning between each iterations.
- Mean and standard deviation calculation.

Benchmarks:

- **TestDFSIO**: **HDFS I/O**, write/read 100 GB of data (10 files), 10 maps & 1 reduce.

► **Write**: **40 MB/s**, Throughput **6 MB/sec**, Average IO rate **7 ± 2 MB/sec**.
► **Read**: **56 MB/s**, Throughput **9 MB/sec**, Average IO rate **11 ± 6 MB/sec**.
► **Results are in agreement with previous studies on small clusters.**

- **TeraSort**: **HDFS I/O + MR**, generate/sort huge volume of data, 135 maps & 32 reduces, replication factor = 1.

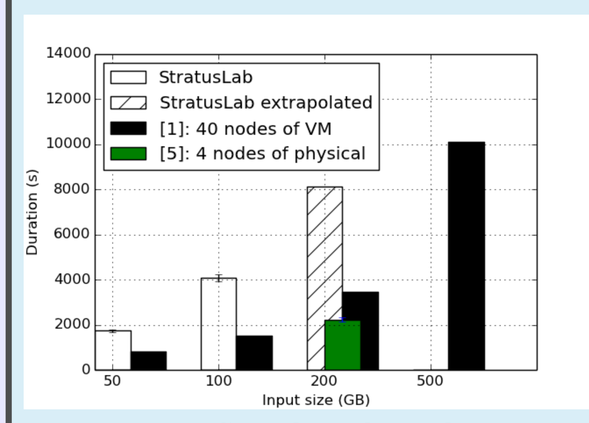

TeraSort benchmark.

► **TeraSort benchmark under performs on the virtual cluster.**

Algorithm comparison

Analysis of meteorological data (see [8]): 101 maps & 1 reduces (default value).

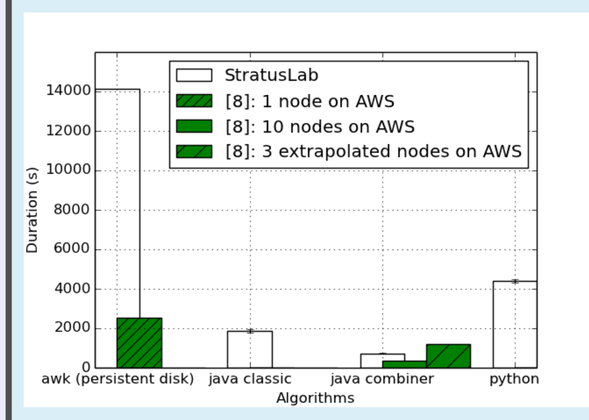

MapReduce job.

► **Java is the most efficient method.**
► **Virtual cluster does not under perform in comparison with EMR AWS.**
► **Streaming adds a performance overhead (factor 2 of time).**

References

[1] Cacheiro et al., Leveraging EGI Fed. Cloud Infra. for Hadoop analytics, IberGrid (2014)
[2] Hortonworks, Hortonworks Data Platform: Installing HDP Manually (2014)
[3] Joshi, Hadoop Perf. Tuning Guide (2012)
[4] Marketplace: https://marketplace.stratuslab.eu/marketplace/metadata
[5] Red Hat, Exploring the next generation of Big Data solutions with Hadoop 2 (2014)
[6] SlipStream: http://sixsq.com/products/slipstream.html
[7] StratusLab: http://www.stratuslab.eu/
[8] White, Hadoop: The Definitive Guide, Sebastopol: O'Reilly Media, Inc. (2010)