THE INAF/OPENPOWER FOUNDATION MoU: STATUS

A. Bulgarelli (INAF) - OpenPOWER for Physical Science Workgroup chair and member of the OpenPOWER Technical Steering committee

Contribution and support from:
+ all people involved in the openpower@inaf.it mailing list
+ all authors of the charter (A. De Rosa, L. Graziani, R. Schneider, U. Becciani, V. Fioretti, L. Clavien, C. Carniel)
+ all people of the OpenPOWER for Physical Science working group
+ CERN people (F. Carminati, J. Apostolakis, M. Bandieramonte, A. Gheata, S. Vallecorsa)
+ IBMers (A. Negro, L. Sangiorgi, R. Ronchese)
+ E4 (D. Gregori, M. Ligabue, C. Gianfreda)
+ INAF people that are working on Power system: A. De Rosa, N. Parmiggiani, L. Baroncelli, A. Aboudan, G. Gianella

Thanks to INAF Direzione Scientifica (R. Smareglia and F. M. Zerbi - primary representative) and OPF Technical Steering Committee for their support.

Outline

- OpenPOWER
- Technical Steering Committee activities
- Workgroup on Physical Science
- Work in progress on Power platform

OPENPOWER

OpenPOWER foundation

- **INAF** becomes a member of the **OpenPOWER Foundation** in 2015 (Academic member)
- The OpenPOWER Foundation was founded in 2013 as an open technical membership organization around the **CPU Power and related technologies, until data centers**.
- The Foundation proposes a **collaborative** model that goes beyond the simple producer-customer relationship: both hardware technology and related software (based on Linux) are "**open**".
  - This allows new manufacturers to deliver Power products without depends on IBM.

- **Open hardware (many vendors)**
- **Open source (linux)**
- **The format is a 'foundation', a place where is possible to discuss with developers**

D.A.V.I.D.E. Supercomputer, developed by E4 Computer Engineering and installed at CINECA, ranked on November TOP500 and Green500 Lists

D.A.V.I.D.E. (Development of an Added Value Infrastructure Designed in Europe), the energy-aware, Peta-flops Class High Performance Cluster, based on OpenPOWER servers and featuring liquid cooling and an innovative technology for monitoring and capping the power consumption, enters November 2017 TOP500 and Green500 lists.

DENVER - November 14, 2017 – At the Supercomputing Conference 2017, E4 Computer Engineering announced that D.A.V.I.D.E. entered for the second time the prestigious TOP500 and Green500 lists, ranking respectively position #440 and #18. D.A.V.I.D.E., a multi-node cluster based on IBM POWER8 servers and NVIDIA Tesla P100 data center GPUs, features two relevant innovations that enhance whole system capabilities in terms of performance and energy-efficiency. D.A.V.I.D.E. entered the TOP500 and GREEN500 list in June 2017 in its air-cooled version, while the current version features liquid cooling and an innovative technology for monitoring and capping the power consumption.

OpenPOWER @ INAF

- Web site and mailing list: openpower@inaf.it
  - http://www.ict.inaf.it/twiki/bin/view/OpenPOWER/WebHome
- The **OpenPOWER workgroup for Physical Science**
  - Member of the Technical Steering Committee
- **You can register yourself at** https://openpowerfoundation.org **and follow the work of the working groups.**

OPF Workgroup products (2017)

- Pending BOD approval:
  - OpenPOWER Ready WG: OpenPOWER Ready Definition and Criteria V2.0 – public review draft
  - 25GIO WG: Advanced Accelerator Adapter Electro-mechanical Specification – Final
- WG Specifications Completed this year:
  - Compliance WG: OpenPOWER CAPI 2.0 Accelerator Compliance
  - Compliance WG: OpenPOWER FSI Protocol Compliance Test Harness and Test Suite (TH/TS) Specification
  - Memory WG: OpenPOWER Memory Bus V1.0
  - Accelerator WG: PSL/AFU Interface CAPI 2.0 v1.0
  - System SW WG: 64bit ELF V2 ABI
  - FSI WG: Open FSI Specification v1.0.0
  - Compliance WG: OpenPOWER ISA Compliance Definition
  - Compliance WG: OpenPOWER CAPI 1.0 Accelerator Compliance Test Specification
- WG Notes Completed this year:
  - Personalized Medicine WG: Porting and Benchmarking of BWAKIT pipeline on Open POWER architecture
  - WG for Physical Sciences: Collecting Use Case Information: Method and Template for Physical Science

Technical Steering Committee responsibilities

- **Recommending** the charter of **new Work Groups** to the Board; including scope of work and confidentiality level;
- **Resolving technical conflicts** within and between Work Groups;
- Evaluating requests for **inter-project collaboration** and recommending the chartering of new Work Groups to facilitate such collaboration;
- **Monitoring the technical progress of Work groups**;
- **Approving OpenPOWER Standards Final Deliverables**;
- Establishing OpenPOWER Standards Final Deliverable quality standards, technical best practices and release dates;
- Approving changes by Work Groups to OSS Communities and OSS Licenses in the Work Group's charter, and under section 3.5 of the Intellectual Property Rights Policy;
- Regularly reporting on progress of Work Groups to the Board.

THE OPENPOWER FOR PHYSICAL SCIENCE WORKGROUP

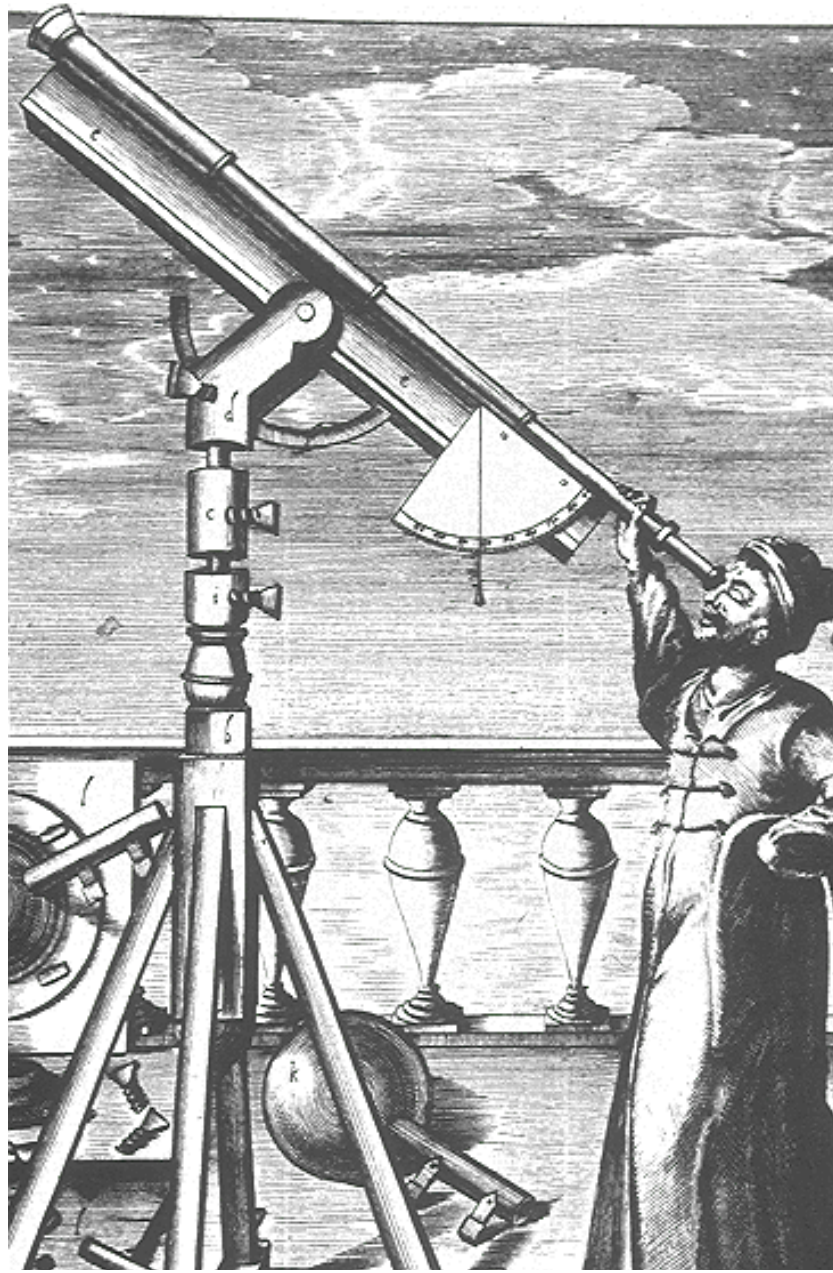

Motivations

- A dedicated **workgroup for Physical Science needs** within the **OpenPOWER Foundation** is started:
  - The proposed workgroup aims at addressing the ICT challenges of **Physical Science projects**.
  - A forum between ‘scientists’ of different fields around common ICT problems
- A forum between 'scientists', and ’hw/sw developers’/‘vendors’:
  - a direct connection with hw/sw developers
  - a direct and different connection with the market

OPF for Physical Science WG: a **forum** of 'scientists' and ‘(technological) developers’ at the same level around a technological solution **(Power architecture and Linux)**.

COLLABORATION WITH OPENPOWER

- **May 2016**: **OpenPOWER:** setup of OPF WG for Physical Science charter subcommittee. Not-INAF key persons for the startup of the WG: C. Bekas (IBM Labs), C. Carniel, P. Altoe, L. Clavien, D. Pleiter (JSC)
- **May 2016** charter submitted to OPF members
- **August 2016** BOD charter approval
- **October 2016**: OPF for Physical Science WG KO meeting
- **May 2017**: first release of the WG Note: “Collecting Use Cases Information”
- **November 2017**: participation of Machine Learning WG. Common use cases?

PARTICIPATION + Intellectual Property Rights (IPR) + CONFIDENTIALITY MODE

- Participation mode: **public**
  - Participation in the WG shall be open to Eligible and non-Eligible people **(i.e. non OPF people,** as defined in the OPF IPR Policy)
    - there are no restrictions on their participation or **feedback/contribution** to provide **Use Cases**
- **IPR (Intellectual Property Rights)**
  - **Our deliverables: notes**
  - The notes can be publicly available
- **Not confidential**
  - contributions and feedbacks are not subject to any requirement of confidentiality

WG MEMBERS (CURRENT)

Ashworth, Mike	Hartree Centre, Science & Technologies Facilities Council	
Bulgarelli, Andrea	INAF - Istituto Nazionale di Astrofisica	
Clavien, Lionel	Groupe T2i SA	
Gregori, Daniele	E4 Computer Engineering SPA	
Pleiter, Dirk	ForschungszentrumJuelichGmbH	
ronchese, Raniero	IBM	
Bastl, Pavel	Institute of Physics, public research institution of the Academy of Sciences of the Czech Republic	
Blemings, Hugh	ISTO	
Bozzo-Rey, Mehdi	IBM	
Breuner, Sven	ThinkParQ Gmbh	
Brown, Jeff	IBM	jeffdb@us.ibm.com
Carniel, Cecilia	IBM	
Chen, Yunju	Wistron Corporation	

De Cesare, Giovanni	INAF - Istituto Nazionale di Astrofisica	
Fioretti, Valentina	INAF - Istituto Nazionale di Astrofisica	
Magugliani, Fabrizio	E4 Computer Engineering SPA	
McKenney, Shawn	Emmion	
Mehta, Kiran	Individual Member	
Smareglia, Riccardo	INAF - Istituto Nazionale di Astrofisica	
Thompson, Lance	IBM	
Zanella, Riccardo	Cineca	

WORKGROUP EXPECTED DELIVERABLES

PROJECTS

- (1) **Current and future Physical Science projects use cases, requirements, common workflows and reference solutions**. Collection of use cases and related requirements. Based on these requirements, identification of common workflows and possible reference solutions in collaboration with other OPF WGs.
- (2) **Scientific software frameworks and libraries.** Identification of widely used software frameworks and libraries used in the Physical Science (based on the experience of WG members), status of the porting to OpenPOWER solutions. Based on the contribution of WG members, execution of performance tests. Usage of proprietary software already developed or available in participating scientific groups.

DELIVERABLES

- **WP1 (method)**. The writing of a Note (or creation of wiki pages) on **methodologies for describing use cases**, to document use cases in a way useful for people that are not an expert in the domain area of the Physical Science projects. This help to link the activities of this WG to those of other WGs and to manage in an effective way the requirements inception.
- **WP2 (Use Cases)**. The writing of Note(s) (or creation of wiki pages) identifying specific **use cases of Physical Science Projects** (from Big Science projects as well as projects at laboratory level), the area (e.g. experimental or computational science), **the workflow and related requirements**. The selection of use cases will depend on the experience and availability of the WG members.
- **WP3 (Software)**. The writing of Note(s) (or creation of wiki pages) that lists the **main available scientific software frameworks and libraries** (both proprietary and open source) for Physical Science projects identified within WP2, and t**he status of the porting** of them on OpenPOWER platforms. The selection of framework and libraries will depend on the experience and availability of the WG members.
- **WP4 (Solutions): Those Notes will identify common workflows and possible platforms (HW and SW)**

FIRST DELIVERABLE

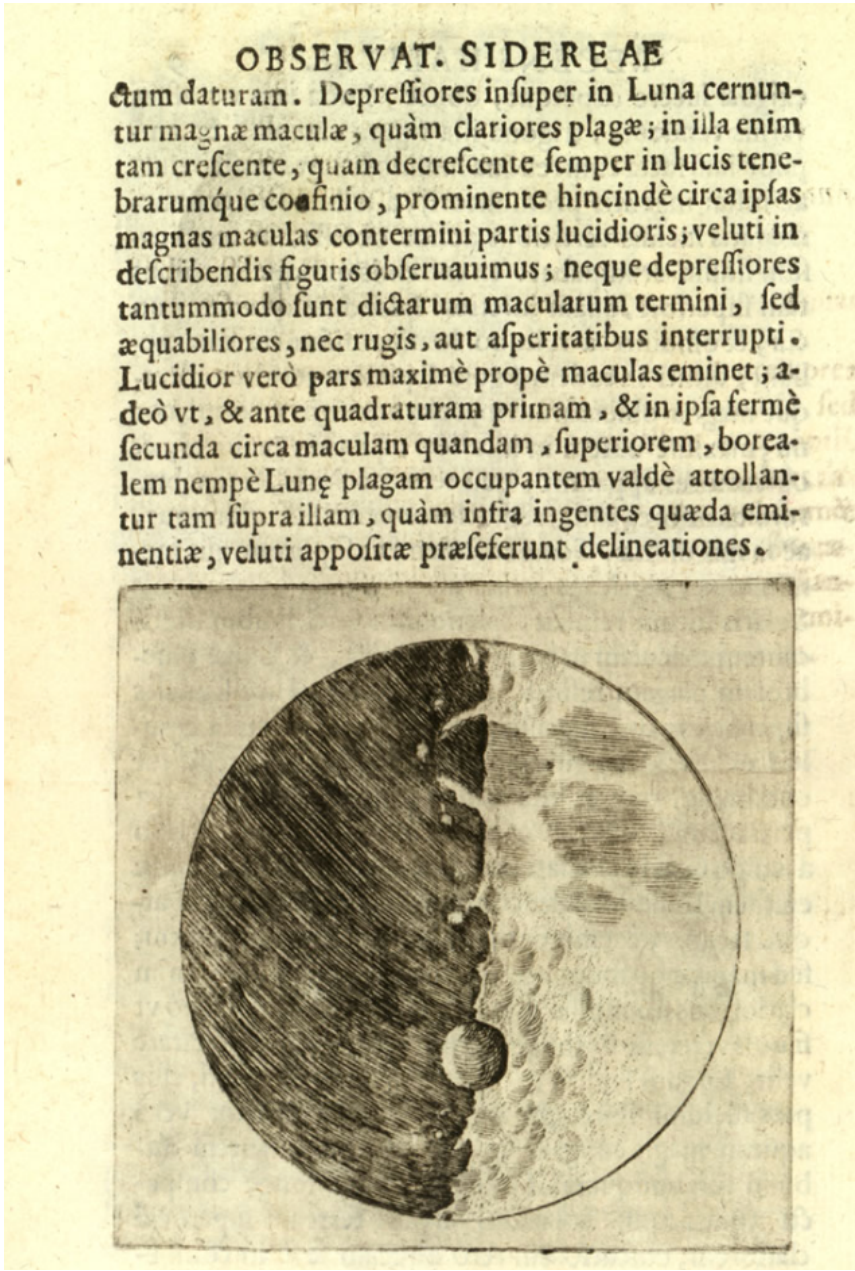

OBSERVAT. SIDEREAE

ctum daturam. Depreſſiores inſuper in Luna cernuntur magnæ maculæ, quàm clariores plagæ; in illa enim tam creſcente, quam decreſcente ſemper in lucis tenebrarumque confinio, prominente hincindè circa ipſas magnas maculas contermini partis lucidioris; veluti in deſcribendis figuris obſeruauimus; neque depreſſiores tantummodo ſunt dictarum macularum termini, ſed æquabiliores, nec rugis, aut aſperitatibus interrupti. Lucidior verò pars maximè propè maculas eminet; adeò vt, & ante quadraturam primam, & in ipſa fermè ſecunda circa maculam quandam, ſuperiorem, borealem nempè Lunę plagam occupantem valdè attollantur tam ſupra illam, quàm infra ingentes quæda eminentiæ, veluti appoſitæ præſeferunt delineationes.

FIRST NOTE

- The main purpose of this Note is to **define a method and a template to collect Physical Science use cases from scientists and research engineers working on Physical Science projects** in the context and within the scope of the OpenPOWER Foundation for Physical Science Workgroup.
- An effective method (shared between all stakeholders) could contribute to
  - **understand the workflow**
    - help Physical Science projects maintain costs within a chosen envelope;
  - starting with **user expectations**
    - map the functionalities to the **scientific requirements**;
  - remove possible **misunderstanding** between the scientific community and ICT stakeholders.

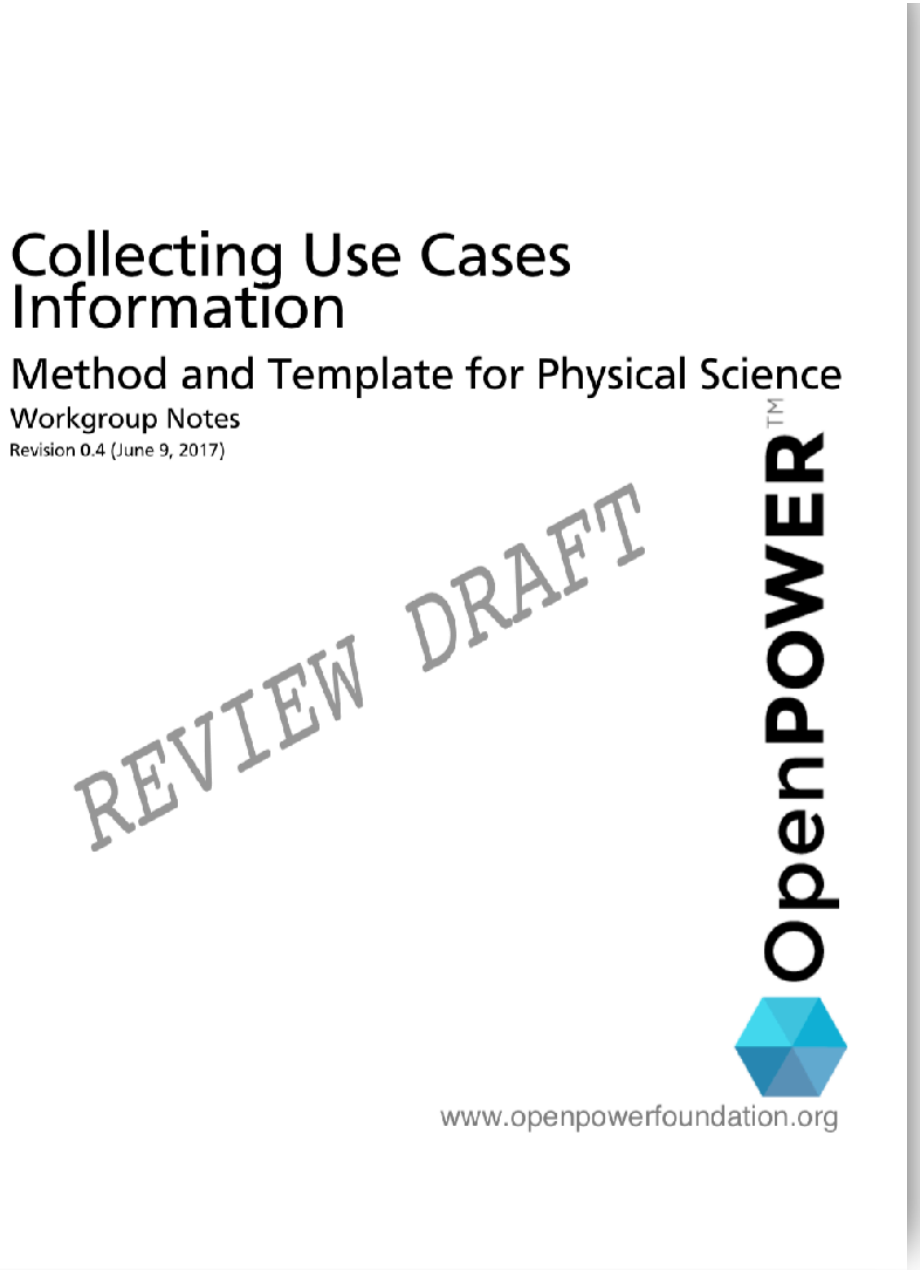

FIRST NOTE/2

- The **stakeholders** are:
  - **physical Science community**: stakeholders that are actively involved in the scientific project, are affected by its outcome or can influence its outcome, and that should provide requirements:
    – scientists;
    – research engineers.
  - **ICT community**: stakeholders that could provide a solution;
  - **WG participants**: members of both communities that should provide support for collecting, analyzing and categorizing requirements.

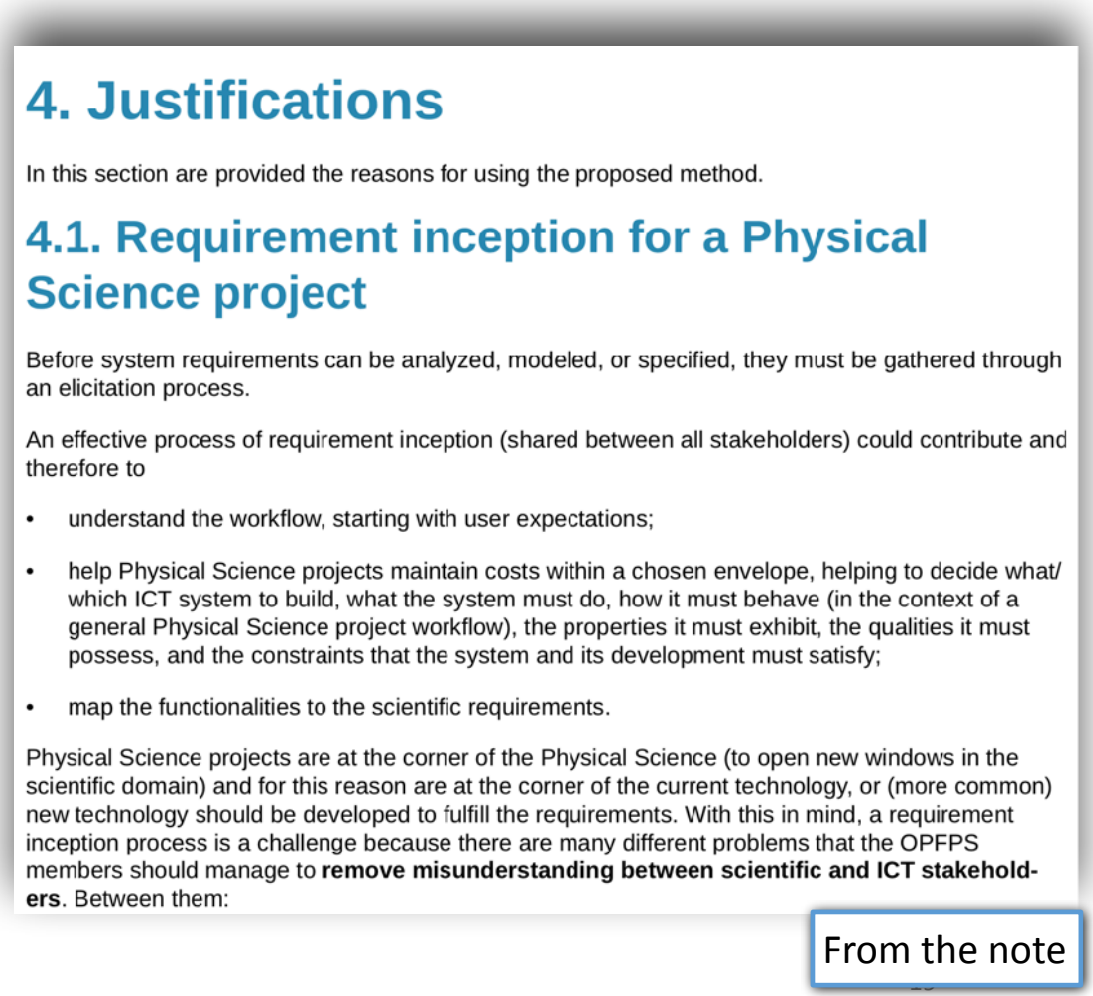

4. Justifications

In this section are provided the reasons for using the proposed method.

4.1. Requirement inception for a Physical Science project

Before system requirements can be analyzed, modeled, or specified, they must be gathered through an elicitation process.

An effective process of requirement inception (shared between all stakeholders) could contribute and therefore to

- understand the workflow, starting with user expectations;
- help Physical Science projects maintain costs within a chosen envelope, helping to decide what/ which ICT system to build, what the system must do, how it must behave (in the context of a general Physical Science project workflow), the properties it must exhibit, the qualities it must possess, and the constraints that the system and its development must satisfy;
- map the functionalities to the scientific requirements.

Physical Science projects are at the corner of the Physical Science (to open new windows in the scientific domain) and for this reason are at the corner of the current technology, or (more common) new technology should be developed to fulfill the requirements. With this in mind, a requirement inception process is a challenge because there are many different problems that the OPFPS members should manage to **remove misunderstanding between scientific and ICT stakeholders**. Between them:

From the note

METHOD: STEP 1

- **Description of the scientific problem and related scientific requirements**. In this context, the WG asks scientific stakeholders to provide an overview/**summary of their problem** (with a short description of the scientific objectives), and to **provide scientific requirements** that could be **measurable**; this helps to identify the scientific requirements for which a possible hardware/ software solution has the greatest impact. **In this way the users will understand the real value of the proposed solution;**

Scientific problem analysis	
Name of the **scientific project**	
Scientific **objectives**	Write some short statements about the objectives of the scientific project for which the scientific stakeholders want to resolve a problem. This section is useful to put the scientific project under analysis in the general context. Define objectives, not requirements.
Scientific requirements and problem definition	Write a **scientific problem statement**. The statement should include: • What is the problem for the scientific project? • Who is affected by it? • What is the impact of this problem on the scientific project? Identify here the **scientific requirements** that the solution can contribute to fulfill. The scientific requirements should be **measurable.** *The problem of …* *affects …* *and results in … (list scientific requirements that are not fulfilled/ that is too expensive to fulfill/ that must be fulfilled)*
Known **challenging** and	Sometimes there are challenges and risks behind the problem. It could be important to find factors that

METHOD: STEP 2

- **Definition of use cases and identification of system requirements (functional and non- functional)** that are used as input to define a solution.

System requirements of the solution (a product/hardware-software solution)	
Outline of the **scope**	Define solution system boundaries, e.g. with a context diagram or in a textual form. Details are in the following sections of the table.
Benefits of a **solution**	The resolution of these problems is also the resolution of scientific requirements, i.e. the increasing of real value for the scientific community. This contributes to define new measurable to compare different solutions. Some questions: • Is there a proposed or a current (in-operation) solution? • Why the current (if exists)/proposed solution is it not profitable? Poor design? Too expensive? • What are key benefits to find a new solution for this problem? Example. *The current solution is... (if present). Benefits of a solution/new solution that create a new system to address the problem include: ... The solution could have a direct impact on the following scientific requirements: ...*
Available technologies of the context	List currently available technologies used in the context of the considered scientific problem.
Assumptions	Assumptions made
Description of the **workflow**/main functional and non-functional **requirements**	It is possible to describe the solution in term of a workflow, a list of functional requirement (list the major features or user capabilities unique to the product) or a use case and a list of non-functional requirements. This Note does not specify "how to" write requirements.

COLLECTING USE CASES

- First use case provide by OpenLAB for geantV
- Planck and AGILE use cases in progress. Other use cases proposed.

Scientific problem analysis	
Name of the **scientific project**	GEANTV
Scientific **objectives** (*Write some short statements about the objectives of the scientific project for which the scientific stakeholders want to resolve a problem. This section is useful to put the scientific project under analysis in the general context. Define objectives, not requirements.*)	*The goal of the GeantV project is the development of a community supported, open-source, generic, particle transport code for HEP (High Energy Physics) integrating both detailed and fast simulation models that fully exploits advanced features of modern CPU architectures, including CPU accelerators.* *The GeantV initiative sets ahead a roadmap starting with a sustained R&D to understand feasibility and limitations of applying from coarse to fine-grained parallelism and other optimisation techniques to particle transport simulation. The main goal of the project is to deliver a production-quality program that improves the current single thread performance of the simulation by a significant factor (between 2 and 5). The program of work aims at producing a usable beta version by 2019 (Run 3) provided the required resources are available and considering the R&D work needed considering the novelty of the techniques used and the complexity of the problem to solve.*
Scientific requirements and problem definition *Write a **scientific***	*The most demanding High Energy Physics applications in terms of computing is the simulation of the passage of particles in matter (transport) by Monte Carlo methods in order to simulate the detector response. This application is indispensable for the design of the detector in the planning phase and for evaluating its efficiency during data taking. Simulation*

ACTIVITIES ON POWER SYSTEM

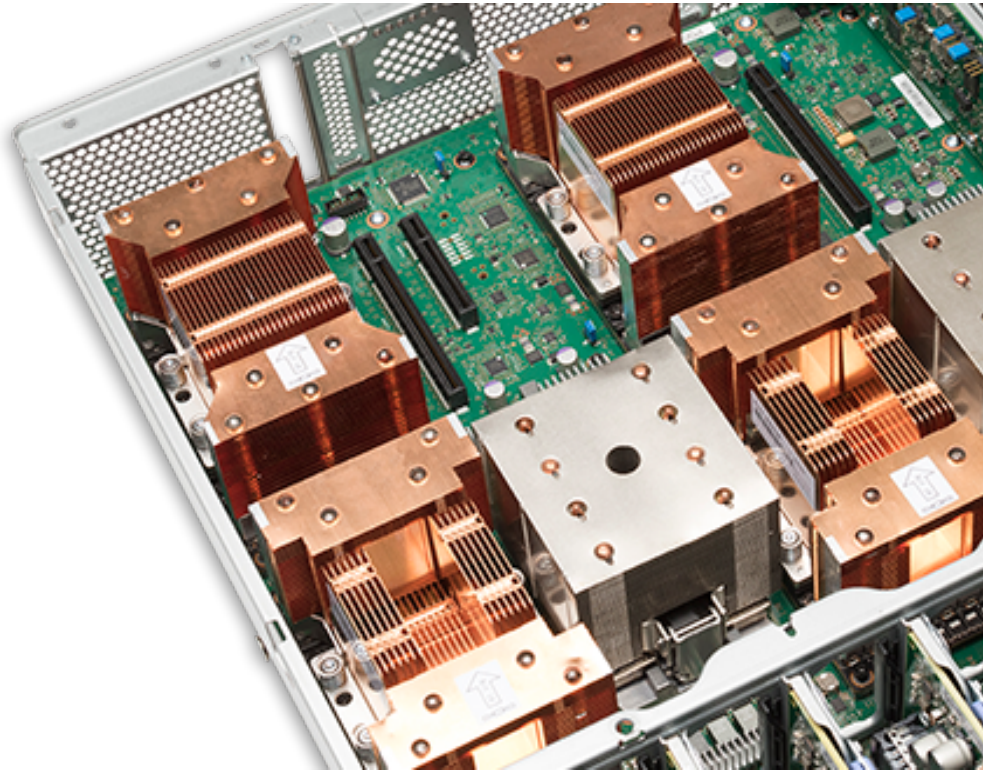

AVAILABLE POWER MACHINES

- **IBM S824L** (Power8) + GPU NVIDIA K40 + Nallatech FPGA
  - 196 CPUs (2 sockets)
- **IBM S822LC** for HPC (Power8+ with NVlink) + 2 GPU NVIDIA K100
  - Ubuntu 16.04.3 LTS
  - 128 CPUs (2 sockets)
  - in comodato d'uso gratuito
- **The following activities will be reported to the OPFPS: some of theme will be delivered as official OPF Notes**
- **For INAF: we are thinking how to share these resources**

Porting of GEANTV on Power with OpenLAB

- Collaboration with CERN/OpenLAB (U. Becciani) **in the context of the INAF/OpenLAB MoU** (see Sofia Vallecorsa talk)
- Work in progress: master thesis of Leonardo Baroncelli
- Collaboration with IBM Labs on Altivec (to be defined)
- "Scalar" porting almost finished. Next task: trying out different compilers, "SIMD" porting with Altivec.

2. AREAS OF COOPERATION

The main goals of the activities in the context of this MoU are supporting the following points:

1. Explore innovative technologies applied to specific products developed at CERN: Geant V, and other tools that eventually will be evaluated by the parties.
2. Cooperate in the porting activities on the platform made available by INAF.
3. Promote and disseminate the joint outcomes of these activities.

Commitments of INAF-ICT:

- Inform the INAF research community about all the material, guides, training courses that may be interesting for this collaboration.
- Promote CERN Openlab in the participation at the OpenPower Foundation.
- Promote the participation of CERN Openlab in the Openpower for Physical Science workgroup leaded by INAF.
- Enable CERN Openlab participation at INAF workshops and meetings as appropriate.
- Cooperate in the porting of relevant programs (e.g. GeantV) in HPC platforms made available to INAF-ICT researchers for this collaboration, and Power8+/Power9 platforms eventually made available at INAF.
- Disseminate the collaboration and promote the CERN Openlab activities as appropriate.

IBM S822LC provided by IBM/E4

MACHINE LEARNING ACTIVITIES

- PowerAI, a tool freely available on Power machines (for S882LC for HPC)
- **Master thesis** of G. Giannella
- **Work in progress** of N. Parmiggiani for AGILE and CTA
- Comparison test with **E4** (planned)

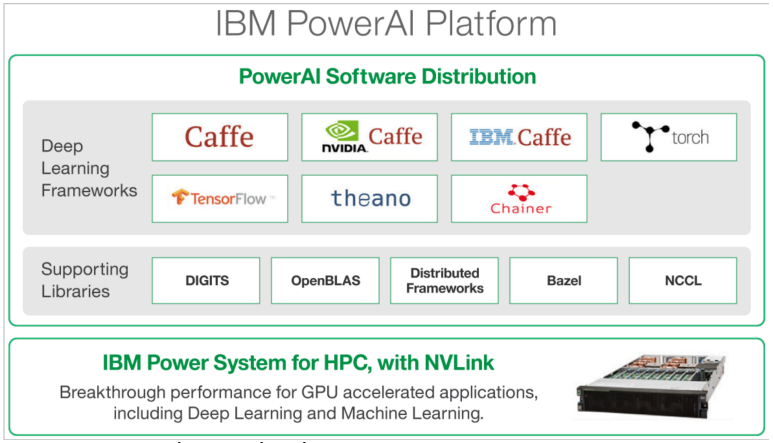

Alma Mater Studiorum – Università di Bologna

Scuola di Scienze
Dipartimento di Fisica e Astronomia
Corso di Laurea magistrale in Astrofisica e Cosmologia

Utilizzo delle Deep Neural Networks per la discriminazione tra eventi gamma e background nella missione spaziale e-ASTROGAM

Tesi di laurea

Presentata da:
Giovanni Giannella

Relatore:
Chiar.mo Prof. Cristian Vignali

Correlatore:
Dott. Andrea Bulgarelli

Dott.ssa Valentina Fioretti

Sessione I
Anno accademico 2016-2017

Porting of ASI/AGILE Science Tools on Power system

- No problems at compile time (gcc 5)
- We have to test xlc and pgi compilers
- Scientific validation on Power
- Simulations of gamma-ray skies

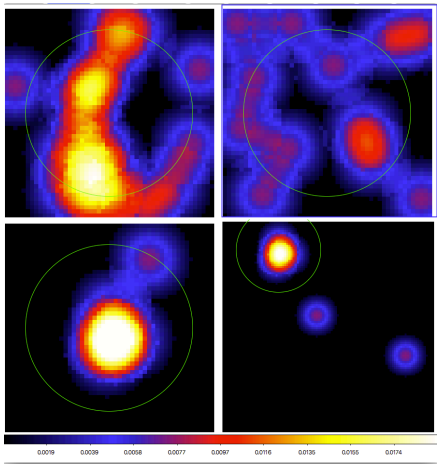

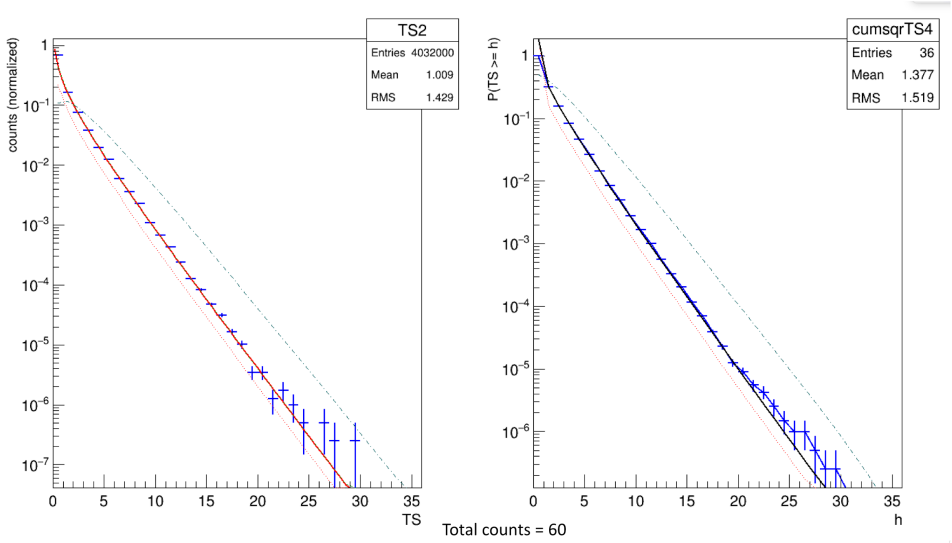

Porting of ESA/Planck software on Power system

- Adriano De Rosa
- Status of the porting 50%
- Performance test will be done with the collaboration of E4

Mixing Intel and Power for data analysis pipeline

- Use case: AGILE real-time analysis pipeline
- We have shared resources between Intel and Power systems with **slurm** resource manager/job scheduler

Jobs | Partitions | Reservations | Burst Buffers | Visible Tabs

JobID	Partition	UserID	Name	State	Time Running	Node Count	NodeList
8335038	large	rt	CI__EB01_FB05_H0025_2AGL0137	RUNNING	01:20:16	1	merlin
8335039	large	rt	CI__EB01_FB05_H0025_2AGL0119	RUNNING	01:20:07	1	merlin
8335040	large	rt	CI__EB01_FB05_H0025_10700I02	RUNNING	01:19:58	1	merlin
8335041	large	rt	CI__EB01_FB05_H0025_2AGL0116	RUNNING	01:19:50	1	merlin
8335042	large	rt	CI__EB01_FB05_H0025_2AGL0187	RUNNING	01:19:38	1	merlin
8335043	large	rt	CI__EB01_FB05_H0025_10338I00	RUNNING	01:19:31	1	merlin
8335044	large	rt	CI__EB01_FB05_H0025_2AGL0055	RUNNING	01:19:22	1	merlin
8335045	large	rt	CI__EB01_FB05_H0025_12639I00	RUNNING	01:19:13	1	merlin
8335046	large	rt	CI__EB01_FB05_H0025_2AGL0080	RUNNING	01:19:04	1	merlin
8335047	large	rt	CI__EB01_FB05_H0025_2AGL0073	RUNNING	01:18:54	1	merlin
8335048	large	rt	CI__EB01_FB05_H0025_2AGL0052	RUNNING	01:18:45	1	merlin
8335049	large	rt	CI__EB01_FB05_H0025_2AGL0014	RUNNING	01:18:35	1	merlin
8335050	large	rt	CI__EB01_FB05_H0025_2AGL0150	RUNNING	01:18:26	1	merlin
8335051	large	rt	CI__EB01_FB05_H0025_2AGL0022	RUNNING	01:18:17	1	merlin
8335052	large	rt	CI__EB01_FB05_H0025_11708I01	RUNNING	01:18:08	1	merlin
8335053	large	rt	CI__EB01_FB05_H0025_2AGL0117	RUNNING	01:17:59	1	merlin
8335054	large	rt	CI__EB01_FB05_H0025_2AGL0173	RUNNING	01:17:49	1	merlin
8335055	large	rt	CI__EB01_FB05_H0025_11601I00	RUNNING	01:17:38	1	merlin

Big Data Analytics with Spark

- We are starting an experience on Big Data Analytics with Spark in the context of the real-time analysis of astrophysical data, e.g. CTA Real-Time Analysis
- Power systems have an optimised version of the framework
- A **dedicated PhD is ongoing with the University Of Modena and Reggio Emilia**

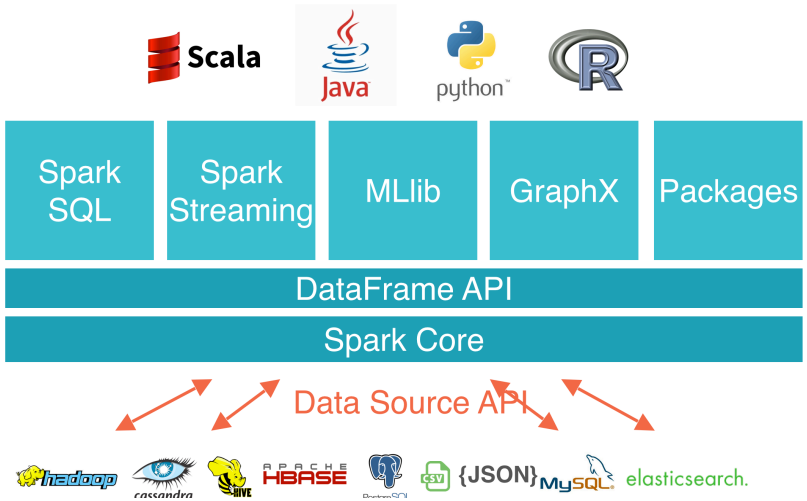

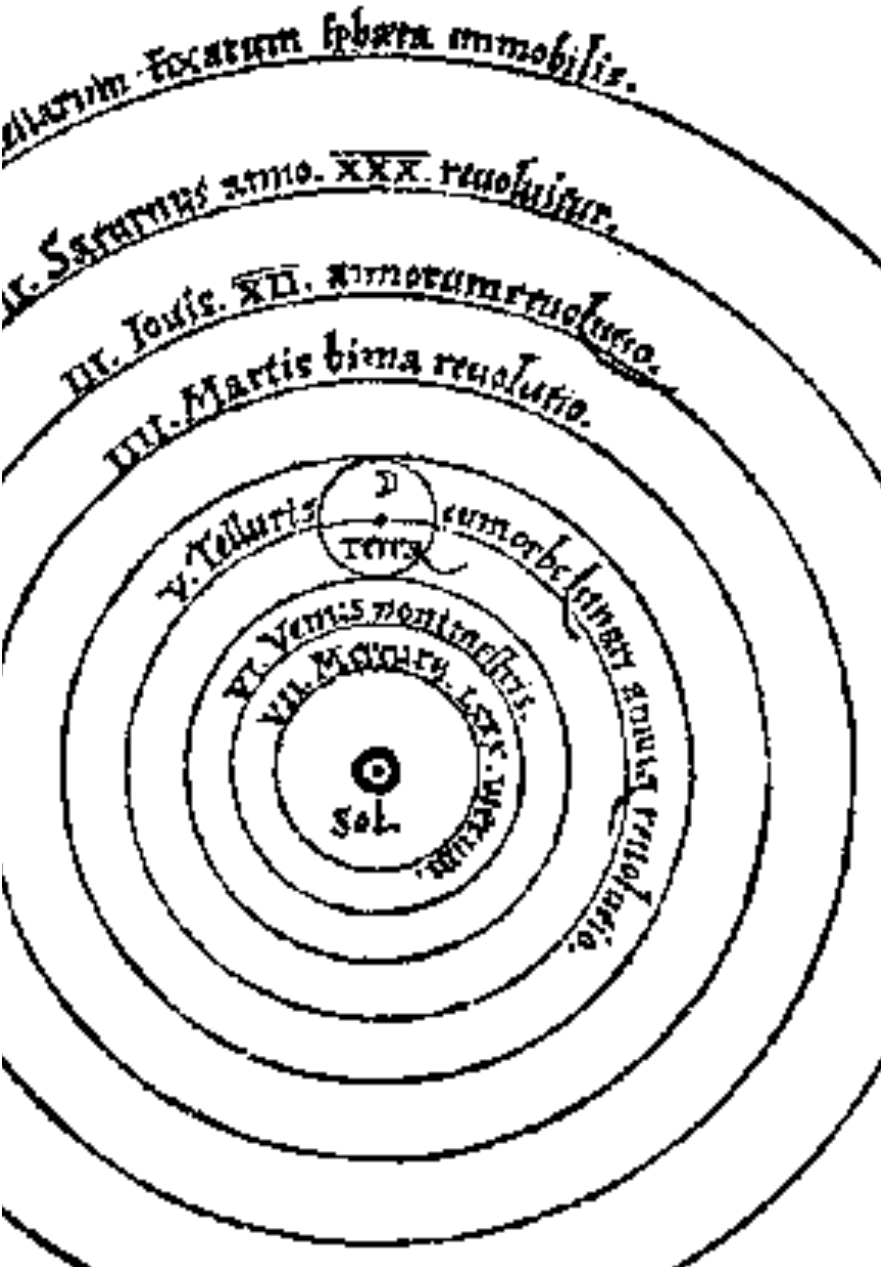

CONCLUSIONS

- Workgroup: the purpose is to **focus OPF members on "Physical Science" use cases**
  - first deliverable: done
  - **We need to collect USE CASES: first remote meeting in December**
- **In parallel, activities on Power system are ongoing**
  - **Porting of scientific code of astrophysical projects**
  - **Machine Learning**
  - **Big Data Analytics**
- **Porting of GeantV in the context of the MoU INAF/OpenLab-CERN**