Big-Data-Oriented Computing on a Mid-Size Data Center: the ADHOC Infrastructure

Dr. Augusto Tortora, Prof. Massimo Brescia, Dr. Stefania Conte, Dr. Paolino Guida, Prof. Guido Russo, Dr. Bernardino Spisso

Speakers : Dr. Augusto Tortora and Prof. Guido Russo

Department of Physics 'E. Pancini', University of Naples Federico II

E-mail: augusto.tortora@unina.it, guido.russo@unina.it

Outline

- Speaker's Info
- Research Topics
- Introduction
- Objectives
- Infrastructure
  - Implementation
  - Storage
  - Networking
  - Computation
  - Security
  - Power and Cooling
  - The Control Room
- Example of applications
- Conclusions and future perspectives

Dr. Eng. Augusto Tortora

Mechanical Engineer, Ph.D. in Industrial Engineering, specialized in Mechanical Engineering with a doctoral thesis on applications of Artificial Neural Networks to energy and propulsion systems.

Since July 2021, researcher at the Department of Physics "E. Pancini", I have developed skills in IT infrastructures, Power and Cooling, and network plant applied to Data Center.

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
Università degli Studi di Napoli Federico II
INFN Istituto Nazionale di Fisica Nucleare
INAF ISTITUTO NAZIONALE DI ASTROFISICA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Prof. Guido Russo

Physicist, full professor at the Department of Physics "E. Pancini" of University of Naples, Italy, with previous experiences at ESO, ESA, NASA on scientific archives.

Over 300 publications on international journals.

The our Research group

Our research group is a multidisciplinary team that encompasses expertise from physics, computer science, engineering, and humanities. It includes representatives from Universities and Research Institutes (CNR, INFN, INAF).

Our main research topics include:

- Development of HPC/HTC Data Center infrastructures;
- International experiments in nuclear physics and astrophysics;
- Preservation and digitization of cultural heritage.

The Astrophysical Data HPC Operating Center (AD HOC) at the Department of Physics

- 33 Racks
- 37 LCP refrigerating units water-air
- 80 servers, 200 GPUs, ~15.000 cores
- Storage: on-line (10 PB) and long-term (2 PB)

TIER II class (~2 PetaFlops)

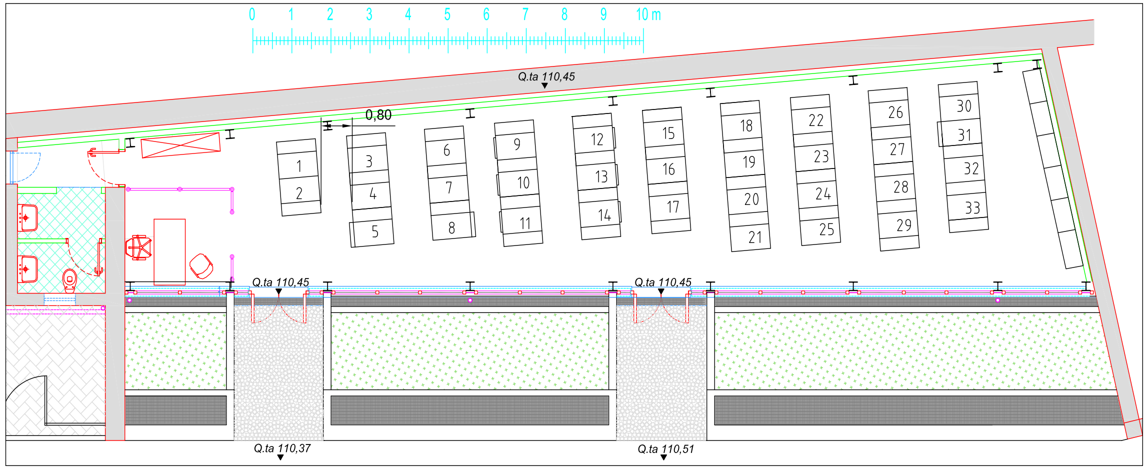

The ADHOC Data Center History

1. Project «SCoPE», start (2006-2009)
2. Project «RECAS», follow-up (2012-14)
3. Project «IBiSCo», follow-up (2016-18)
4. Project «ICSC», follow-up (2022-25)
5. Project «STILES», last (2023-25)

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
Università degli Studi di Napoli Federico II
INFN Istituto Nazionale di Fisica Nucleare
INAF Istituto Nazionale di Astrofisica
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Objectives

- demonstrate that for large unstructured datasets, it is possible to leverage a mid-size Data Center for astrophysical analyses.

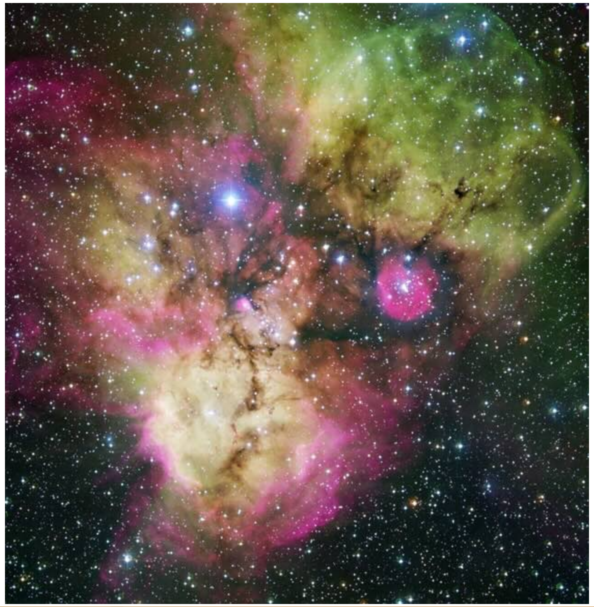

Requirements: Users

- **astrophysical scientists:** Archiving and analysis of images/spectra from ground based instruments (ELT, SKA, Rubin-LSST, VST), space-borne telescopes (ESA Euclid, JWST). Machine/Deep Learning paradigms for detection/classification of sources and cosmology;

- **matter properties physicists:** developing and testing alternative jet fuels by introducing a state-of-the-art mathematical models for accurate estimations of the fuel consumption in jet engines;

- **high energy physicists:** Experiments (ATLAS, BelleII), splitting data in two steps:
  1. sample of "interesting" events to be selected;
  2. Experimental confirmation of theoretical studies through extraction and analysis of several dozen of Terabytes of data;

General Requirements

- Data sets must be sequentially read and processed in a specific order
- Mandatory high speed networking
- Processing parallelization through a massive exploitation of GPUs
- Constraints imposed by limited budget

Implementation

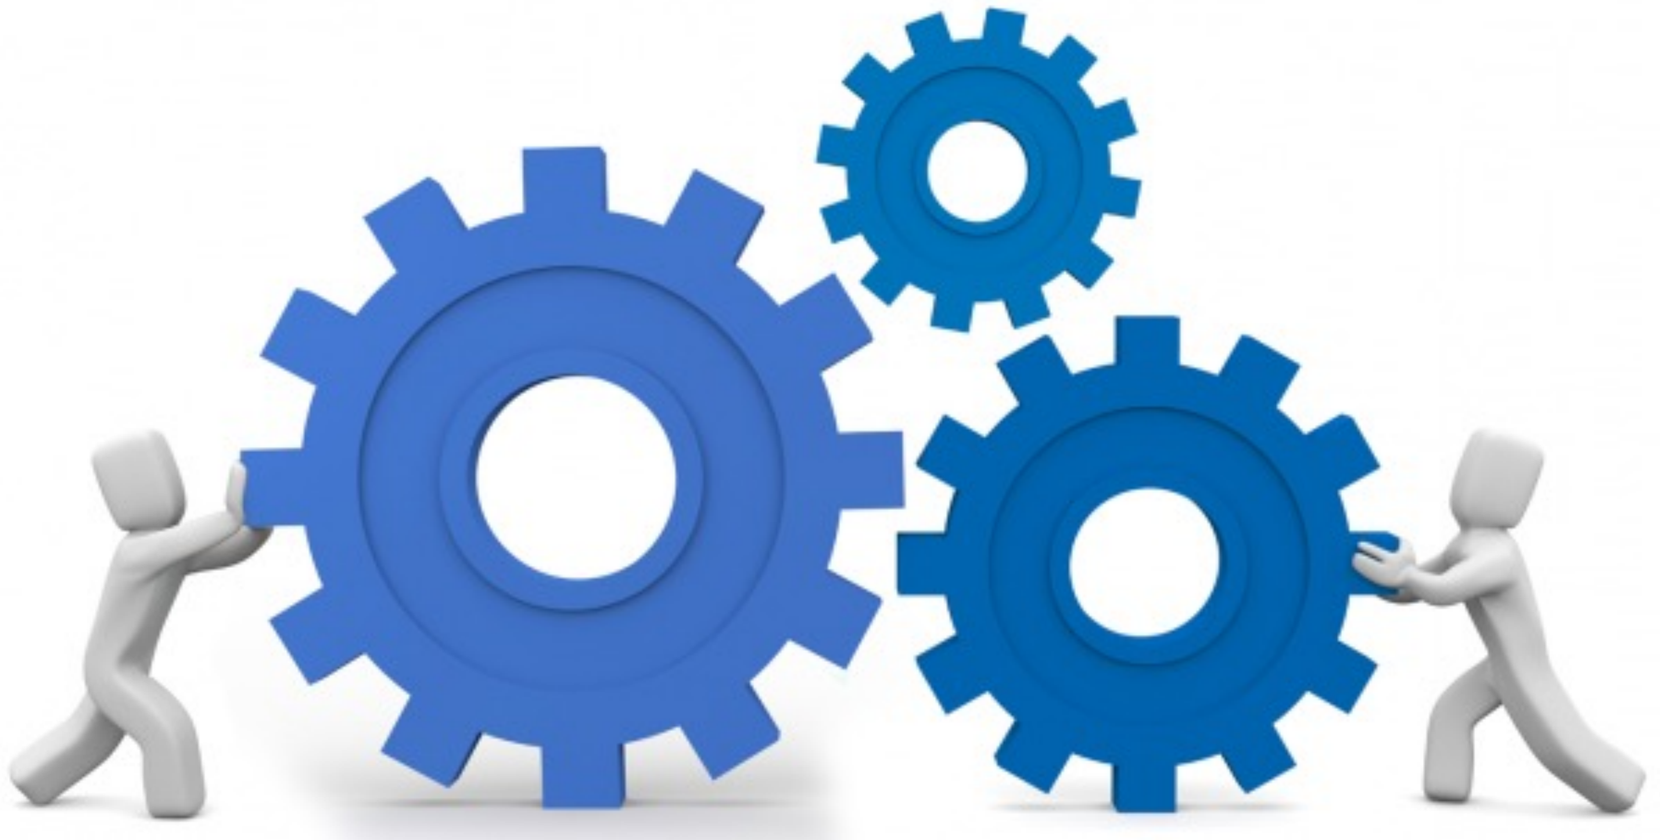

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
Università degli Studi di Napoli Federico II
INFN Istituto Nazionale di Fisica Nucleare
INAF Istituto Nazionale di Astrofisica
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Storage

- 500 SATA disks, 22 TB each
- 100 NVMe disks, 7.5 TB each (raw capacity)
- 12 enclosures, with several of these groups connected to a single FC-AL controller
- Two server every 2.5 PBytes as a front-end
- Each server is accessible via multiple networking options: 2x25 GbE, 2x100 GbE, 2x200 Infiniband

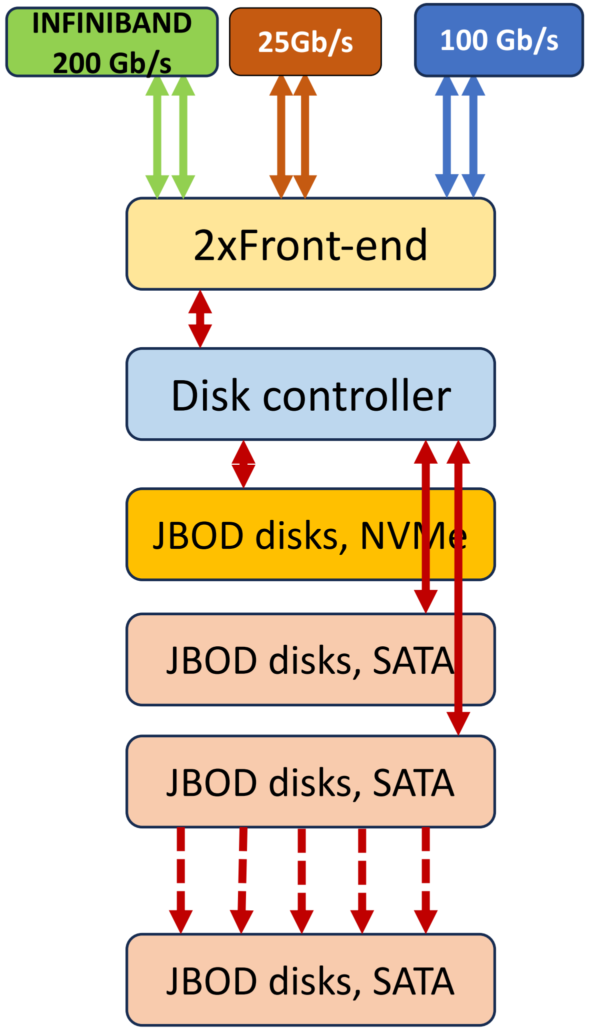

Networking

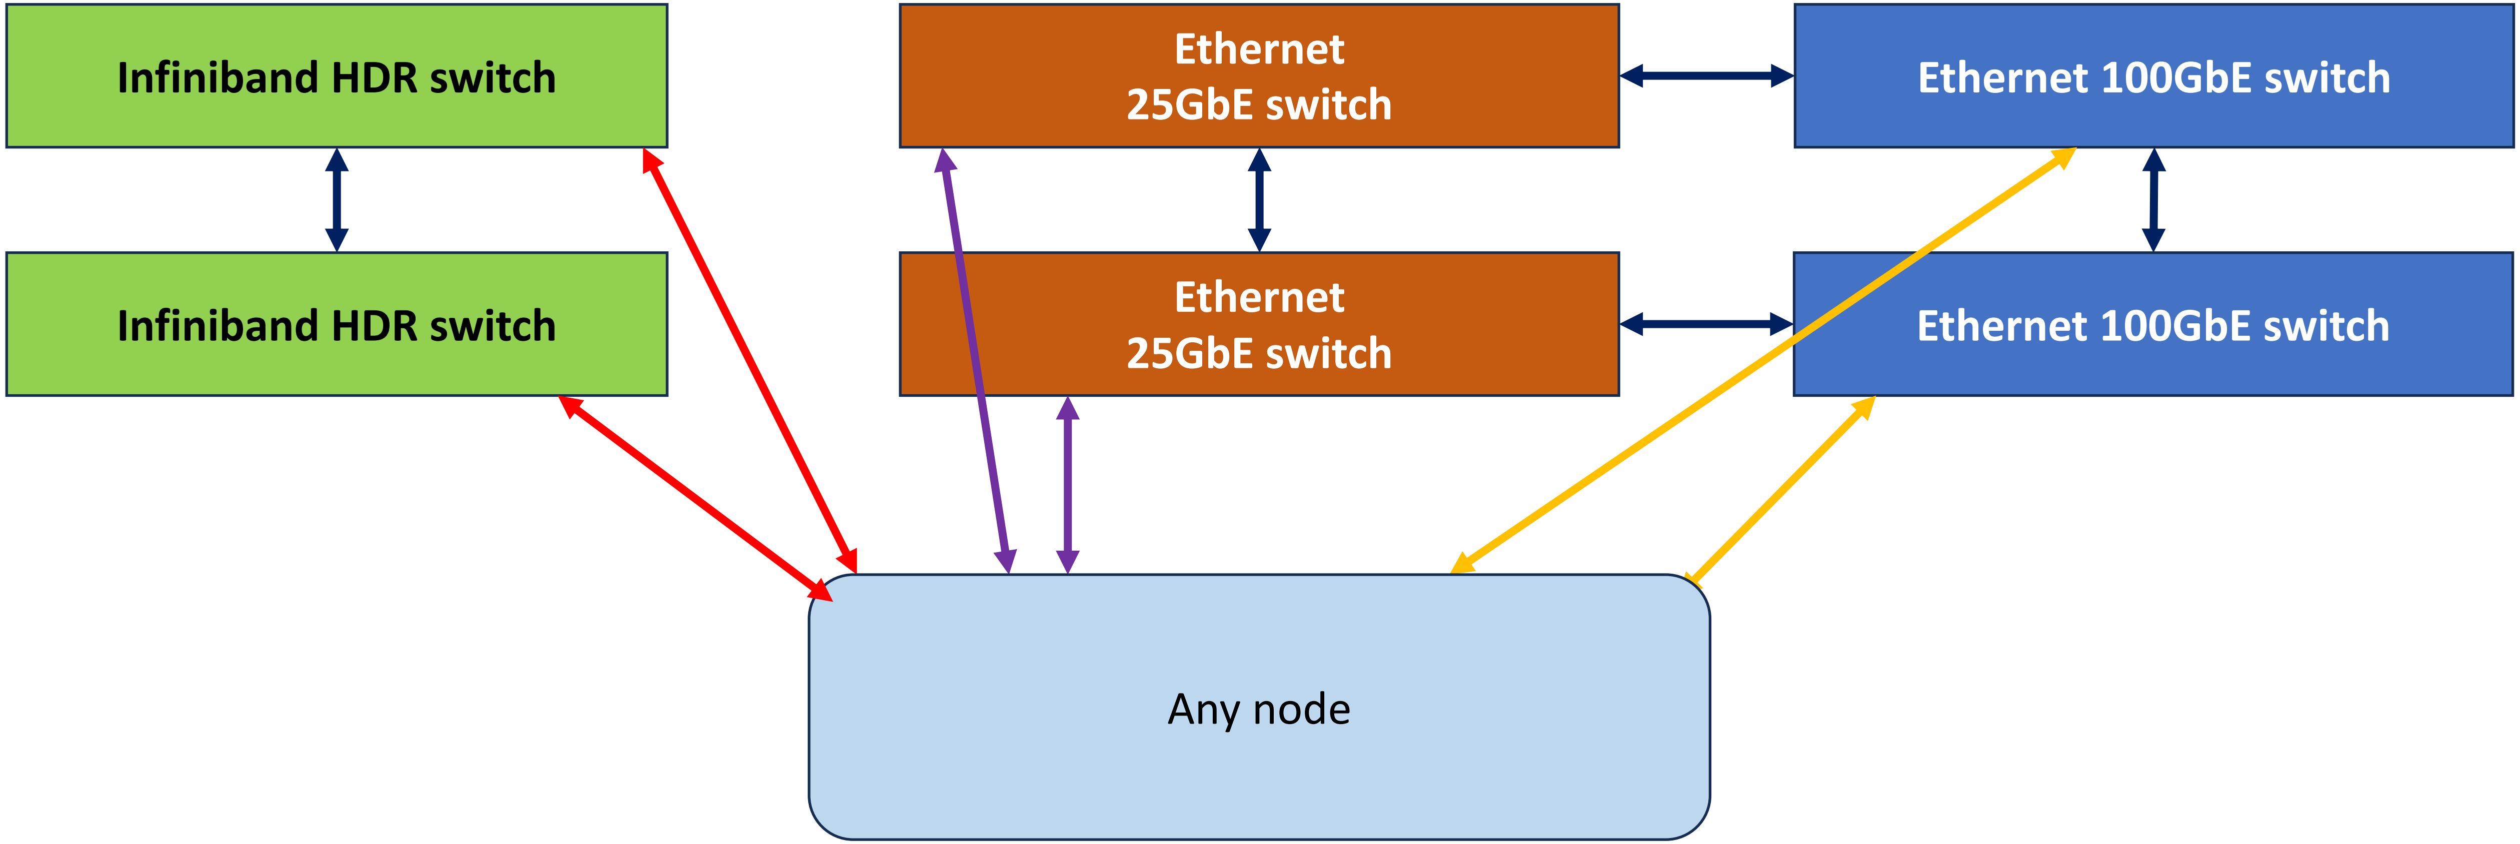
Infiniband HDR switch
Infiniband HDR switch
Ethernet 25GbE switch
Ethernet 25GbE switch
Ethernet 100GbE switch
Ethernet 100GbE switch
Any node

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
Università degli Studi di Napoli Federico II
INFN Istituto Nazionale di Fisica Nucleare
INAF Istituto Nazionale di Astrofisica
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Computing

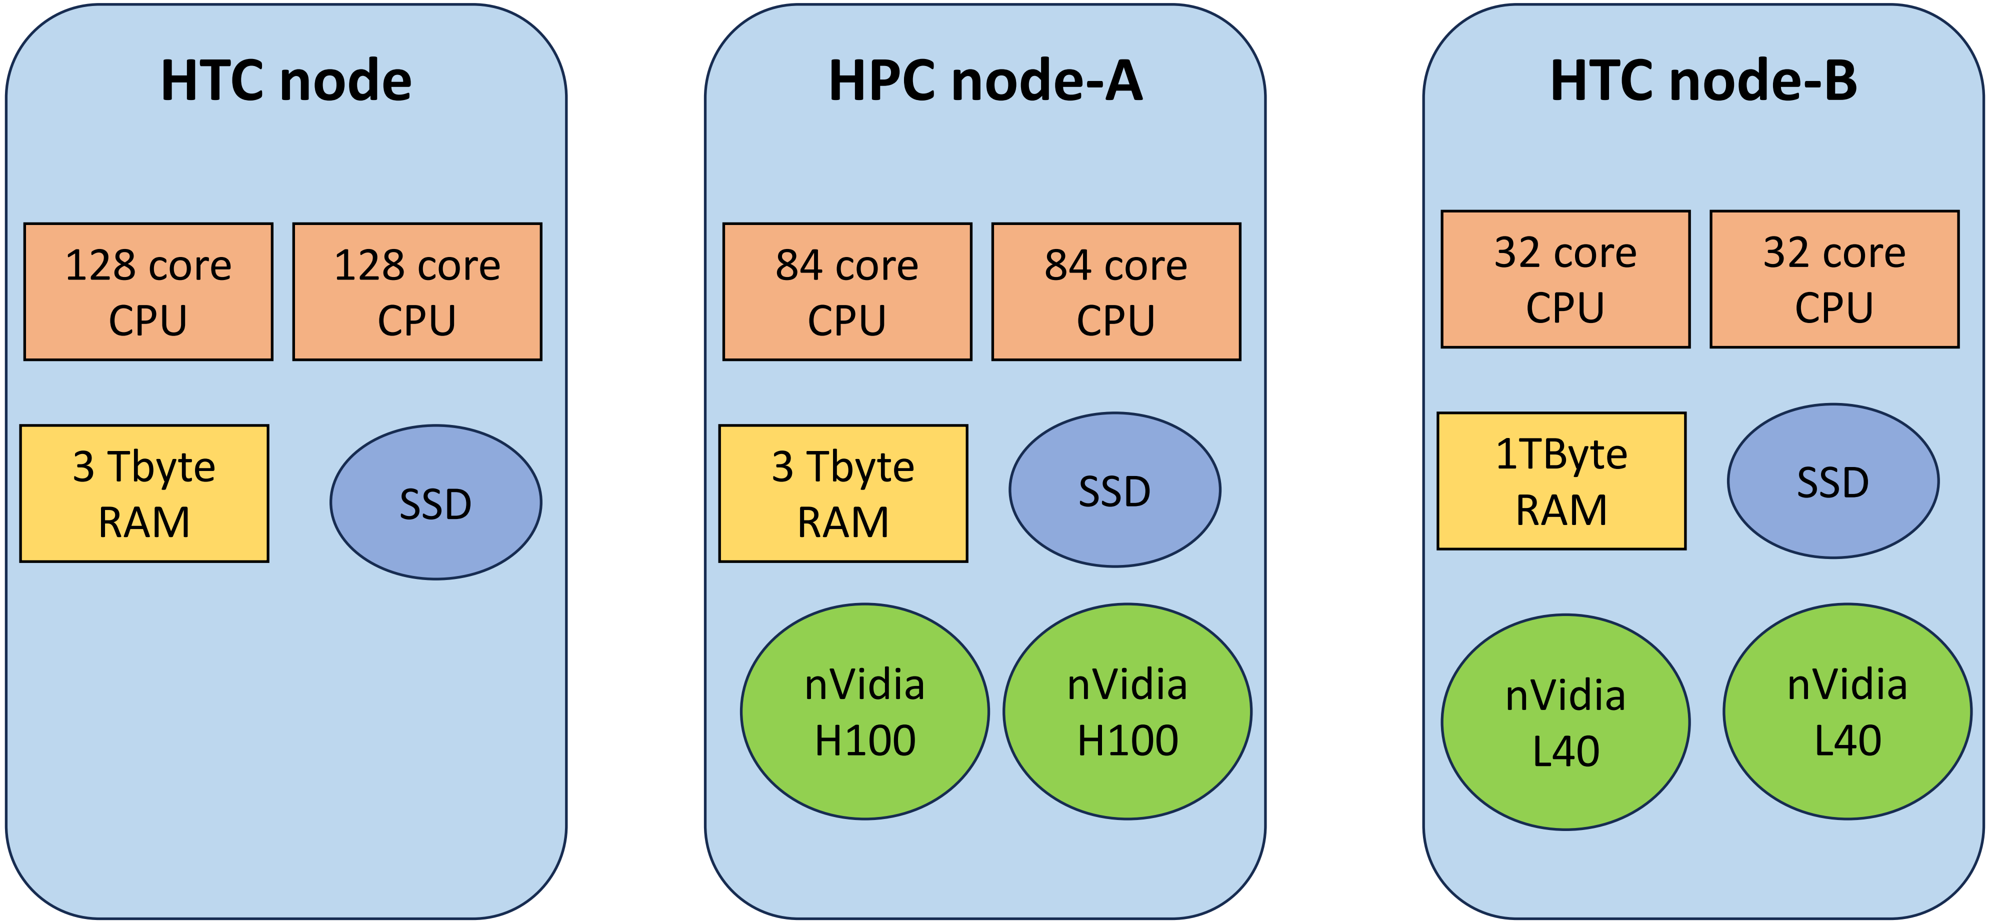

SECURITY

- 2 firewalls, with strict access rules, using public IP addresses only for the User Interfaces servers.
- All accesses are logged, with disk space for 4 years of the operation before the circular buffer starts to be overwritten.
- Moreover, the physical access is strictly controlled via fingerprint and NFC smartcard, according to our ISO 9001 certification, again with a very long log history.
- Moreover, 10 IP cameras are present with a NVR to check anomalous accesses.

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
Università degli Studi di Napoli Federico II
INFN Istituto Nazionale di Fisica Nucleare
INAF ISTITUTO NAZIONALE DI ASTROFISICA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Power and Cooling

- Monitoring system for PUE management
- Using of Rittal metered PSM
- 2 420KW UPS
- 37 LCP refrigerating units
- 30 °C inside the Rack under full load conditions
- 2 Chillers of 400 thermal KW

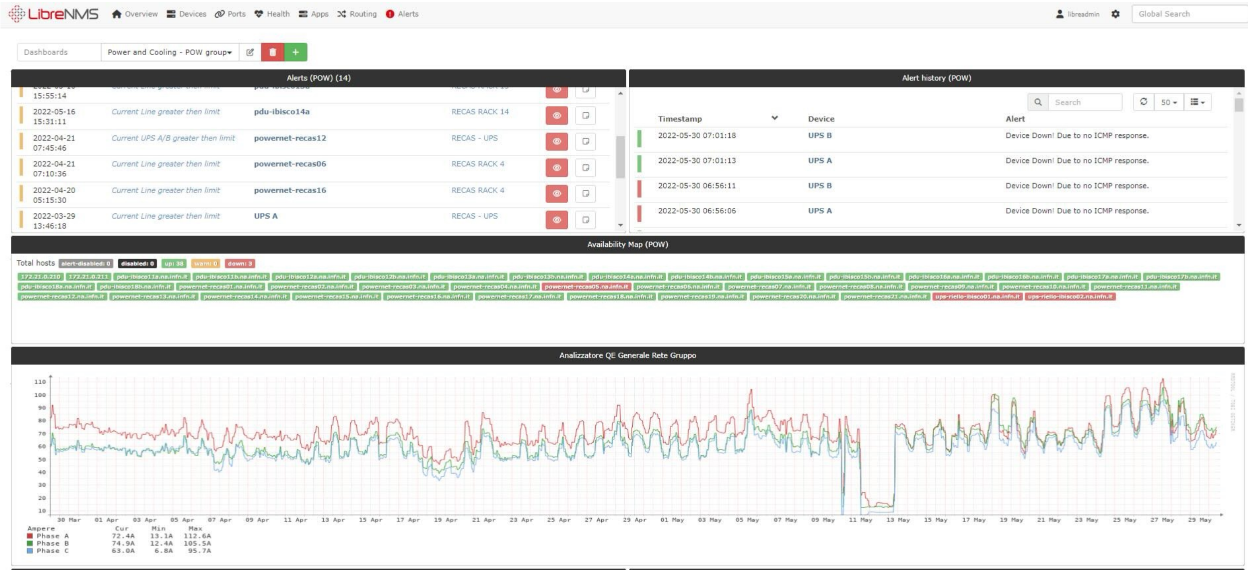

Control Room

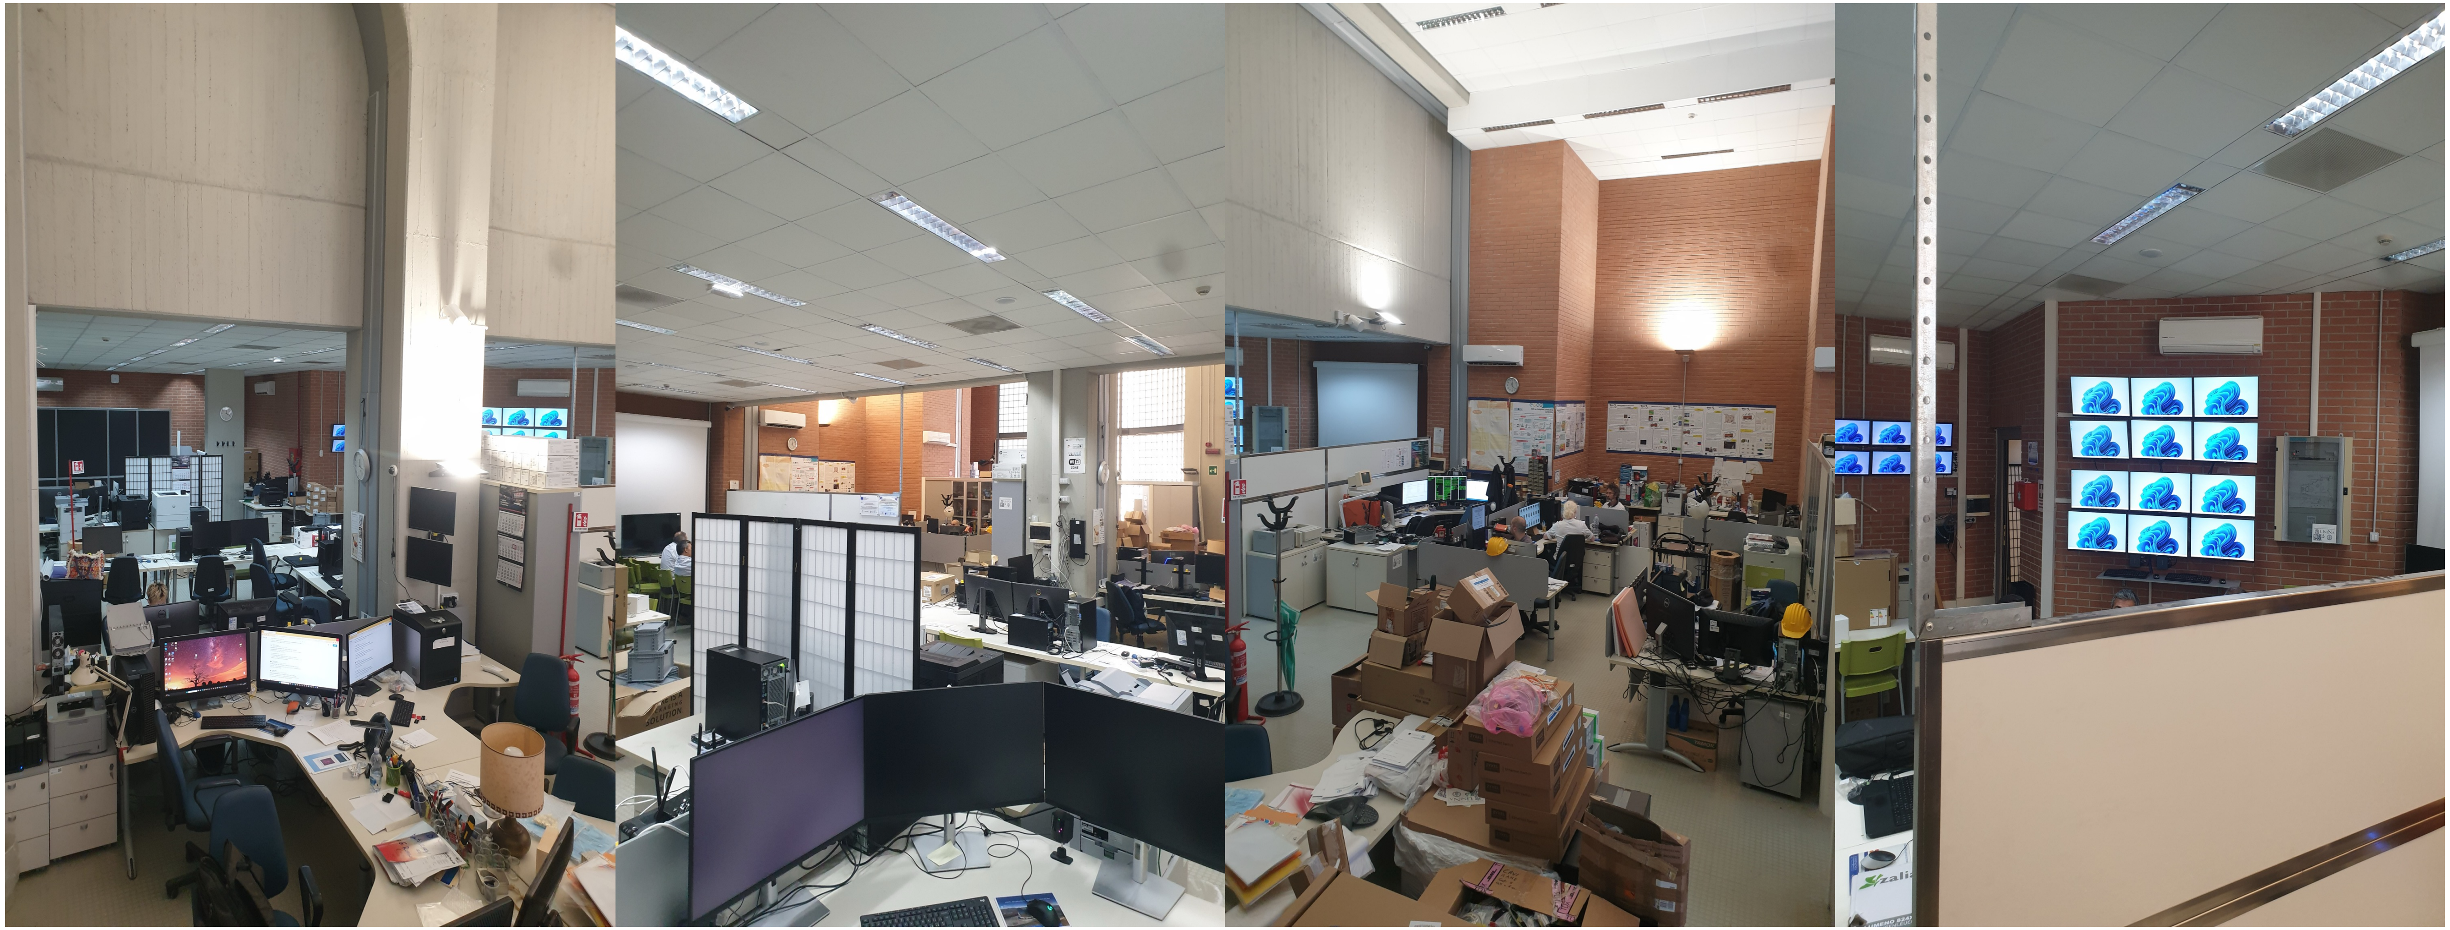

Example of application

• V. Rubin Observatory LSST (Legacy Survey of Space and Time) Project 1/2

- Rubin is a 8m class telescope characterized by a wide field of view and a high resolution;
- The camera observes in six filters covering wavelengths 320-1050 nm.
- About 90% of the observing time is dedicated to uniformly observe about 800 times 18000 deg2 of the sky (adding up on all 6 bands), during 10 years and will produce a co-added map up to magnitude 27.5 in r band.
- Maps of the entire sky visible in the southern hemisphere every few nights (20 TB of data will be collected each night); Fast detection of variable sources within 60sec from observation;
- Database obtained (of about 300 Petabytes) include: about 37 billion observations of 20 billion galaxies and 17 billion stars, 6 million solar system objects;

Example of application

- **V. Rubin Observatory LSST (Legacy Survey of Space and Time) Project 2/2**

- Analysis of crowded stellar fields (Galactic Bulge, Globulars, Magellanic Clouds) through specific software tools like the PSF Daophot/Allframe suite;
- This software requires parallel computing paradigms as well as machine and deep learning methods;
- The aim is characterizing stellar clusters and homogeneously deriving key parameters e.g., age, distance, reddening, metallicity, etc.) for globular and open clusters in the entire survey.

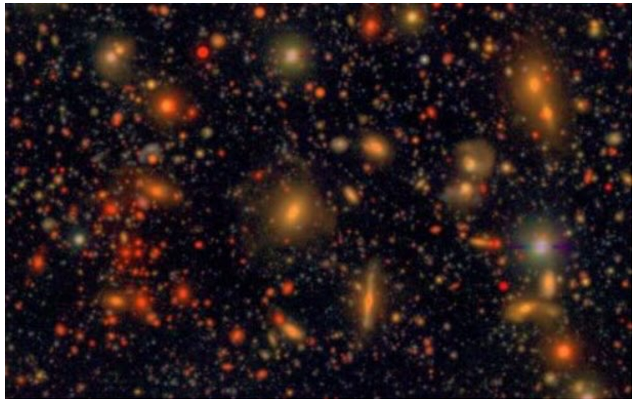

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
Università degli Studi di Napoli Federico II
INFN Istituto Nazionale di Fisica Nucleare
INAF ISTITUTO NAZIONALE DI ASTROFISICA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Example of application

- **Strengthening the Italian leadership in ELT and SKA (STILES) Project 1/2**

ELT:

- ELT is an optical/nearIR telescope with a 39m primary mirror, the largest of its kind ever built or planned;
- ELT is built by ESO, and will be located atop Cerro Armazones, a ~3000m peak in the Chilean desert.
- ELT is designed to exploit the full power of Adaptive Optics that removes atmospheric disturbances so as to reach the full resolution obtainable from the mirror.

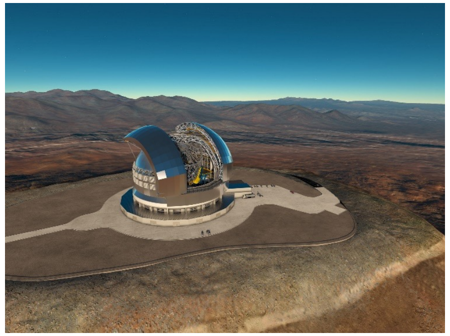

Example of application

• Strengthening the Italian leadership in ELT and SKA (STILES) Project 2/2

SKA OBSERVATORY:

- will comprise two radio interferometers spread over two continents.
- 512 stations with 256 low frequency antenna (50-350 MHz) array (SKA-Low) will reside in Western Australia, distributed over an area of 65 km in diameter;
- The mid-frequency dish array (SKA-Mid) (350 MHz-15 GHz) will be hosted in South Africa's Karoo region (197 dishes distributed over a region of 150 km in diameter).

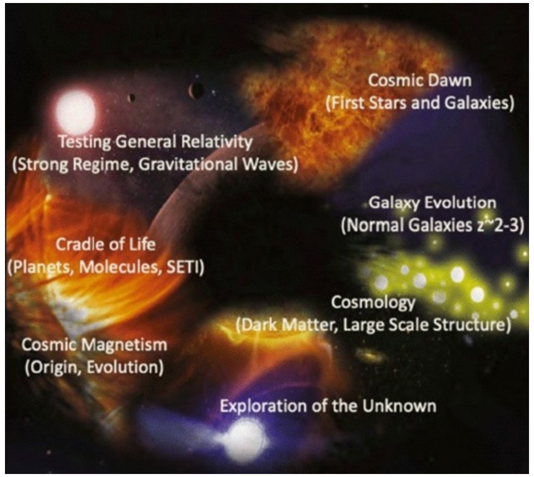

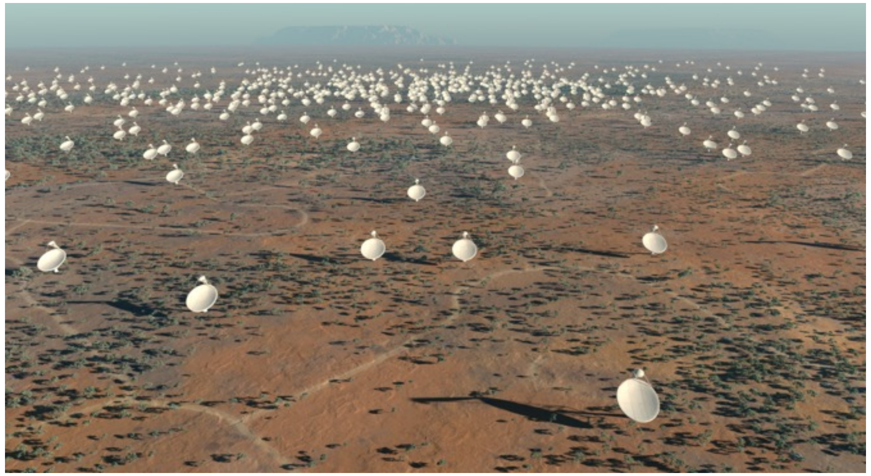

Example of application

Objectives:

- The computing and storage support offered by our infrastructure will achieve a synergy between ELT and SKA projects, since multi-frequency, multi-messenger approach is now recognized as a pillar of modern astronomy;
- Implement innovative data-mining techniques (generically referred to as Machine/Deep Learning);

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
Università degli Studi di Napoli Federico II
INFN Istituto Nazionale di Fisica Nucleare
INAF ISTITUTO NAZIONALE DI ASTROFISICA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Conclusions and Future Perspectives

Our goal is to provide a computing center, mainly devoted to astrophysical applications, built up on an existing infrastructure, improving all energy-consuming devices with new powerful yet not-so-much expensive hardware.

A simplified storage architecture with 10 Pbyte capacity and an efficient cooling system have been realized. The Data Center is operational, and a few applications have been illustrated.

In the future, our goals are to increase the computational power, and to maintain the Data Center operations with minimal downtime (TIER II).

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
Università degli Studi di Napoli Federico II
INFN Istituto Nazionale di Fisica Nucleare
INAF Istituto Nazionale di Astrofisica
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Dr. Eng. Augusto Tortora
augusto.tortora@unina.it

Università degli Studi di Napoli Federico II - Dipartimento di Fisica "Ettore Pancini"

Prof. Guido Russo
guido.russo@unina.it