

DELIVERABLE D1.2.3

HIGH-LEVEL VIEW OF THE INFRASTRUCTURE

FINAL RELEASE

TeRABIT

APRIL 2026

Mauro Campanella (INFN-GARR)

Silvia Calegari, Giacinto Donvito, Claudio Grandi (INFN)

Massimo Carboni, WP leader; Claudia Battista, Direttrice (GARR)

Mirko Cestari, Giuseppa Muscianisi, Debora Testi (CINECA)

Giorgio Bolzon, Stefano Salon (OGS)

Eccetto dove diversamente specificato, il contenuto di questo sito è
concesso in licenza con Attribuzione Creative Commons - 4.0 internazionale.



Finanziato
dall'Unione europea
NextGenerationEU



Ministero
dell'Università
e della Ricerca



Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA



terabit

Table of Content

EXECUTIVE SUMMARY	2
1. HIGH LEVEL VIEW OF THE THREE INFRASTRUCTURES	1
1.1 The current status of the three RIs	1
1.1.1. GARR-T.....	1
1.1.2. PRACE-Italy.....	5
1.1.3. HPC-BD-AI.....	7
2. CURRENT INTEROPERABILITY AND RI FEDERATION EFFORT	9
3. REFERENCES	11

List of Tables

Table 1 - Distribution of the HPC Bubbles on the HPC-BD-AI infrastructure	8
Table 2: total HPC bubbles nodes and interfaces (Terabit, DARE and ICSC)	8

List of Figures

Figure 1- The GARR-X/T Network RI	1
Figure 2- The GARR-T Network evolution.....	2
Figure 3 - The GARR-T Network evolution in Sardinia.....	3
Figure 4: GARR-T backbone capacity growth versus users' requirements (black).....	3
Figure 5- The GARR-T Network evolution architecture.....	4
Figure 6- PRACE-Italy Tier1 node Galileo100	5
Figure 7: GAIA HPC Cloud total resources.....	6
Figure 8 - HPC-BD-AI sites, connected by the GARR network.....	7
Figure 9: UrganSHAKE system Overview	10



**Finanziato
dall'Unione europea**
NextGenerationEU



EXECUTIVE SUMMARY

The deliverable reports the status of the three research infrastructures at the end of the projects, after 40 months from its start.

Each infrastructure has evolved according to plans, exceeding, for all three, the foreseen growth and technical upgrade, as a consequence of tenders' results and of the synergy with the ICSC national centre and other PNRR projects. The document reports a summary of the interoperability and shared services effort to the objective of strengthening a federation of the three infrastructures, including agreements and interfaces to allow workflows communication for authenticated users. The architecture has constantly been jointly discussed with the ICSC experts to ensure interoperability between the two projects.

The infrastructures will continue to operate and evolve to serve research and education users.

1. HIGH LEVEL VIEW OF THE THREE INFRASTRUCTURES

The high-level view of the infrastructures at the start of the project has been provided by D1.2.1 and at the middle of the project by D1.2.2. The following sections report the situation in April 2026 and a final section on federation effort.

The GARR network continues to serve the other two RIs for their high-capacity connectivity to research and education in Italy and in other countries and to general Internet. More information on the detailed architecture for each RI can be found in "D1.2.4.3 Capacity planning". Additional architectural information of the RIs' development in TeRABIT are provided in D2.4.1.2, D3.26 and D4.8.

1.1 The current status of the three RIs

1.1.1. GARR-T

The topology of the GARR network at the end of 2023 is reported in Figure 1. It can connect user's sites at a capacity ranging from 100Mbps to 400Gbps.




Figure 1- The GARR-X/T Network RI

In Figure 1 the links from continental Italy to Sardinia are provided by leased circuits. The main results of the WP2 activity can be summarized as:

- Extension of the fiber optic backbone to reach a total capacity of over 36 Tbps
- In Sardinia, the realization of a dark fibre ring of about 1,100 km, connected to the mainland via fiber spectrum, with 4 submarine circuits of 400 Gbps each
- Support for the Einstein Telescope: the new network in Sardinia is the technological pillar that supports the Italian candidacy for the large gravitational wave observatory
- in Sicily and the South of Italy: strengthening connectivity in the areas of Southern Italy for research without borders

The latest map of the GARR-T fiber infrastructure is shown Figure 2. TeRABIT upgraded the GARR-T network in Sicily and the connectivity of the Sardinia Island using resilient spectrum in the two fibers (orange lines) and additional fibers inside the island.




Figure 2- The GARR-T Network evolution

The positive synergy with ICSC allowed to upgrade other parts of Italy to the technological architecture of GARR-T. Figure 3 details the new topology of GARR-T in Sardinia and shows the additional fiber that connects the Sos Enattos mine site, where the Einstein Telescope Research Infrastructure is proposed to be hosted to detect gravitational waves. The latter connectivity has been funded separately by the region.



Figure 3 - The GARR-T Network evolution in Sardinia

Figure 4 shows the evolution of the backbone Capacity (green line) compared to total end-user granted access capacity (black line). The PNRR projects allowed the GARR-T network to scale adequately to be prepare to support the expected growth of users' requirements.




Figure 4: GARR-T backbone capacity growth versus users' requirements (black)

The GARR-T architectural model is confirmed and it is based on in-house management of the whole infrastructure, from fibre to services. A high-level view of its components is depicted in Figure 5.





Figure 5- The GARR-T Network evolution architecture

The GARR-T network, as of today, has an international connectivity of 800 Gbps with GÉANT, the European Research and education backbone and features a 1.6 Tbps link between the INFN data centre in Bologna and CERN in Geneva.

1.1.2. PRACE-Italy

The PRACE-Italy infrastructure is hosted in Casalecchio di Reno (Bologna, Italy), at CINECA and consists of the Galileo 100 Tier1 HPC node. Galileo 100 main components are depicted in Figure 6.



Cloud	HPC	Storage
OpenStack Cloud partition	564 computing nodes (CPU/GPU, 0.5 to 2TB SSD)	Ceph Block, Object
7000 vCPUs	28.000 CPUs	22PB

Figure 6- PRACE-Italy Tier1 node Galileo100

Galileo 100 is offering both cloud and HPC services and it is connected to the GARR network. The service offering includes high-end GPUs (Nvidia V100) for big-data and machine learning applications, which were utilized in 2022 for about 92% of their total capacity. The TeRABIT project had the objective of a significative growth of the capabilities of the node. At the same time, the PNRR project MEET (Monitoring Earth's Evolution and Tectonics) of the Italian National institute of Vulcanology and Geophysics had very similar objectives. Funding from the projects and CINECA were united to optimize the tender result. The opportunity gave also the possibility to review the final architecture and upgrade strategy.

it was deemed more effective and closer to users' requirement development to create a new node, separate from Galileo100, named GAIA.

GAIA HPC Cloud is an infrastructure combining cloud flexibility with High Performance Computing power. Hosted at CINECA Dama Technopole in Bologna, it is the largest HPC cloud for research in Italy. It supports data processing, analysis and data management services, sensitive workloads, and Kubernetes-based automation. GAIA proves particularly suitable for handling AI workloads, such as fine-tuning and inference, applicable across multiple scientific domains.

Based on OpenStack virtualization tool, it includes 500 Dell servers with over 130,000 physical cores and 256 NVIDIA GPUs. The storage includes 10 PB raw Ceph and 50 PB full-flash VAST system, suitable for HPC-and-cloud synergy.

GAIA operates under an Infrastructure as a Service (IaaS) model. CINECA manages hardware and virtualization, while users control their environments. The system empowers researchers with scalable, customizable, and high-performance resources.




Figure 7: GAIA HPC Cloud total resources

The GAIA node contributes to the CINECA Multiprotocol datalake and it is fully integrated in the CINECA HPC ecosystem, favoring the convergence between traditional HPC simulation and cloud-based use e.g., based on GPUs.

1.1.3. HPC-BD-AI

The HPC-BD-AI infrastructure continues to consist of a distributed set of sites offering cloud services for computation, including storage. The infrastructure leverages the INFN DataCloud which offers a set of Cloud-native, open services and federated resources, exploitable with dynamic solutions covering the IaaS, PaaS and SaaS levels.

INFN DataCloud has the functionality to distribute workloads across multiple locations, with the possibility to federate with both public and private clouds using *de facto* standard protocols.

The HPC-BD-AI infrastructure has been upgraded with both ICSC and TeRABIT funds. The final topology is shown in Figure 8, together with the resources of other partners. HPC Bubbles provided by TeRABIT are installed according to what shown in Table 1.



Figure 8 - HPC-BD-AI sites, connected by the GARR network

Table 1 - Distribution of the HPC Bubbles on the HPC-BD-AI infrastructure

	CPU Nodes	GPU Nodes	FPGA Xilinx Nodes	Storage Nodes
Bari	0	4	0	18
CNAF	16	21	2	36
Milan Bicocca	0	0	2	0
Naples	10	0	2	0
Padua	6	6	0	0
Pisa	8	0	0	0
Rome 1	8	0	0	0
Turin	6	6	0	0
TeRABIT Total	54	37	6	54

The concept of HPC Bubbles has been adopted also by ICSC and the DARE project (Digital lifelong pRevEntion) to provide more HPC bubbles resources to users. The total number of offered resources is listed in Table 2.

	CPU Node	GPU Node	FPGA Xilinx Node	FPGA Terasic Node	Storage Node
BA	24	6	0	0	32
CNAF	26	30	2	2	52
MIB	0	0	2	2	0
NA	18	1	2	0	8
PD	6	6	0	0	0
PI	20	0	0	0	0
RM1	12	0	0	0	0
TO	14	6	0	0	0
LNGS	0	6	0	0	12
CT	12	0	0	0	8
LNF	12	0	0	0	0
LNFESA	8	6	0	0	6
LNL	4	0	0	0	0
MI	4	0	0	0	0
TOTALE	160	61	6	4	118

Table 2: total HPC bubbles nodes and interfaces (Terabit, DARE and ICSC)

2. CURRENT INTEROPERABILITY AND RI FEDERATION EFFORT

The three infrastructures can currently interoperate and exchange information through the GARR-T, as the default connectivity provider. GARR-T provides the RIs also connectivity with all research education sites in Italy, the worldwide NRENs using GÉANT and general internet.

The federation effort focused on the UrgentSHAKE pilot use case, focused on the creation of a service to perform an urgent analysis of damages caused by an earthquake. The development of the use case is done through a joint discussion and agreed decisions with experts from ICSC Spoke0.

To create the seamless secure landscape for HPC services the components considered are:

- High-capacity connectivity. This is offered by GARR to the project partners.
- Federated authentication systems for users. The authentication systems are implemented using different architectures and technologies in the three RIs. The eduGAIN and IAM federation are examples of how a user identity may cross the borders between organisations, domain and services, leading to a federated identity environment. For the use case, the interaction between the three RIs has been only partially automated.
- Access control systems to resources, machine actionable. The effort links to the authorization part of the identity. this component has also been partially automated.
- Monitoring system: the system has been developed by GARR and reported in D2.4.1.2. The system is highly adaptable to other RIs and systems.
- Accounting system. In a federated environment account is paramount to keep track not only of the use of resources by a specific identity, but also to account all the resources used by groups of identities, as an example.
- Cybersecurity agreements and controls.

The TeRABIT project team engineered a specific workflow system for UrgentSHAKE, with a single user interface that provides a common entry point to all the agreed federated HPC services of the RIs.

The service requires a communication interface between the RIs, best based on the cloud IT services provided by all RIs. The complete system is shown in Figure 9.




Figure 9: UrgentSHAKE system Overview

The documents D3.2.6, D4.8 describe more details of the first proof of concept of the federation of the resources provided by GARR, CINECA and INFN, showing the possibility to interoperate the resources of the three infrastructures.

This use case demonstrates that all the technical components are available and can be used to create users' services on top of an agile RIs federation.

The AAI components are under constant developments and there is still an evolving to reconcile the security needs of infrastructures and research agility, however significant progresses have been achieved for greater automation and control.

3. REFERENCES

CINECA	https://www.cineca.it/en
DARE Project	https://www.fondazionedare.it/en/risultati-di-progetto-ad-accesso-libero/
GAIA Cloud	https://www.hpc.cineca.it/systems/hardware/gaia/
GARR	http://www.garr.it/en/
GÉANT	http://www.geant.org
INFN	https://home.infn.it/en/
INFN-Cloud	https://www.cloud.infn.it/
ICSC	https://www.supercomputing-icsc.it/en/icsc-home/
MEET project	https://meet.ingv.it/
OGS	https://www.ogs.it/en