The INFN scientific computing infrastructure: present status and future evolution

Tommaso Boccali, Gianpaolo Carlino, Luca dell'Agnello on behalf of C3S

INFN and scientific computing

The **Italian Institute for Nuclear Physics** (INFN) is the Italian research agency dedicated to the study of the fundamental constituents of matter and the laws that govern them. It conducts theoretical and experimental research in the fields of **sub-nuclear, nuclear and astroparticle physics**. Since its foundation in 1951, the activities it fosters require larger and larger computing and storage resources, due to the increasing complexity of the experiments.

INFN has also been the seminal institution for the research network in Italy, now handled by GARR.

INFN has several facilities: most of them host, at least, one computing center, the main being the Tier-1 at CNAF and the 10 Tier-2s in other sites (see fig. on the right).

INFN recently organized a survey of its computing centers, varying in size from the WLCG Tier-1 to small local facilities.
Questions included **current resource deployment, infrastructure capabilities, effort in support.**

29 answers were collected, for sites ranging from few tens to more than 15000 hosted CPU cores.

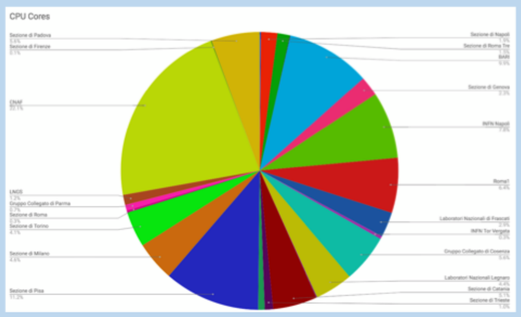

Cores Distribution

A deep analysis on the collected data is ongoing, but some features already stand out with clarity.

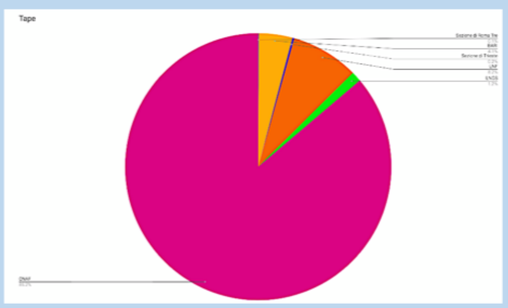

Tape Distribution

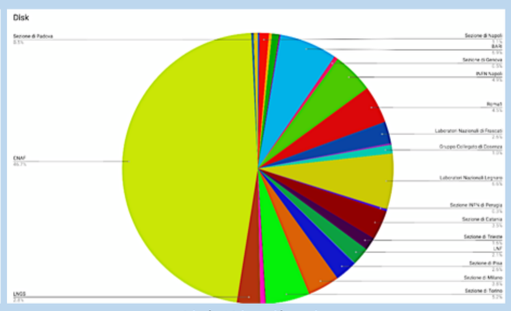

Disk Distribution

Cost, expandability and optimizations

We divided sites into **Big/Medium/Small** centers. Eventually, the large ones correspond to official WLCG Tier-1/2 sites.

At infrastructure level, our sites in 2017 used a total of **~3.4 MW** of electricity.

Resource Type	Unit of Measurement	Number (2017)
CPU	Cores	70000
Disk	TB	57000
Tape	TB	97500

Total INFN resources

The partitioning into Big/Medium/Small sites helps understanding the different cost patterns.
The cost per unit resource is evidently smaller for Big Sites, which suggests economies of scale are important.

Size	Site	Size	Site	Size	Site
Big	CNAF	Medium	Napoli	Small	Milano Bicocca
	Torino		Roma3		LNS
	Pisa		Genova		Pavia
	Milano		Cosenza		Roma2
	Legnaro		Trieste		Perugia
	Padova		Frascati - KLOE		Tor Vergata (2 siti)
	Roma1 (Tier-2)		Parma		Ferrara
	LNF		Gran Sasso		Roma1 (non Tier-2)
	Bari				Firenze
	Napoli (Tier-2)				
	Catania				

Overall, taking into consideration only Moore's law and eventual additional space into the data centers, by 2026 resources could increase up to a factor 7 (but other infrastructural constraints could be a limiting factor).

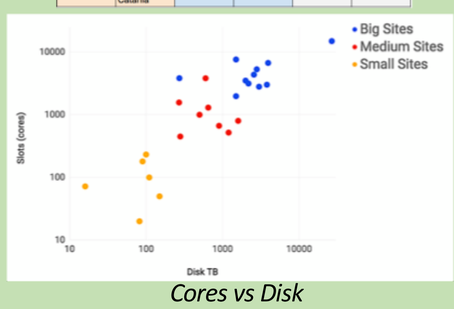

Cores vs Disk

Small / medium sites are often co-located with University departments, with a positive influence of the real INFN costs in both power and manpower costs.

		Big Sites	Medium Sites	Small Sites
Answering sites	29	11	8	10
Total number of racks	504	340	97	67
Number of Servers	4072	3307	509	256
Number of slots (cores)	67993	57218	10121	654
Disk storage (TB)	57332.5	50364.5	6000	968
Tape storage (TB)	97470	87975	9495	0
Site FTEs	92.25	65.3	16.8	10.15
Technical Support FTEs	61.05	47.2	10	3.85
FTE/Slot	0.0014	0.0011	0.0017	0.0155
FTE/Disk_TB	0.0011	0.0009	0.0017	0.004
Available power (kW)	10359.6	8676	1240	443.6
Computing centers space (m^2)	5248.7	4119	702	427.7
IT total space (m^2)	2798	2057	427	314
Estimated expense for power (Eur)	5039250	4452750	306000	280500

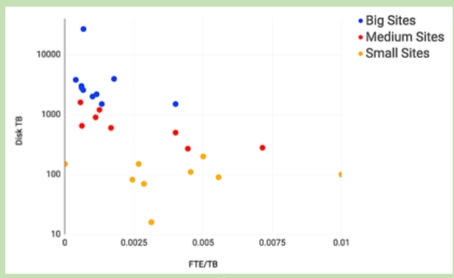

FTE/TB vs TB

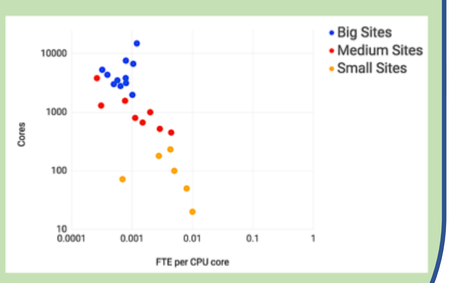

FTE/Cores vs Cores

Future directions & opportunities

INFN is committed to continuing to support its scientific community and cope with its future requirements.
R&D activities towards HL-LHC have started, in multiple directions:

- Evolution of Computing towards **Clouds** (*Indigo-DataCloud, HNSciCloud, EOSCpilot, EOSC-hub*)
- Solutions for **Exabyte level storage**, including caches and optimized access (*XDC*)
- **Use of GPUs** via Cloud Interfaces (*DEEP-Hybrid DataCloud, CNAF-GARR-PISA-BARI distributed Cache, …*)
- **Data lake** (*ESCAPE, XDC*)

In the framework of the project Copernicus, an integration with ASI (Italian Space Agency) and CNR (National Research Council) resources is ongoing. In general, collaborations with other Italian funding agencies are expected in order to evolve to a multi-disciplinary national computing infrastructure.

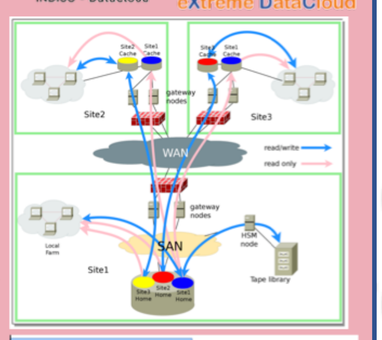

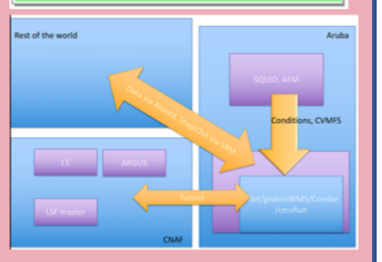

INFN is exploring non conventional, innovative ways to increase the computing resources offered by its Tier-1 center at CNAF (→ **oral "*INFN Tier-1: a distributed site*"**):

- Elastic expansion to other INFN Computing centers → Bari (caching based on GPFS/AFM)
- Elastic expansion to Commercial Cloud Providers using a CNAF-developed tunneling technology
  - → **Aruba**
  - → **Microsoft**
- Extension of the farm to the **nearby CINECA PRACE Center**
  - Using a single dark fiber pairs allowing for up to 1.2 Tb/s
  - No need for caches or special technical tricks

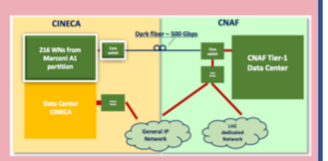

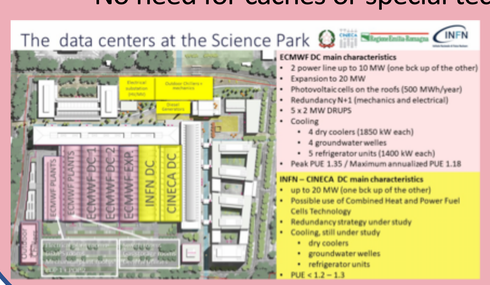

For the highly demanding computing requirements of HL-LHC and also to take into account the data lake approach, we are planning an evolution of the Tier-1 with a **new data center (20 MW for IT) in Bologna**, shared with CINECA, and located close to the premises of ECMWF facility.

Towards the data lake

INFN is actively planning R&D activities and tests in order to plan its future evolution.

In the framework of XDC, ESCAPE and collaborations with other European centers, INFN is performing tests on possible data lake solutions.

Distributed caching: an attempt to have a Xrootd-based caching system using distributed endpoints, with smart placing policies. Currently being set up between the INFN Tier-1, Pisa and Bari.

Testing of new networking equipment: test a private 100 Gbit/s connection between 3-4 INFN sites, to be used only for internal "data lake" transfers.

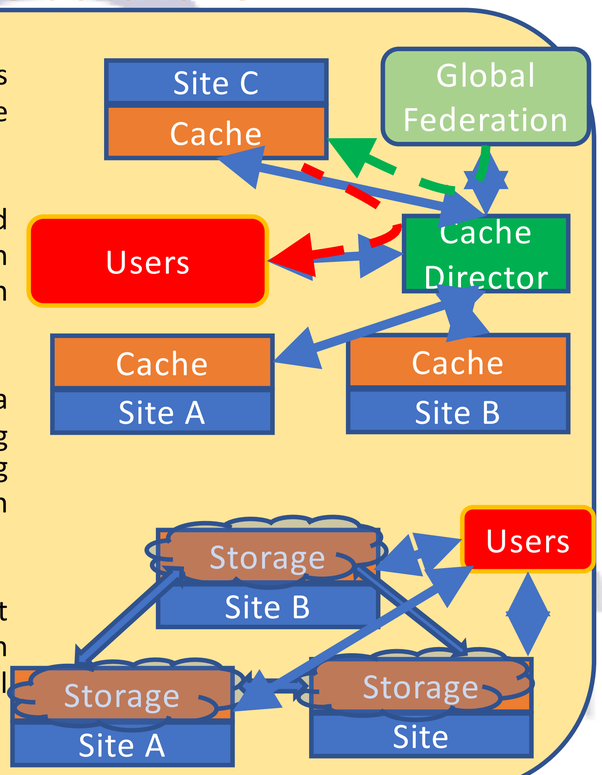

How to organize a data lake in Italy? Still open question, a few possible implementations:

1. The main Data Center in Bologna, other sites serving as Compute Nodes (with caches)
2. A single logical Data Center, physically distributed with co-location of CPUs
3. Compute nodes logically distributed per region

The IBISCo project, recently submitted to the Italian call "Research and Innovation 2014-2020 National Operative Program», aims at enhancing the already existing computing infrastructure in Southern Italy that could form the logical data center envisaged in the 2nd option.