

Cloud@CNAF Management and Evolution

C. Duma¹, A. Costantini¹, D. Michelotto¹ and D. Salomoni¹

¹INFN Division CNAF, Bologna, Italy

E-mail: ds@cnafe.infn.it

Abstract.

Cloud@CNAF is the cloud infrastructure hosted at CNAF, based on open source solutions aiming to serve different use cases present here. The infrastructure is the result of the collaboration of a transversal group of people from all CNAF functional units: networking, storage, farming, national services, distributed systems. If 2016 was for the Cloud@CNAF IaaS (Infrastructure as a Service) based on OpenStack [1], a period of consolidation and improvement, 2017 was an year of consolidation and operation ended with an extreme event - the flooding of the DataCenter, when an aqueduct pipe located in the street nearby CNAF went broke. This event caused down of the entire DataCenter, including the Cloud@CNAF infrastructure. This paper presents the activities carried out throughout 2018 to ensure the functioning of the center cloud infrastructure, that saw the its migration from CNAF to INFN-Ferrara, starting to the re-design of the entire to cope with the limited availability of space and weight imposed by the new location, to the physical migration of the racks and remote management and operation of infrastructure in order to continue to provide high-quality services for our users and communities.

1. Introduction

The main goal of Cloud@CNAF [1] project is to provide a production quality Cloud Infrastructure for CNAF internal activities as well as national and international projects hosted at CNAF:

- Internal activities
 - Provisioning VM for CNAF departments and staff members
 - Provisioning of VM for CNAF staff members
 - Tutorial and courses
- National and international projects
 - Providing VMs for experiments hosted at CNAF, like CMS, ATLAS, EEE
 - testbeds for testing the services developed by projects like the OpenCityPlatform & INDIGO-DataCloud

The infrastructure made available is based on OpenStack [2], version Mitaka, with all the services deployed using a High-Availability (HA) setup or in a clustered manner (for ex. for the DBs used). During 2016 the infrastructure has been enhanced, by adding new resources, compute and network, and its operation has been improved and guaranteed by adding the monitoring part, improving the support, automating the maintenance activities.

Thanks to this enhancement, Cloud@CNAF was able to offer high reliable services to the users and communities who rely on such infrastructure.

At the end of 2017, on November 9th early at morning, an aqueduct pipe located in the street nearby CNAF, broke as documented in Ref. [3]. As a result, a river of water and mud flowed towards the Tier1 data center. The level of the water did not exceed the threshold of safety of the waterproof doors but, due to the porosity of the external walls and the floor, it could find a way into the data center. Both electric lines failed at about 7.10AM CET. Access to the data center was possible only in the afternoon, after all the water had been pumped out. As a result, the entire Tier1 data center went down, including the Cloud@CNAF infrastructure.

2. The resource migration

Some weeks after the flooding, we decided to move the Cloud@CNAF core services in a different location in order to recover the services we provided for community and experiments. Thanks to a strong relationship, both University of Parma/INFN-Parma and INFN-Ferrara proposed to host our core machinery and related services. Due to the geographical proximity and the presence of Point of Presence (PoP) GARR, we decided to move the Cloud@CNAF core machineries to the INFN-Ferrara location.

Unfortunately, INFN-Ferrara was not able to host all the Cloud@CNAF resources due to a limited power availability and weight. For such reason, we decided to carry on an important activity aimed at re-designing the new infrastructure. In order to do that, we selected the services and the related machinery to move to the new - temporary - location to fit the maximum power consumption and weight estimated for each of the two rooms devoted to host our services (see Table 1 for details).

3. Re-design the new infrastructure

Due to the limitations described in Table 1 we were pushed to re-design the Cloud@CNAF infrastructure by using (only) three racks to host Cloud@CNAF core services (see Table 1 for the list of services). Among these three racks, the first one hosted the storage resources, the second hosted Openstack controller and network services, together with the GPFS cluster and other services. The third rack hosted Ovirt and Openstack compute nodes and some other ancillary services. Rack 1 and 2 have been connected by 2x40Gbps through our Brocade VDX switches and Rack 1 and 3 have been connected by 2x10Gbps through PowerConnect switches.

Moreover, Rack 1 is connected to PoP GARR with 1x1Gbps fiber connection. A complete overview of the new infrastructure and related resource location is shown in Figure 1. As depicted by Figure 1 and taking into account the limitations described in Table 1, we were able to limit the power consumption up to 13,79kW in respect to Room 1 (limit 15kW) and up to 5.8kW (limit 7kW) in respect to Room 2.

The whole migration process (from the design to the reconfiguration of the new infrastructure) took almost a business week and after that the Cloud@CNAF infrastructure and related services were up and running, able to serve again different projects and communities.

Thanks to the experience and the documentation provided, in June 2018 - after the Tier1 returned in its production status, Cloud@CNAF has been migrated back in less than three business days.

4. Cloud@CNAF evolution

Starting from the activity carried out in 2016 related to the improvements done at the infrastructure level [1], in 2018 (after the return of the core infrastructure services due to the flooding) the increase of the computing resources, in terms of quality and quantity, continued in order to enhance both the services and the performance offered to users.

Thanks to such activity, during the last year the Cloud@CNAF saw a growth on the number of users and use cases implemented in the infrastructure, in particular the number of projects

increased up to 87 which means a total consumption of 1035 virtual CPUS, 1766TB of RAM, with a total of 267 virtual machines (see Figure 2 for more details).

Among others, some of the project that used the cloud infrastructure are:

- HARMONY - under the TTLab coordination, is a project aimed at finding resourceful medicines offensive against neoplasms in hematology
- EEE - Extreme Energy Events - Science inside Schools (EEE), is a special research activity about the origin of cosmic rays carried out with the essential contribution of students and teachers of high schools,
- USER Support - for the development of experiments dashboard and the hosting of the production instance of the dashboard, displayed on the monitor present on the CNAF hallway,
- DODAS - for Elastic Extension of Computing Centre batch resources on external clouds,
- Services devoted to EU projects like DEEP-HDC [4], XDC [5] and many more.

5. Conclusions and future work

Due to a damage in the aqueduct pipe located in the street nearby CNAF, a river of water and mud flowed towards the Tier1 data center causing the shutdown of the entire data center. For such reason, the services and related resources hosted by Cloud@CNAF went down. To cope with this problem, we decided to temporarily migrate the core resources and services of Cloud@CNAF to INFN-Ferrara and to do this a complete re-design of the entire infrastructure was needed to tackle the limitations in terms of power consumption and weight imposed by the new location. Due to the joint effort of all the CNAF people and the INFN-Ferrara colleagues we were able to re-design, migrate and make operational the new Cloud@CNAF infrastructure and related hosted services in less than a business week. Thanks to the experience and the documentation provided, in June 2018 - after the Tier1 returned in its production status, Cloud@CNAF has been migrated back in less than three business days. Even taking into account the above described problems, we were able to maintain and evolve the Cloud@CNAF infrastructure, giving the possibility to the old and new users to continue their activities and carry out their results. In the next year, new and challenging activities are planned, in particular the migration to the OpenStack Rocky version.

	Rack1	Rack2	Rack3	Room1 (Max)	Room2 (Max)	Tot
Power consumption (kW)	8,88	4,91	5,8	13,79 (15)	5,8 (7)	19,59
Weight (Kg)	201	151	92	352 (400Kg/mq)	92 (400Kg/mq)	444
Occupancy (U)	9	12	10	21	10	31

Table 1. Power consumption weight and occupancy for each Rack

Rack1	Rack2	Rack3
VDX	VDX	PowerConnect x2
EqualLogic	Cloud controllers	Ovirt nodes
Powervault	Cloud networks	Compute nodes
	Gridstore	DBs nodes
	Other services	Cloud UI

Table 2. List of resources and services hosted per Rack

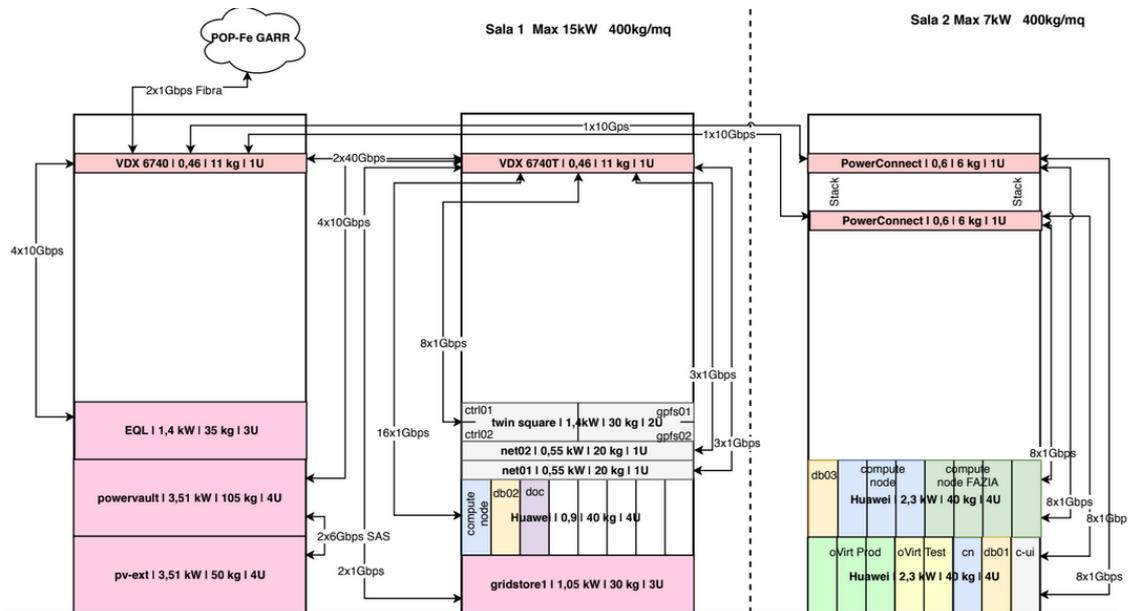


Figure 1. The new architecture of the Cloud@CNAF developed to cope the limitations at INFN-Ferrara.

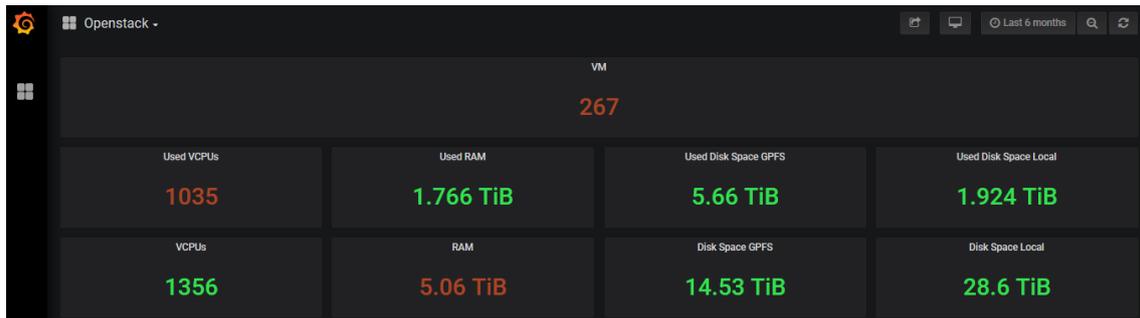


Figure 2. Cloud@CNAF monitoring and status

6. References

- [1] Cloud@CNAF - maintenance and operation, C. Duma, R. Burchi, A. Costantini, D. Michelotto, M. Panella, D. Salomoni and G. Zizzi, CNAF Annual Report 2016, <https://www.cnaf.infn.it/Annual-Report/annual-report-2016.pdf>
- [2] Web site: <https://www.openstack.org/>
- [3] The flood, L. dell'Agnello, CNAF Annual Report 2017, <https://www.cnaf.infn.it/wp-content/uploads/2018/09/cnaf-annual-report-2017.pdf>
- [4] Web site: <https://deep-hybrid-datacloud.eu/>
- [5] Web site: www.extreme-datacloud.eu