

# Running Big Data on the EGI Federated Cloud

Javier Lopez Cacheiro, Álvaro Simón, Juan Villasuso,  
Esteban Freire, Roberto Rosende, Ivan Díaz, Alejandro  
Feijoo, Pablo Rey, Carlos Fernandez

**CESGA**

Boris Parak

**CESNET**

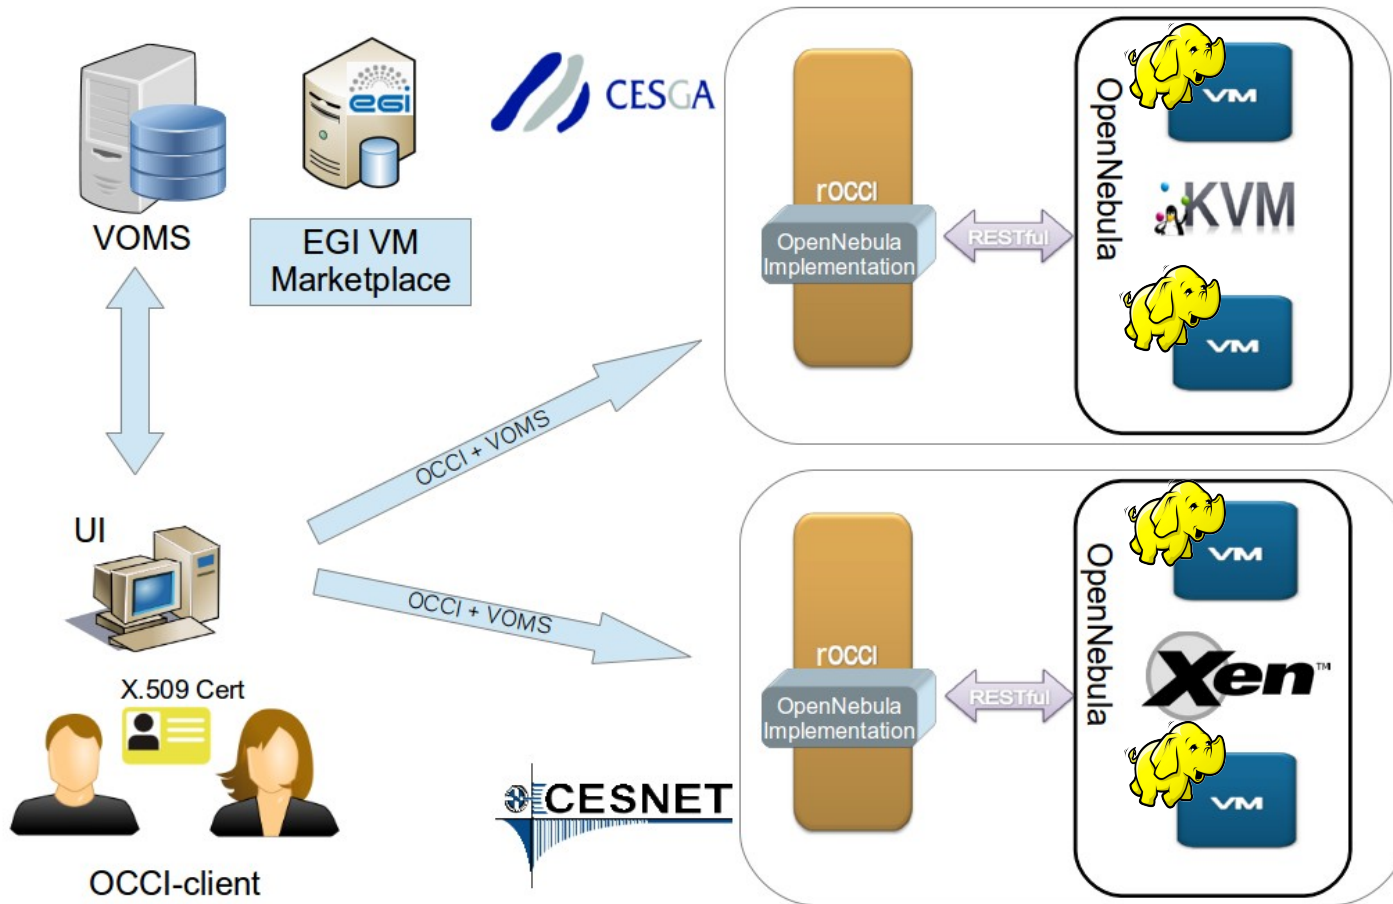
- Introduction
- Methodology
- Results
- Conclusions
- Links

- Motivation
  - Big Data is an emerging technology that could benefit from existing developments in cloud.
  - It represent a real benchmark of the Fedcloud infrastructure under heavy usage conditions.
  - Run Big Data calculations over federated clouds is attracting the interest of the research community.

- The aim of our work:
  - To check the suitability of EGI Fedcloud infrastructure to run Hadoop through a series of real-world benchmarks...

- To deploy a Hadoop cluster inside the Fedcloud infrastructure. In three steps:
  - ***Cluster startup***: obtaining the necessary resources.
  - ***Hadoop configuration***: configuring the Hadoop cluster.
  - ***Benchmark execution***: Data load and run MapReduce jobs.

- Cluster startup:
  - **Create a customized VM** image template to run Hadoop. (SL6 in our case).
  - **Configure Firewall** settings to allow inter-cluster communication.
  - **Modules.** Modules package is used to simplify software deployment.
  - **Java.** We use Oracle Java JDK 1.6.0\_33 because it offers better performance.
  - **Hadoop.** VM image already includes Hadoop v1.0.3 and v1.0.4



- Hadoop configuration:
  - Hadoop software is already available from VM images.
  - Our Hadoop cluster will consist of one *master* and a variable number of *slave* nodes. **N+1** VMs
  - **Master** node will run: *namenode*, *secondarynamenode* and *jobtracker* services.
  - **Slaves** will run just one *tasktracker* and one *datanode* services.
  - We are using small instances. 1 CPU, 1GB of RAM.



Parameter	Type	Value
<i>fs.inmemory.size.mb</i>	core-site	200MB
<i>io.file.buffer.size</i>	core-site	128KB
<i>mapreduce.task.io.sort.factor</i>	core-site	100
<i>mapreduce.task.io.sort.mb</i>	core-site	100
<b><i>mapred.tasktracker.map.tasks.maximum</i></b>	mapred-site	<b>1</b>
<b><i>mapred.tasktracker.reduce.tasks.maximum</i></b>	mapred-site	<b>1</b>
<b><i>mapred.reduce.tasks</i></b>	mapred-site	<b>0.95 x num. slaves</b>
<i>dfs.datanode.du.reserved</i>	hdfs-site	1GB
<b><i>dfs.block.size</i></b>	hdfs-site	<b>1MB / 64MB</b>
<b><i>dfs.replication</i></b>	hdfs-site	<b>3</b>
<b>HADOOP_HEAPSIZE</b>	hadoop-env	<b>512MB</b>

- Benchmark execution:
  - We selected two use cases which are representative of a broad range of Hadoop cluster jobs:
    - **Enciclopædia Britannica** 1911 data set word count (176MB). Datablock size **1MB**.
    - **Wikipedia** data set word count (41GB). Datablock size **64MB**.



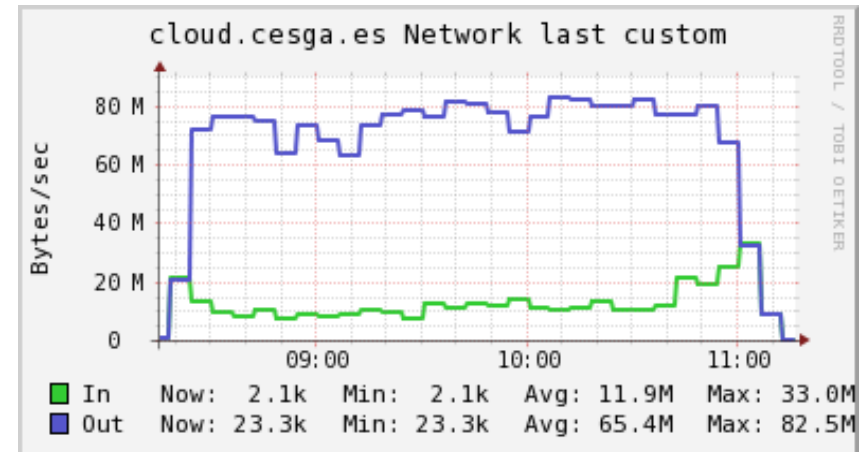
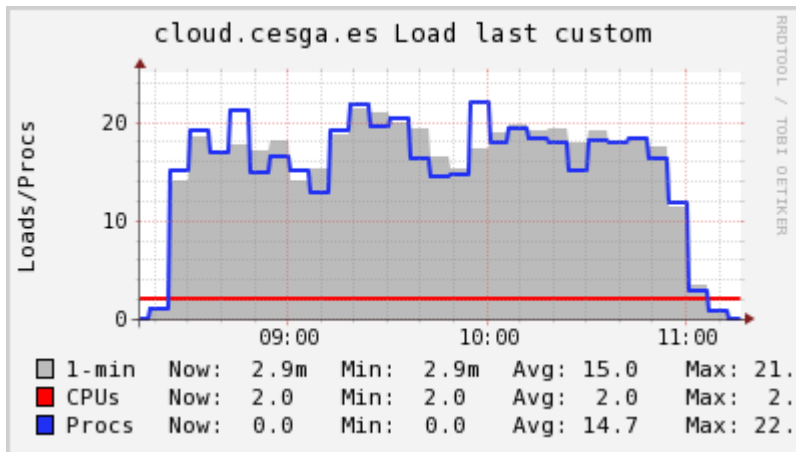
**WIKIPEDIA**  
The Free Encyclopedia



- **Benchmark execution:**
  - The benchmarks were run both in a federated cloud deployment (EGI Fedcloud) and in one-site-only cluster deployment.
  - Benchmarks are run for different cluster sizes that ranges from **10** to **101** Vms.
  - All the measurements where repeated at least 5 times.

- Benchmark results:
  - We are instantiating from 10 to 101 VMs so it can also serve as a stress test of the Fedcloud rOCCI framework.
  - After some initial VM deployment tests we discovered that most of the Fedcloud RPs were misconfigured or did not support rOCCI.
  - 101 node cluster startup time ranged from **2.5** to **3** hours.
  - The time involved to configure and start the Hadoop cluster was really short (less than **1** minute in the 101 nodes test).

- OpenNebula frontend load and network usage during 101 nodes deployment.



- Benchmark times for the Enciclopædia Britannica use case:

Cluster size	Federated		One-site-only	
	put (s)	wordcount (s)	put (s)	wordcount (s)
10 (CESGA)			47 ± 2	169 ± 1
10 (CESNET)			9.5 ± 0.2	160 ± 1
21	235 ± 12	108 ± 2	37 ± 1	97 ± 1
31	189 ± 4	90 ± 2	36 ± 2	78 ± 2
41	190 ± 11	81 ± 4	33 ± 2	71 ± 3
51	133 ± 7	73 ± 3	33 ± 1	66 ± 2
101	94 ± 7	67 ± 22	33 ± 4	52 ± 5

\*The put time: time to load the files into the Hadoop HDFS filesystem.



\*The wordcount time: time to perform a simple word count mapreduce job.

- Benchmark times for the Wikipedia use case:

Cluster size	Federated	Federated	One-site-only	One-site-only
	put (s)	wordcount (s)	put (s)	wordcount (s)
51	19190	1001 ± 39	6705	1347 ± 117
101	13208	705 ± 14	5665	725 ± 18

\*The put time: time to load the files into the Hadoop HDFS filesystem.

\*The wordcount time: time to perform a simple word count mapreduce job.

- These results show the suitability of a federated cloud infrastructure like EGI Fedcloud to run Big Data analytics.
- • Small VM instances are good enough to run Hadoop (assuming an appropriate configuration tuning).
- It was tested rOCCI performance under heavy usage.
- About future work it is possible to offer Hadoop cluster on-demand through a simple web frontend.
  
- • There were some disappointing results about cluster startup (in order of 3h for the 101 node cluster).
- EGI Fedcloud does not provide (yet) a workload or broker service.
- Only a few Resource Provides are able to provide the required resources.



- Apache Hadoop: <http://hadoop.apache.org/>
- rOCCI API: <https://github.com/gwdg/rOCCI>
- Hadoop Fedcloud deployment tools:  
<https://github.com/grid-admin/hadoop>

# Questions?

