INFN and the evolution of distributed scientific computing in Italy

Davide Salomoni[1], *Ahmad* Alkhansa[1], *Marica* Antonacci[2], *Patrizia* Belluomo[3], *Massimo* Biasotto[4], *Luca Giovanni* Carbone[5], *Daniele* Cesini[1], *Diego* Ciangottini[6], *Vincenzo* Ciaschini[1], *Alessandro* Costantini[1], *Alessandra* Doria[7], *Giacinto* Donvito[2], *Doina Cristina* Duma[1], *Federica* Fanzago[8*], *Nadina* Foggetti[2], *Federico* Fornari[1], *Emidio Maria* Giorgio[9], *Alessandro* Italiano[2], *Giada* Malatesta[1], *Barbara* Martelli[1], *Diego* Michelotto[1], *Lucia* Morganti[1], *Jacopo* Gasparetto[1], *Gianluca* Peco[10], *Carmelo* Pellegrino[1], *Andrea* Rendina[1], *Massimo* Sgaravatto[8], *Francesco* Sinisi[1], *Daniele* Spiga[6], *Vincenzo* Spinoso[2], *Bernardino* Spisso[7], *Stefano* Stalio[11], *Lucio* Strizzolo[12], *Sergio* Traldi[8], *Marco* Verlato[8] and *Enrico* Vianello[1]

[1]INFN, CNAF, Viale Berti Pichat 6, 40127 Bologna, Italy
[2]INFN, Sezione di Bari, Via Giovanni Amendola 173, 70126 Bari, Italy
[3]INFN, Sezione di Catania, Via Santa Sofia 62, 95123 Catania, Italy
[4]INFN, Laboratori Nazionali di Legnaro, Viale dell'Università 2, 35020 Legnaro PD, Italy
[5]INFN, Sezione di Milano-Bicocca, Piazza della scienza 3, 20126 Milano, Italy
[6]INFN, Sezione di Perugia, Via Alessandro Pascoli 23 c, 06123 Perugia, Italy
[7]INFN, Sezione di Napoli, Strada Comunale Cinthia, 80126 Napoli, Italy
[8]INFN, Sezione di Padova, Via Marzolo 8, 35131 Padova, Italy
[9]INFN, Laboratori Nazionali del Sud, Via Santa Sofia 62, 95123 Catania, Italy
[10]INFN, Sezione di Bologna, Viale Berti Pichat 6, 40127 Bologna, Italy
[11]INFN, Laboratori Nazionali del Gran Sasso, Via Giovanni Acitelli 22, 67100 L'Aquila, Italy
[12]INFN, Sezione di Trieste, Galleria Padriciano 99, 34149 Trieste, Italy

Abstract. INFN has been running a distributed infrastructure (the Tier-1 at Bologna-CNAF and 9 Tier-2 centres) for more than 20 years which currently offers about 150000 CPU cores and 120 PB of space both in tape and disk storage, serving more than 40 international scientific collaborations. This Grid-based infrastructure was augmented in 2019 with the INFN Cloud: a production quality multi-site federated Cloud infrastructure, composed by a core backbone, and which is able to integrate other INFN sites and public or private Clouds as well. The INFN Cloud provides a customizable and extensible portfolio offering computing and storage services spanning the IaaS, PaaS and SaaS layers, with dedicated solutions to serve special purposes, such as ISO-certified regions for the handling of sensitive data. INFN is now revising and expanding its infrastructure to tackle the challenges expected in the next 10 years of scientific computing adopting a "cloud-first" approach, through which all the INFN data centres will be federated via the INFN Cloud middleware and integrated with key HPC centres, such as the pre-exascale Leonardo machine at CINECA. In such a process, which involves both the infrastructures and the higher level

* Corresponding author: federica.fanzago@pd.infn.it

services, initiatives and projects such as the "Italian National Centre on HPC, Big Data and Quantum Computing" (funded in the context of the Italian "National Recovery and Resilience Plan") and the Bologna Technopole are precious opportunities that will be exploited to offer advanced resources and services to universities, research institutions and industry. In this paper we describe how INFN is evolving its computing infrastructure, with the ambition to create and operate a national vendor-neutral, open, scalable, and flexible "datalake" able to serve much more than just INFN users and experiments.

1 Introduction

INFN is the Italian institute dedicated to the study of the fundamental constituents of matter and the laws that govern them [1]. Its theoretical and experimental research activities spread from subnuclear to astroparticle physics and are undertaken in close collaboration with the Italian universities in the context of an international framework. Since its foundation, in 1951, INFN was conceived as a geographically distributed institute with a collaborative work environment. Currently INFN integrates 4 national laboratories and about 30 divisions, associated groups and other centres.

Addressing the computing needs of its experiments and communities is one of the tasks that INFN has to accomplish. INFN has always considered crucial to own, operate and develop a reliable computing infrastructure, distributed over its facilities. In particular, in the early 2000s, to support in primis the LHC experiments, a computing infrastructure consisting of a large Tier1 centre in Bologna and 9 smaller Tier2s was set up.

While the development of scientific computing within INFN was originally driven by the needs of its own physics communities, being the forefront of computing in research seeded many projects with a much broader scope.

INFN has therefore been evolving its infrastructure exploiting Grid and Cloud technologies, always in the context of the Open Science paradigm, creating the High-Performance Computing/Big Data/Artificial Intelligence infrastructure (HPC-BD-AI). This is now the largest distributed research and academic computing infrastructure.

The computing needs of the experiments in various disciplines are constantly increasing. The “High-Luminosity LHC” experiments and the next generation of many other experiments present unprecedented needs for computing. Moreover users are asking for more flexibility in selecting the most suitable tools and in creating dynamic environments, with increasingly distributed, personalised and integrated solutions. To cope with these challenges, it is necessary to foster and simplify the utilisation of all viable technologies, to be able to pervasive use geographically distributed storage, to be more efficient, elastic and resilient.

In this paper we describe how INFN is evolving its organisation and its infrastructure to effectively tackle the new challenges, and also exploiting new opportunities. The overall goal is to implement a transparent integration between distributed computing and storage resources covering the Edge, Cloud and traditional HPC domains. This integrated architecture is essential for achieving scalable and efficient processing of large amounts of data to flexibly support different use cases.

2 INFN and the Grid

INFN in the early 2000s was one of the main actors in several national and European projects with the goal of building a distributed Grid based computing and storage infrastructure.

In the context of the Worldwide LHC Computing Grid (WLCG) project [2], ten main international centres were selected. In Italy, this was the Tier-1 at Bologna CNAF. 9 additional Tier-2 centres were then added at Padova-LNL, LNF, Turin, Milan, Pisa, Rome, Naples, Bari, Catania. These computing centres are all still operational. Their size has increased by about 100 times since their creation: the INFN distributed computing infrastructure now provides 150.000 CPU cores, 120 PB of storage disk and 120 PB of storage on tape. They are connected through dedicated network links provided by the GARR consortium [3].

The INFN Distributed Computing federation delivers the LHC experiments O(7-20)% of their computing budgets. It also supports several non LHC (e.g nuclear and astrophysics) experiments. The resources assigned to these communities are about 10-20% of the total.

Fig. 1. Map of INFN facilities and LHC tiers.

3 INFN Cloud

The Grid model cannot easily meet all use cases of all user communities. In particular, small experiments experienced some problems in using the Grid based infrastructure because of some issues: strict policies to access and use the distributed resources, limited support for interactive activities, and required knowledge of the Grid middleware tools, etc. For these reasons, several INFN sites have deployed local cloud infrastructures. These local Clouds usually differ between themselves with respect to their hardware, software, and cloud middleware.

To optimise the use of available resources and expertise in the institute, INFN decided to implement a national cloud for research. The fundamental idea was to realise a federation of existing distributed infrastructures, able to extend to other private or commercial resource providers, and to provide to final users a dynamic set of ready-to-use services based on specific use cases. This was implemented by leveraging the experience gained in several national and European projects where INFN participated.

INFN Cloud was officially made available to users in March 2021 [4]. It is composed of a core backbone that connects the two biggest INFN data centres, one at CNAF, Bologna and the other in Bari, and a set of geographically distributed and federated sites. The two backbone sites, connected at high speed, are used to host the core services that require high availability.

INFN Cloud also integrates some ISO certificate regions used where sensitive data are stored and processed, as required by some life science and health projects [5].

A site can join the INFN Cloud infrastructure accepting the Rules of Participation, after the approval of the management board. Such rules define how these resources shall be provided, the expected SLAs, etc. in compliance with INFN, national and European laws.

The implementation of INFN Cloud is based on three fundamental pillars. The first is an open-source vendor independent infrastructure. The second is a dynamic resources orchestration that allows matching between services required by users and the available federated resources. The matchmaking is implemented by the Indigo PaaS orchestrator service, developed in the INDIGO-DataCloud European project. The third pillar is a unique and standard mechanism for the authentication and the authorization at all the cloud levels, realised by the INDIGO Identity and Access Management (IAM) service, also developed during the INDIGO-DataCloud project.

3.1 The INFN Cloud portfolio and services

INFN Cloud provides users with a dynamic and extensible portfolio of ready to use services spanning from the instantiation of simple VMs up to complex platforms already configured to implement specific workflows. Services are instantiated by the users as needed. All services are described through an Infrastructure as Code paradigm: they are represented through TOSCA templates (a standard language for describing the topology of applications in cloud). Services are implemented using a “lego-like” approach, composing and reusing simple building blocks. Ansible roles (to manage the automated configuration of virtual environments), and Docker containers (to encapsulate high-level application software and runtime) are among the technologies used in the implementation.

Services can be deployed by users via the dashboard web interface or command line. The dashboard hides from the user all the details about the cloud infrastructure and the TOSCA templates. Users have only to select the service they want to deploy and fill the form provided by the dashboard with the parameters for the service configuration. Then they can monitor and manage the deployment in a simple way.

INFN Cloud also provides some centrally managed services: object storage accessed through the S3 protocol, an image repository implemented through Harbor, and a Jupyter notebook.

4 New computing challenges and computing model

In the next few years, the computing needs of experiments are going to dramatically increase. Linear estimates foresee that considering the increased needs for computing at High-Luminosity LHC and the expected speed-up factor from technology advancements, by 2028 we will require 20x more computing power than our projected resources. Similar challenges are relevant for several other experiments in other scientific areas. To cope with these levels of requests, infrastructural and technological changes are also needed.

Also thanks to the funds provided by some projects of the "National Recovery and Resilience Plan" (NRRP), as discussed in section 5, INFN is renewing its infrastructure, to be ready for the High Luminosity-LHC era. The on-going interventions include increasing data centres capacity, the use of more compact computing, lowering the power usage effectiveness (PUE), and the update of the network infrastructure.

For what concerns the hardware, software and services, INFN is trying to foster and simplify the utilisation of any viable technologies: GPUs, FPGA, and Quantum when available.

By evolving its distributed infrastructure considering a "cloud datalake" model, INFN wants to make its infrastructure more flexible and scalable. INFN will keep owning a solid baseline of computing resources on stable sites, but it must be possible to easily integrate, even temporarily, any available computing resource of external providers: commercial clouds, other sciences' resources and HPC systems. The goal is to abstract from physical machines and to form a national pool of resources and high-level services.

INFN has already demonstrated the capability to execute LHC workflows on HPC systems exploiting HPC CINECA resources and on commercial clouds (e.g. ARUBA) in the context of some grant projects started in 2019 [6,7].

Another important optimization is related to storage and data management. The idea is to have a pervasive use of geographically distributed storage, "the datalake", reducing the number of replicas of data and deploying optimised caches.

This model, where CPU and storage are no longer coupled together, requires a dynamic and performant network to support intensive data processes and to take advantage of all the resources, wherever they are.

INFN Cloud can be considered the initial seed of such a national datalake for research and beyond, building on existing/renewed/new e-Infrastructures. It can therefore be considered the base of the evolution of the INFN distributed computing vision.

5 Opportunities for services and infrastructure improvements

The issues described in the previous sections are really challenging, but there are some opportunities that can be exploited.

The first one is the "Bologna Tecnopolo", which will become the national hub for supercomputing, competitive at international level. It is already hosting the European Centre for Medium-Range Weather Forecasts (ECMWF) and the Leonardo pre-exascale supercomputing. The INFN CNAF Tier-1 data centre will also soon be relocated here. In its final configuration the Tecnopolo will provide for INFN-CNAF and CINECA an infrastructure up to 23 MW ready for post-exascale and for the next generation of scientific experiments.

The other big opportunity to exploit is the Italian National Recovery and Resilience Plan (NRRP).

In the context of the NRRP, a project for the implementation of a National Research Center for high performance, big data and quantum computing, called ICSC, was approved and funded with 430 million euros. The project, led by INFN, involves 51 partners: research institutions, Universities, and private companies. ICSC will become a pervasive initiative throughout Italy. It foresees a Hub & Spoke organisation: 10 thematic spokes for technology developments and software applications, and one Infrastructure spoke [8].

To complement ICSC, INFN also proposed the TeRABIT project ("Terabit network for Research and Academic Big data in Italy") [9]. This project was approved and funded with 41 million euros. It aims to create a distributed, hyper-networked, hybrid Cloud-HPC computing environment by integrating the distributed INFN infrastructure with HPC resources of PRACE-Italy (CINECA [10]), through a high-speed network provided by GARR and complement the ICSC National Center. In this project INFN wants to augment its infrastructure with "HPC bubbles": clusters of CPUs-GPUs and FPGAs with fast storage for high performance computing. The aim is to realise an "Edge-Cloud Continuum" computing, exploiting AI technologies, which allows users to process big data in a more flexible way. HPC bubbles will be integrated among them, with cloud services and also with traditional HPC systems as Leonardo.

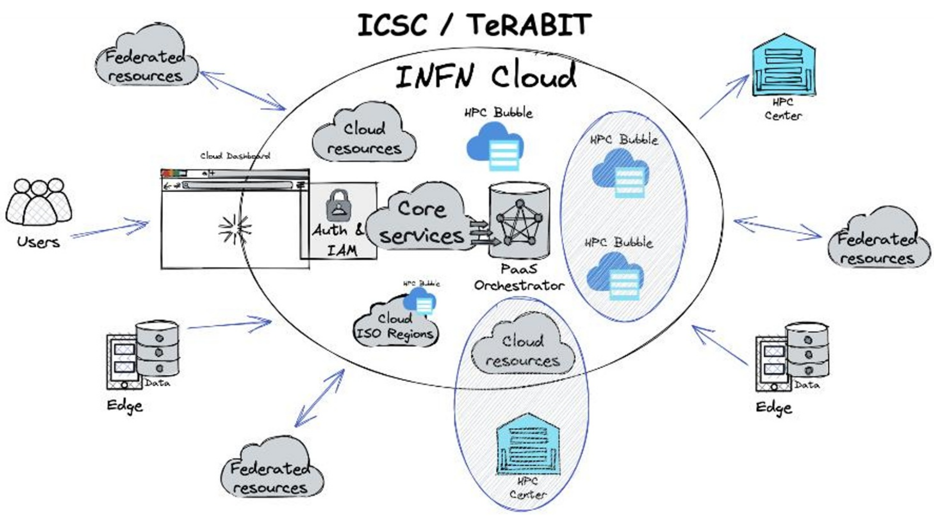

Fig. 2. The Edge-Cloud continuum computing.

6 The INFN DataCloud working group

To cope with these challenges, INFN decided to reorganise its computing management structure. In the new structure, activities are organised into working groups.

One of them is the infrastructure working group, called DataCloud, that is responsible for managing and evolving the INFN distributed infrastructure towards a cloud datalake model as described above.

The DataCloud activities are organised in 7 work packages (WPs).

WP1 is responsible for the operations. This includes the management and the updates of the infrastructure, its monitoring and accounting, the operations of the ISO certified regions, the coordination of the sites.

WP2 is responsible for the support, both for end users and site managers. It is also responsible for training and communication activities.

WP3 is responsible for the definition of sustainability models and for the evolution of the datalake architectural model, in collaboration with WLCG and other related projects or initiatives.

WP4 deals with security and policies. This includes the operations of a Security Operation Center (SOC), the definition of security policies to access the resources, the definition and updates of the rules of participation.

WP5 manages the development and maintenance of the middleware and tools deployed in the infrastructure. This includes the PaaS orchestrator and the IAM software.

WP6 is the R&D work package, where prototypes to address specific use cases are implemented and validated in testbeds. Activities in this WP include the implementation of prototypes for software distribution, on top of the CernVM File System (CVMFS), and the validation and customization of the WLCG data management tools to enable a national datalake.

WP7 is responsible for the integrated management system and for the management of legal compliance issues. Among its activities are ISO certification, data protection, and the support for the definition of the data management plans for projects using the infrastructure.

7 Conclusions

In the coming decade, computing and storage needs of LHC and experiments of other areas will be higher than what existing model can guarantee. For this reason, INFN is working on solutions to cope with this challenge by adopting the new cloud datalake computing model and exploiting the usage of high-performance resources and also integrating commercial providers. INFN has the ambition to create and operate an open, scalable and flexible cloud datalake that should become the computing environment not only for INFN core activities, but also for industrial research in Italy and beyond at an international level.

References

1. INFN, https://home.infn.it/en/
2. Worldwide LHC Computing Grid, https://wlcg.web.cern.ch/
3. Consortium GARR, https://www.garr.it/
4. INFN Cloud, https://www.cloud.infn.it/
5. A.Retico et al., *Enhancing the impact of Artificial Intelligence in Medicine: A joint AIFM-INFN Italian initiative for a dedicated cloud-based computing infrastructure*, Physica Medica Volume **91** (2021)
6. T. Boccali et al., *Enabling CMS Experiment to the utilization of multiple hardware architectures: a Power9 Testbed at CINECA*, Journal of Physics: Conference Series **2438** (2023), ACAT-2021
7. T.Boccali et al., *Extension of the INFN Tier-1 on a HPC system,* EPJ Web of Conferences **245**, 09009 (2020)
8. ICSC, https://www.supercomputing-icsc.it/en/icsc-home/
9. TeRABIT project, https://www.terabit-project.it/
10. Cineca, https://www.cineca.it/en