Cloud activities @ INFN-Bari

People involved

- *Marica Antonacci*
- *Stefano Nicotri*
- *Roberto Valentini*
- *Vincenzo Spinoso*
- *Domenico Diacono*
- *Giacinto Donvito*
- *Alessandro Italiano*

Cloud layer

- OpenStack [production]
  - PRISMA cloud
  - HTCondor cluster
- PROXMOX [production]
  - General purpose services
- CloudStack [tested]
  - Evaluated for PRISMA
- OVirt [tested]
  - Evaluated as alternative for PROXMOX

Cloud, underlying hardware

Prisma Cloud	*HTCondor cluster*	*Calcolo Farm*
600 CPU 3TB of RAM 110 TB - R3 10G W/L	1216 CPU 3TB of RAM 6 TB - R3 1G Local	264 CPU 0,35TB of RAM 13 TB vi iSCSI 1G L

PRISMA cloud activities

PRISMA cloud activities

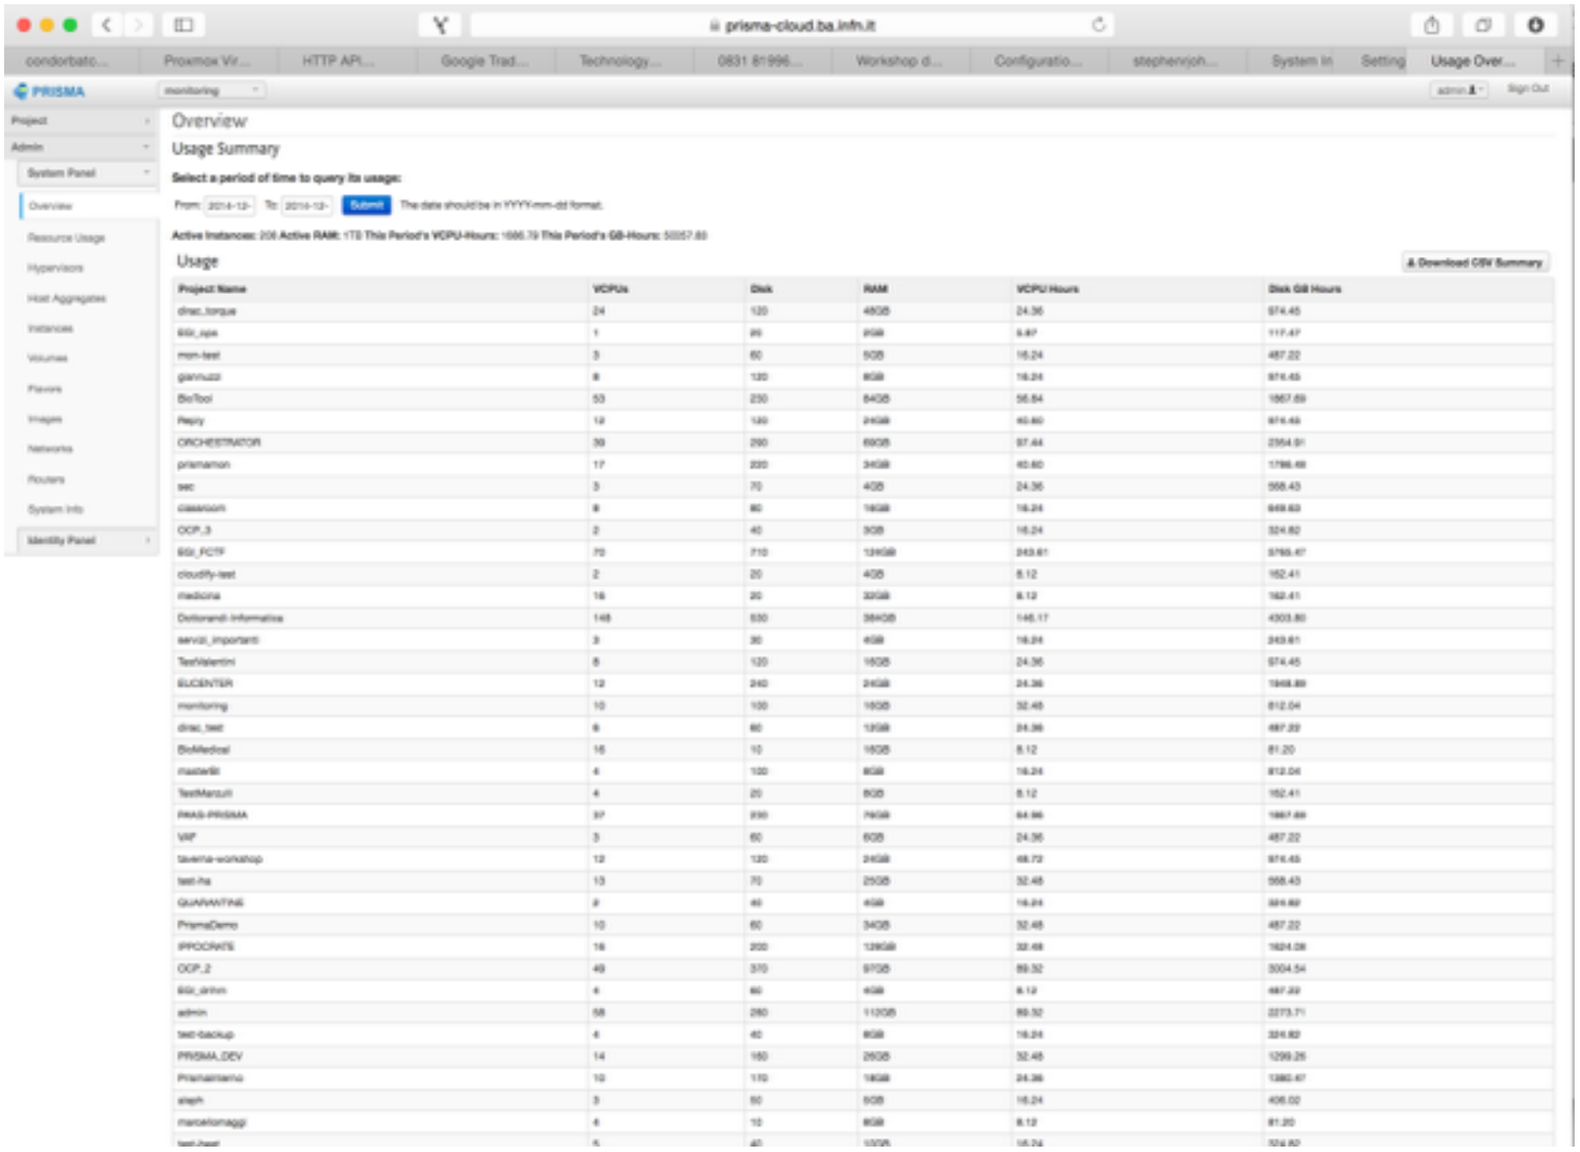

PRISMA cloud activities

Project Name	VCPUs
dirac_torque	24
EGI_ops	0
[illegible]	3
giannuzzi	8
BioTool	53
Reply	12
ORCHESTRATOR	39
prismamon	17
sec	3
classroom	8
OCP_3	2
PrismaInterno	10
cloudify-test	2
[illegible]	16
Dottorandi-Informatica	148
[illegible]	3
TestValentini	32
EUCENTER	12
monitoring	10
dirac_test	6
BioMedical	16
[illegible]	4
TestMarzulli	4
PAAS-PRISMA	37
VAF	3

Project Name	VCPUs
taverna-workshop	12
test-ha	13
[illegible]	2
[illegible]	12
[illegible]	16
[illegible]	57
[illegible]	4
admin	58
test-backup	2
PRISMA_DEV	14
EGI_FCTF	68
aleph	3
marcellomaggi	4
test-heat	5
AlexisCorso	16
condor-test	4
sismic	20
amiderha	8
PRISMA-BARI-Services	52
NetworkAnalysis	22
BioInf	16
uniba_matematica	4
Dirac	22
latraccia	8
BioVel	32

PhD students can requests computing resource.

EGI Federated CLoud Task Force

PRISMA cloud activities

- *BioVel:*
  - Biovel Portal: it is a browser-based interface that contains everything necessary to run users' workflow without user having to install anything on their own computer.
  - Biovel OpenRefine: it is a powerful tool for working with messy data, cleaning it, transforming it from one format into another, extending it with web services, and linking it to databases.
  - BioSTIF: it is a web map application designed for the interactive visualisation of spatio-temporal data.
- *Drihm:*
  - *Run Hydrology and Hydraulic models on the cloud*
  - *This VO use windows VM contextualised with cloudbase*
- *DCH-RP: Long term preservation in the cloud.*
  - *Sul nostro sito abbiamo un'istanza della HAPPI VM:*
  - *SCIDIP-ES HAPPI supports the archive manager and curator to capture and manage part of the Preservation Descriptive Information (PDI).*
- *CERN PoC: Run jobs on the EGI FedCloud through Vac/VCycle*
- *Chipster/Elixir:*
  - *VM Chipster: a user-friendly analysis software for high-throughput data. It contains over 300 analysis tools for next generation sequencing (NGS), microarray, proteomics and sequence data*

PRISMA Cloud consolidation

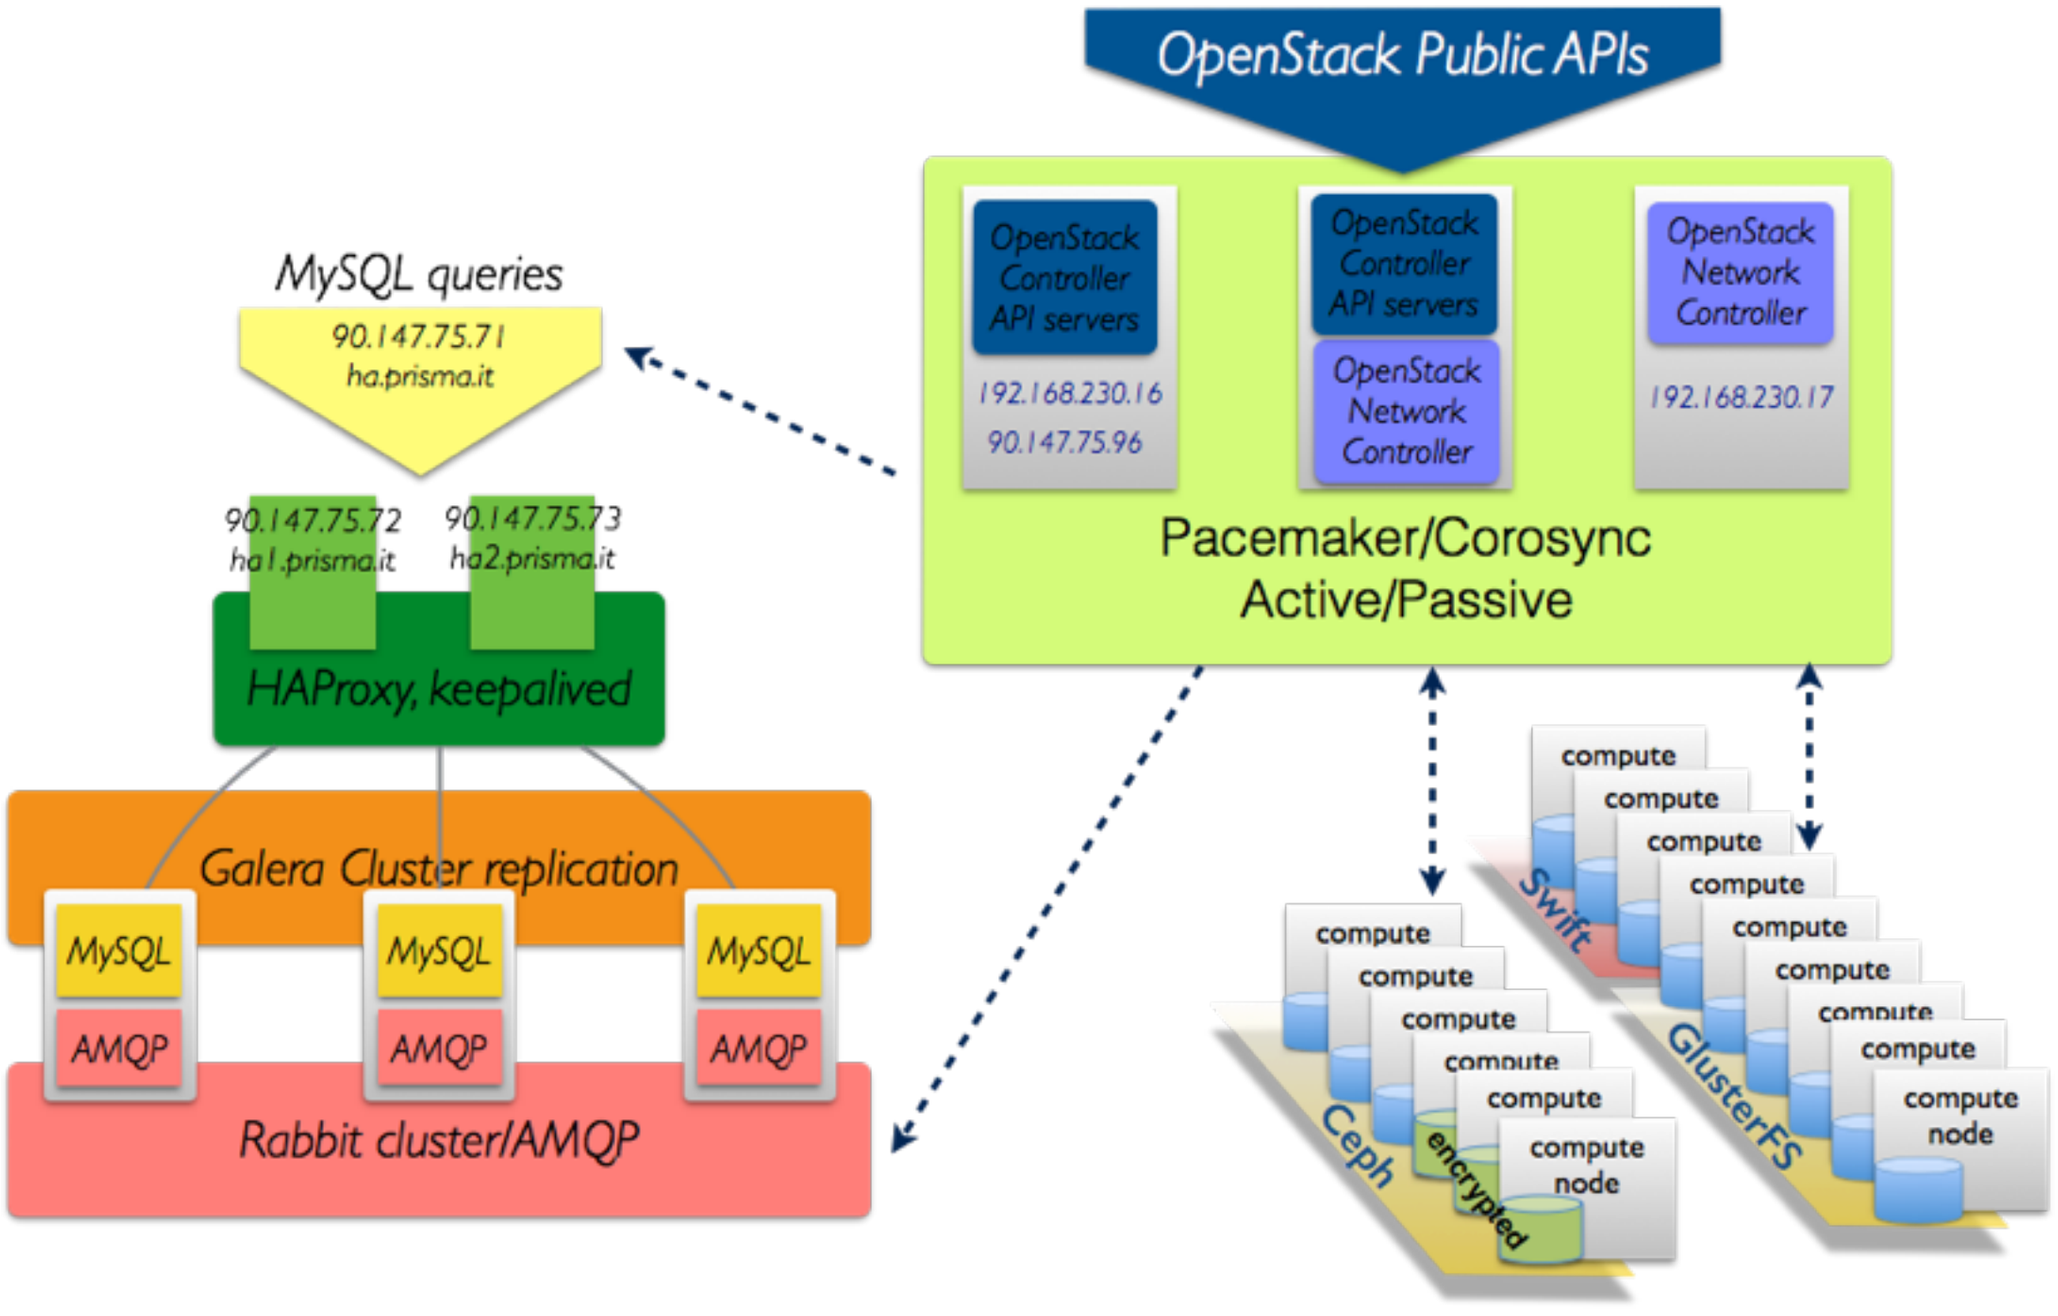

PRISMA cloud features

- *Network*
  - *GRE and FLAT network*
  - *All endpoint use a secure connections*
  - *Keystone v3 in production. Support for domain*
- *Storage*
  - *CEPH as Cinder backend*
  - *SWIFT standard deployment, used mainly for image backup*
- *Resource allocation*
  - *Quality of Service using HostAggregation with different overbooking policies*
- *Orchestration*
  - *Gained experience with HEAT and Puppet. Not yet a full integration with Ceilometer*

Multi-Region activities

- *Bari è una delle tre sedi (insieme a LNGS e PD) che è stata coinvolta fin dall'inizio nella sperimentazione del set-up di una cloud multi-region*
- *Le attività svolte e in corso:*
  - *installazione e configurazione dei vari componenti*
  - *debugging e troubleshooting*
  - *manutenzione (aggiornamenti, applicazione di patch, etc.)*
  - *supporto per l'integrazione delle nuove regioni*
  - *definizione di un'architettura di monitoring basata su Zabbix*
- *I test in corso a Bari:*
  - *set-up di keystone per abilitare le API V3 che consentono in particolare l'utilizzo dei domain (in Icehouse non è possibile far coesistere gli endpoint V2 e V3)*
  - *realizzazione di un repository centralizzato per le immagini virtuali (PoC)*

Multi-Region deployment

- *Keystone server + mysql (nodo del cluster Galera)*
  - *contribuisce all'implementazione dell'identity service distribuito geograficamente (BA,LNGS,PD)*
- *HAProxy (load balancer dei keystone server distribuiti) - con IP inserito in DNS HA*
- *Swift proxy + storage nodes (release Juno)*
  - *contribuisce all'implementazione dell'object storage distribuito geograficamente (BA,LNGS,PD)*
- *Istanza di Zabbix per il monitoraggio dei servizi delle varie sedi*
- *Openstack region "bari":*
- *2 host fisici: 1 controller+network node; 1 compute node*
- *servizi abilitati:*
  - *Neutron: configurazione GRE - plugin ovs, VPNaaS, LBaaS*
  - *Glance. Default backend: swift distribuito*
  - *Cinder. cinder-volume e cinder-backup con driver RBD*
  - *Nova. Instanze su cephFS*
  - *Heat.*
  - *Dashboard (Horizon): http://bari-region-ctl.ba.infn.it/horizon*

HTCondor cluster

- *We have exploited OpenStack in order to deploy a HTCondor cluster*
- *OpenStack infrastructure completely deployed and managed by Puppet*
  - *Controller[All-In-One], Swift cluster, Compute Nodes*
  - *Puppet config under version controll*
- *Condor cluster16K slots with 0,7K VM managed by Puppet*
- *We have planed to test swift as storage backend for batch jobs*

General purpose cloud features

- General purpose services have been deployed using PROXMOX
  - PROXMOX main features:
    - KVM/Container, CentralManagement, Backup/Restore, VM HA, Network, Storage
  - Running services:
    - IMAP, SMTP, DNS, LDAP, KERBEROS, RADIUS, ZABBIX, DCC, AFS, OwnCloud, OpenVPN, DHCP, WebServers and several user VMs
- Storage high availability
  - each Storage device has a target iSCSI on two networks
  - two 10GB switches provide connectivity
  - each PROXMOX HV is connected with both switches
  - multipaths-tools provides storage failover