Advice on INAF Future Engagements in Computational Astrophysics

G. Brunetti (INAF)
G. Bodo (INAF)
K. Dolag (LMU Munich and MPA Garching, Germany)
K. Heitmann (Argonne National Laboratory, USA)
L. Mayer (CTAC, Inst. for Computational Science, Univ. of Zurich, Switzerland)
A. Possenti (INAF)

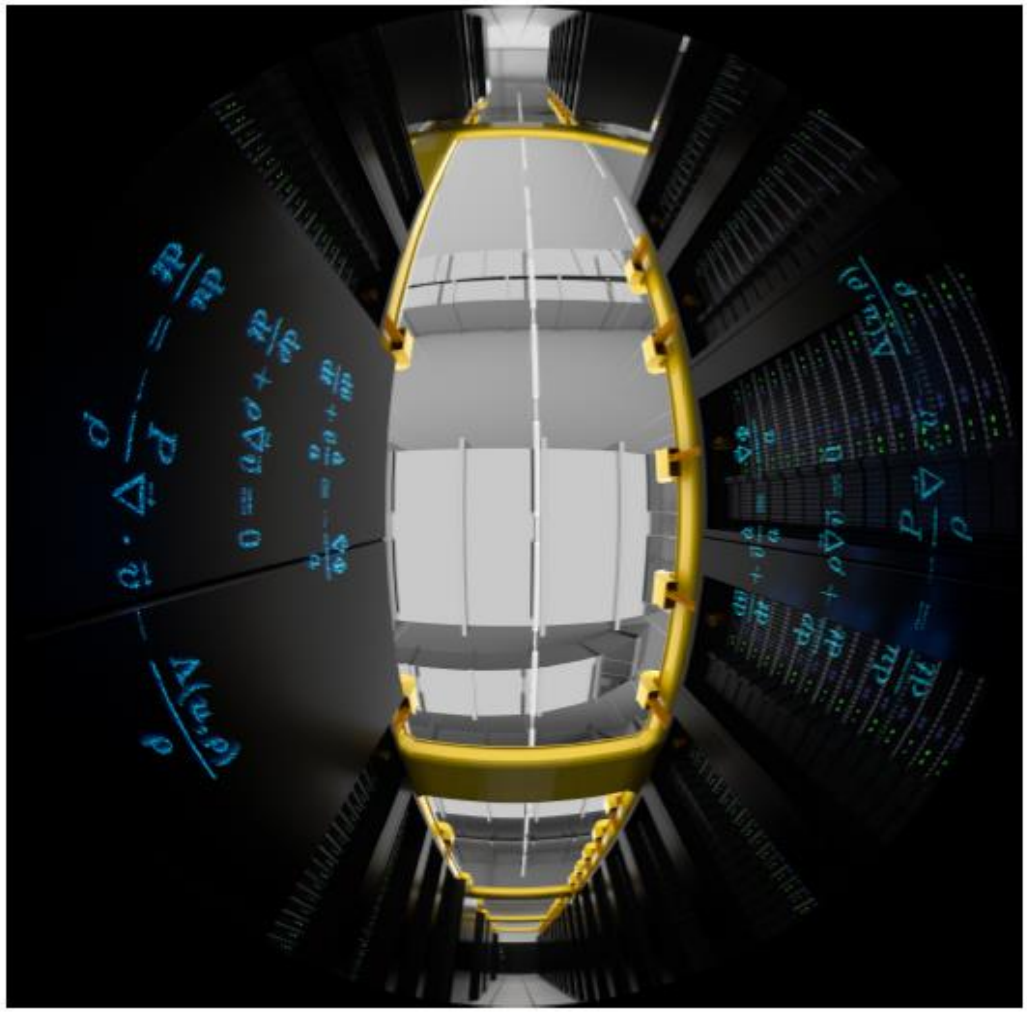

Advice on INAF Future Engagements in Computational Astrophysics

G. Brunetti (INAF)
G. Bodo (INAF)
K. Dolag (LMU Munich and MPA Garching, Germany)
K. Heitmann (Argonne National Laboratory, USA)
L. Mayer (CTAC, Inst. for Computational Science, Univ. of Zurich, Switzerland)
A. Possenti (INAF)

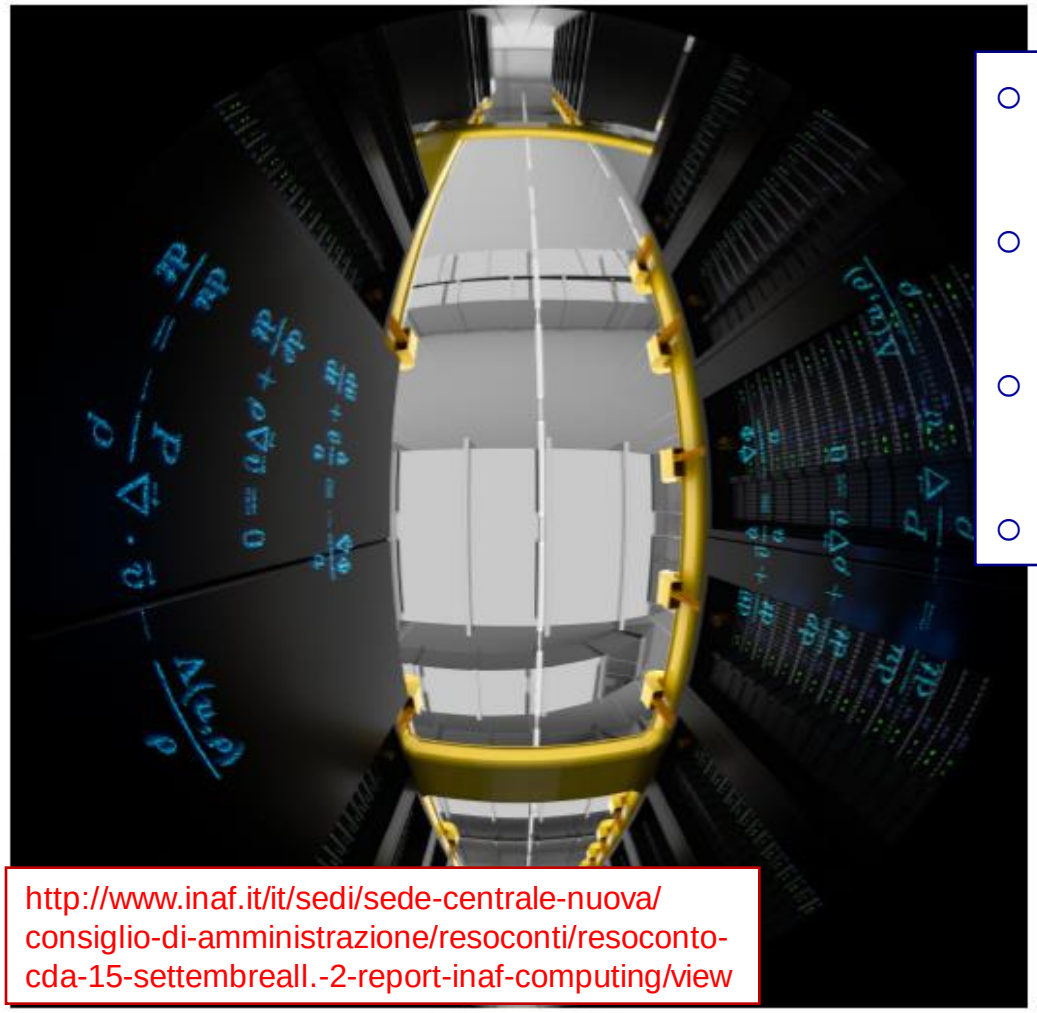

http://www.inaf.it/it/sedi/sede-centrale-nuova/consiglio-di-amministrazione/resoconti/resoconto-cda-15-settembreall.-2-report-inaf-computing/view

FOCUS : Calcolo Critico

- General perspectives of & advice on the intervention on CC
- Optimization of the overall architecture
- Suggestions on the configuration(s) of the infrastructure(s)
- Organization

CRITICAL IT INVESTMENT

Is a large investment necessary/useful in critical IT ?

(critical IT activity ? : 5+ kcores, 5+ PB)

HPC : service (access to computing facilities) does not require necessarily a large investment within INAF (no strong oversubscription observed in CINECA etc calls), **unless a '*change of vision*' is adopted: i.e. structural investment is combined with research (R&D) with the ambitious aim to place INAF among the most competitive communities on the roadmap towards exascale** etc..

CRITICAL IT INVESTMENT

Is a large investment necessary/useful in critical IT ?

(critical IT activity ? : 5+ kcores, 5+ PB)

HPC : service (access to computing facilities) does not require necessarily a large investment within INAF (no strong oversubscription observed in CINECA etc calls), **unless a '*change of vision*' is adopted: i.e. structural investment is combined with research (R&D) with the ambitious aim to place INAF among the most competitive communities on the roadmap towards exascale** etc..

Need synergies and critical mass

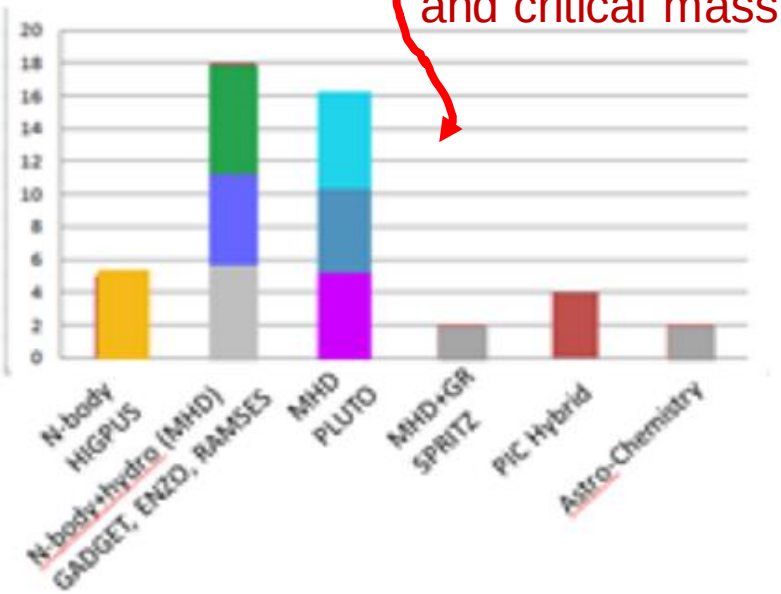

CRITICAL IT INVESTMENT

Is a large investment necessary/useful in critical IT ?

(critical IT activity ? : 5+ kcores, 5+ PB)

HPC : service (access to computing facilities) does not require necessarily a large investment within INAF (no strong oversubscription observed in CINECA etc calls), **unless a '*change of vision*' is adopted: i.e. structural investment is combined with research (R&D) with the ambitious aim to place INAF among the most competitive communities on the roadmap towards exascale** etc..

PROJECTS : clearly INAF is currently involved in **large observational projects** facing strong challenges in the field of data management and analysis. **These existing projects** (SKA precursors [LOFAR, MeerKAT], Euclid, Rubin, SKA RC) provide a first/initial critical mass and **motivate a roadmap of investments in critical IT and a parallel investment to strengthen competences** (computer science and technology) **in the community**.

CRITICAL IT INVESTMENT

Is a large investment necessary/useful in critical IT ? (critical IT activity ? : 5+ kcores, 5+ PB)

HPC :

PROJECTS : clearly INAF is currently involved in **large observational projects** facing strong challenges in the field of data management and analysis. **These existing projects** (SKA precursors [LOFAR, MeerKAT], Euclid, Rubin, SKA RC) provide a first/initial critical mass and **motivate a roadmap of investments in critical IT and a parallel investment to strengthen competences** (computer science and technology) **in the community**.

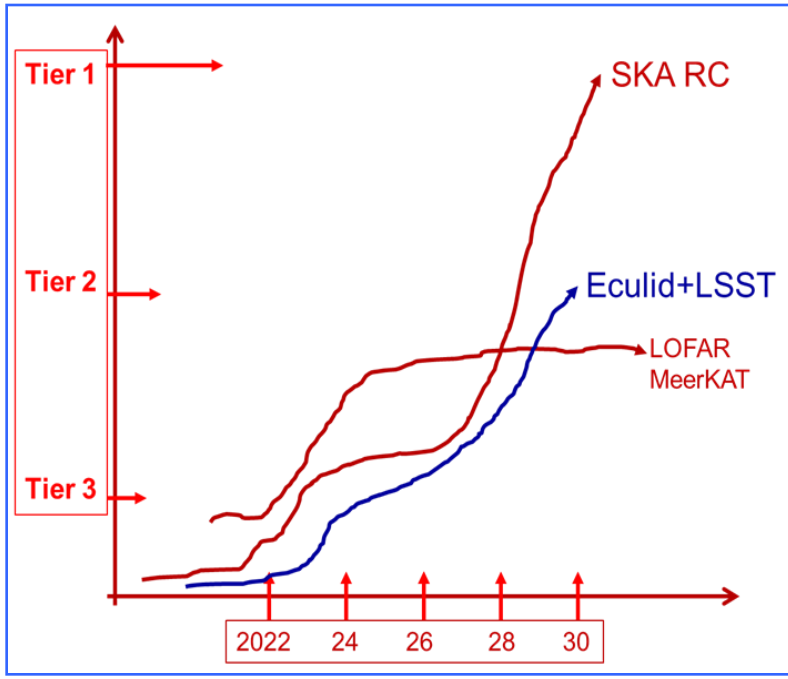

CRITICAL IT INVESTMENT

Is a large investment necessary/useful in critical IT ? (critical IT activity ? : 5+ kcores, 5+ PB)

HPC :

PROJECTS : clearly INAF is currently involved in **large observational projects** facing strong challenges in the field of data management and analysis. **These existing projects** (SKA precursors [LOFAR, MeerKAT], Euclid, Rubin, SKA RC) provide a first/initial critical mass and **motivate a roadmap of investments in critical IT and a parallel investment to strengthen competences** (computer science and technology) **in the community**.

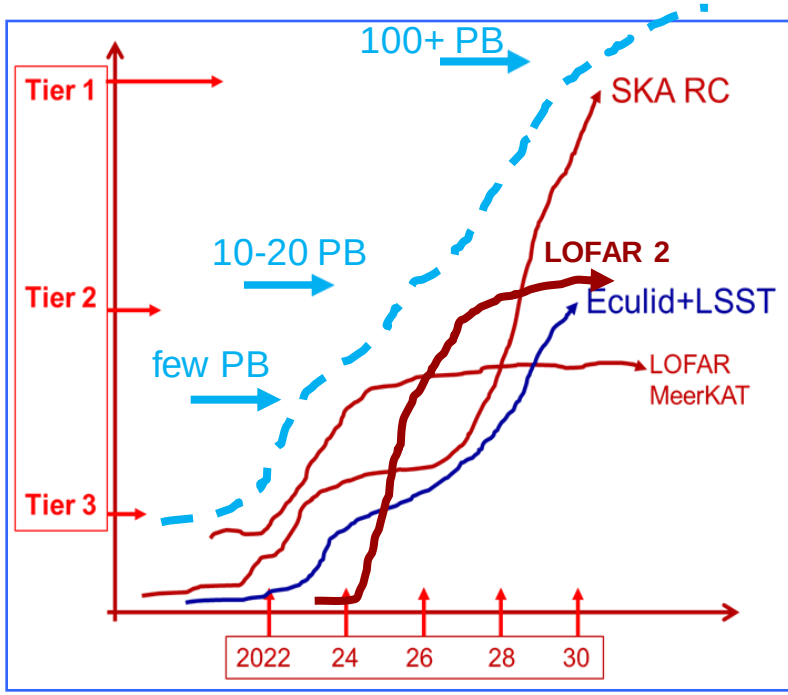

es, an involvement at critical level in LOFAR 2.0 KP will require 10-15 PB in 2024+ and significant computing resources

CRITICAL IT INVESTMENT

Is a large investment necessary in critical IT ? (critical IT activity ? : 5+ kcores, 5+ PB)

HPC : service (access to computing facilities) does not require necessarily a large investment within INAF (no strong oversubscription observed in CINECA etc calls), **unless a '*change of vision*' is adopted: i.e. structural investment is combined with research (R&D) with the ambitious aim to place INAF among the most competitive communities on the roadmap towards exascale** etc..

PROJECTS : clearly INAF is currently involved in **large observational projects** facing strong challenges in the field of data management and analysis. **These existing projects** (SKA precursors [LOFAR, MeerKAT], Euclid, Rubin, SKA RC) provide a first/initial critical mass and **motivate a roadmap of investments in critical IT and a parallel investment to strengthen competences** (computer science and technology) **in the community**.

GRASPING FUTURE OPPORTUNITIES : the existence of a large computing infrastructure in INAF **will allow INAF to grasp opportunities in the future**. Of course large projects with critical IT needs are not born "every day"… yet INAF might be in a better position to negotiate for .. e.g., *ground base segments or data centers in the future, in-kind contributions for large projects*.

MISSION of CRITICAL IT INVESTMENT

Clear **vision** on critical computing timely/necessary for the growth of the community in the field of computer science and for a strong (effective) role of INAF in the international projects and science programs with large computational needs.

The vision must identify what the **mission** of the critical IT investment is and provide a model/roadmap to operationalize it.

MISSION of CRITICAL IT INVESTMENT

Clear vision on critical computing timely/necessary for the growth of the community in the field of computer science and for a strong (effective) role of INAF in the international projects and science programs with large computational needs.
The vision must identify what the mission of the critical IT investment is and provide a model/roadmap to operationalize it.

LARGE INTERNATIONAL PROJECTS :

The effective role of INAF in large observational projects and their scientific exploitation will also depend on **IT resources** and on **the competences to develop novel IT algorithms for data analysis and manage Big Data**.

MISSION of CRITICAL IT INVESTMENT

Clear **vision** on critical computing timely/necessary for the growth of the community in the field of computer science and for a strong (effective) role of INAF in the international projects and science programs with large computational needs.
The vision must identify what the **mission** of the critical IT investment is and provide a model/roadmap to operationalize it.

LARGE INTERNATIONAL PROJECTS :

The effective role of INAF in large observational projects and their scientific exploitation will also depend on **IT resources** and on **the competences to develop novel IT algorithms for data analysis and manage Big Data**.

Type 1 : critical IT support for research programs **within INAF** aiming at the scientific exploitations of big data from international facilities or advanced numerical simulations (EUCLID, HPC, LSST, radio)

Type 2 : critical IT support for offering shared computational and archiving resources **within large international projects**, e.g. models of in-kind contribution etc

MISSION of CRITICAL IT INVESTMENT

Clear **vision** on critical computing timely/necessary for the growth of the community in the field of computer science and for a strong (effective) role of INAF in the international projects and science programs with large computational needs.
The vision must identify what the **mission** of the critical IT investment is and provide a model/roadmap to operationalize it.

LARGE INTERNATIONAL PROJECTS :

The effective role of INAF in large observational projects and their scientific exploitation will also depend on **IT resources** and on **the competences to develop novel IT algorithms for data analysis and manage Big Data**.

Type 1 : critical IT support for research programs **within INAF** aiming at the scientific exploitations of big data from international facilities or advanced numerical simulations (EUCLID, HPC, LSST, radio)

Type 2 : critical IT support for offering shared computational and archiving resources **within large international projects**, e.g. models of in-kind contribution etc

IT resources & organization should allow INAF to grasp new opportunities.

MISSION of CRITICAL IT INVESTMENT

Clear **vision** on critical computing timely/necessary for the growth of the community in the field of computer science and for a strong (effective) role of INAF in the international projects and science programs with large computational needs.
The vision must identify what the **mission** of the critical IT investment is and provide a model/roadmap to operationalize it.

LARGE INTERNATIONAL PROJECTS :

The effective role of INAF in large observational projects and their scientific exploitation will also depend on **IT resources** and on the **competences to develop novel IT algorithms for data analysis and manage Big Data**.
IT resources & organization should allow INAF to grasp new opportunities.

THEORY & SIMULATIONS:

Prepare the community for pre-exascale/exascale supercomputers

- **Cosmology** (connection with large projects)
- **Plasma astrophysics** (connection with radio, CTA, Athena,..)
- **Grasp opportunities from new architectures**

HOW .. ?? INGREDIENTS..

1. INFRASTRUCTURE(s) : archiving + computing power

2. COMPUTER SCIENCE DIVISION : researchers/R&D

ORGANIZATION : efficient National coordination

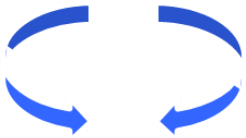

SYNERGIES : PROJECTS, PARTNERS (CINECA, TECNOPOLO…)

Infrastructure(s) :

One “combined” solution in a centralized location

- the **diversity of the needs** within the communities (HPC and OBS Projects) are as large as across the communities : **splitting in different (large) infrastructures is not convenient**.
- **centralized** infrastructure is more **cost efficient** (hardware as well as system engineers and cooling systems) than distributing machines in different places; **operating the infrastructure** (even a fraction of it !) **would require a large technical/system investment within INAF** (no strong/ready expertise is available @INAF…)
- **Roadmap** :
  - **incremental HPC** : it must attract and match the needs of a large community (several existing groups)
  - **incremental Critical IT for Data Analysis** : starts from SKA precursors (2022+)… Euclid (2024+)… full SKA RC
  - **dynamical** : it can evolve matching the needs of Projects with **Critical IT elements** as soon as they appear on the horizon (challenge in selection of…. @INAF DS+)

Hybrid + Partitions + Management

Infrastructure(s) :

One “combined” solution in a centralized location + distributed satellite systems (if needed).

- the **diversity of the needs** within the communities (HPC and OBS Projects) are as large as across the communities : **splitting in different (large) infrastructures is not convenient**.
- **centralized** infrastructure is more **cost efficient** (hardware as well as system engineers and cooling systems) than distributing machines in different places; **operating the infrastructure** (even a fraction of it !) **would require a large technical/system investment within INAF** (no strong/ready expertise is available @INAF…)

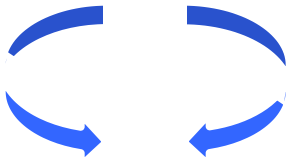

- **distributed** (limited number) **satellite systems** (Tier3-) dedicated to **specific tasks** may be useful for particular projects and contribute to/strengthen critical mass/competences in groups/projects (es ERC Grants, .. specific develop tasks in Lprojects…)

Computer Science Division :

The formation of a **division for computational astrophysics** in INAF is necessary to optimize the return on the large investment that is planned in critical IT infrastructures, and the coordination between infrastructures and large projects with critical IT needs. Personnel (about **15 FTE**) geographically distributed.

Mission :

- **research (R&D):** carry out **advanced research** in the field of computational astrophysics (innovative codes and architectures, .. porting etc) to impact main INAF scientific lines (simulations, .. pipelines for DATA)
- **Infrastructure(s):** coordinate numerical and technological **experiments** aimed at the optimization & design of the architecture of computing infrastructure(s). **Support** of the activities within the major projects and programs using the critical IT infrastructure(s) at INAF.
- **training & fertilization:** training actions within INAF also in connection with external bodies (e.g. Universities and supercomputing centers). **Support service** for INAF researchers for the use of very large (Tier 0) facilities abroad.

Computer Science Division :

The formation of a **division for computational astrophysics** in INAF is necessary to optimize the return on the large investment that is planned in critical IT infrastructures, and the coordination between infrastructures and large projects with critical IT needs. Personnel (about **15 FTE**) geographically distributed.

Mission :

- **research (R&D):** carry out **advanced research** in the field of computational astrophysics (innovative codes and architectures, .. porting etc) to impact main INAF scientific lines (simulations, .. pipelines for DATA)
- **Infrastructure(s):** coordinate numerical and technological **experiments** aimed at the optimization & design of the architecture of computing infrastructure(s). **Support** of the activities within the major projects and programs using the critical IT infrastructure(s) at INAF.
- **training & fertilization:** training actions within INAF also in connection with external bodies (e.g. Universities and supercomputing centers). **Support service** for INAF researchers for the use of very large (Tier 0) facilities abroad.

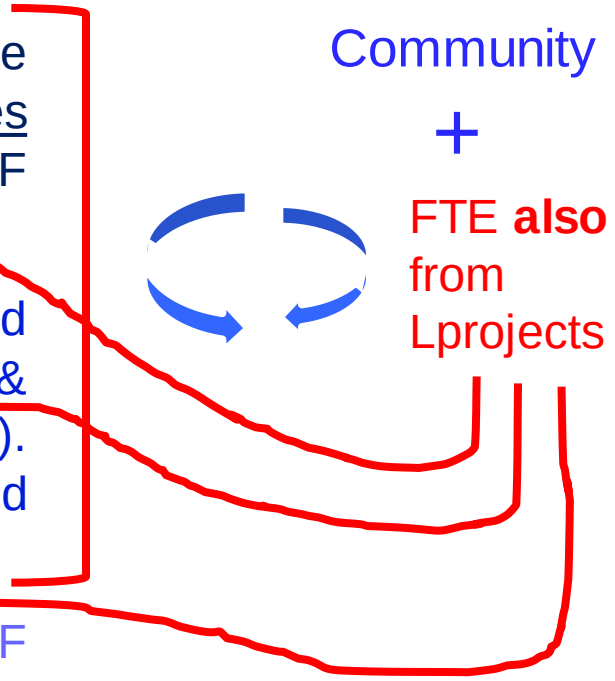

Management

3 Players : Projects+HPC, Infrastructure(s), Computer Science Division

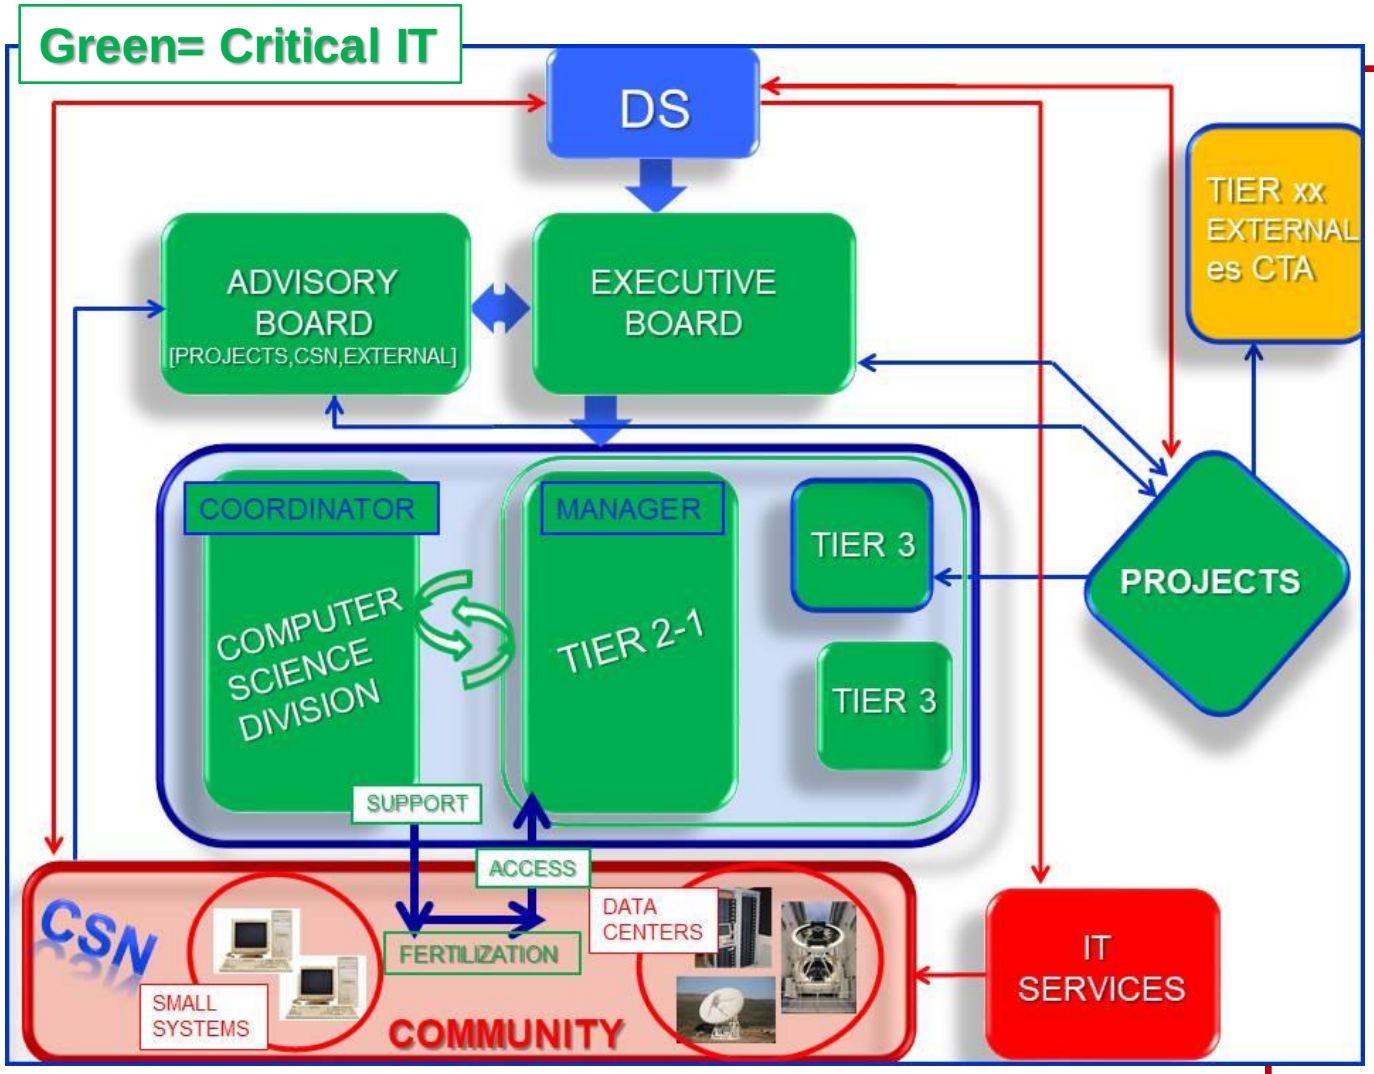

EXECUTIVE BOARD

- **Conflicts** between critical IT needs of (established) INAF Projects+HPC & management decisions should be **minimized** An **executive board** (President /Chair + voting members) offers a **good balance** and allows a **clear decision-making flows**.
- Provides sci-tech roadmap and decides on the priorities. Supervise management and coordination of infrastructure(s) (Manager) and of the computer science division (Coordinator). Hardware acquisition and plans for the enrolment of the personnel of division.
- A delegate of the **executive board** should keep the relations with the international and national partners and institutions
- We suggest a **delegate** from **Projects** and one from INAF **Directors** in the **executive board**

APPOINTMENT LETTER

It is clear that we are facing a perspective intervention of a quite large size and a radical change in the way computing is handled in our community. For this reason the Governance Body of INAF charged the Science Director to set up a committee to gain a broad scientific view on the program and some specific advice on it.

Beside any other suggestion you may provide us, INAF asks your advice on the following:

- **General perspectives of the intervention.** Is it appropriate ? Is it timely ? Is it in condition to be sufficiently transformational to mark a difference in producing outstanding scientific results ? Are there alternatives or significant changes to be considered ? What best roadmap would you see to match the needs of the community and the timescale of large programs and facilities ?
- **Optimization of the overall architecture.** Is the MFC linked to smaller specialized nodes an appropriate architecture for the purposes of the intervention ? Are there minimum or maximum investment thresholds you would recommend ?
- **Specific suggestions on the MFC configuration(s)** given its foreseeable use in the diverse scientific engagements of INAF. Are there particularly innovative/promising technologies on which to invest in the reference time frame of the roadmap ?
- **Organization.** Which would be an efficient management model and decision-three flow for the MFC, its possible connected nodes and related R&D activities ? How to best organize the user interface ?

We are aiming to receive a compact advisory paper (possibly within 10 pages, unless the committee considers a larger format more appropriate). The paper will be received by the Science Directorate and presented to INAF Council.

WORK PLAN : MAY - JULY 2021

1. basic survey of the needs/desiderata of the main projects

- current radio facilities (ALMA,SRT+EVN/VLBI,LOFAR+MeerKAT+...)
- SKA, SKA RC
- HPC-theory
- Euclid
- Gaia
- LSST
- ELT
- ASTRI Mini-Array
- CTA

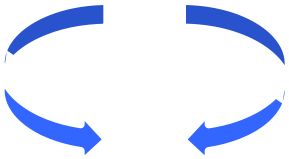

including auditions/talks in several cases

2. INAF plan : appropriate scale ? transformational ?

- Basic proposal for a single/multiple IF
  [multiple vs single IF, sizes of the IFs, configurations+technologies of the IFs]
- Roadmap to match desiderata
- R&D and training actions: define a critical mass

3. Organization

- Efficient management: basic principles
- Role of Projects & connections with CINECA and TECNOPOLO
- Connection with existing offices (es ICT)

MISSION of CRITICAL IT INVESTMENT

A clear vision on critical computing is timely for the growth of the community in the field of computer science and for a strong (effective) role of INAF in the international projects and science programs with large computational needs.

LARGE INTERNATIONAL PROJECTS :

The effective role of INAF in large observational projects and their scientific exploitation will also depend on IT resources and on the competences to develop novel IT algorithms for data analysis and manage Big Data.
A large investment is necessary to support :

- **Roadmap to the SKA**
  - SKA precursors/pathfinders (LOFAR, MeerKAT+, .. MeerKAT, ASKAP, etc)
  - SKA RC
- **Observational Cosmology & Time-domain**
  - EUCLID+LSST

THEORY & SIMULATIONS:

Prepare the community for pre-exascale/exascale supercomputers

- **Cosmology** (connection with large projects)
- **Plasma astrophysics** (connection with radio, CTA, Athena,..)
- **Grasp opportunities from new architectures**

MISSION of CRITICAL IT INVESTMENT

Clear **vision** on critical computing timely/necessary for the growth of the community in the field of computer science and for a strong (effective) role of INAF in the international projects and science programs with large computational needs.

LARGE INTERNATIONAL PROJECTS :

The effective role of INAF in large observational projects and their scientific exploitation will also depend on **IT resources** and on the **competences to develop novel IT algorithms for data analysis and manage Big Data**.
A large investment is timely to support current Lprojects:

- **Roadmap to the SKA**
  - SKA precursors/pathfinders (LOFAR, MeerKAT+, .. MeerKAT, ASKAP, etc)
  - SKA RC
- **Observational Cosmology & Time-domain**
  - EUCLID + LSST

THEORY & SIMULATIONS:

Critical Projects & Programs :

Projects & HPC play a central role in the vision of the investment for critical IT in INAF: **they represent the reason for the investment itself**. The success of the operation depends on how well the projects are supported.

Evaluation of needs :

Type 1 : critical IT support for research programs in INAF aiming at the scientific exploitations of big data from international facilities or advanced numerical simulations (EUCLID, HPC, LSST, radio)

Type 2 : critical IT support for offering shared computational resources within large international projects, e.g. part of distributed resources for archive or computing power (radio)

Critical Projects & Programs :

Projects & HPC play a central role in the vision of the investment for critical IT in INAF: **they represent the reason for the investment itself**. The success of the operation depends on how well the projects are supported.

Evaluation of needs :

Type 1 : critical IT support for research programs in INAF aiming at the scientific exploitations of big data from international facilities or advanced numerical simulations (EUCLID, HPC, LSST, radio)

Type 2 : critical IT support for offering shared computational resources within large international projects, e.g. part of distributed resources for archive or computing power (radio)

Critical IT needs:

- **Tier 2-3 (radio), Tier 3...1 (SKA RC), Tier 2 (obs cosmo..), Tier 2 (HPC)**
- **1...5 PB/yr (radio), ...50+ PB/yr (SKA RC), 10 PB (obs cosmo..), few PB (HPC)**
- **mid + fat + superfat (5 TB) nodes**
- **CPU + GPU**

Infrastructure(s) :

One combined solution in a centralized location would much better allow INAF to capture all the different needs than individual, local solutions **(+distributed satellite systems if needed)**

- "commercial cloud" is **not** an option (see report)
- the **diversity of the needs** within the communities are as large as across the communities !
- centralized infrastructure is more **cost efficient** (hardware as well as system engineers and cooling systems) than distributing machines in different places
- Size & Architecture & Partitions :
  - 2023-27 : reaching **Tier 2** (up to 30,000cores, Pflop, 30+ PB)
    - light CPU nodes : 10-30 % (256 GB RAM)
    - medium CPU nodes : 30-50 % (512 GB RAM)
    - fat CPU nodes : 5-10 % (2 TB RAM)
    - medium GPU nodes : 20-30 % (2-4 GPU/node, 512 GB RAM)
    - super-fat CPU nodes : 1-2 % (5 TB RAM)

Protocols (see report)
Nodes id (projects)
Partitions between type 1 & 2
Possibility to use the full platform

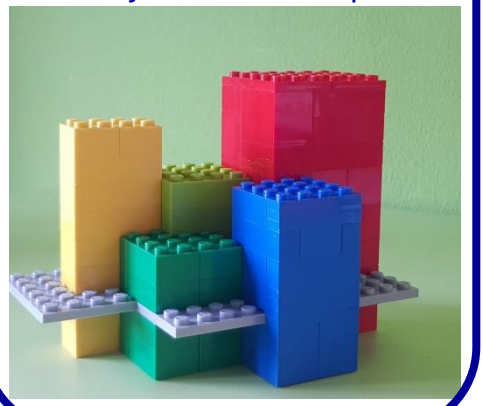

Infrastructure(s) :

One combined solution in a centralized location would much better allow INAF to capture all the different needs than individual, local solutions **(+distributed satellite systems if needed)**

- "commercial cloud" is **not** an option (see report)
- the **diversity of the needs** within the communities are as large as across the communities !
- centralized infrastructure is more **cost efficient** (hardware as well as system engineers and cooling systems) than distributing machines in different places
- Size & Architecture & Partitions :
  - 2023-27 : reaching **Tier 2** (up to 30,000cores, Pflop, 30+ PB)
    - light CPU nodes : 10-30 % (256 GB RAM)
    - medium CPU nodes : 30-50 % (512 GB RAM)
    - fat CPU nodes : 5-10 % (2 TB RAM)
    - medium GPU nodes : 20-30 % (2-4 GPU/node, 512 GB RAM)
    - super-fat CPU nodes : 1-2 % (5 TB RAM)
  - 2028-30 : reaching **Tier 1** (3+ Pflop, 70 PB/yr) + **Tier 2**

Protocols (see report)
Nodes id (projects)
Partitions between type 1 & 2
Possibility to use the full platform

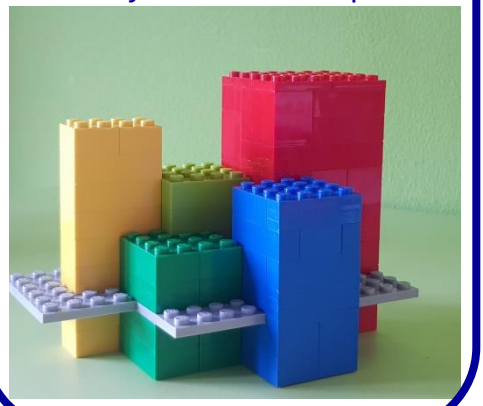

Infrastructure(s) :

One combined solution in a centralized location would much better allow INAF to capture all the different needs than individual, local solutions **(+distributed satellite systems if needed)**

- "commercial cloud" is **not** an option (see report)
- the **diversity of the needs** within the communities are as large as across the communities !
- centralized infrastructure is more **cost efficient** (hardware as well as system engineers and cooling systems) than distributing machines in different places
- Size & Architecture & Partitions :
  - 2023-27 : reaching **Tier 2** (up to 30,000cores, Pflop, 30+ PB)
  - 2028-30 : reaching **Tier 1** (3+ Pflop, 70 PB/yr) + **Tier 2**
- **Location** : convenient site in which **large computers are already operated** by specialized groups of **systems engineers** (Tecnopolo (BO), or similar).
  [Urgent to sign agreements with CINECA etc…]
  [e.s., Tier 1 platform might be acquired along with LEONARDO upgrade…]

Management

3 Players : Projects, Infrastructure(s), Computer Science Division

A management should be able to manage and operate infrastructure(s) to guarantee the critical IT needs of Projects, coordinate a Computer Division (research+services) and its interplay with Projects and the broad community.

Management

3 Players : Projects, Infrastructure(s), Computer Science Division

A management should be able to manage and operate infrastructure(s) to guarantee the critical IT needs of Projects, coordinate a Computer Division (research+services) and its interplay with Projects and the broad community.

- In the current configuration of INAF a management sitting “within” the DS appears the **most natural approach**: it simplifies the coordination of the Critical IT with UTG (projects) and INAF Directors (personnel).

Management

3 Players : Projects, Infrastructure(s), Computer Science Division

A management should be able to manage and operate infrastructure(s) to guarantee the critical IT needs of Projects, coordinate a Computer Division (research+services) and its interplay with Projects and the broad community.

- In the current configuration of INAF a management sitting "within" the DS appears the **most natural approach**: it simplifies the coordination of the Critical IT with UTG (projects) and INAF Directors (personnel).

Limitations :

- What is the decision-making process "within INAF" that should determine which **projects need to be included in the critical IT model** ?
- What is the decision process which sets the **$$funds** for critical IT ?

Management

3 Players : Projects, Infrastructure(s), Computer Science Division

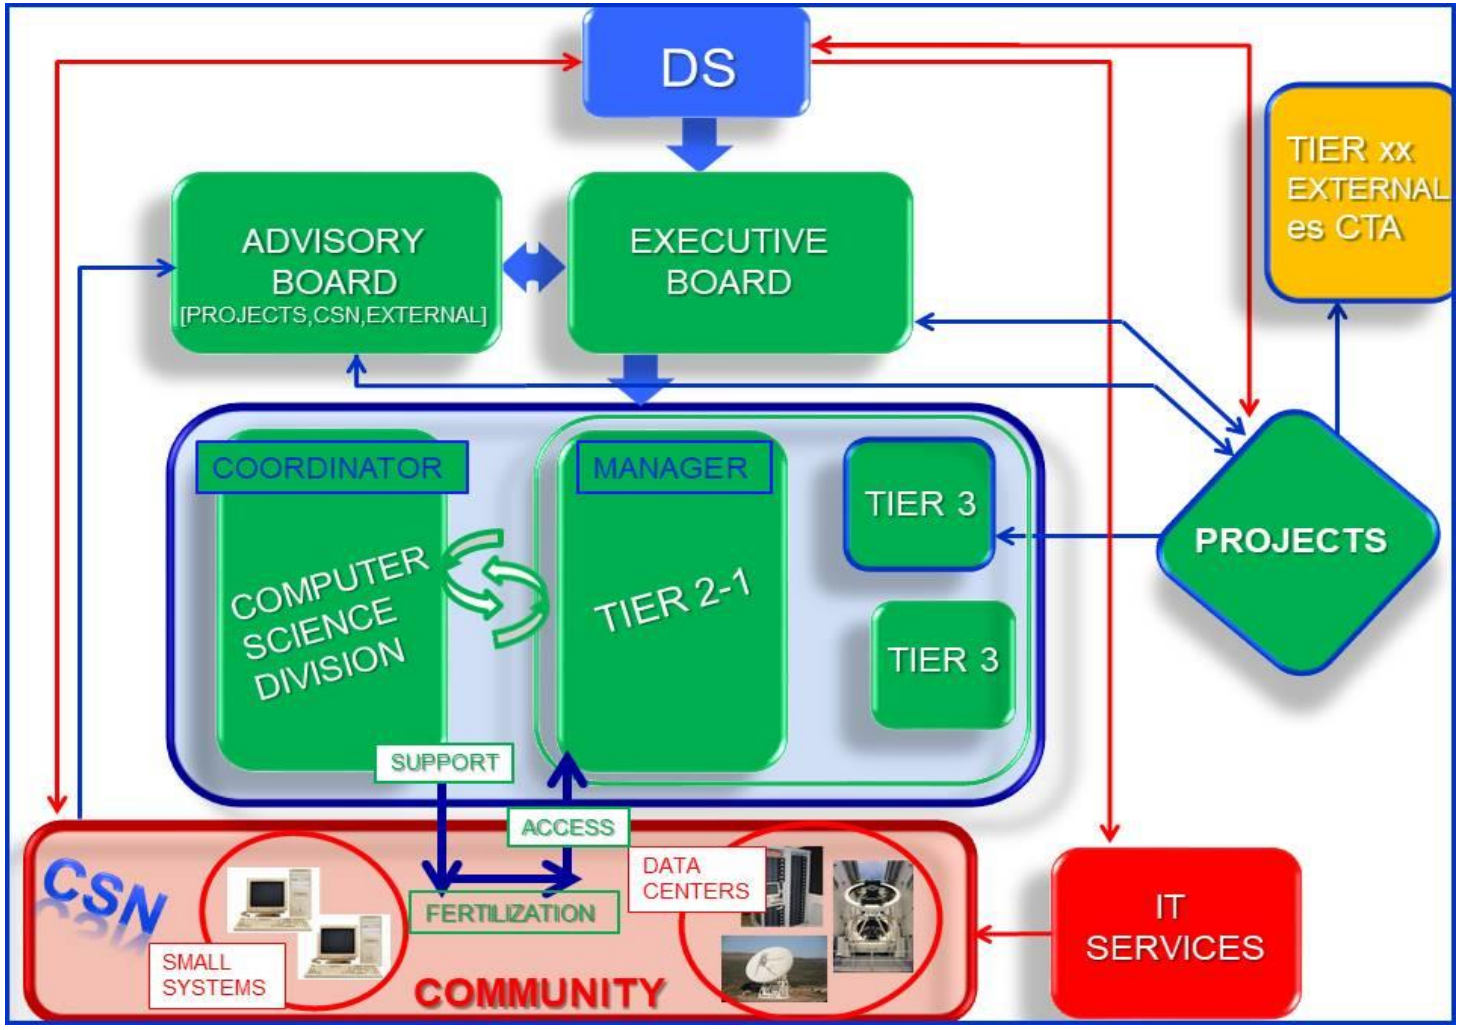

Management

3 Players : Projects, Infrastructure(s), Computer Science Division

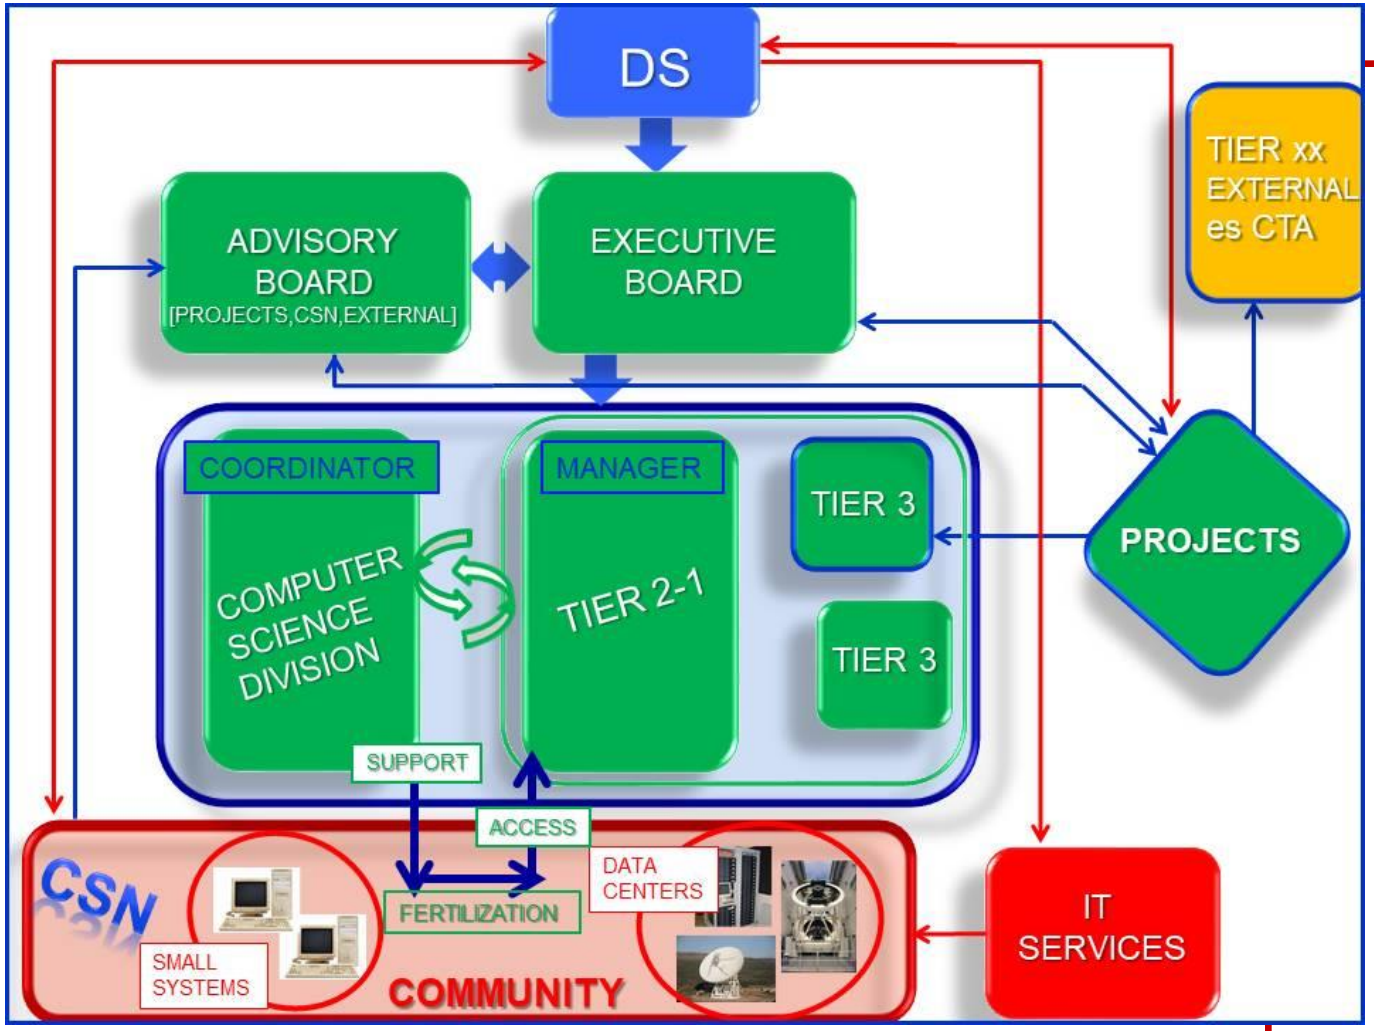

EXECUTIVE BOARD

- **Conflicts** between critical IT needs of (established) INAF Projects and management decisions should be **minimized** An **executive board** (President /Chair + voting members) offers a **good balance** and allows a **clear decision-making flows**.
- Provides sci-tech roadmap and decides on the priorities. Supervise management and coordination of infrastructure(s) (Manager) and of the computer science division (Coordinator). Hardware acquisition and plans for the enrolment of the personnel of division.
- A delegate of the **executive board** should keep the relations with the international and national partners and institutions
- We suggest a **delegate** from **Projects** and one from INAF **Directors** in the **executive board**

ROADMAP

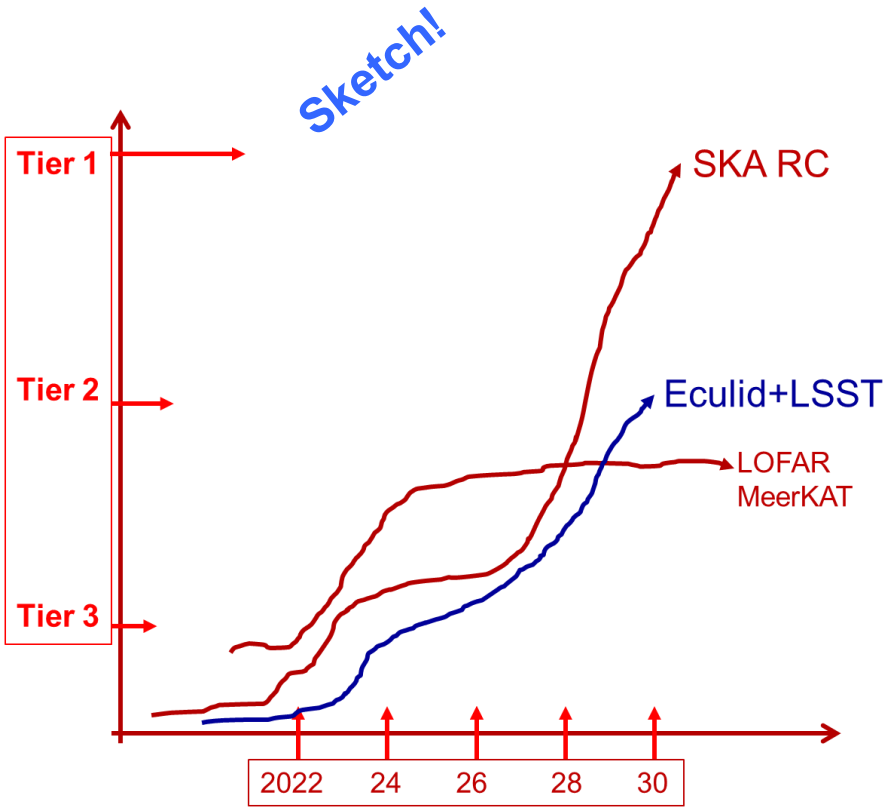

HPC Activity INAF

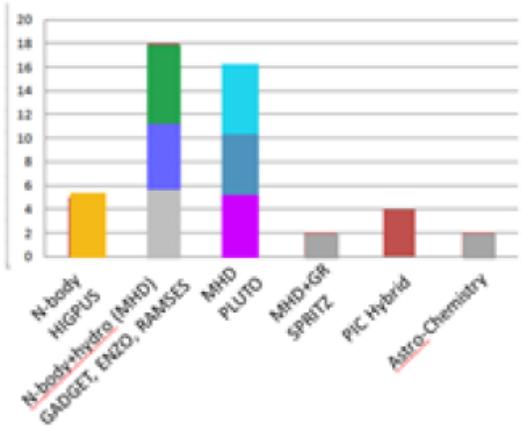

ROADMAP

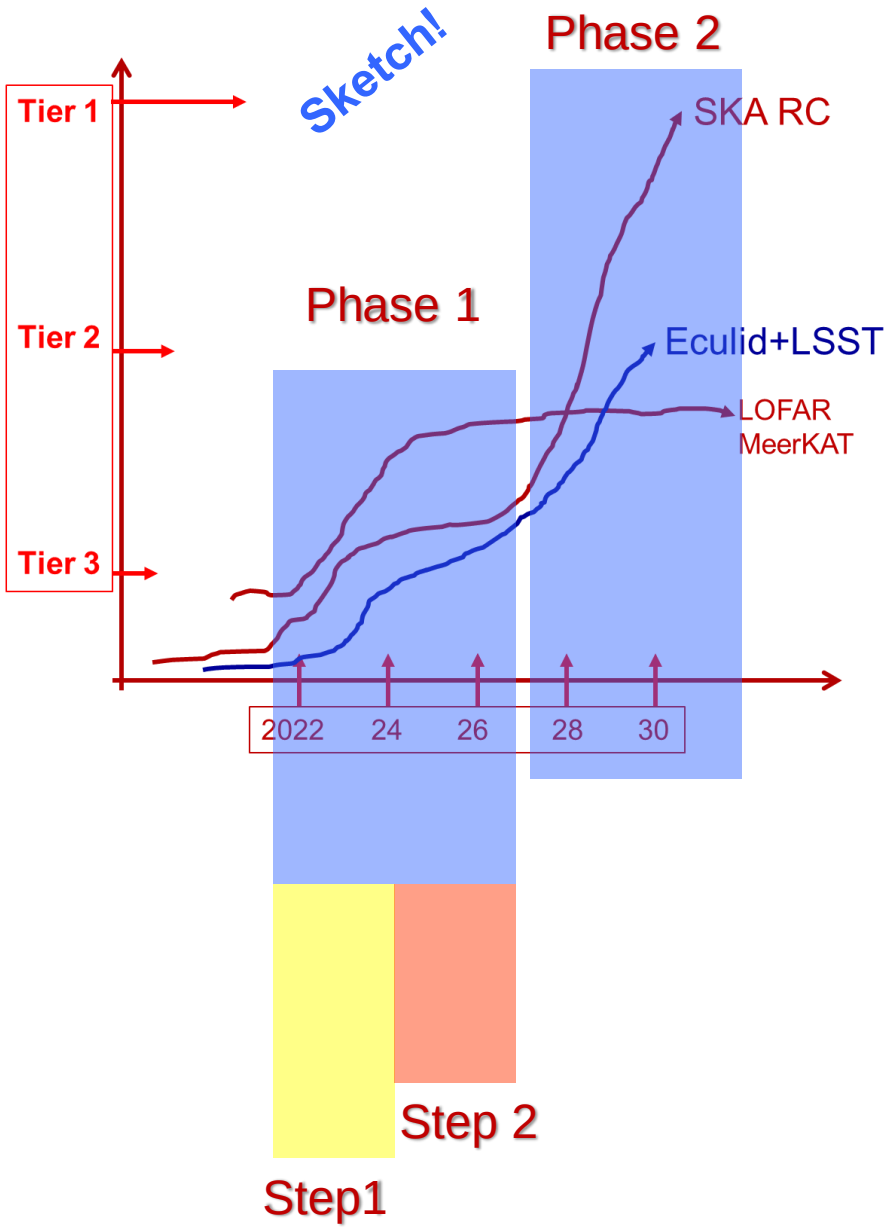

PRIORITIES Step1 (2021-24):

1. **Computer science division**: organize a network of researchers with strong skills in the field of critical IT and with strong connections and interest in the major projects underway and in the HPC activity. Acquisition of new staff, preferably by reinforcing existing groups.

2. Renew MoU with CINECA establishing agreements that combine service and R&D activities in the critical IT field. Agreements for Tecnopolo.

3. First **investment on a platform** that aims at a computing power of about 200-400 Tflops and **several PB/yr storage**.

Step 2 (2025-28)

4. Reach Tier 2 & 30+ PB

5. Complete computer science division

Phase 2 (2028+)

6. SKA RC : Tier 1 & 50+ PB/yr