

# INFN Corporate Cloud - Management and Evolution

C. Duma<sup>1</sup>, A. Costantini<sup>1</sup>, D. Michelotto<sup>1</sup> and D. Salomoni<sup>1</sup>

<sup>1</sup>INFN Division CNAF, Bologna, Italy

E-mail: [ds@cnafe.infn.it](mailto:ds@cnafe.infn.it)

## Abstract.

This paper describes the achievements and the evolution of INFN Corporate Cloud (INFN-CC), the geographically distributed private Cloud infrastructure aimed at providing ICT services starting from the Infrastructure as a Service (IaaS) cloud level based on OpenStack. In particular, the contribution provided by CANF in terms of operations and possible evolution is here described and analysed.

## 1. Introduction

The INFN Cloud Working Group has been active for almost three years within the so called "Commissione Calcolo e Reti" (CCR) in INFN. Its activity being that of testing and acquiring expertise on technologies related to Cloud Computing and of selecting solutions that can be adopted in INFN sites in order to meet the computing needs of the INFN scientific community and more generally to ease information sharing inside and outside INFN. In the recent past, a number of projects related to Cloud Computing started in INFN thanks to the knowledge and expertise that were the outcome of the activity of the Cloud Working Group. A restricted team has been working in the last two years on the deployment of a distributed private cloud infrastructure to be hosted in a limited number of INFN sites. The INFN-CC working group planned and tested possible architectural designs for the implementation of a distributed private cloud infrastructure and implemented a prototype that is described hereafter.

INFN-CC [1, 2] is intended to represent a part of the INFN Cloud infrastructure, with peculiar features that make it the optimal cloud facility for a number of usecases that are of great importance for INFN.

While the INFN Cloud ecosystem will be able to federate heterogeneous installations that will forcibly adopt a loose coupling scheme, INFN-CC tightly couples a few homogeneous OpenStack installations that share a number of services, while being independent, but still coordinated, on other aspects.

The focus of INFN-CC is on resource replication, distribution and high availability, both for network services and for user applications. INFN-CC represents a single, though distributed, administrative domain.

## 2. INFN-CC - a distributed cloud

As already mentioned, INFN Corporate Cloud (INFN-CC) is the INFN geographically distributed private Cloud infrastructure aimed at providing services starting from the IaaS level and it is based on OpenStack that has been deployed in three of the major INFN data-centres

in Italy (INFN-CNAF, INFN-Bari and INFN-LNF). INFN-CC has a twofold purpose: on one hand its fully redundant architecture and its resiliency characteristics make of it the perfect environment for providing critical network services for the INFN community, on the other hand the fact that it is hosted in modern and large data-centres makes of INFN-CC the platform of choice for a number of scientific computing use cases. INFN-CC also deploys a higher PaaS layer, developed within the EU funded project INDIGO-DataCloud [3], in order to provide to the INFN scientific communities not only an easier access solution to computing and storage resources, but also both automatic instantiation and configuration of services or applications used for their everyday work, like batch-system on demand or big data analytics facilities. The PaaS layer, together with the underlying IaaS, is able to provide automatic scalability of the clusters instantiated and fault tolerance in case of a single node or complete site failures.

### *2.1. Architecture and services*

Technically speaking, from the OpenStack [4] point of view INFN-CC is a multi-region cloud composed of different OpenStack installations sharing a set of services that are managed globally while maintaining other services local, as shown in Figure 2.

The available INFN-CC services can be summarized in the following categories:

Ancillary services such as:

- A distributed Percona XtraDB Cluster relies on this network and is the back-end both for the identity service and the image service catalog,
- A distributed DNS dynamically modified, by humans as well as monitoring processes, in order to make clients point only to working endpoints.

Local services, implemented independently on each site, they have a local scope such as compute, volume and network. In particular Compute and Volume services rely on a CEPH [5] back-end. Each site has a CEPH instance with different priority and the CEPH rbd mirror is employed to replicate data across INFN-CC sites for disaster recovery.

Global services, implemented on all sites for high availability, backed by common DBs when needed, they have a global scope and are here listed:

- Openstack Horizon, providing the GUI to access the INFN-resources and services via Web,
- OpenStack Keystone access points, pointing to the above mentioned distributed DBMS, are available on all INFN-CC sites,
- OpenStack Swift relies on the INFN-CC private network and is deployed geographically,
- Openstack Glance relies on CEPH as a storage backend and on the Percona cluster as a catalog, this way it is fully distributed as well.

In order to provide the above mentioned list of services, particular care was made to define the network setup among the INFN-CC sites in order to provide standard connectivity to the VMs and cross site connectivity among the three INFN sites. The connectivity, made by a level 3 distributed private network provided by GARR, allows an easy cloud management and fast data exchange. As shown in Figure ??, in the INFN-CC model the VM networks remain private and do not cross the border of their own “region”. Also public networks are separate and managed locally, except they might benefit of a cross-site DNS domain namespace in order to allow for easy service migration. Hosts on the management networks of the different sites, on the other hand, must be able to intercommunicate, possibly taking advantage of a set of loose firewall rules, in order to speed up the system setup and maintenance. Moreover, a cross-site DNS domain namespace enable to dynamically migrate cloud services, when needed, for high availability.

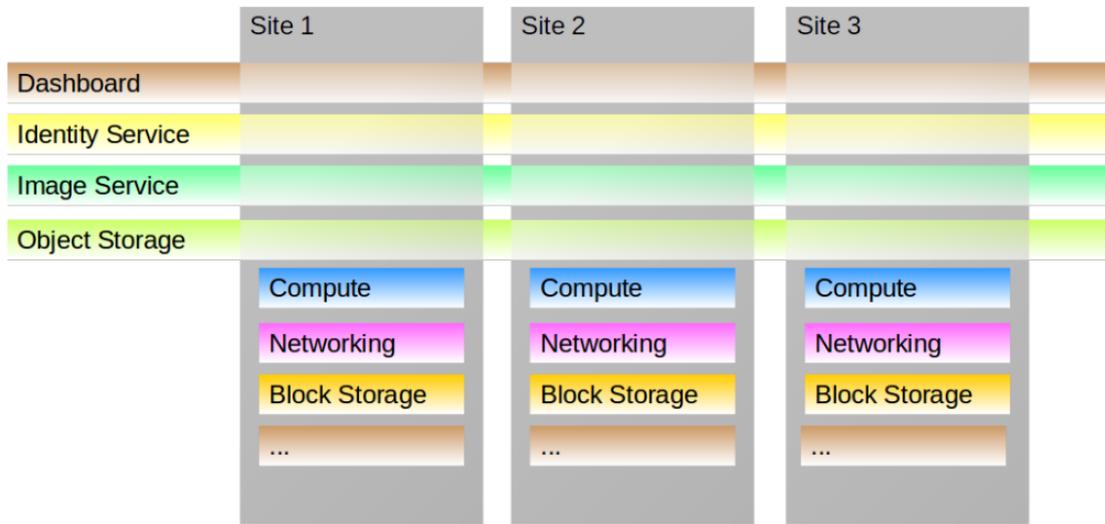


Figure 1. INFN-CC architecture and related services.

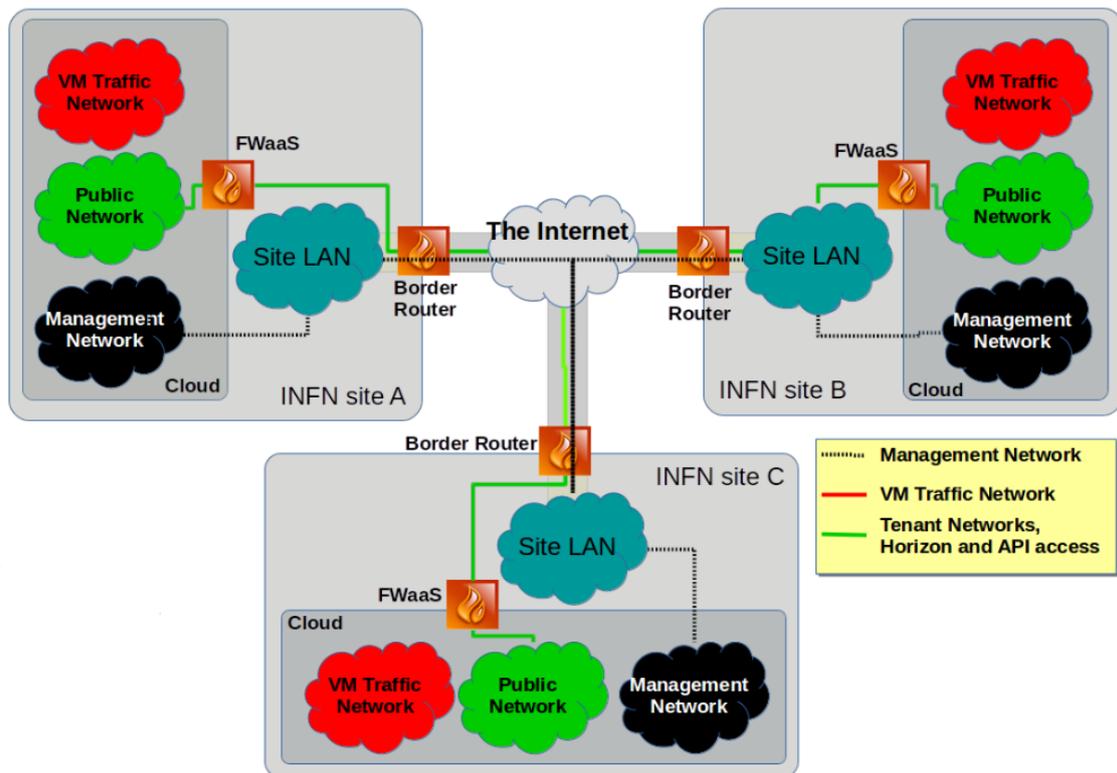


Figure 2. INFN-CC networking and related components.

## 2.2. INFN-CC functionalities

INFN-CC provides some interesting functionalities and features thanks to its geographical distribution among different sites. In particular, we can highlight the following:

- Single point of access to distributed resources, fully exploiting the native functionalities of

OpenStack and with no (or very small) need of external integration tools.

- Single Sign On (SSO) and common authorization platform. User roles and projects are the same throughout the infrastructure, while quotas for projects vary from site to site.
- Secure dashboard and API access to all services for all users. The dashboard, OpenStack APIs and EC2 APIs are available. All services are implemented on top of an SSL layer, in order to secure resource access and data privacy.
- Easy sharing of VM images and snapshots through a common Object Storage deployment. A single image/snapshot database is used by all the project sites. This means that all VM images and snapshots are available in all sites.
- Common DNS name space for distributed resources. DNS HA provides high availability for distributed resources.
- Block device sharing over remote sites; a rough way to implement is through CEPH backend volume backups, faster and more efficient ways are under investigation. The final approach will mainly depend on the WAN bandwidth and latency among the INFN-CC sites.
- Self-service backup for instances and block storage. Backed-up data can be accessed/restored transparently from/to any site. Final users and tenant administrators are responsible for backing up their instances and the attached block devices. Adequate tools, native to OpenStack, are provided. As the backup storage backend, both for instance images and snapshots and for block devices, are replicated and distributed, data backup is transparently available in all the cloud sites and is still available in the case of a site failure.

### *2.3. The management model*

As resources in INFN-CC are so closely coupled and interdependent, they must be managed carefully by expert staff and must always work correctly. For this reasons a limited number of cloud administrator, no matter where they are based, are allowed to administer hosts offering OpenStack services in any INFN-CC site both for normal maintenance and in case of emergency. This approach is eased by the homogeneity of the infrastructure, but requires a trust agreement that breaks the barriers of the single site: remote administrators must be trusted exactly as residents. On the other hand, infrastructure and hardware management is not easily performed from a remote site and should be full responsibility of the local IT staff. Cloud administrators and local IT staff should of course interact for better problem detection and solving. This management model brings issues that exceed the technical and organizational problems of a distributed management team: hosting sites must agree on having external people manage part of their infrastructure as if they were local staff.

## **3. Use cases for the INFN-CC**

The architecture of INFN-CC is particularly fit for a wide range of use cases where a strict relation exists between resources that are distributed over different sites. Most of these use cases are related to the delivery of computing services for the INFN community, be they of local interest for users belonging to a single INFN site or of general interest for the whole community. This does not mean that scientific computing is unfit for the INFN-CC, but often scientific computing environments do not need the high availability features provided by INFN-CC and can take advantage of other cloud deployments.

### *3.1. Scientific computing*

Massive data analysis or simulations do not usually need an environment like that of INFN-CC, but this does not mean that the INFN-CC doors are closed to scientific computing. Tier 3 virtualization, the last mile of data analysis, as well as software development environments are

the first use cases that might take advantage of INFN-CC and use it efficiently. Further use cases might be applicable in the future, according to the available resources and to the project development.

### *3.2. Distributed Web application*

In this use case, a generic distributed web application can use a distributed SQL database (accessed through HAProxy) and a distributed object storage data backend (with almost no single point of failure). The failure of one instance does not affect final users, that are still able to use the application.

## **4. Operations and evolution**

As previously described in the text, the INFN-CC infrastructure is a multi-region cloud composed of different OpenStack installations. The actual configuration of the CNAF region is here described:

- A network node: bare metal resource hosting the OpenStack Networking service that deploys several processes across a number of nodes. These processes interact with each other and other OpenStack services. The main process of the OpenStack Networking service is neutron-server, a Python daemon that exposes the OpenStack Networking API and passes tenant requests to a suite of plug-ins for additional processing
- Two Compute nodes: bare metal resources on which the VM's are actually deployed. Each compute node runs an hypervisor to deploy and run the VM.
- A storage node: bare metal resource hosting a cluster distribution of CEPH object storage that provides interfaces for object-, block- and file-level storage
- A ToR switch

Up to date, CNAF is providing to the INFN-CC the following IaaS resources: 20 VCPUs, 50GB RAM, 50 FIPs and 50TB of volume storage. A set of new resources are expected to be acquired by 2019 and became part of the CNAF region of INFN-CC infrastructure. In the next year, a contribution from CNAF in terms of development and integration of new services is also expected, in particular in the deployment and testing of CEPH distributions.

## **5. Conclusions**

In this contribution, the INFN-CC cloud infrastructure has been briefly presented from different aspects. Architecture, services, maintenance model and possible usecases have been also discussed. In particular, the contribution of CNAF in terms of operations and evolution of both the CNAF region and the the INFN-CC cloud infrastructure at whole have been described. In the next years, the evolution of the services offered by INFN-CC is expected to bring new and changelling usecases. In this respect, CNAF is aiming to contribute in terms of manpower and expertise to improve both the reliability and the quality of the service offered to INFN and worldwide.

## **6. References**

- [1] Web site: <https://indico.cern.ch/event/587955/contributions/2935944>
- [2] Web site: [https://wiki.infn.it/cn/ccr/cloud/infn\\_cc](https://wiki.infn.it/cn/ccr/cloud/infn_cc)
- [3] Web site: [www.indigo-datacloud.eu](http://www.indigo-datacloud.eu)
- [4] Web site: <https://www.openstack.org/>
- [5] <https://ceph.com>
- [6] Web site: <https://deep-hybrid-datacloud.eu/>
- [7] Web site: [www.extreme-datacloud.eu](http://www.extreme-datacloud.eu)