INAF
ISTITUTO NAZIONALE
DI ASTROFISICA

WORKSHOP GARR 2024

NET MAKERS

New Experiments in Astrophysics and Computing Infrastructures for the Italian Community

Ugo Becciani

6 November 2024

INAF

The INAF (National Institute for Astrophysics) consists of 16 research structures, various sections at universities, and its Central Headquarters in Monte Mario, Rome.

Approximately 1,200 research and administrative-management staff form the core of the institution's human resources, which also benefits from several hundred young researchers on fixed-term contracts and university associates.

The INAF community plays a fundamental role both scientifically in many international projects that will shape future developments in the coming years, such as ESO programs, SKA and CTA programs, and many space projects (i.e. Euclid, Gaia etc).

INAF and computing challenges. Main Projects/Objectives

Scientific Themes	Main Projects
Radioastronomy	Square Kilometer Array (**SKA**), Low Frequency Array (LOFAR2.0) e **Meerkat+**
Observational Astrophysics e time-domain	Spatial missions: **Euclid** e **Gaia** Legacy Survey of Space and Time (**Vera Rubin observatory**) Extremely Large Telescope (**ELT**)
High-Energy	**CTA, ASTRI,** FERMI, Dampe, HERD, AMS02, SWGO, etc
Large Scale Simulation	**HPC Theory** (P-GADGET3 -> OpenGADGET, PLUTO, plasma physics simulations, etc..)
Big Data	HPC computing processing, **Management and distribution of large dataset in the Datalake**, High rate analysis

The SKA Regional Centres network

SKAO: two **sites**...

SKA1-MID, Karoo, South Africa:

133 SKA1 + 64 MeerKAT dishes. Max baseline ~150km

Bands: **2** (0.95–1.76 GHz), **5** (4.6–14(24) GHz), **1** (0.35–1.1 GHz)

[© R. Braun]

SKA1-LOW, Murchison, Australia:

130,000 dipoles (512 stations x 256 antennas); 50–350 MHz

~80km baselines; large areal concentration in core

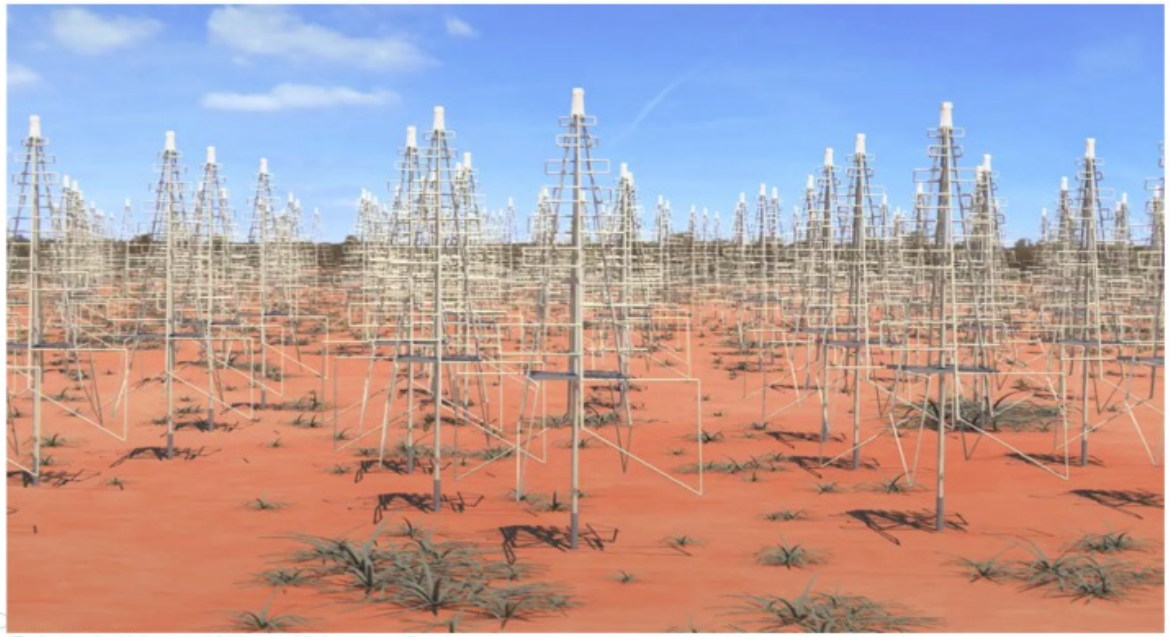

[© R. Braun]

... two kinds of **antennae**

SKAO: some numbers of an unprecedented observatory

Element	SKA1	SKA2
Dishes, feeds, receivers	~200	~2500
Aperture arrays	~130,000	~1,000,000
Signal transport	~1 Pb/s	~10 Pb/s
Signal processing	~exa-MACs	~exa-MACs
High performance computing	~100s peta-flops	~exa-flops
Data storage	Exa-byte capacity	Exa-byte
Power requirements	~10MW	~50MW

[© R. Braun]

Not only antennae…

Computing and data archiving are the key ingredients for extracting the best of the science from antennae, network and receivers

[© SKA organization 2021]

Thus, computing and data archiving are the real «limiting factors» for the capabilities of the new Observatory: SKAO will become more and more efficient with the improvement of those

The SKAO data flow: from the antennae to the Regional Centers

CSP: Central Signal Processor

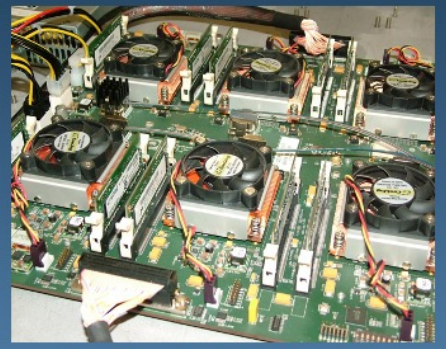

e.g. FPGAs in the ASKAP correlator

5 Tb/s
data buffer of 2 weeks

SDP: Science Data Processor

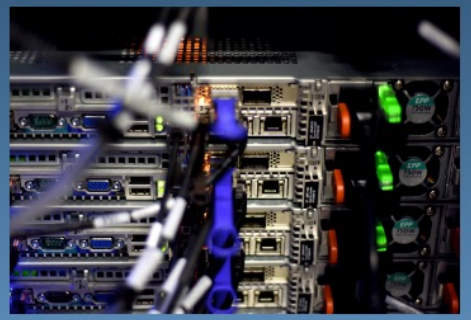

e.g. SDP prototype, Cambridge

5 + 9 Tb/s
data
buffer of
2 minutes

600 PB/yr
data persistence

SRC: SKA Regional Centre network

SRC
SRC
SRC
SRC
SRC

SKA portal

Distributed facilities

USERS

Adapted from Philippa Hartley (SKAO)

it | SRC
SKAO Regional Centre Italy

Some early estimates

Preliminary Plan

		SRC Net v0.1	SRC Net v0.2	SRC Net v0.3	SRC Net v1.0b	SRC Net v1.0
		Jan 2025	January 2026	Sep 2026	Nov 2027	Jun 2028
Deployment (%)		2.00	10.00	15.00	50.00	100.00
Country	Share (%)	Computing (PFLOPS)	Computing (PFLOPS)	Computing (PFLOPS)	Computing (PFLOPS)	Computing (PFLOPS)
Italy	6	**0.04**	**0.21**	**0.32**	**1.05**	**2.10**
Total	100	0.70	3.50	5.25	17.50	35.00

		SRC Net v0.1	SRC Net v0.2	SRC Net v0.3	SRC Net v1.0b	SRC Net v1.0
		Jan 2025	January 2026	Sep 2026	Nov 2027	Jun 2028
Deployment (%)		2.00	10.00	15.00	50.00	100.00
Country	Share (%)	Storage (PB)	Storage (PB)	Storage (PB)	Storage (PB)	Storage (PB)
Italy	6	**1.27**	**6.36**	**9.54**	**31.80**	**63.60**
Total	100	21.20	106.00	159.00	530.00	1060.00

it | SRC
SKAO Regional Centre Italy

The Italian foreseen contributions

personnel: FTE for International Effort (2022+) e.g. (13 March 2024-01 September 2024) ≈ 2.2 FTE

Anticipated FTE for the National Effort (2025+) ≈ 2.5 FTE

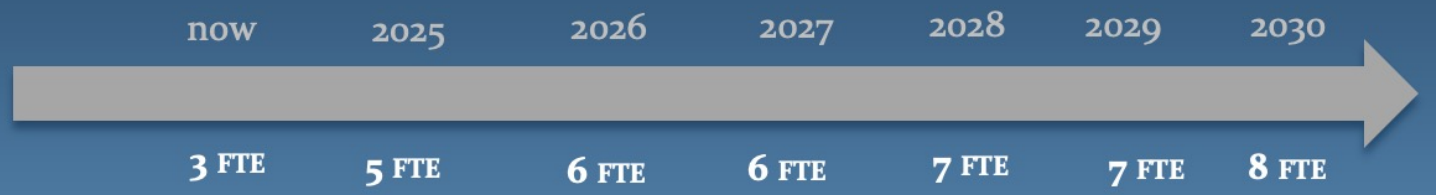

hardware: already funded assets and soon available to SKAnet v. 0.1

Item	Early 2025	End 2025	
CPU power	0.1 PF/s (Tier 3 – dedicated, CPU only)	1.5 PF/s (Tier 3 – dedicated, CPU+GPU)	15 PF/s (Tier 1 – shared, CPU+GPU)
Storage	0.3 PB on-line disk (S3), 1.2 PB Tape	2 PB on-line disk (S3), 5 PB Tape	10 PB Flash (LUSTRE, shared)
Network	10 Gb/s (LAN + WAN)	100 Gb/s (LAN) 100 Gb/s (WAN)	400 Gb/s (LAN) 100 Gb/s (WAN)

LOFAR: the LOw Frequency ARray

10 Countries
24+14+14(16) stations

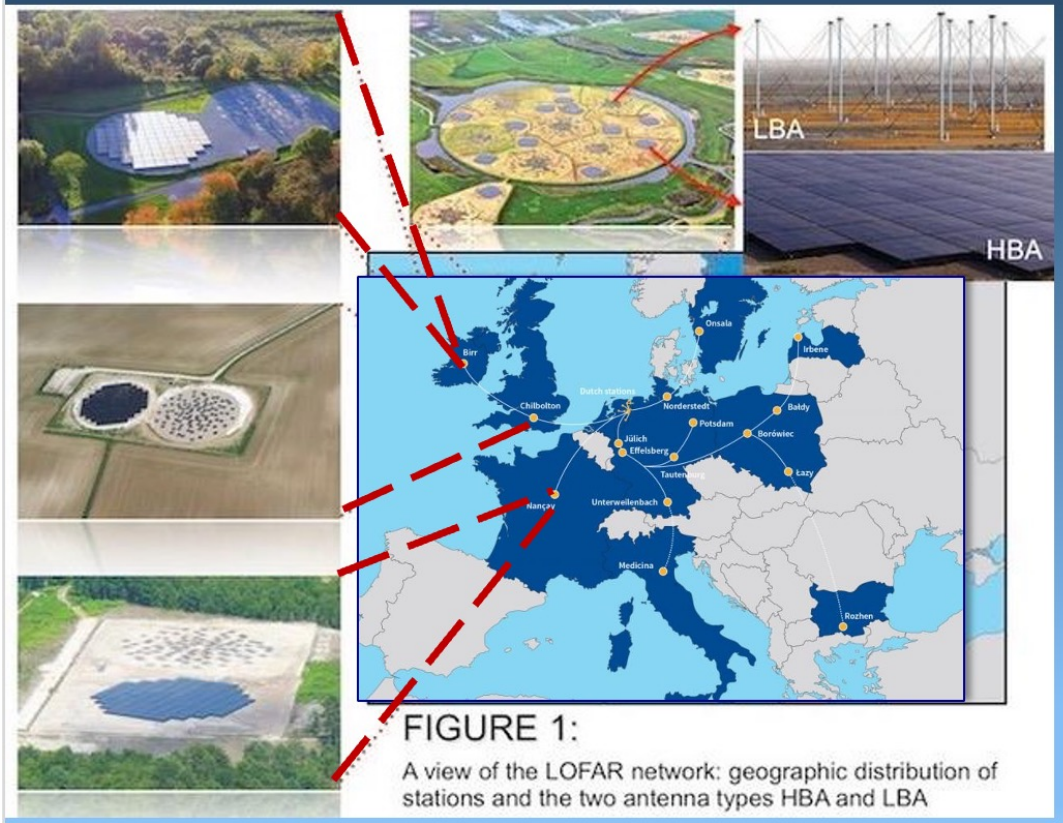

FIGURE 1:
A view of the LOFAR network: geographic distribution of stations and the two antenna types HBA and LBA

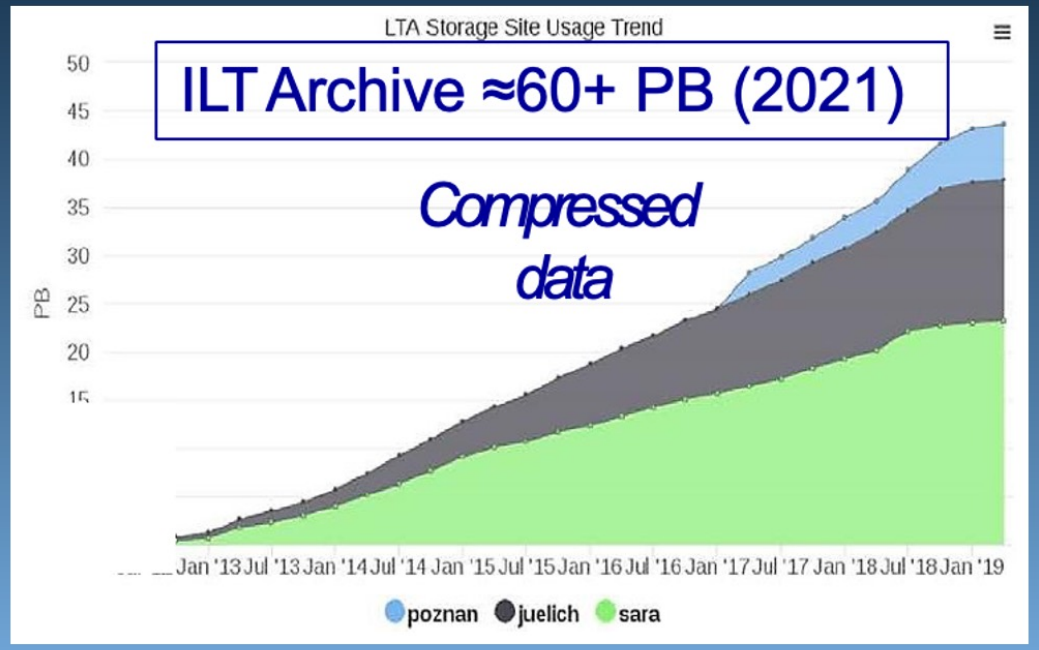

- ✓ 250 Gb/s across the entire network
- ✓ Large FoV, n baselines, n channels, produce typical TB-size datasets
- ✓ Archiving problem and managing Big Data

an Station

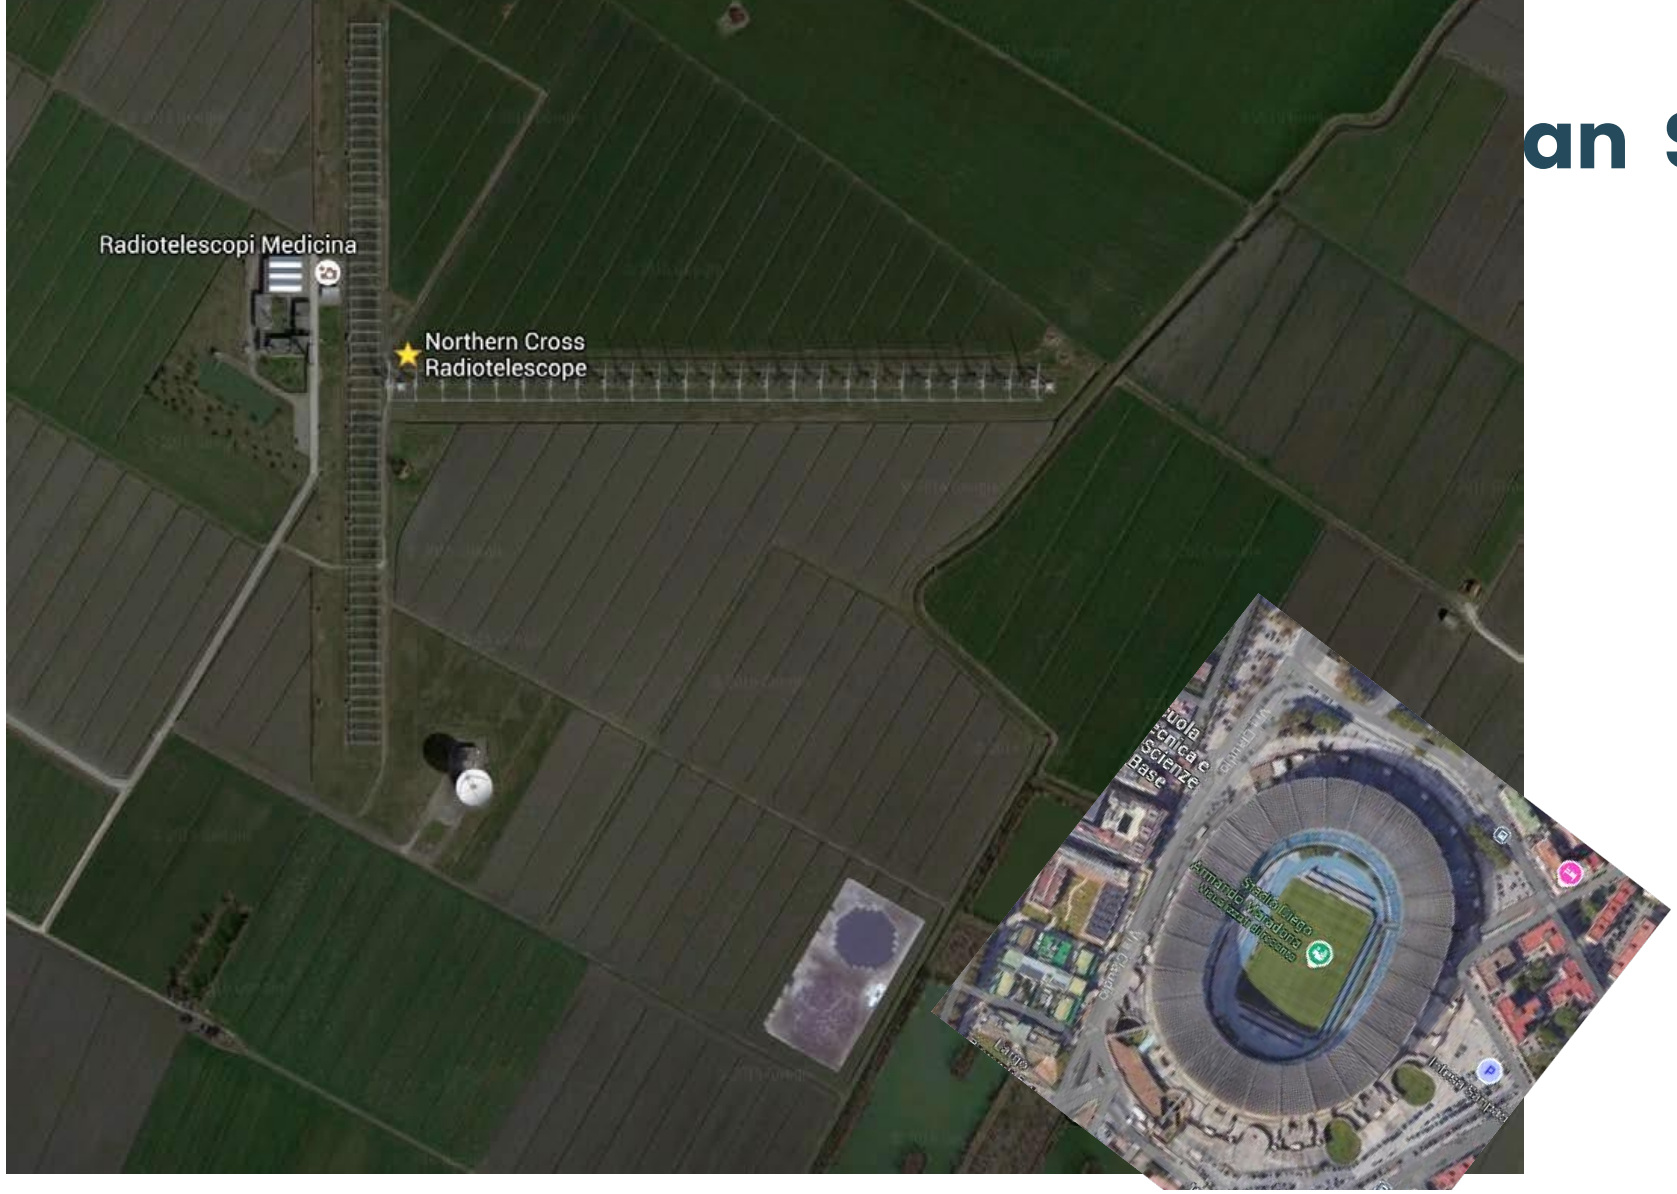

LOFAR 2.0 upgrade (2025+…)

Upgrade :
- Electronics
- Correlator
- Dual beam (es Medicina)

will bring LOFAR into a new regime entering an uncharted territory at the very low frequencies (10-60 MHz)

LOFAR 2 rate ~4 x LOFAR

Large Programs:

12-20 PB/yr

120 PB total

LOFAR ERIC

LOFAR is @ top priorities in PNIR 2021-27

- **PIs/coPIs** of the largest LOFAR 2 surveys: **data flux in IT will increase !!**
- Use of long baselines to avoid confusion limit: **critical for computing**
- Unique synergies with **EUCLID** and **LSST** in 2025-30

- Strong case for BIG data in IT
- Strong case for **proto-SRC**

USC VIII

ICSC
Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

TECNOPOLO MANIFATTURA
DATA VALLEY HUB

CTAO: Cherenkov Telescope Array Observatory

New generation VHE gamma-ray observatory

- Full sky coverage
  - 2 sites, one in each hemisphere (Chile & Spain)
- Wide energy range: 20 GeV - 200 TeV and beyond
  - 3 different telescope sizes (23 - 12 - 4 m diameter)
- 10x improved sensitivity
  - Large number of telescopes over a wide area
- Excellent energy and angular resolution

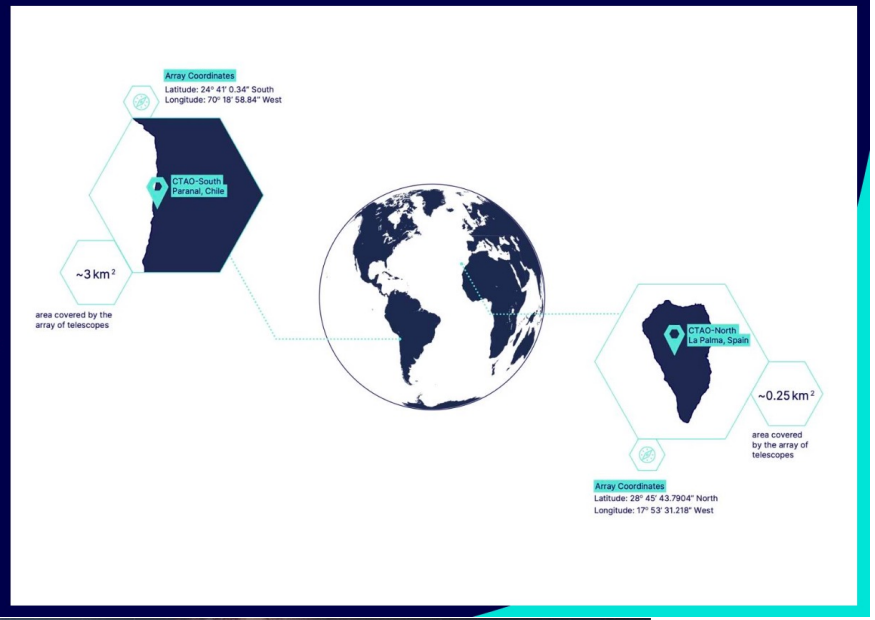

CTAO telescopes

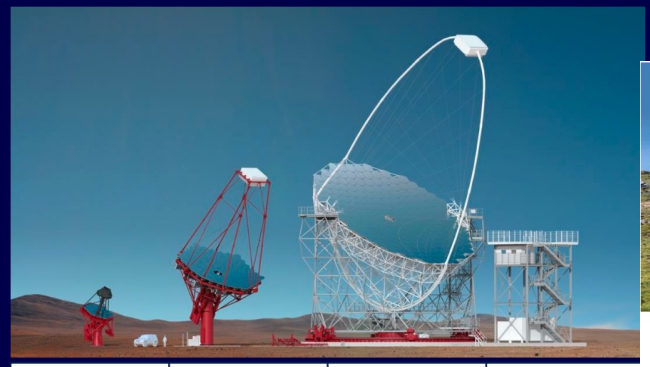

	SST	MST	LST
Mirror ⌀	~4m	~11.5m	~23m
FoV	~9deg	~7.5deg	~4.3deg

CTAO-North, La Palma, Spain

- 4 Large-Sized Telescopes (LSTs)
- 9 Medium–Sized Telescopes (MSTs)

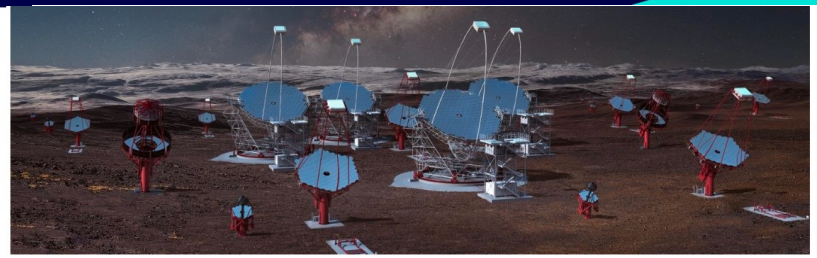

CTAO-South, Chile

- 14 Medium–Sized Telescopes (MSTs)
- 37 Small-Sized Telescopes (SSTs)
- **2 Large-Sized Telescopes (LSTs)**

CTAO raw data rates offsite

2 Gbps network @ 365d/y * 24h/day with 75% efficiency → Max 6PB/y

-> Required Data Volume Reduction ratio: > 10

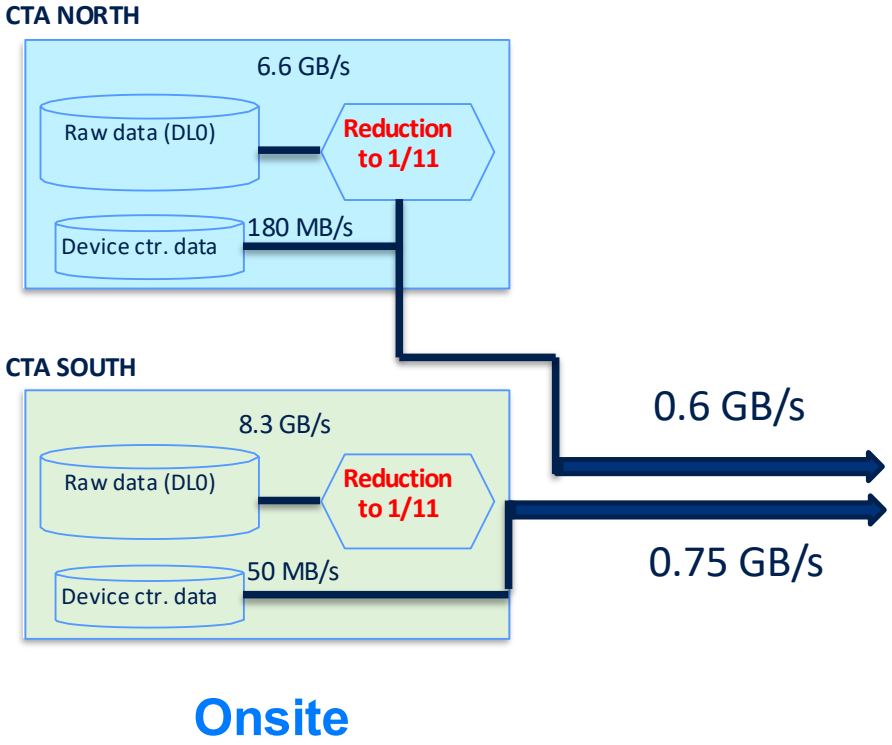

More than 100 PB on disk and 350 PB on tape by 2031

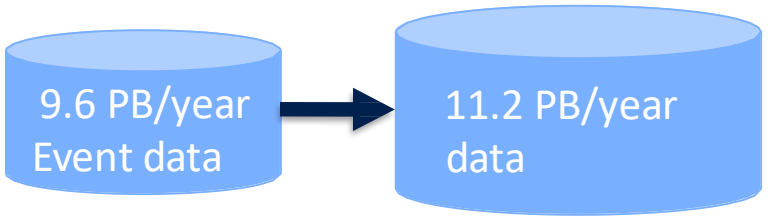

DL0-reduced

The Gaia mission is nearing the end of its operational life: **after 10.5 years of continuous science operations** (more than twice the initial lifetime), from the second half of Jan 2025, data taken with focal plane instruments will no longer be considered part of the scientific data stream (>> reconfiguration of Level-0 -1 and -2 pipelines at primary Data Processing Centers -DPCs- like the **Italian DPCT**).

GAIA MISSION: THE DPCT- A BIG DATA/ HTC FACILITY

INTERNET LINK : from 1 to 10 Gbps via GARR
HOT STORAGE SYSTEM CAPACITY: **2.5+ PB overall raw disk space** distributed between two HP P7400 storage units and **one P8400**.
COMPUTING : **14 servers** HP DL580 G7/G9 with a total of about **600 CPU cores and 4.5TB RAM**.
DEV & TEST: 7 servers HP
DB SERVERS: **3 servers** HP DL580 G7 (**32 cores**, **256MB RAM each**) based on **Oracle RAC technology** (**DBMS Oracle**).
3 LEVELS BACKUP : L1 on primary storage array, L2 on disks (StoreOnce 6600) and L3 on tape libraries (HP ESL G3).
HPC INTERCONNECTION: access to HPC super computer at CINECA for dedicated processing.

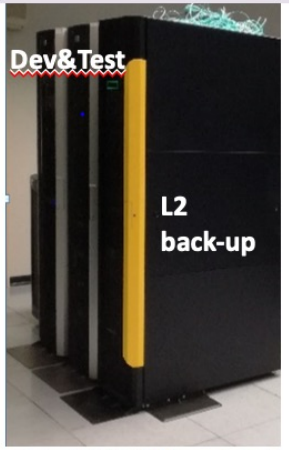

Direct link to HPC systems at CINECA via dedicated MOU

Vera Rubin Telescope

The goal of the Vera Rubin Observatory project is to conduct the 10-year Legacy Survey of Space and Time (LSST). **LSST will deliver a 500 petabyte set of images and data products** that will address some of the most pressing questions about the structure and evolution of the universe and the objects in it. The Rubin Observatory LSST is designed to address four science areas:

- Probing dark energy and dark matter.
- Taking an inventory of the solar system.
- Exploring the transient optical sky.
- Mapping the Milky Way.

Vera Rubin Telescope

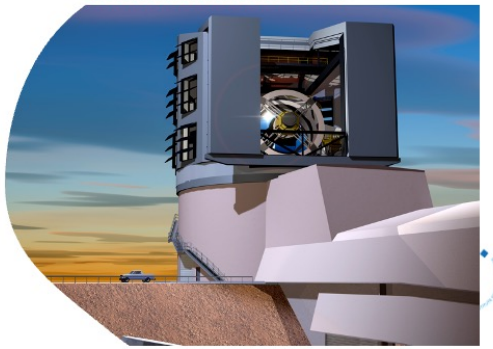

Astrophysical Data HPC Operation Center

Designed, developed and hosted by Dept. of Physics
Regulated by a next MoU between INAF and Dept.

Our ICT @ Università degli Studi di Napoli Federico II

Data/processing Science services:
bio- geo- astro- informatics

Data/processing services for astrophysics:
ELT, SKA, VST, ALMA, Rubin-LSST

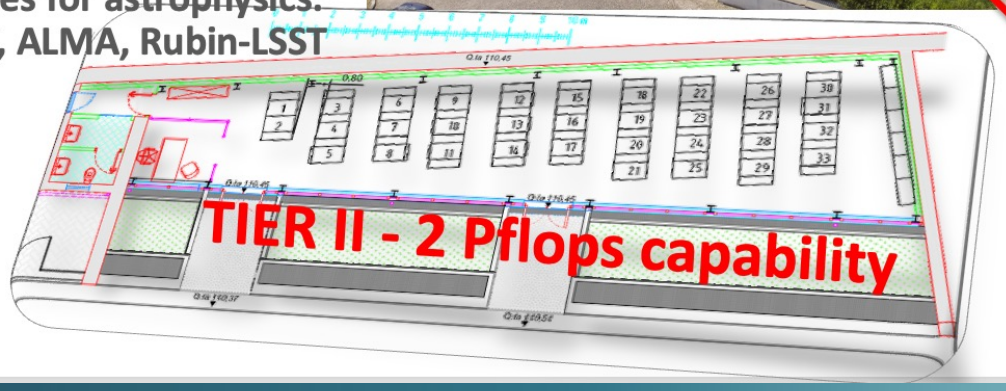

Vera Rubin Telescope

Italian Astrophysical Community HPC Facilities for Vera Rubin and other experiments

20 PB of Data Storage (raw)
52 HPC/HTC multi-core servers
46 GPUs H100 (80 GB/device)
20 GPUs L40 (48 GB/device)

Astrophysical Data HPC Operation Center

6PB @ Bologna 2PB @ Napoli

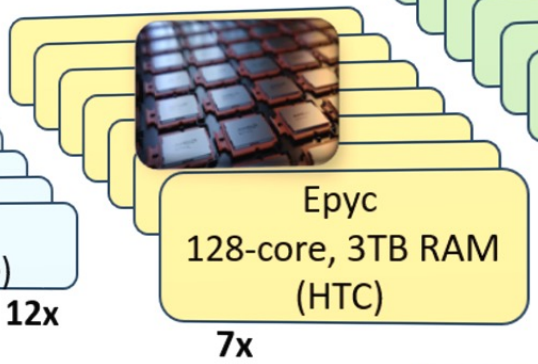

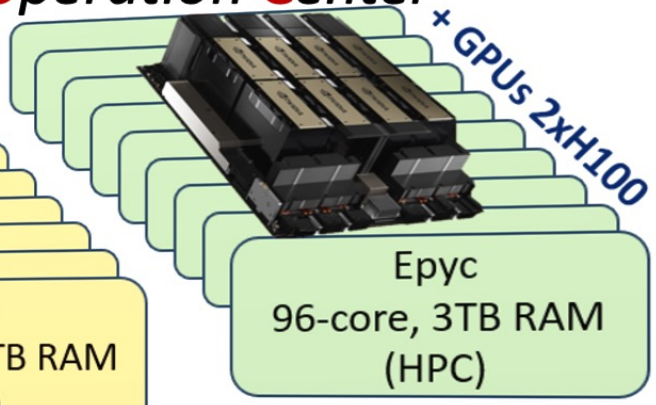

2x NAS - 72 TB

Printer A3 col 3x | Printer A4 b/n 3x | Plotter A1

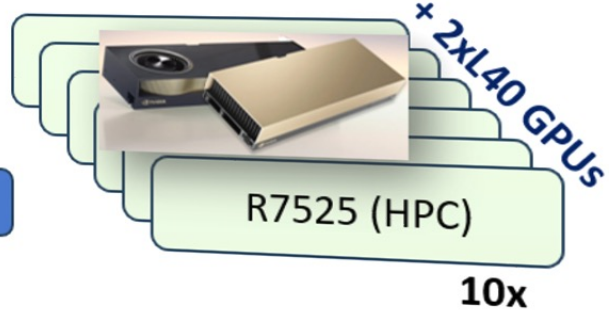

INAF Short-Term Computational Infrastructure

Tier-1 System @ technopole

The New HPC Tier 1@ CN will be a system with Data Centric nodes and GPU nodes. Globally 5+15 PF.

The total investment will be aprox 15 Million of Euros (4 Meuros from INAF).

INAF participate to this project **with dedicated resources (expected aprox 4-6 PF, 10 Pbytes SSD)**

INAF TIER-3 System @ technopole

System for LT storage and associated hardware for use in SKA-precursors, Path-finders, and related development environments.

.

➔ **STORAGE LONG TERM**: Tape Library LTO

➔ **STORAGE medium/high speed - online storage**


LEONARDO
cefla
GRUPPO ICM


LEONARDO


ECMWF Data Center
CINECA & INFN Exascale Supercomputer Center
Conference and Education Center
Conference and Education Center
Innovation Center
Competence Center Industry 4.0
Civic Protection and Agency for Development and Innovation
Meteo National Agency
University Centers
Biobank and Life Science
ENEA Center

INAF computing infrastructure at the CN-HPC

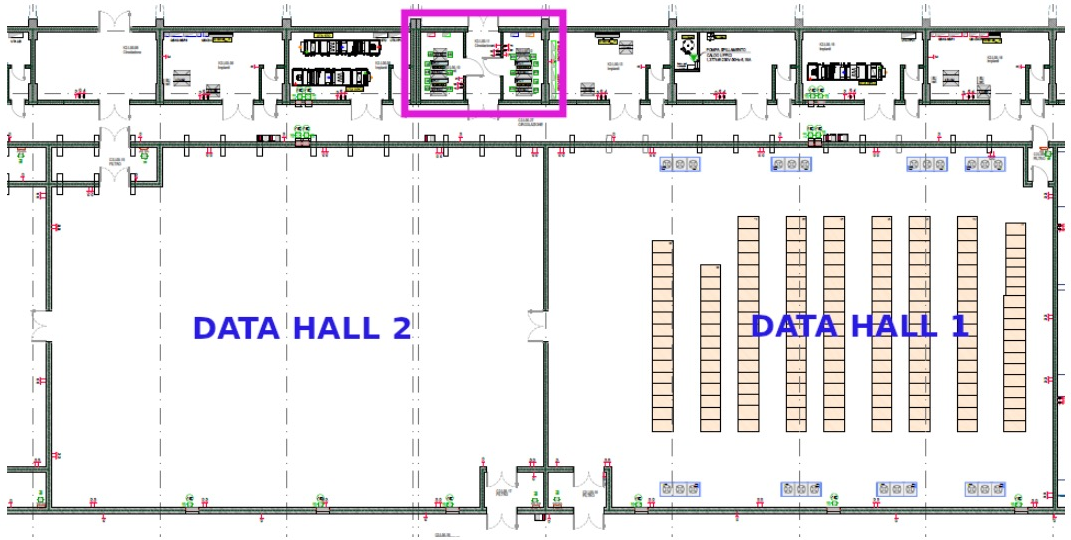

Figure 2. Technopole data centre floor plan. Leonardo layout is depicted in Data Hall 1 (on the right) and the procured System will be hosted in Data Hall 2 (on the left). The network room is indicated by a purple box at the top of the figure.

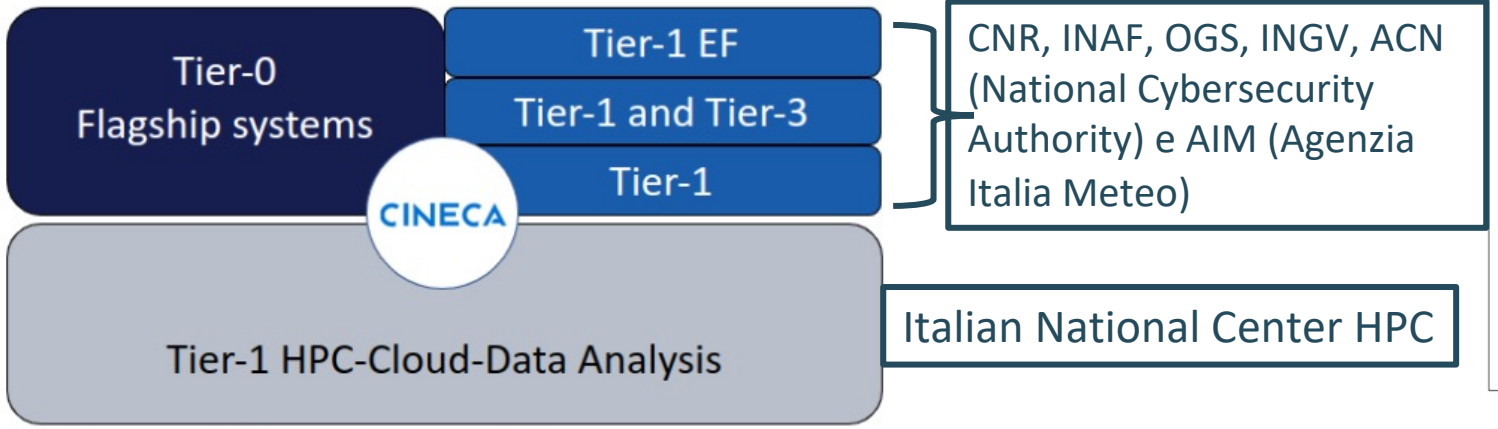

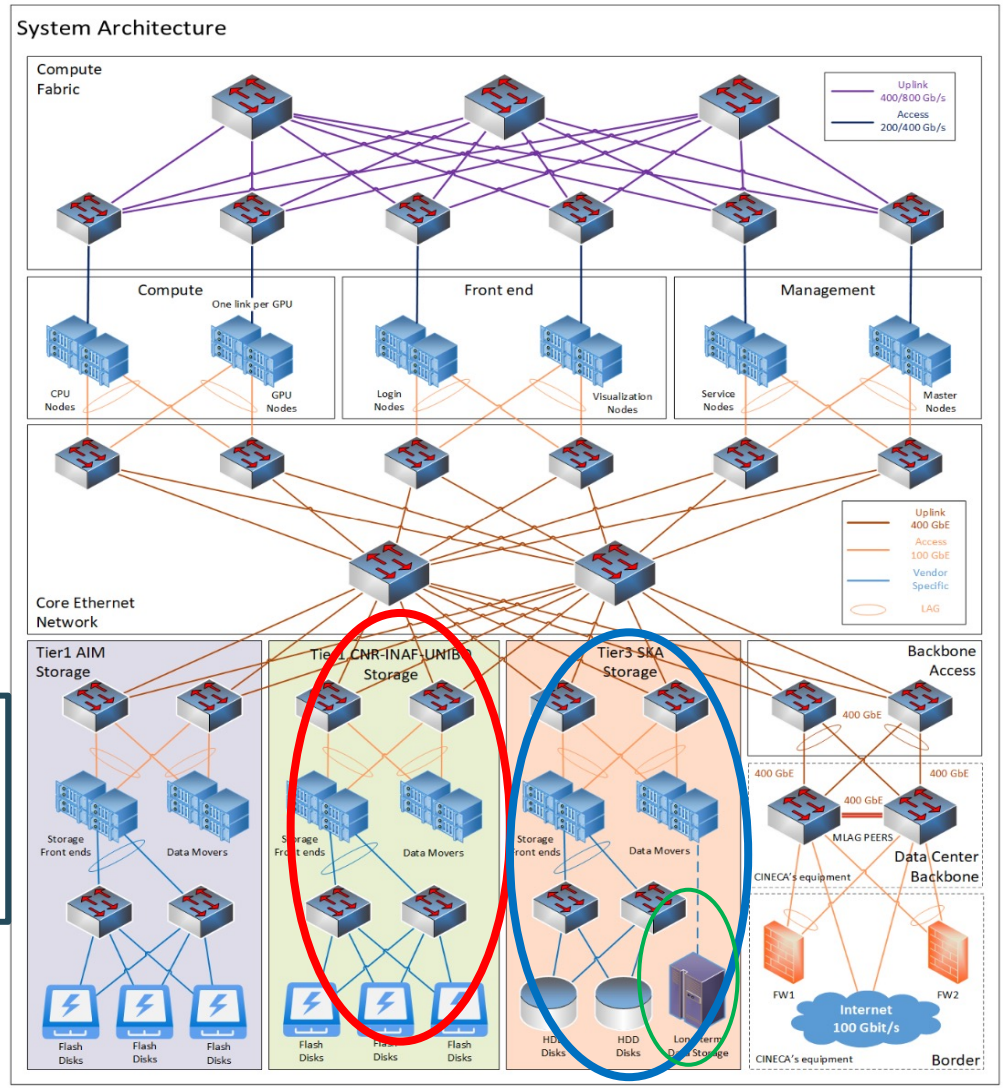

Figure 7: Reference design of the system architecture.

TIER 1 – The Data Centric partition

The CPU partition must provide at least a HPL performance of **4 PFlops**.

The CPU **must be based on x86_64 architecture**: ***at least 56 cores***, a theoretical peak performance (FP64) of at least 3.5 TFlops.

The nodes must be equipped with at least ***3 GBytes per core and not less than 512 GBytes.***

At least 25 compute nodes of the CPU partition must feature 1 TByte (fat node mainly for INAF)

Expected 5-8 nodes for INAF

TIER 1 – The Booster partition

The GPU partition must provide at least a HPL performance of **15 PFlops.**

The GPUs must provide at least a ***theoretical peak performance (FP64) of 60 TFlops.***

The **node** must be equipped with **2 CPUs, + 4 GPUs**.

Node memory at least 512 GBytes of DDR5 memory and it must be greater equal to the sum of all GPU memories installed in the node.

Expected about 15-20 nodes for INAF

INAF e l'Infratsruttura di calcolo attuale e in prospettiva a breve termine

Long Term Data Storage (LTS)

Tape library with a net capacity of **at least 5 Pbytes (*expected aprox 8-10 PB*) upgradable** and scalable to reach **at least 300 PBytes capacity (probably 400 Pbytes).**

On-line Storage: A capacity-oriented HDD-based storage with **3-5 PBytes of net space available** (read throughput of 30 GB/s write throughput of 15 GB/s.)

The LTS must also foresee a **backup technology (self-checking and self-recovery),** must be also become a logical partition of the Tier-3 SKA Online Storage.

Integrated native policy manager capable of managing the data movement between disks and tapes.

Availability of policies for automated data movement and archiving between different layers (e.g., flash, disk, tape, etc.).

NET MAKERS

LA COMUNITÀ CHE INNOVA LA RETE

INAF
ISTITUTO NAZIONALE
DI ASTROFISICA

High-Performance Computing, AI and Big Data analysis. Future challenges in astrophysics

Shaping the Future