

L'esperienza di Cloud Federata del **GARR**

Giuseppe Attardi, Federico Ruggieri

WS INAF - Bologna, 30 Novembre 2017



The GARR network

- More than **15.000 km** of GARR owned fibers
 - ~**9.000 Km** of backbone
 - ~**6.000 Km** of access links
- About **1000 user sites** interconnected
- > **1 Tbps** aggregated access capacity
- > **2 Tbps** total backbone capacity
- **2x100 Gbps** IP capacity to GÉANT
- Cross border fibers with ARNES (Slovenia), SWITCH (Switzerland).
- > **100 Gbps** to General Internet and Internet Exchanges in Italy
- **NOC and engineering** are in-house, in Rome.



DATA, HPC & HTC Centres



HPC: CINECA

HTC: INFN, RECAS, ENEA, GARR, etc.

All sites connected to the GARR network with optical fibres from 10 to 100 Gb links.

Infrastruttura Hardware



Federated Cloud

Objectives

- Facilitate **transition towards cloud computing**
- Allow **resource sharing**, maintaining **control of use**
- Exchange **best practices** on management and use
- Evolve towards **native cloud applications**
- Expand **catalogue of cloud applications**

GARR Commitments

- Architecture Design
- Ready to use OpenStack Distro
- OpenSource Code Base (git.garr.it)
- Upgrades and Maintenance
- Solution for multiple tenancy
- Federation and Delegation
- Federated Authentication
- Asset Management

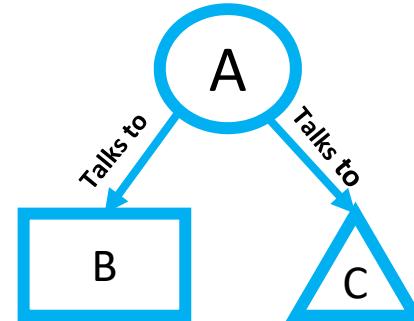
Training on Cloud Computing

- Hands-on Workshop on Federated Cloud Deployment
- Editions:
 - May 2017. WS GARR 2017
 - June 2017. 9 countries from Eastern Europe
 - October 2017. Università Napoli II



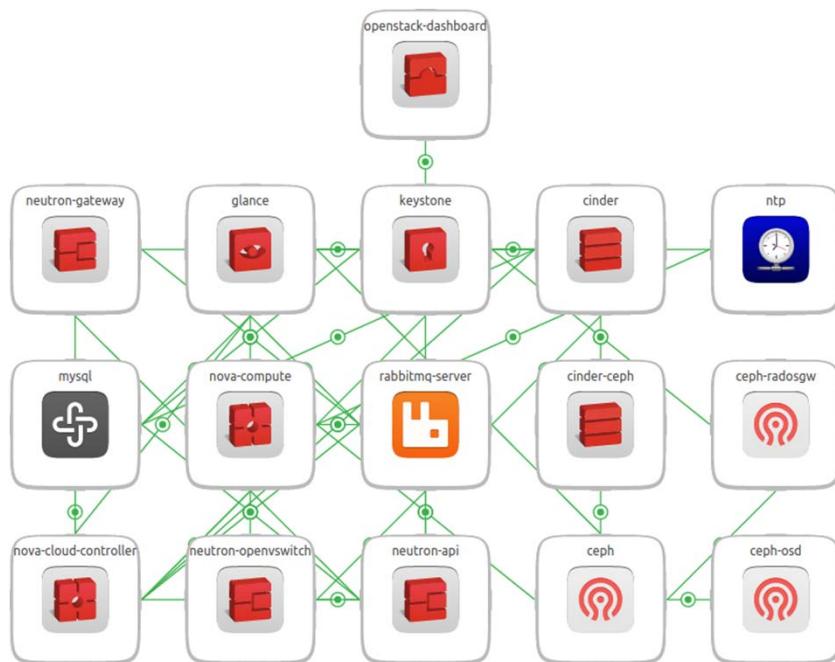
Declarative Modeling

- Describe **what** you want, not **how** to do it
 - Workflow Engine computes the **differences** between **current** and **desired state**
 - **Generates execution plan** to produce the desired model
-
- App A requires:
 - X GB memory and Y CPU
 - N GB storage
 - talking with B and C
 - An URL endpoint
 - To run locally, close to B

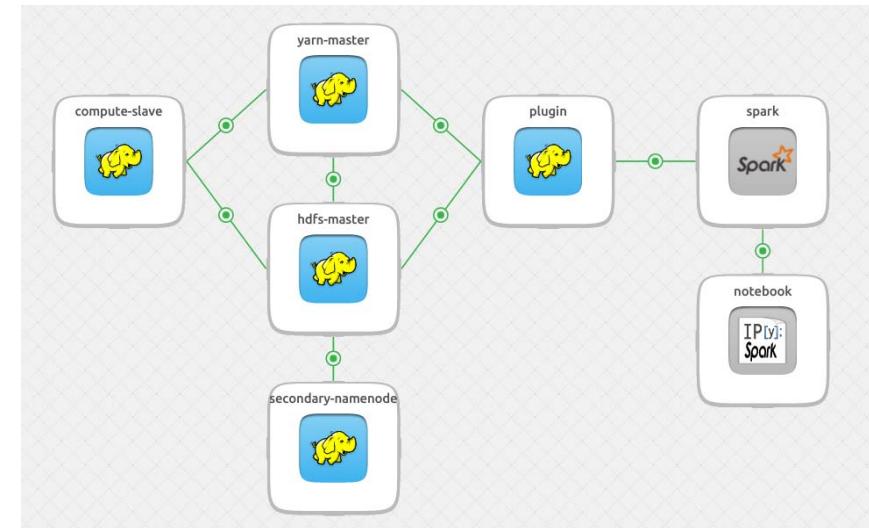


A Single Automation Tool for Platform & Application Deployment

Platform Deployment: OpenStack

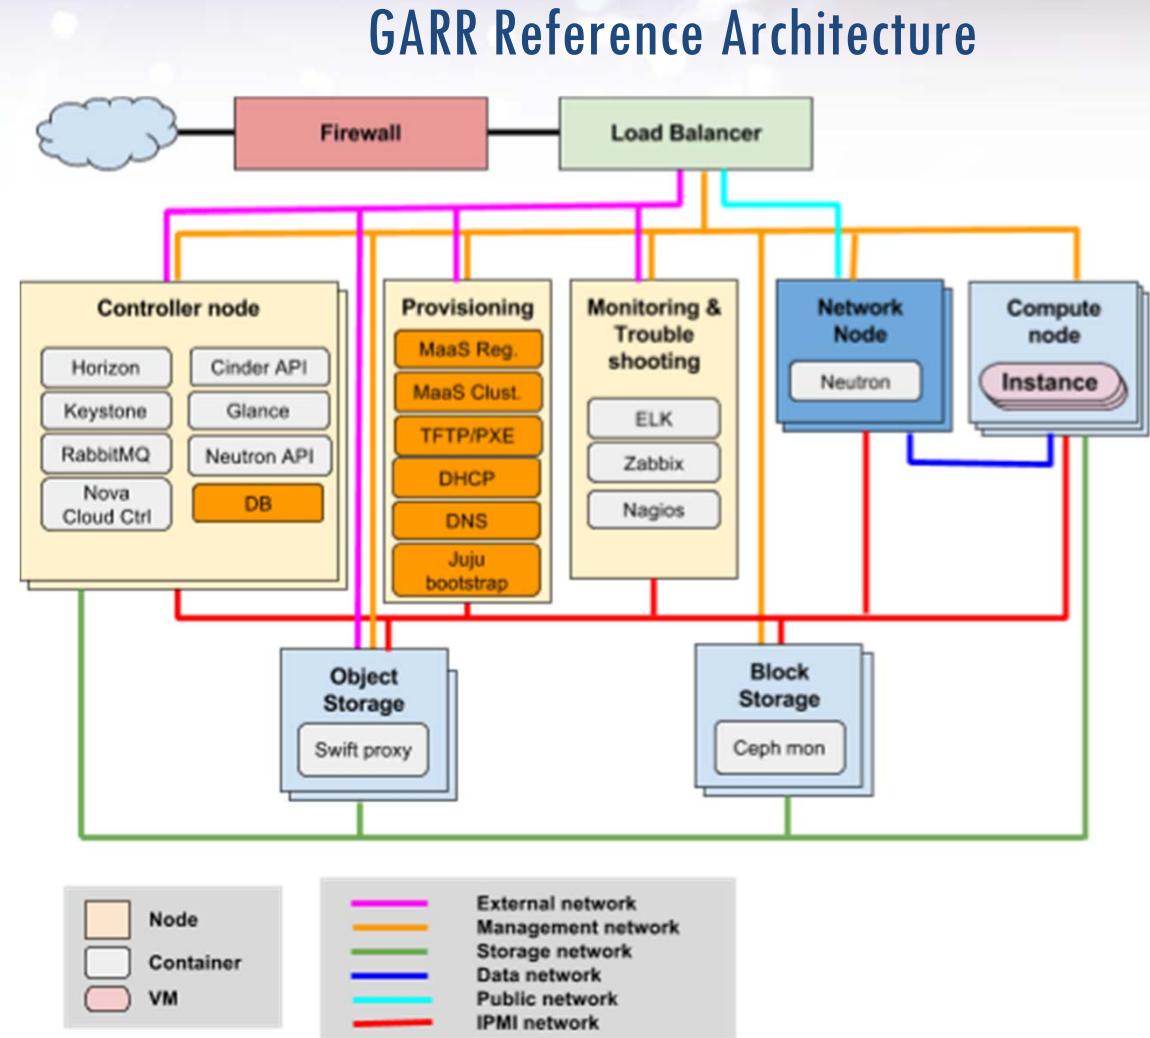


Application Deployment: Big Data Analytics

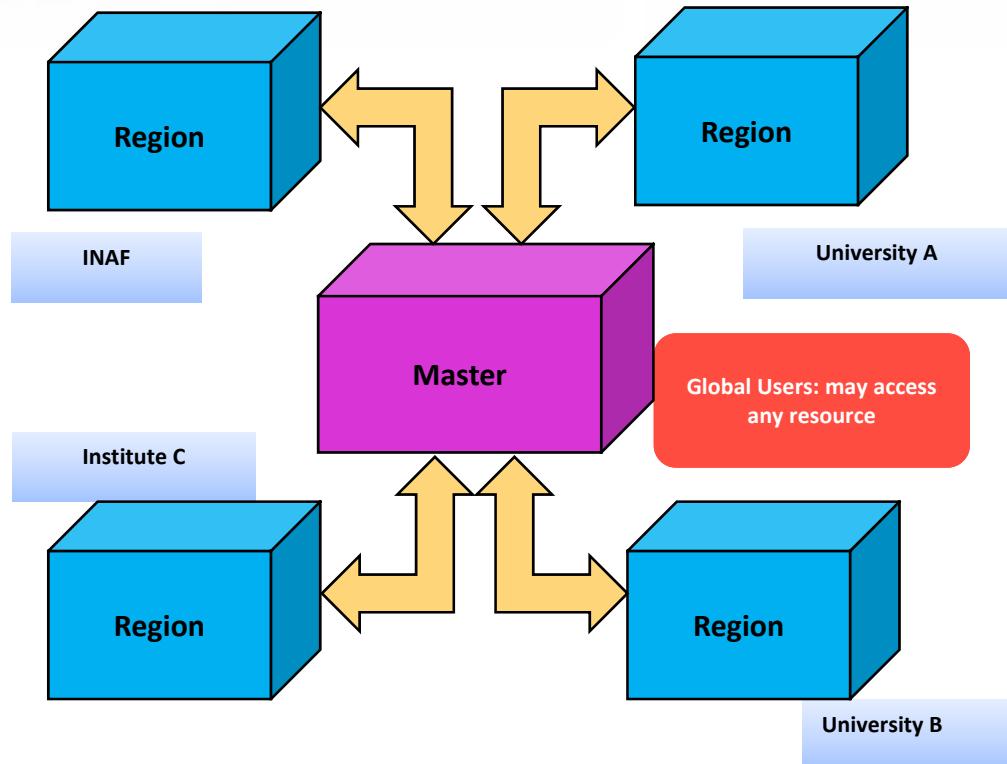


Federation with OpenStack

- Widely used and well supported Cloud Computing software
- Over 45.000 developers world wide
- Complex to manage
- Designed a Reference Architecture:
 - Declarative modeling
 - Easy to configure and to replicate
 - Managed with automated orchestration tools



Federated Cloud Architecture



- **Federated Region Deployment**
 - Simple procedure
 - From predefined model
 - **Time to deploy from scratch: a few hours**
- **Federated Authentication**
 - SAML2 (Idem, EduGain)
 - OIDC (Google)
 - Single user account over whole federation
- **Delegated Administration**
 - Resources controlled through quotas
 - Region Administrator
 - Virtual Datacenter Administrator

Cloud Developer Community

- Build a community of users and developers:
 - <https://cloud.garr.it/community/>
- Build a shared Catalogue of services
- Examples built by GARR:
 - Moodle as a Service
 - Jupyter Notebooks as a Service

Deployment as a Service (DaaS)

Example: Deploying/Scaling Moodle in the Cloud

Juju / attardi@local / default

4 applications | 3 machines

< mysql

This application has been marked to be destroyed on next deployment.

Units 1

Configure

Relations

Expose Off

Change version cs:mysql-56

Search the store

Logout

The screenshot shows the Juju graphical interface. On the left, there's a sidebar for the 'mysql' application, which has been marked for destruction. It lists one unit, no relations, and exposes nothing. On the right, two nodes are shown: a MySQL node (with an orange 'm' icon) and a PostgreSQL node (with a blue 'p' icon). They are connected by a line, indicating a relation between them.

Jupyter Notebook Server

- Experiment live with Machine Learning and GPUs

jupyter MNIST in Keras (autosaved)

Control Panel Logout

| Python 3 O

File Edit View Insert Cell Kernel Help

Code CellToolbar

Build the neural network

Build the neural-network. Here we'll do a simple 3 layer fully connected network.

The diagram illustrates a 3-layer fully connected neural network. It consists of three layers of nodes: an input layer with 10 nodes (labeled 0 through 9), a hidden layer with 512 nodes, and another hidden layer with 512 nodes. The network is fully connected between adjacent layers, with arrows indicating the flow of information from one node in a layer to every node in the next layer. Below the network, there is a small grayscale image of a handwritten digit, likely a sample from the MNIST dataset.

App Deployed on AWS (external cloud)

The screenshot displays two cloud management interfaces side-by-side.

Juju UI (Left): Shows the deployment status of two applications: mysql (1 unit) and mediawiki (1 unit). A message indicates "You have placed all of your units". The default environment contains 1 unit, trusty, with 1x3.5GHz, 3.75GB, and 8.00GB storage. It also lists 0 containers and 1 Root container for mysql/0.

AWS EC2 Dashboard (Right): Shows the EC2 Dashboard with the "Launch Instance" button highlighted. The Instances section lists four running instances:

Name	Instance ID	Instance Type	Availability Zone
	i-00f1ec62d532cea8d	m3.medium	us-east-1b
	i-02d632e9b1d7b85...	m3.medium	us-east-1a
juju-controller...	i-033284f4eacc5497b	t2.medium	us-east-1a
	i-09d3b4dd6ed8799...	t2.micro	us-east-1b

Active Services

- VM
 - Virtual machines
- Virtual Datacenter
 - Set of resources autonomously managed
- Deployment as a Service (DaaS)
 - Self provisioning of ready to use application packages
 - (WordPress, IdP, Moodle, Spark, ML, etc.)

Status

- **Resources**
 - ~9000 vCPU
 - 10 PB Storage
- **Usage**
 - Over 700 users
 - Over 1200 VM
- **Guarantees**
 - Service Continuity
 - Data Protection

Usage

[Download CSV Summary \(?start=2016-11-01&end=2017-03-22&format=csv\)](#)

Project Name	VCPUs	Disk	RAM	VCPU Hours <small>?</small>	Disk GB Hours <small>?</small>	Memory MB Hours <small>?</small>
isti	368	2.8TB	656GB	164333.47	1322184.30	279763673.66
unipa	336	7.5TB	1.2TB	693206.14	16207815.52	2692199747.8
ws2017ipv6	128	2.5TB	256GB	14955.82	299116.37	30629516.41
lns-prj1	54	543GB	71.5GB	118279.27	1629777.45	189620873.70
garrdemo318	32	20GB	32GB	16914.14	10571.34	17320083.68
INFN-FI	32	40GB	16GB	601.92	771.89	309779.80
wsosadmin	27	540GB	54GB	4510.02	90200.37	9236517.59
GEANT	25	481GB	48.5GB	29681.28	591502.48	60615623.66
infn-vlabs	24	480GB	48GB	8052.53	161050.61	16491582.24
demo	17	340GB	34GB	9877.99	197559.85	20230128.54
SSSUP	16	80GB	16GB	12651.09	63258.04	12954892.40
garrdemo125	8	100GB	128GB	13429.14	167864.31	220023110.77

New: Container Platform Architecture

- Automated platform deployment on bare metal, AWS or other clouds
- Automated workload deployment
- Distributed storage system Ceph
- Storage cluster for sharing big data
- Docker containers managed by Kubernetes



GPU-based Artificial Intelligence Platform

- GPUs on cloud servers with pass-through
- Ready to use with fully loaded with most popular Open Source Deep Learning libraries
- According to Jerome Huang, CEO of NVIDIA:

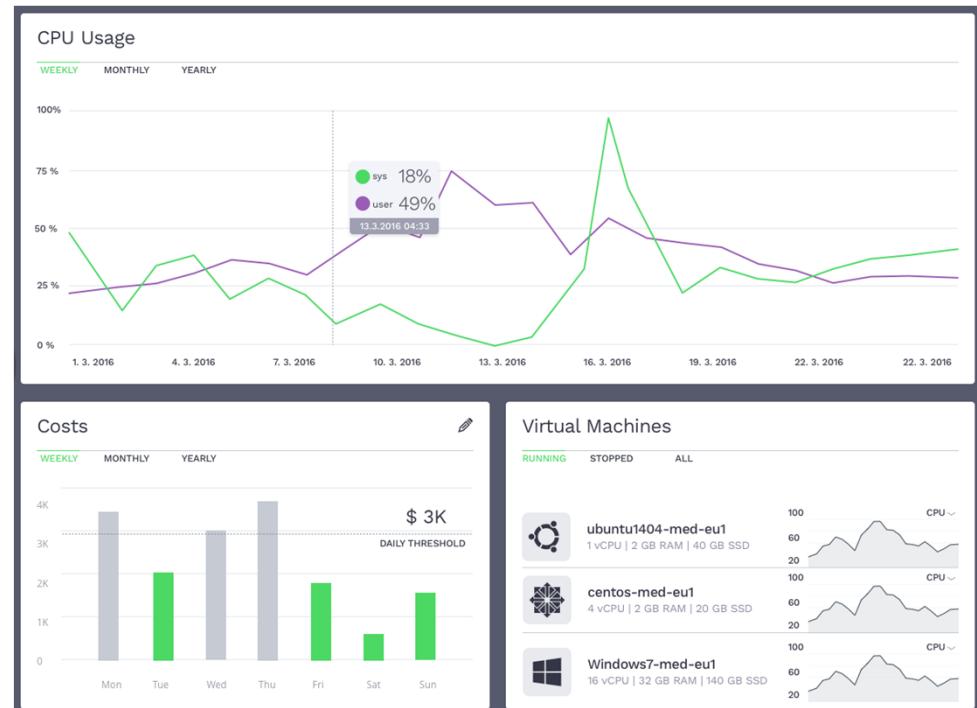
“The combination of deep learning, big data, and GPU computing makes ours the most revolutionary time in computer science”

- **Server with:**
 - 2 Xeon ES5-2698
 - 512 GB RAM
 - 2 x 800 GB SSD
 - 4 Nvidia GPUs Volta V100
 - **Deep Learning frameworks**
 - **Registry of Containers**
 - **Repository of annotated dat**
-
- **Accessible to researchers on one condition:**
give back training data and code for using them



Billing/Accounting

- Our own addition to OpenStack
- Provides detailed reporting on usage of every resource:
 - CPU
 - Disk (read/write)
 - Bandwidth
- Domain/Region Administrators can
 - Control usage and costs
 - In real time
 - Set limits on usage





THANK YOU !