RED ESPAÑOLA DE SUPERCOMPUTACIÓN

Infraestructuras Científicas y Técnicas Singulares

SPANISH SUPERCOMPUTING NETWORK

RESOURCES AND ACCESS

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Sergi Girona
RES Coordinator

RES: HPC Services for Spain

Infraestructuras Científicas y Técnicas Singulares

- The RES was created in 2006.
- It is coordinated by the Barcelona Supercomputing Center (BSC-CNS).
- It is part of the Spanish “Map of Unique Scientific and Technical Infrastructures” (ICTS).

RES: HPC Services for Spain

RES is made up of 12 institutions and **13 supercomputers**.

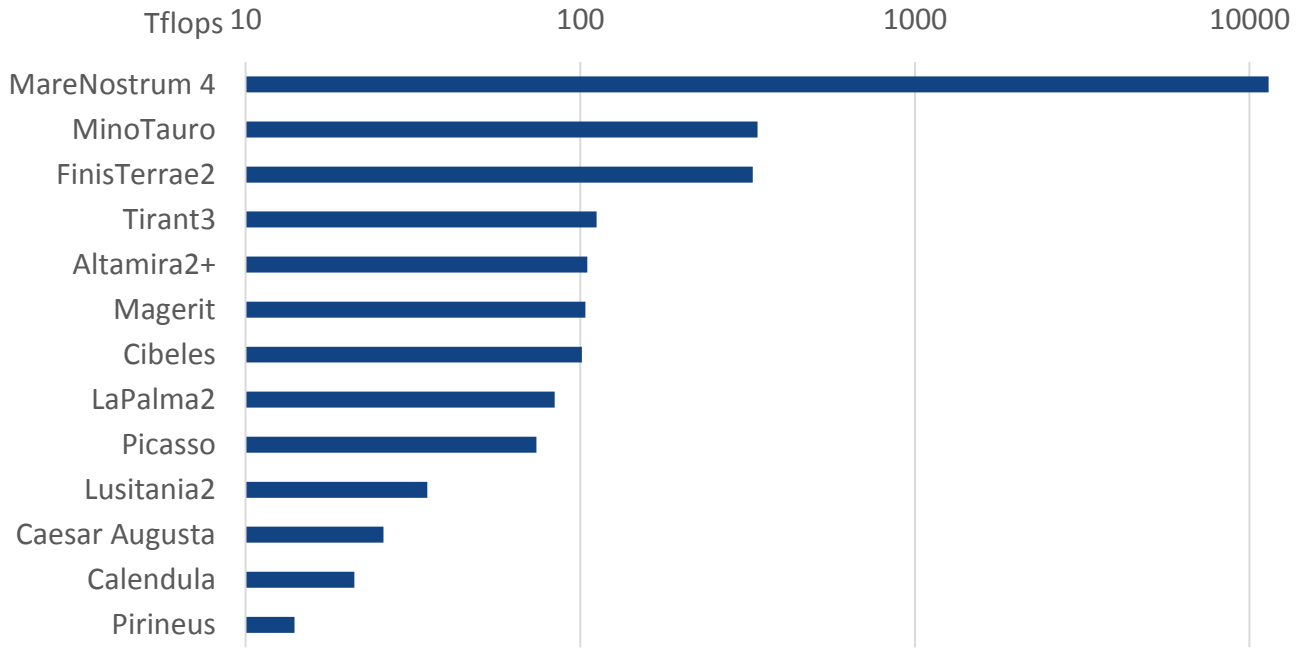

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
CESGA
Centro de Supercomputación de Galicia
Instituto de Astrofísica de Canarias
CSUC
Universidad Autónoma de Madrid
FCSC
Instituto Universitario de Investigación Biocomputación y Física de Sistemas Complejos
Universidad Zaragoza
CénitS
UC Universidad de Cantabria
Universidad de Málaga
Universitat de València
CeSViMa

RES supercomputers

BSC (MareNostrum 4) 165888 cores, 11400 Tflops
Main processors: Intel(R) Xeon(R) Platinum 8160
Memory: 390 TB
Disk: 19 PB

UPM (Magerit II) 3920 cores, 103 Tflops
Main processors : IBM Power7 3.3 GHz
Memory: 7840 GB
Disk: 1728 TB

UMA (Picasso) 4016 cores, 74Tflops
Main processors: Intel SandyBridge-EP E5-2670
Memory: 22400 GB
Disk: 720 TB

UV (Tirant 3) 5376 cores, 111,8 Tflops
Main processors: Intel SandyBridge-EP E5-2670
Memory: 32 GB
Disk: 14 + 10 TB

CSUC (Pirineus) 1344 cores, 14,3 Tflops
Main processors: Intel Xeon X7542 with 6 cores
Memory: 61400 GB
Disk: 112 TB

CénitS (Lusitania 2) 420 cores, 34,89 Tflops
Main processors Intel Sandybridge Xeon
Memory: 10 GB
Disk: 328 TB






RES supercomputers

BSC (MinoTauro) 1300 cores, 339 Tflops
Main processor: 39x 2 Intel Xeon E5-2630 v3, 61x 2 Intel Xeon E5649
Memory: 20 TB

CESGA (FinisTerrae 2) 7712 cores, 328,3Tflops
Main processor: Intel Xeon E5-2680v3
Memory: 40 TB
Disk: 960 TB

UC (Altamira 2+) 5120cores, 105 Tflops
Main processor: Intel SandyBridge
Memory: 15,4 TB

UZ (Caesaraugusta) 3072 cores, 25.8 Tflops
Main processor: AMD Opteron 6272, 2.1 GHz (Interlagos)
Memory: 256 GB RAM memory

FCSCL (Caléndula) 2800 cores, 21,12 Tflops
Main processor: IntelE5450
Memory: 3520 GB
Disk: 6 TB

UAM (Cibeles) 4480 cores, 105 Tflops
Main processor: Intel SandyBridge
Memory: 8960 GB
Disk: 300 TB

IAC (LaPalma) 4032 cores, 83,85 Tflops
Main processor: Intel SandyBrigde
Memory: 8064 GB
Disk: 60 TB



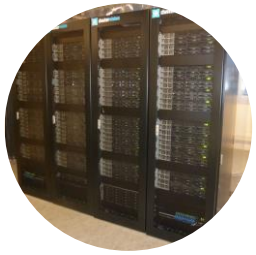


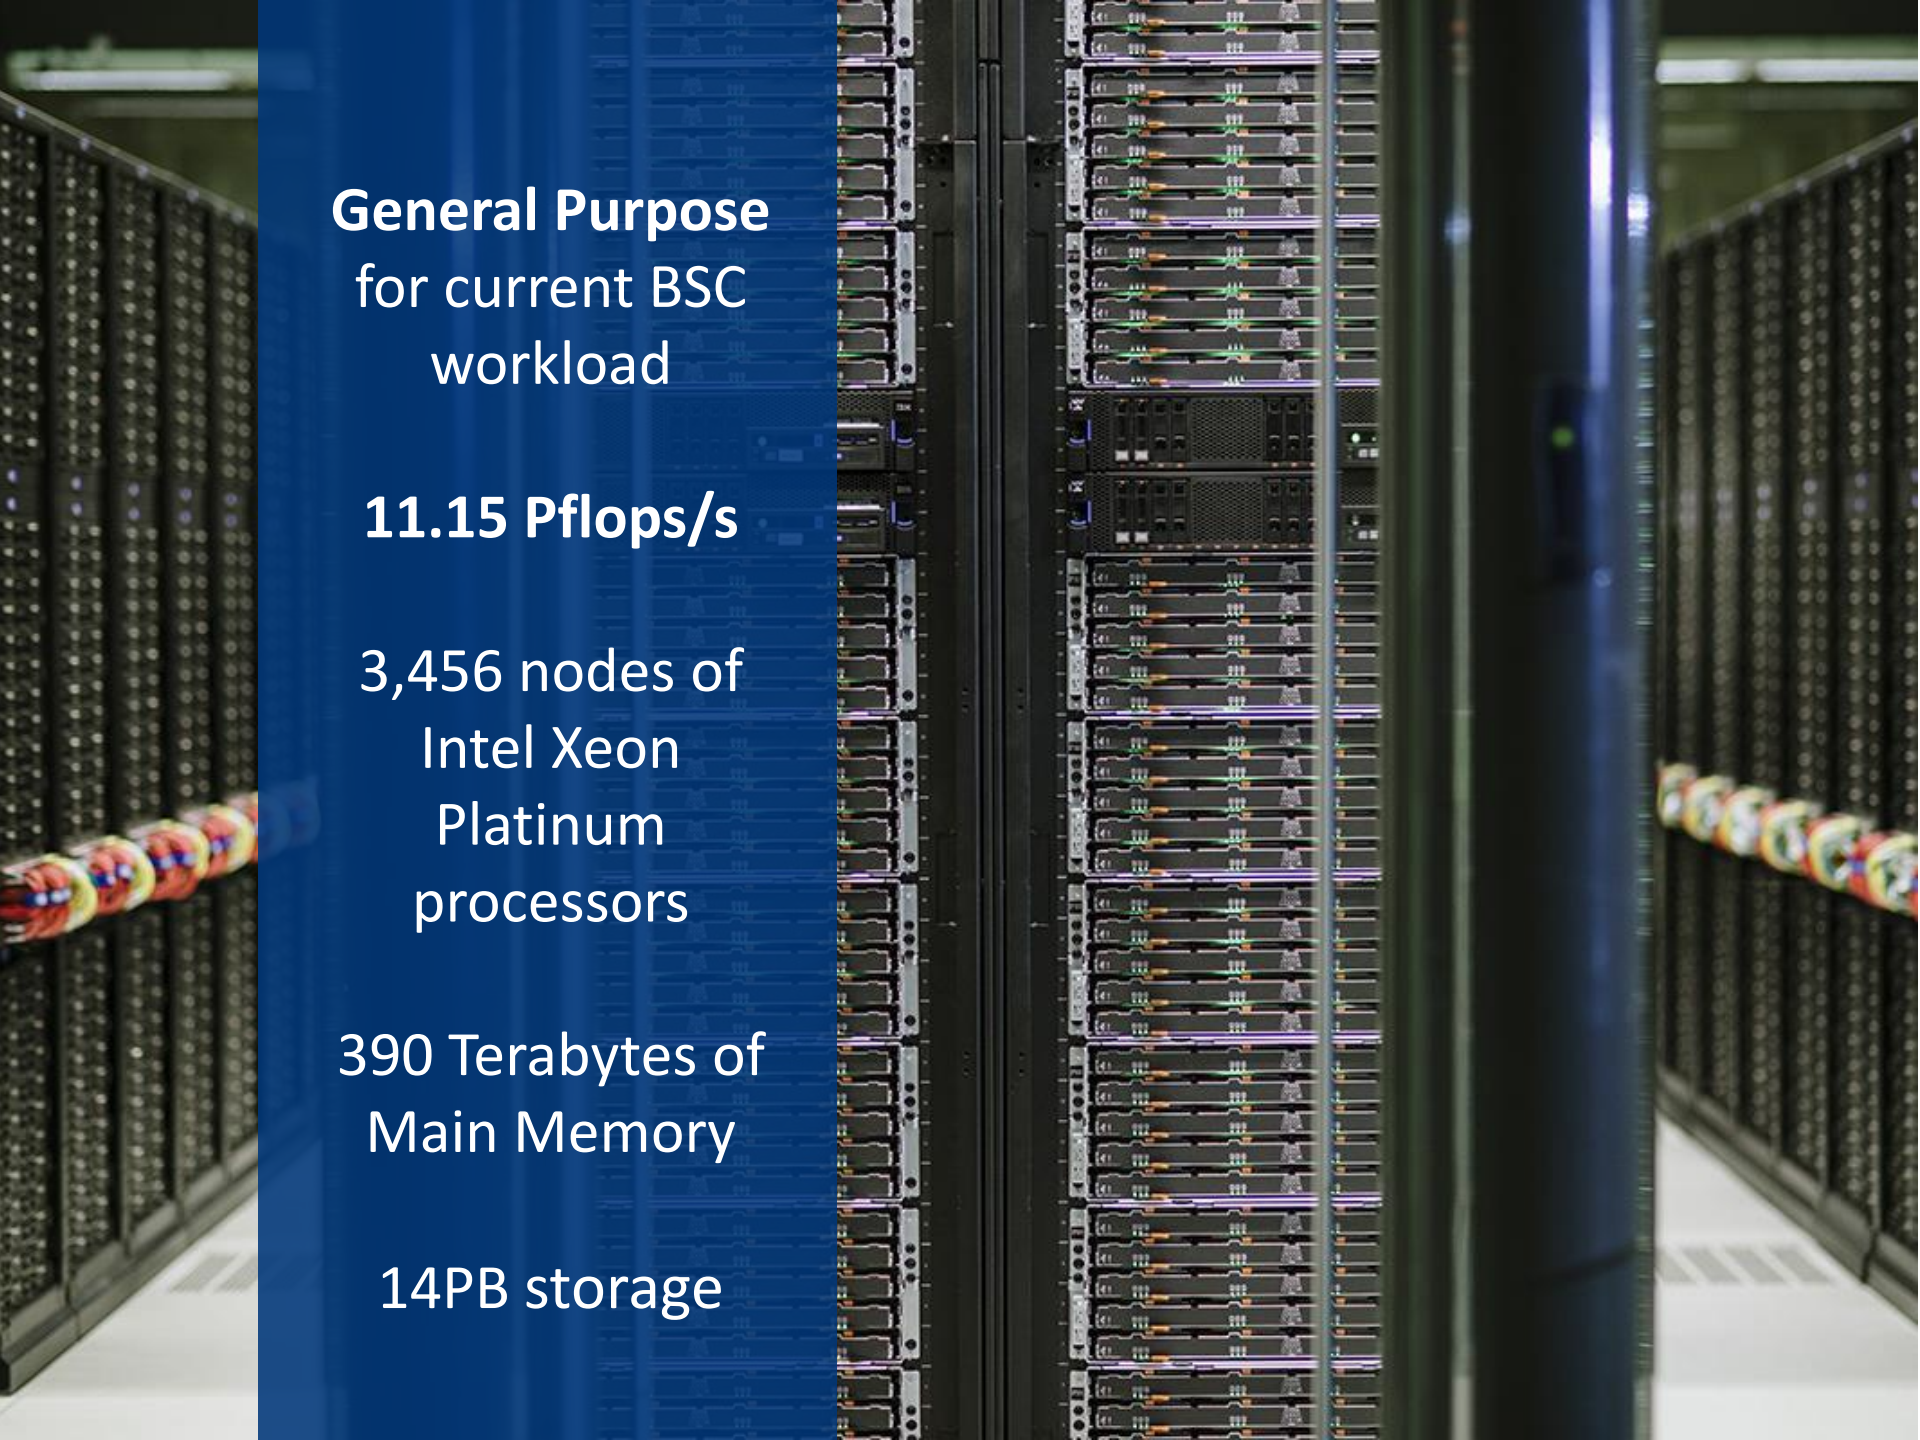

General Purpose for current BSC workload

11.15 Pflops/s

3,456 nodes of Intel Xeon Platinum processors

390 Terabytes of Main Memory

14PB storage

General Purpose for current BSC workload

11.15 Pflops/s

3,456 nodes of Intel Xeon Platinum processors

390 Terabytes of Main Memory

14PB storage

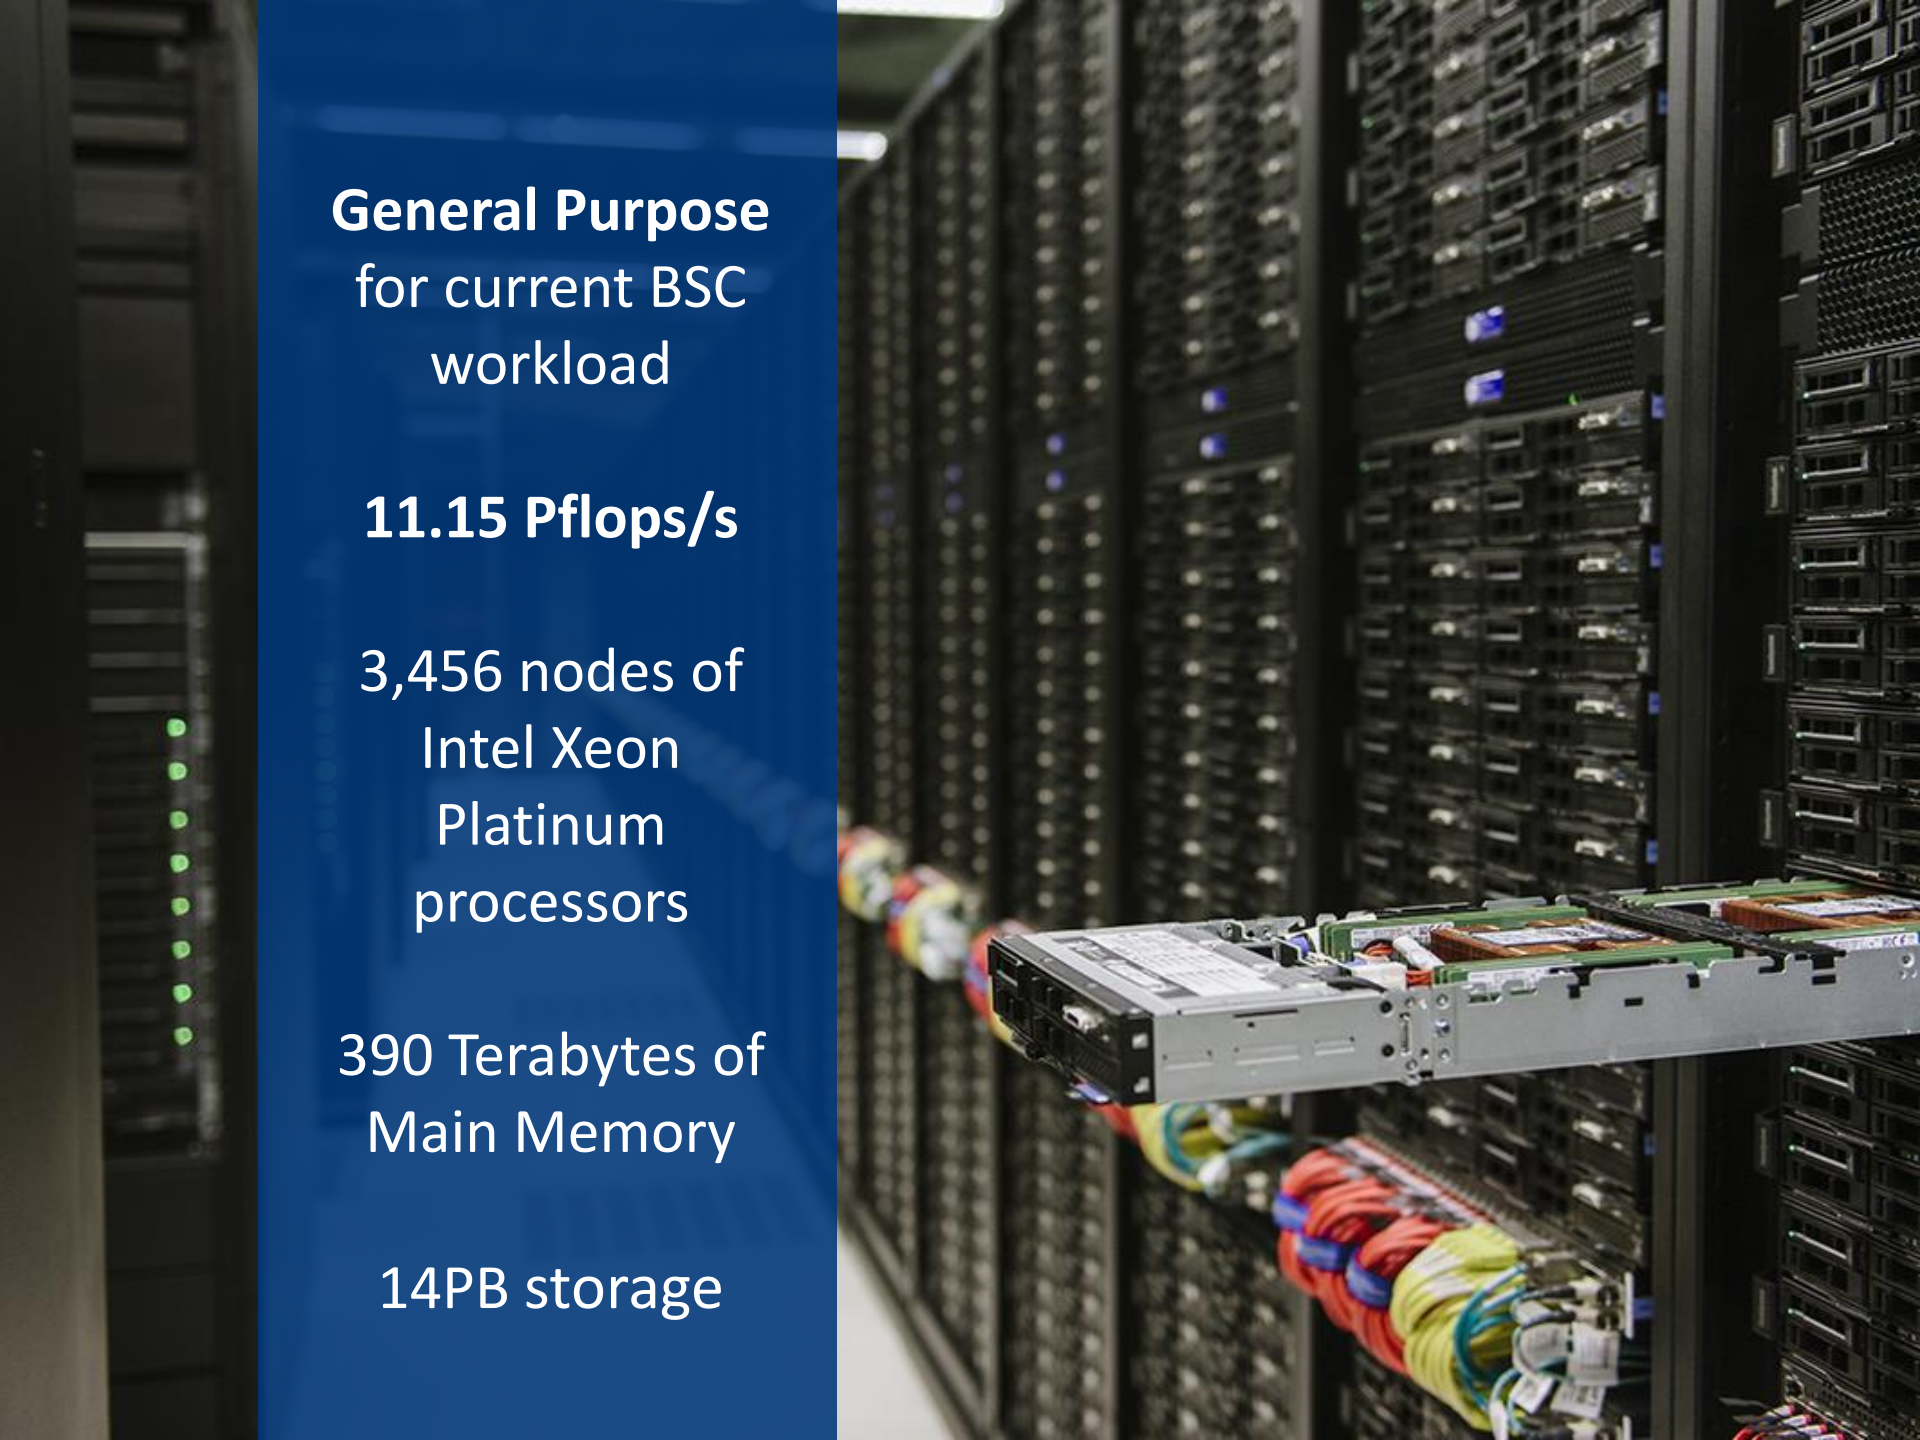

General Purpose for current BSC workload

11.15 Pflops/s

3,456 nodes of Intel Xeon Platinum processors

390 Terabytes of Main Memory

14PB storage

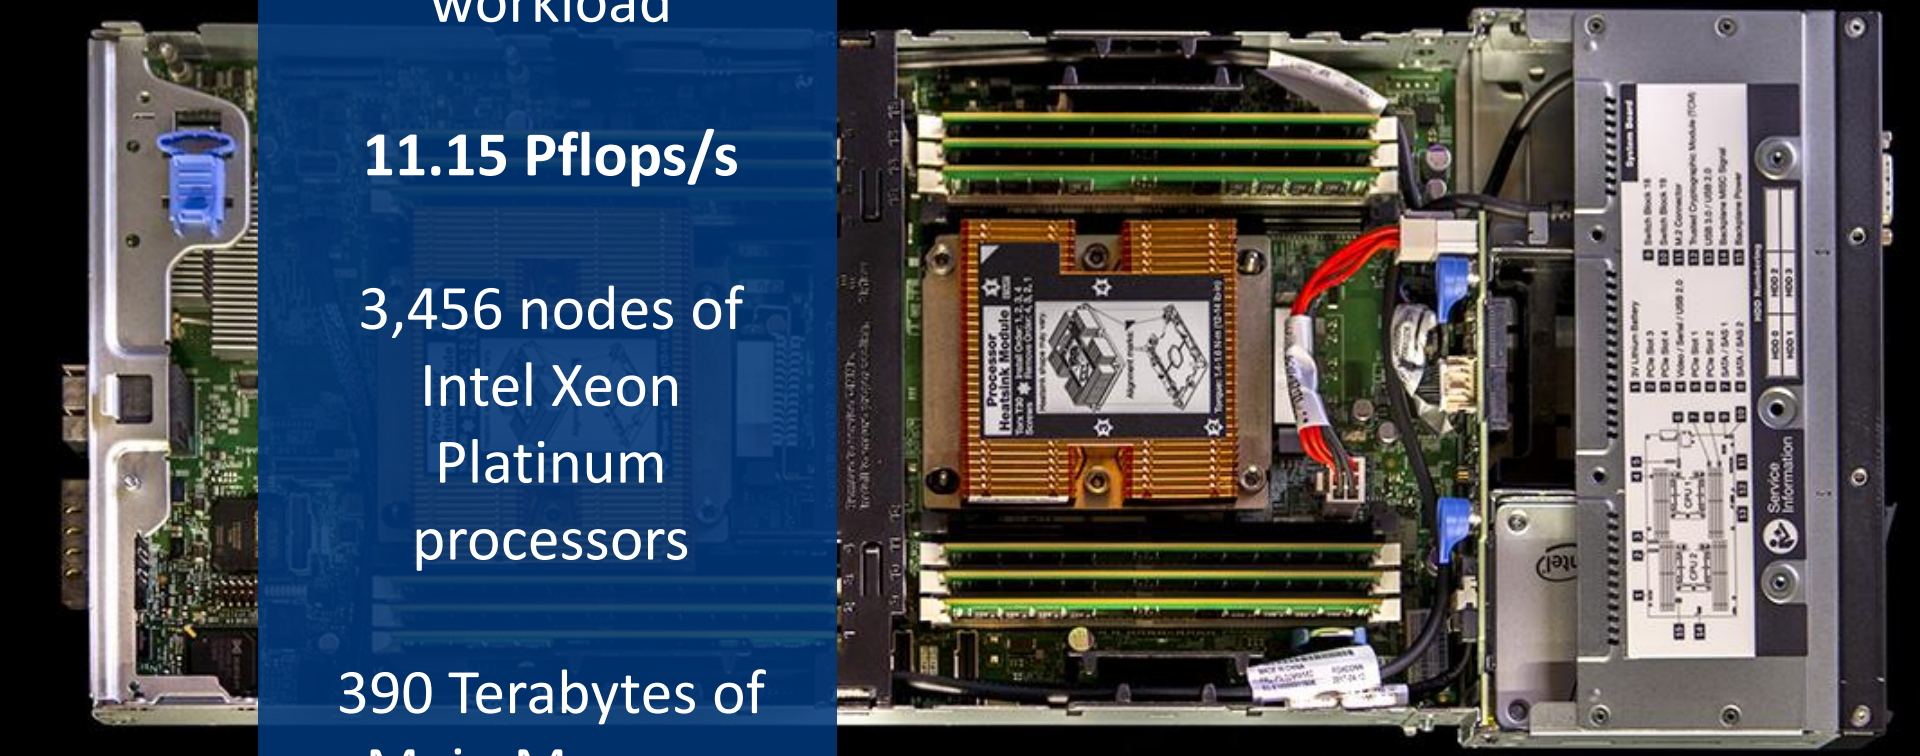


Interconnected
with OmniPath
network


Interconnected
with OmniPath
network

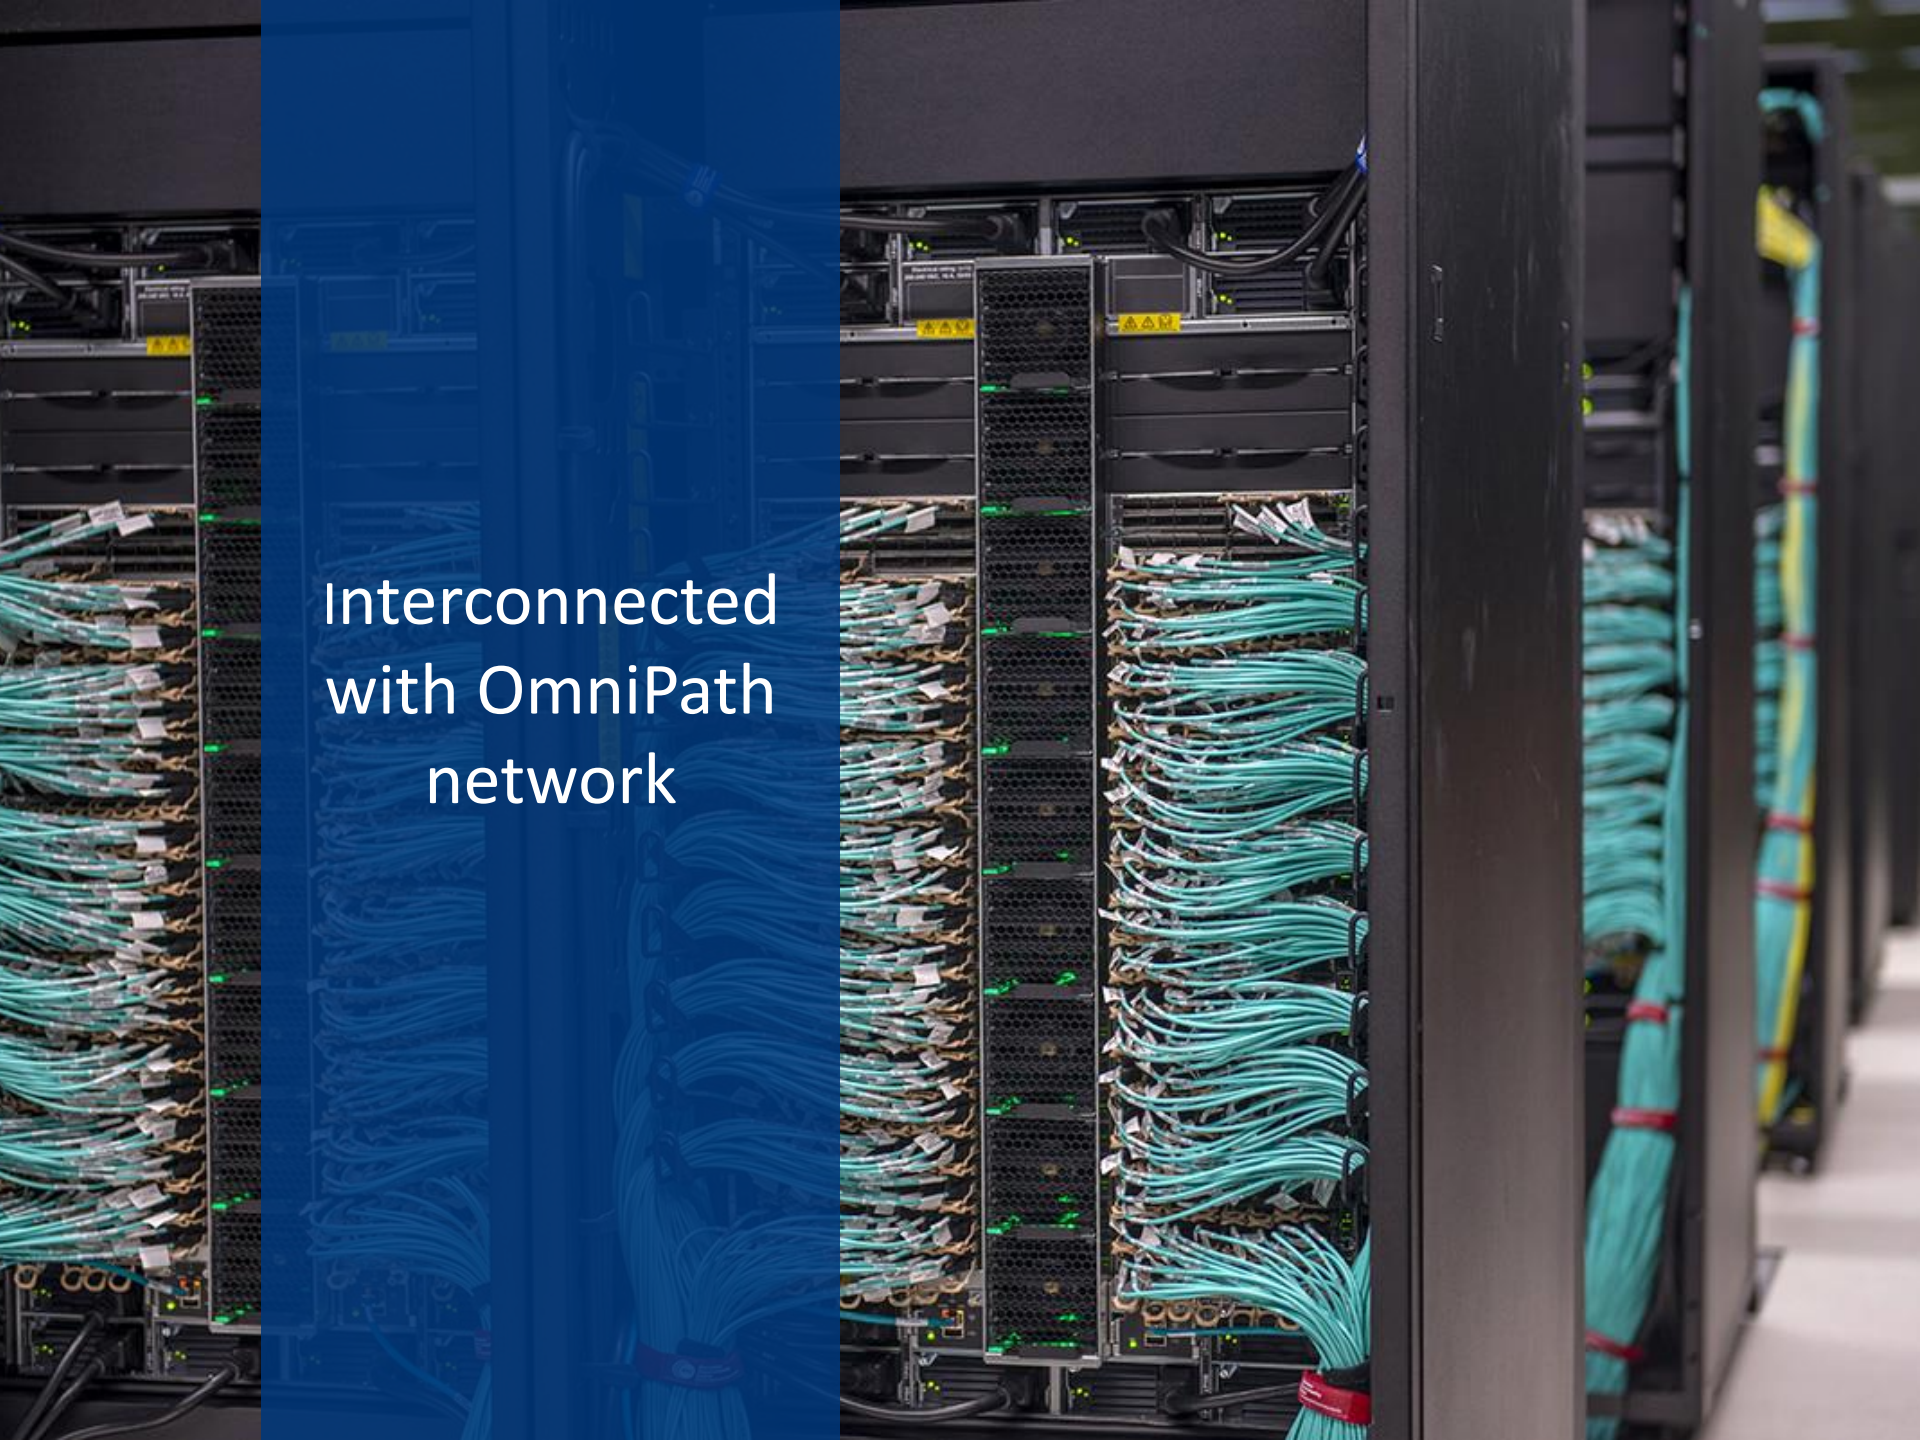
Interconnected
with OmniPath
network

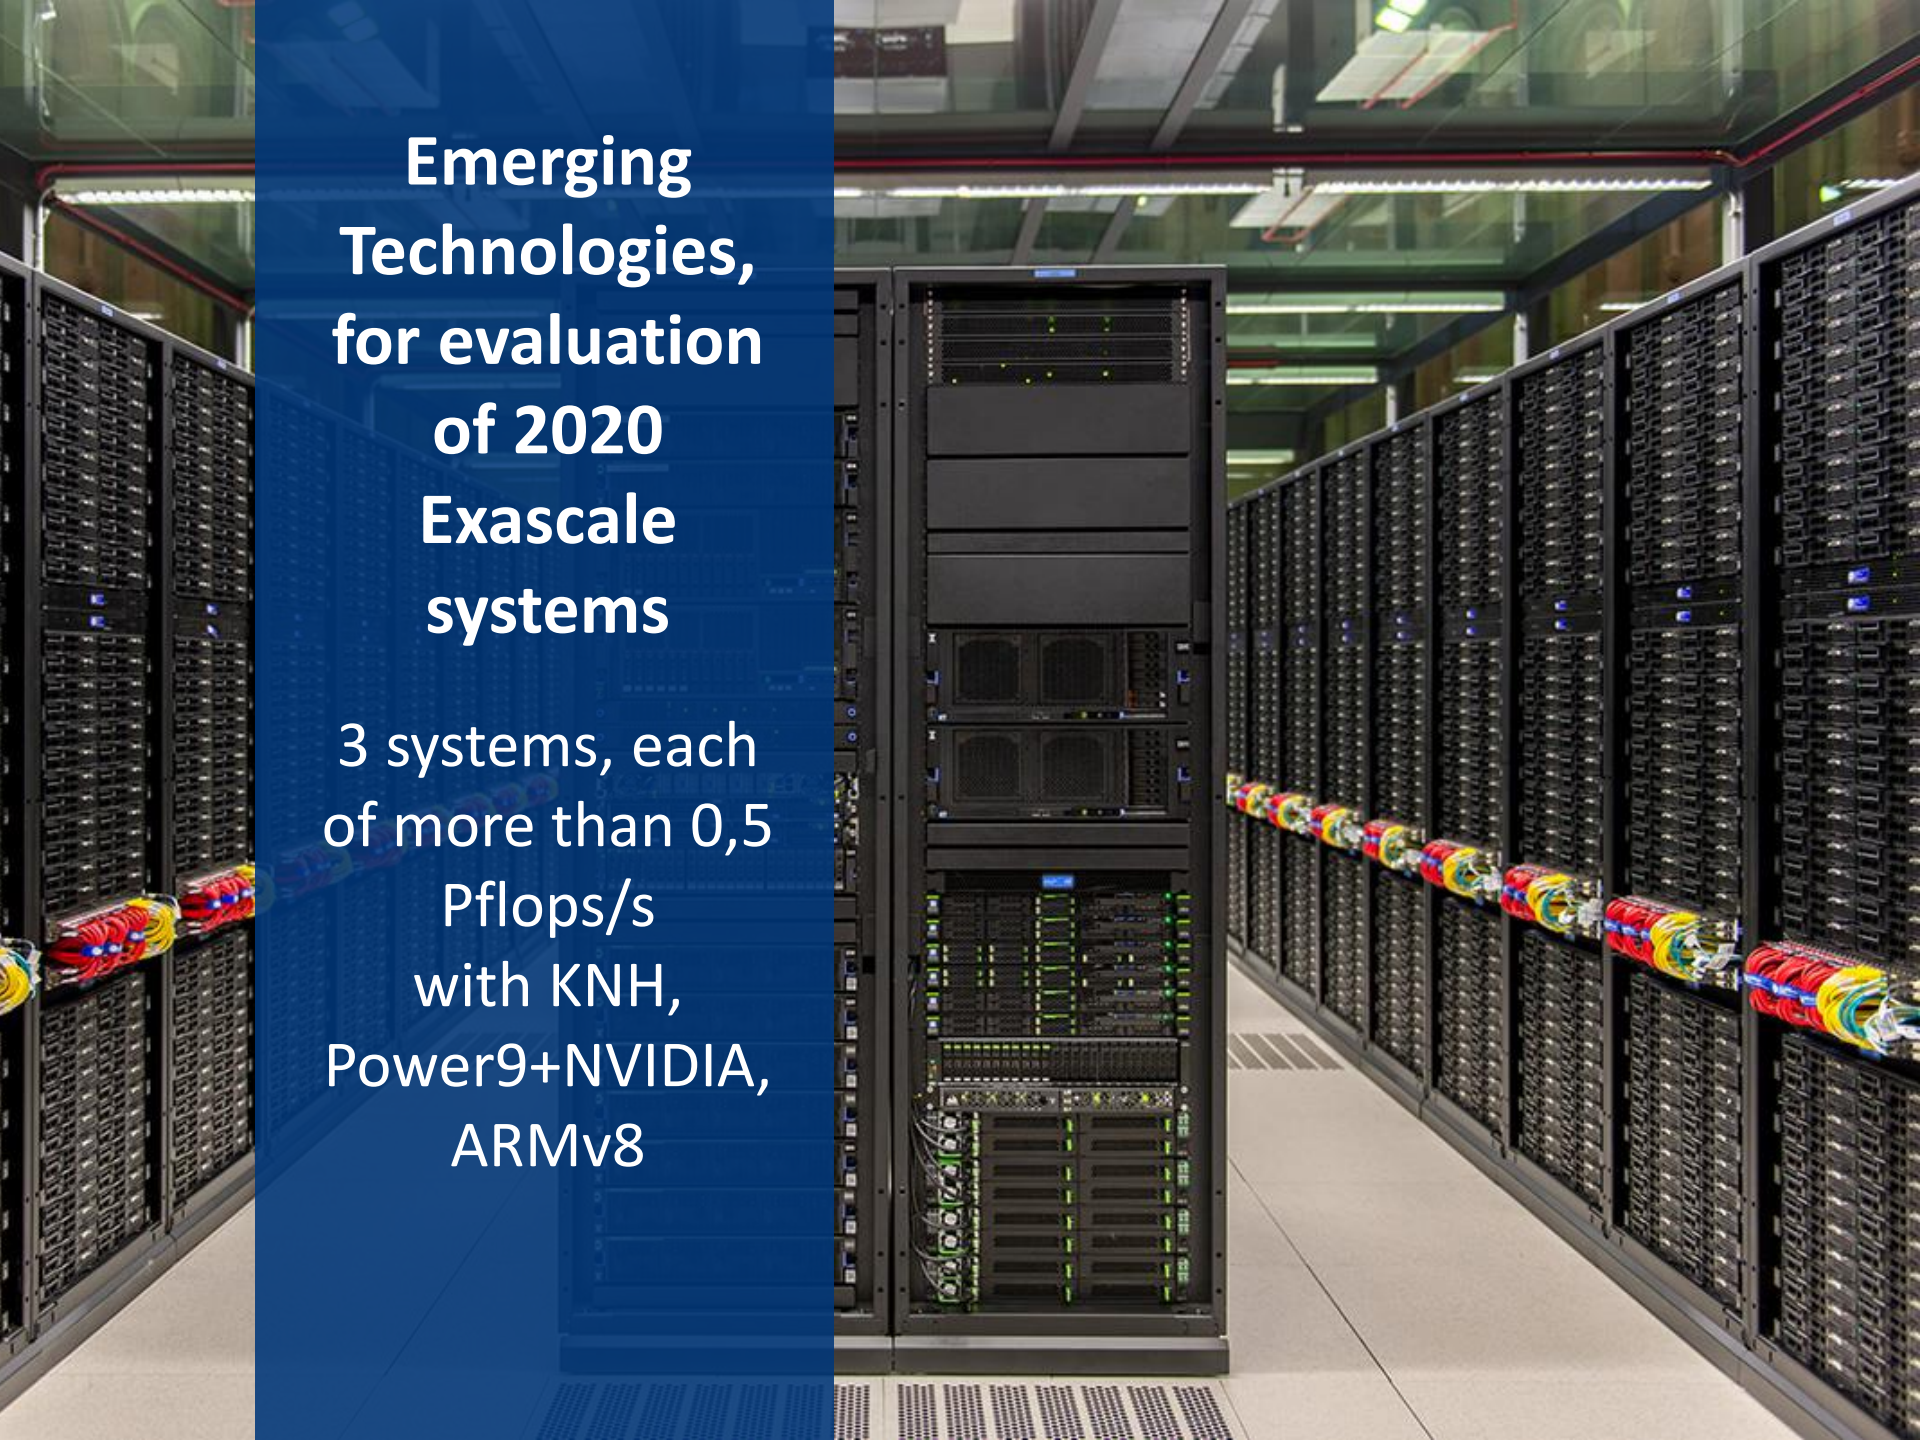

Emerging Technologies, for evaluation of 2020 Exascale systems

3 systems, each of more than 0,5 Pflops/s with KNH, Power9+NVIDIA, ARMv8

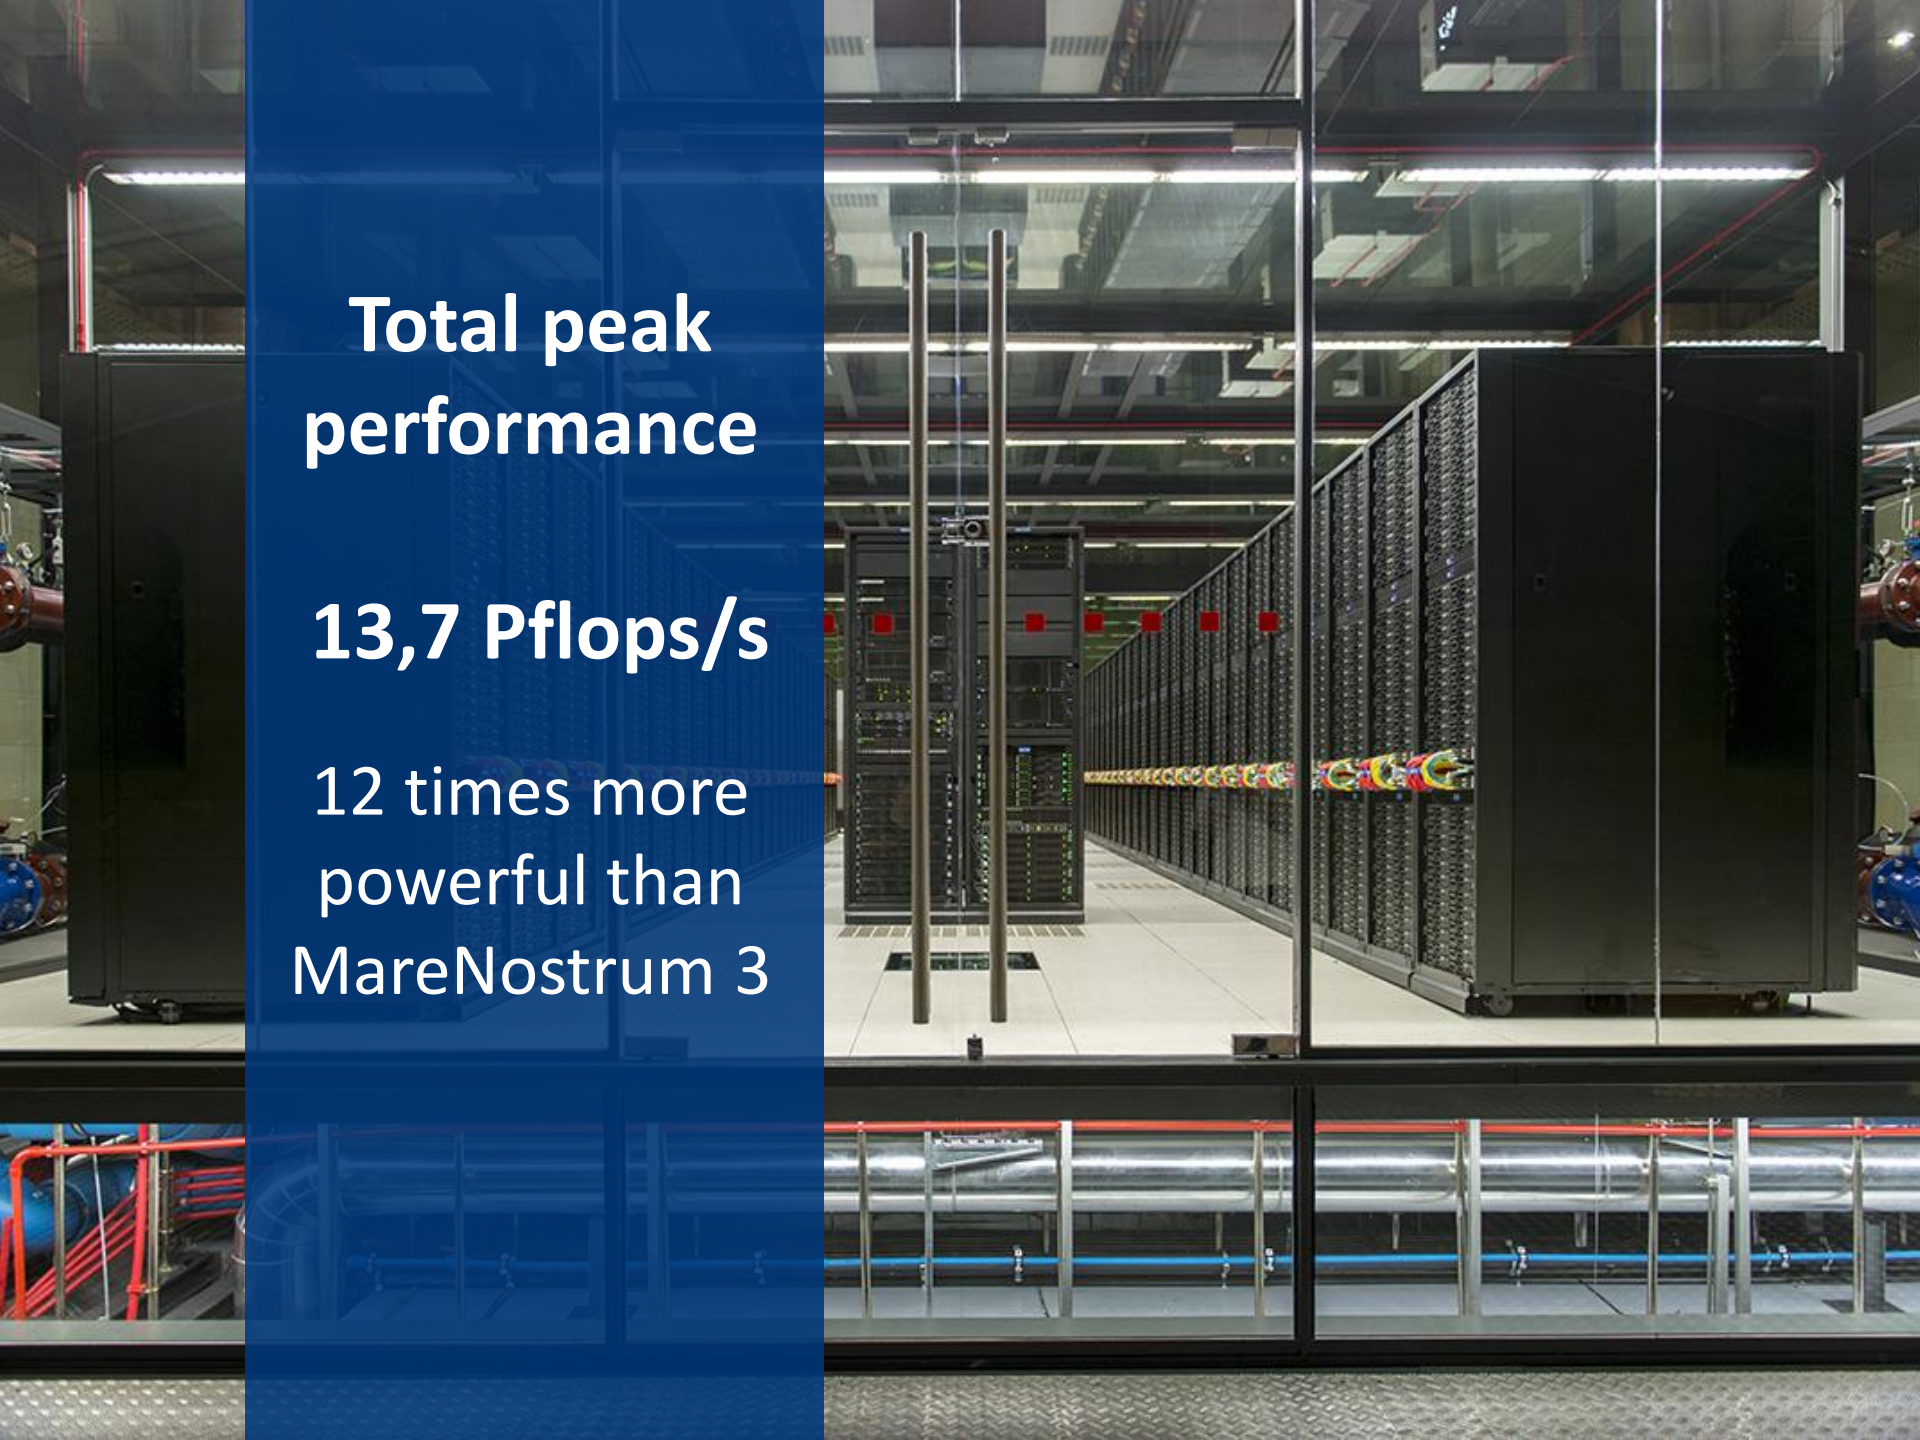

Total peak performance

13,7 Pflops/s

12 times more powerful than MareNostrum 3

RES: Big Data and storage

Storage components in MareNostrum 4:

- **Disk storage capacity of 14 Petabytes**
  - 7 x ESS GL6
    - 2 IBM Power System 822L
    - 6 DCS3700 JBOD expansions
  - 2x Metadata block
    - 2 IBM Power System 822L
    - 2 Flash System V900

Long-term storage in BSC (Active Archive):

- Not directly accessible from HPC Machines, but can be used from any HPC Machine through a batch system:
  - 5.7 PB GPFS Filesystem
  - Robot SL8500 (Tapes): capacity 6 PB

BSC infrastructure roadmap

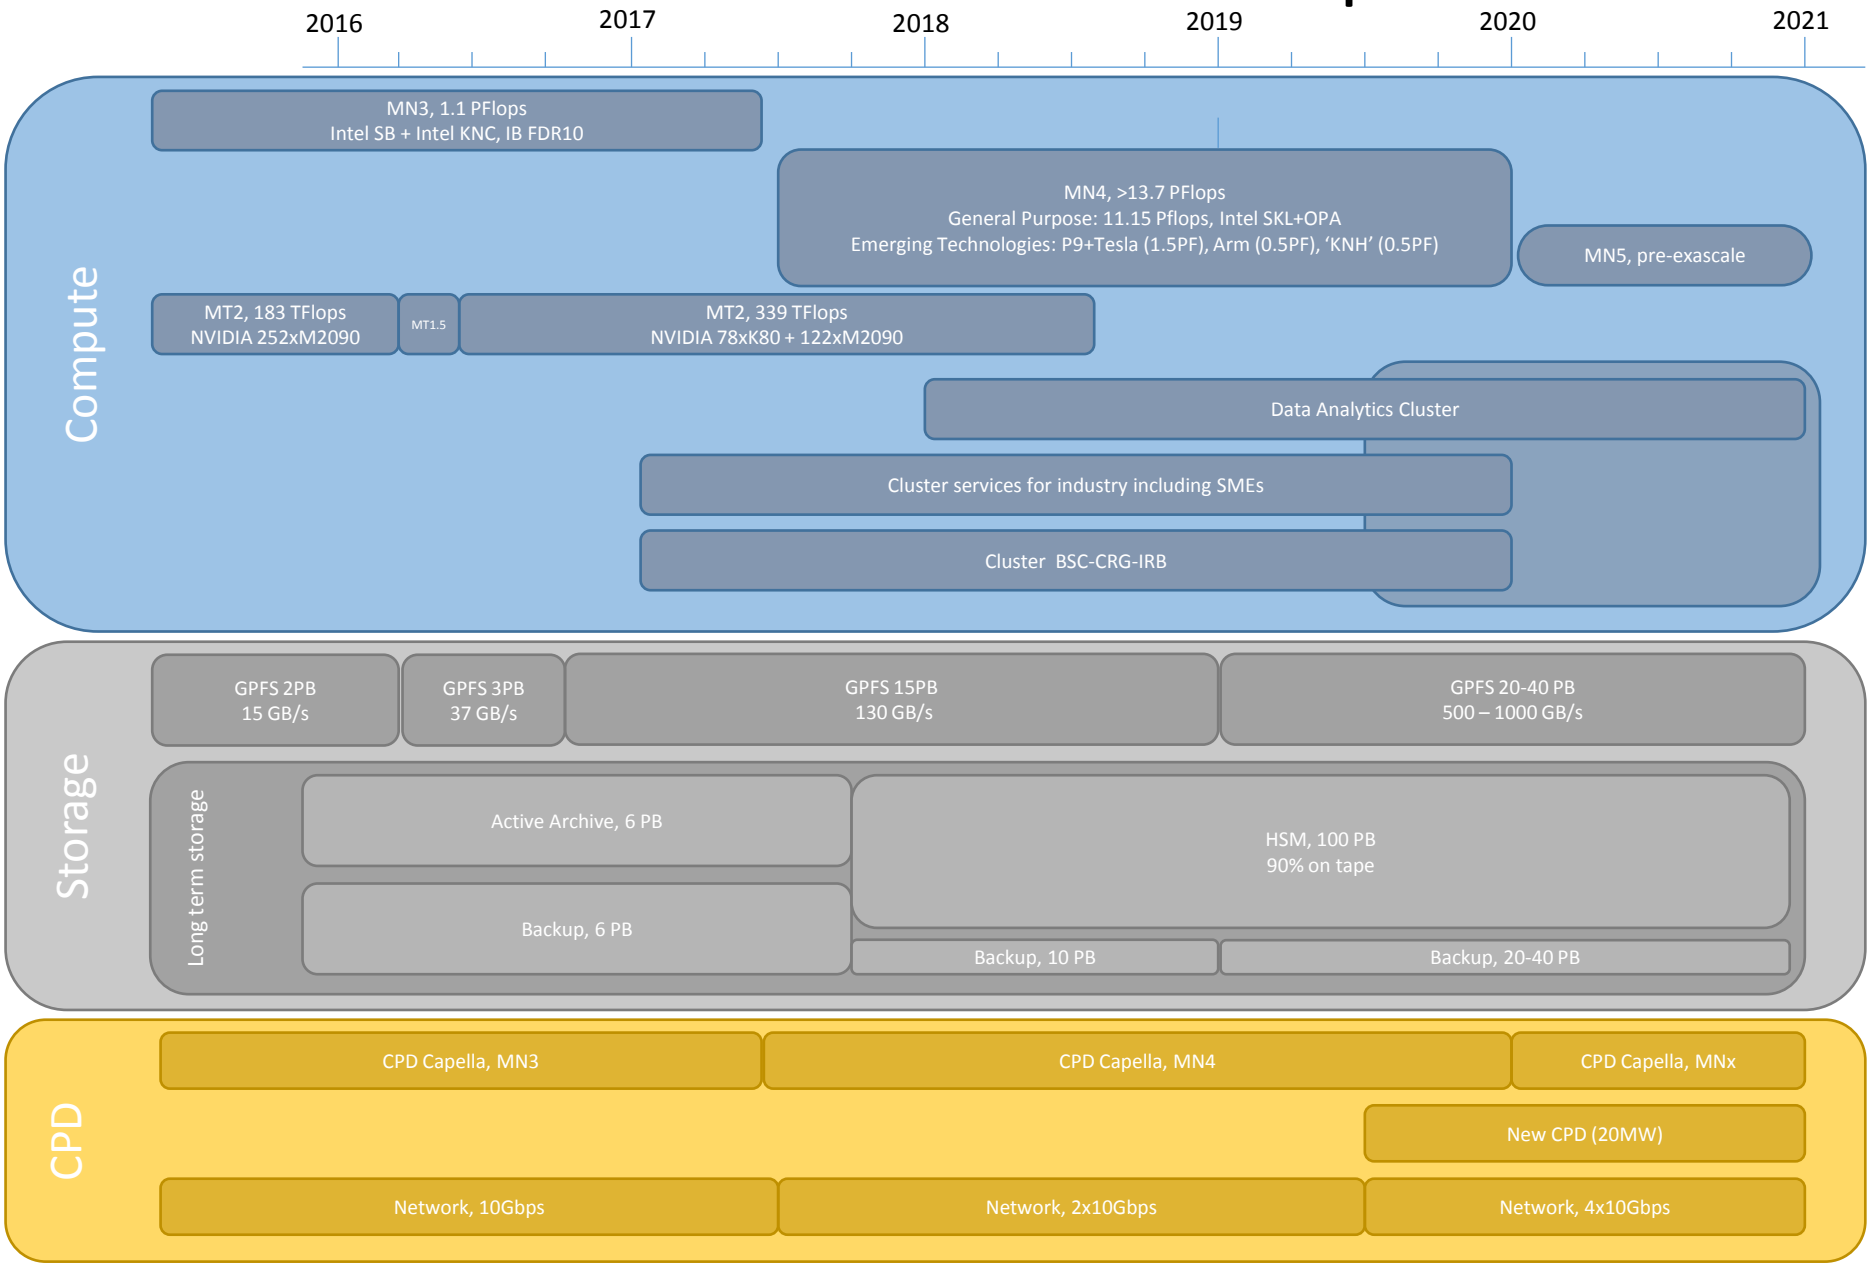

RES: HPC Services for Spain

- **Objective**: manage high performance computing technologies to promote the progress of excellent science and innovation in Spain.
- It offers HPC services for **non-profit R&D** purposes.
- Since 2006, it has granted more than **1,000 Million CPU hours** to 2,473 research activities.

Hours granted per area

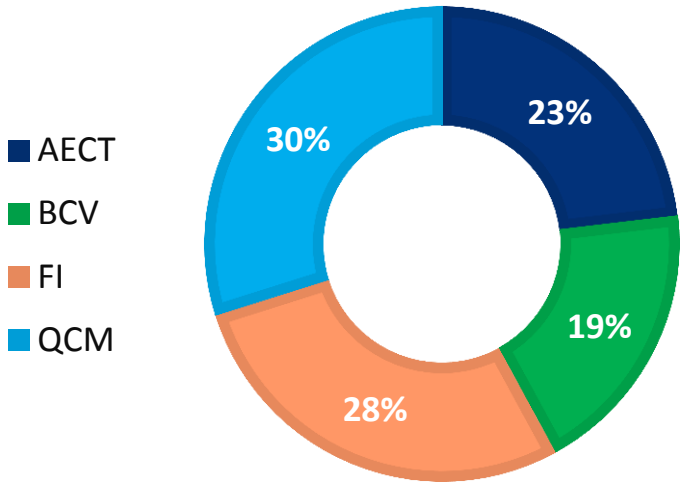

Research areas

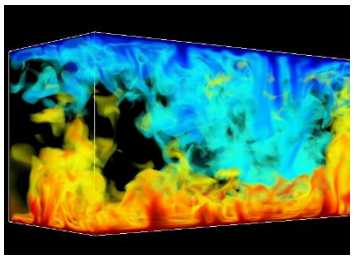

Mathematics, physics and engineering

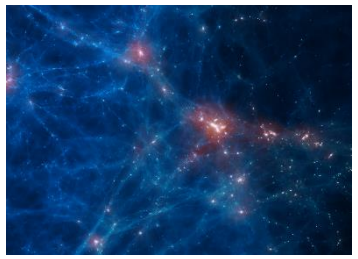

Astronomy, space and earth sciences

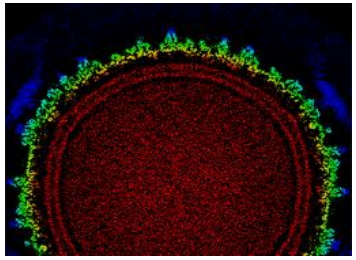

Life and health sciences

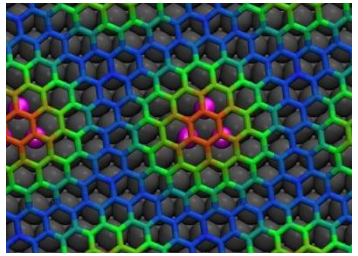

Chemistry and materials sciences

RES: Big Data and storage

RES Working group in Data Services (ongoing):

- **Collect information** about the needs of data services in different scientific areas
- **Identify the available resources** in the RES related to data storage and Big Data
- **Stablish a model** to provide data services: access model, evaluation process, data curation
- **Look for future agreements** and collaborations with other initiatives

How to apply?

- RES resources are **open for Open R&D**:
  - Computing time: CPU hours and local storage
  - Technical support: application analysis, porting of applications, search for the best algorithm… to improve performance and ensure the most effective use of HPC resources.
  - Free of cost at the point of usage
    - Spin-offs free access for 3 years.
- Three **open competitive calls** per year.

Period	Deadline for applications	Starting date
P1	January	1st March
P2	May	1st July
P3	September	1st November

Next deadline: May 2018

Resources granted: computing power

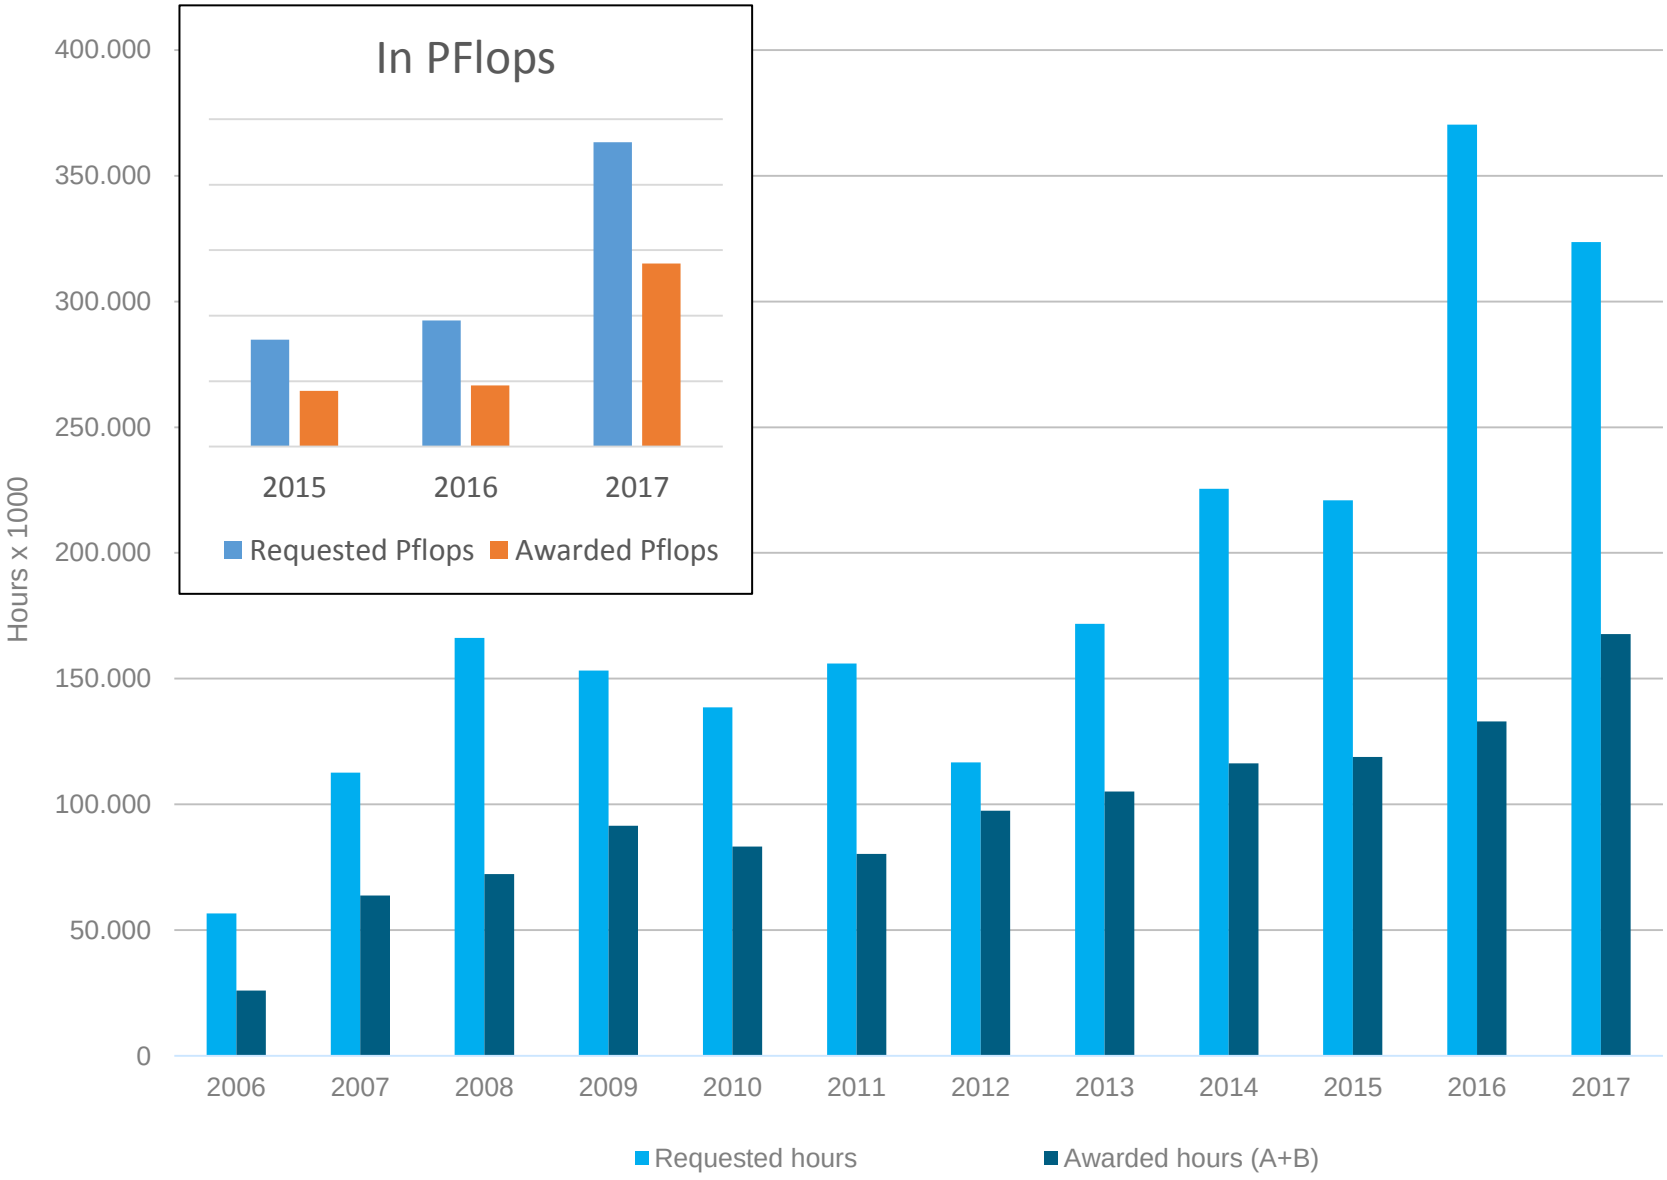

Resources granted: disk storage

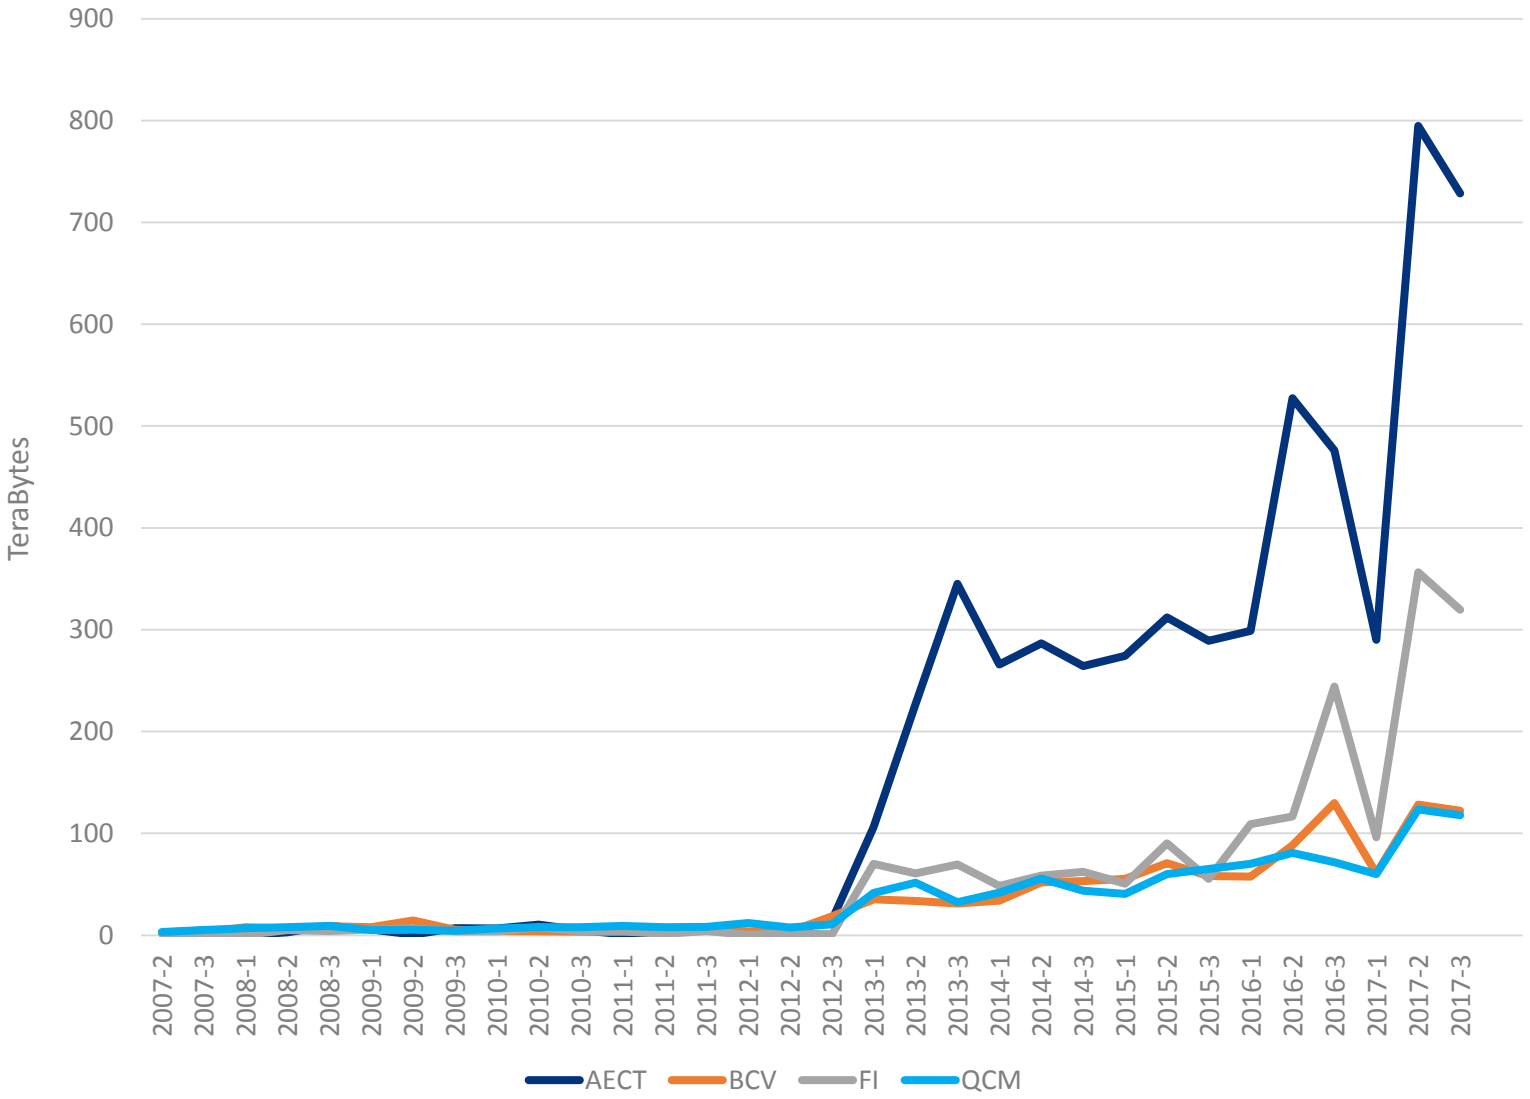

Proposal evaluation

Application

Technical experts panel

- Members appointed by RES nodes

Scientific experts panel

- Astronomy, Space and Earth Sciences
- Life and Health Sciences
- Mathematics, Physics and Engineering
- Chemistry and Materials Science and Technology

- Technical appropriateness to HPC architecture (10%)

- Scientific interest (20%)
- Relevance of calculations in the research project (30%)
- Scientific credentials and experience in HPC (20%)
- Supercomputation needs (20%)

Proposal evaluation

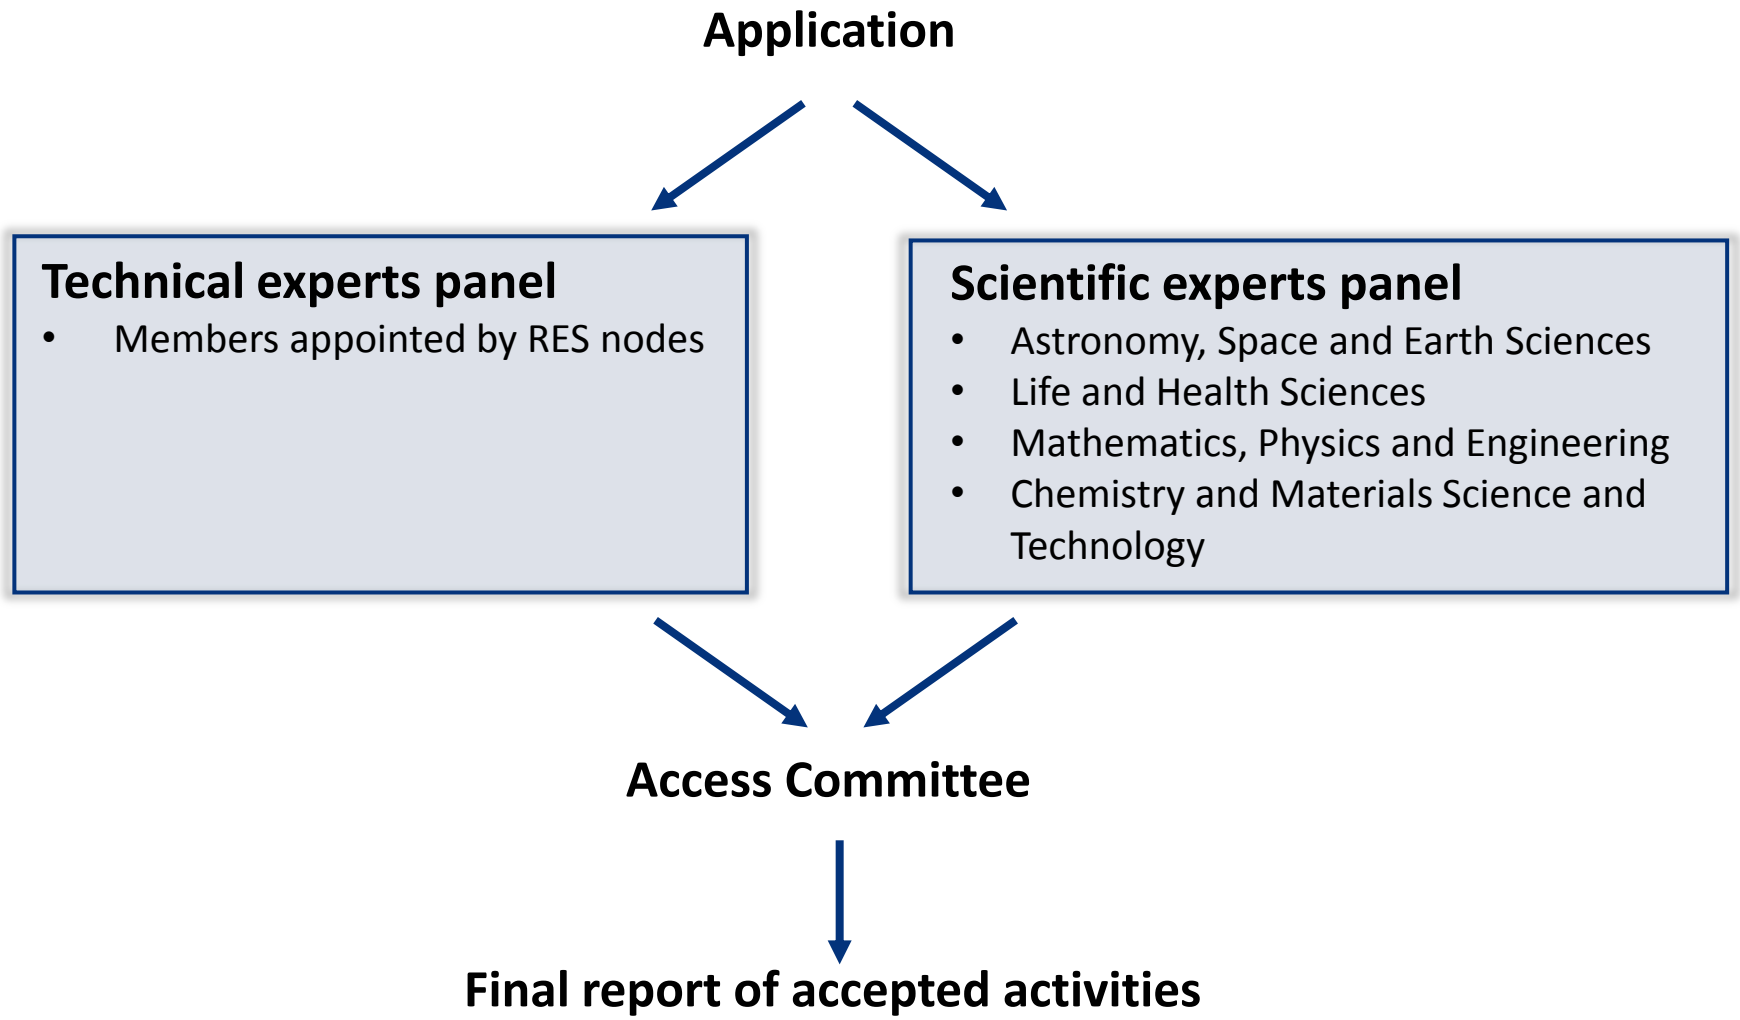

RES Users' Committee

- **CURES aims to provide advice and feedback to RES coordinators:**
  - Promotes optimal use of high performance computing facilities
  - Shares information about users' experiences
  - Voices user concerns
- You can **contact CURES** through RES intranet:

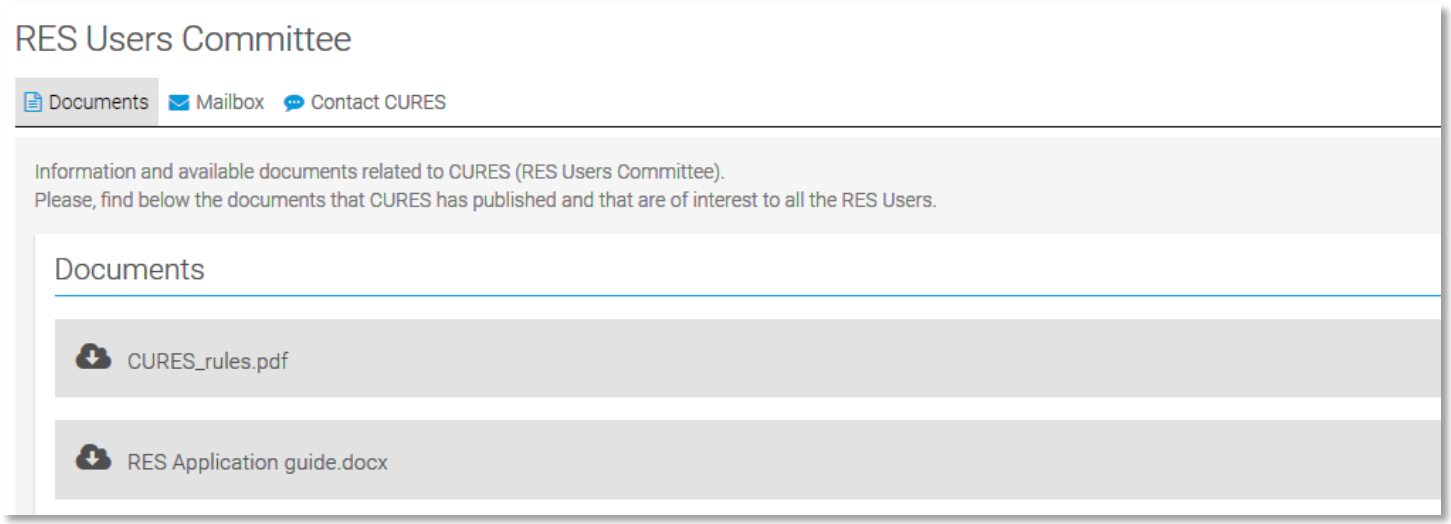

RES events: scientific seminars

The RES promotes scientific seminars which address supercomputing technology applications in specific scientific areas. These events are mainly organized by RES users and are open to the entire research community.

In 2017:

- ✓ 5 scientific seminars
- ✓ More than 300 attendees

Annual call for proposals

http://www.res.es/en/events

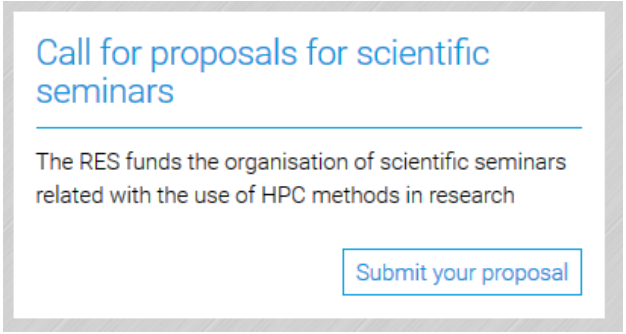

RES events: technical training

These workshops are organized by the RES nodes and aim to provide the knowledge and skills needed to use and manage the supercomputing facilities.

PRACE

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

PATC Courses:

BSC is a PRACE Advanced Training Centre

https://www.bsc.es/education/training/patc-courses

RES events: RES Users' Meeting

20 September 2018 - Valencia

The agenda includes:

- Information about RES and PRACE
- Parallel scientific sessions
- Poster session
- Evening social event

Services

- **Big Data & ML 4 HPC:**
  - Installation & maintenance of Big Data & ML tools/stacks.
  - Develop necessary tools to adapt Big Data clusters in HPC envs.
- **Advising (and best practices):**
  - Code development.
  - Data management and formatting.
- **Collaboration with researches:**
  - Applied Learning Methods.
  - Big Data Frameworks.
  - Data-Center Optimization.
  - Data-Centric Architectures.
  - Internet of Things and Stream Processing.

Applications

- **Hadoop.**
- **Spark.**

- **Cassandra.**
- **Hive.**

- **TensorFlow.**
- **Caffe.**
- **Theano.**
- **… (Sonnet, Lasagne, Scikit-Learn, Keras, PyTorch).**

- **Virtually anything you need (and request).**

BIG DATA Tools and Software at BSC-CS

PyCOMPSs/COMPSs

- Programmatic workflows
  - Standard sequential coordination scripts and applications in Python or Java
  - Incremental changes: Task annotations + directionality hints
- Runtime
  - Exploitation of inherent parallelism
  - DAG generation based on data dependences: files and objects
  - Tasks and objects offload
- Platform agnostic
  - Clusters
  - Clouds, distributed computing

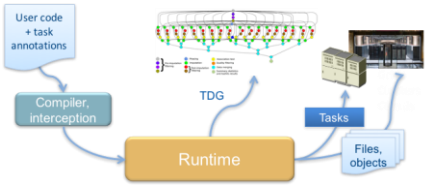

dataClay

- dataCLay: platform that manages **Self-Contained Objects** (data and code)
- Platform features:
  - Store and retrieve objects as seen by applications
  - Remote execution of methods
  - Add new classes
  - Enrich existing classes: With new methods and With new fields

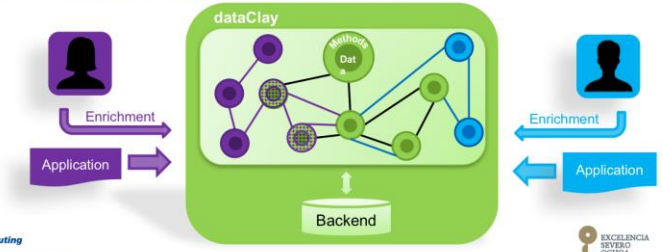

Hecuba

- Set of tools and interfaces that aim to facilitate an efficient and easy interaction with non-relational data-bases
- Currently implemented on Apache Cassandra database
  - However, easy to port to any non-relational key-value data store
- Mapping of Python dictionaries into Cassandra tables
  - Both consist on values indexed by keys
  - Only Python data type supported right now
- Redefinition of Python iterators

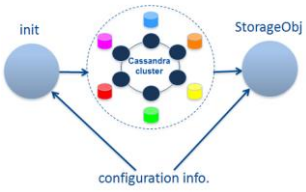

spark4mn

- Spark deployed in MareNostrum supercomputer
- Set of commands and templates
  - Spark4mn
    - sets up the cluster, and launches applications, everything as one job.
  - spark4mn_benchmark
    - N jobs
  - spark4mn_plot
    - Metrics

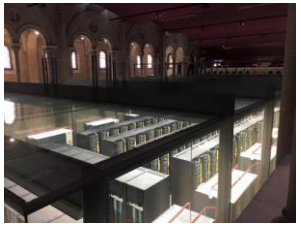

BIG DATA Applications

Tiramisu

- Goal: to exploit the representations learnt by CNNs
- Input: sets of images
  - For each set of images an activation set is extracted using deep learning toolkits (Caffe)
- Tiramisu performs next cognitive step → Data Mining and Knowledge Discovery on top of Deep Learning
  - Operations with the activation sets to derive new activation sets
  - Enables unsupervised Image clustering
  - Easy to use by data scientists
  - BSC development on top of PyCOMPSs

Case of study: Respiratory system simulator

- **Qbeast**: D8-tree distributed engine
- Prototype implemented on top on key-value data store: Cassandra
  - Managed by Hecuba
- Peer-to-peer architecture
- Linear scalability
- Enable to visualize results at simulation time
- Queries parallelized with PyCOMPSs

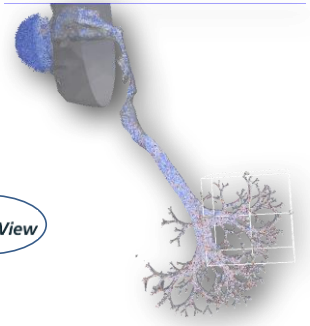

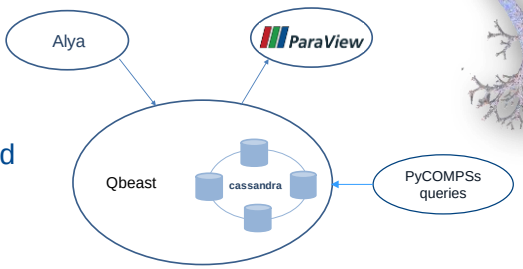

Guidance

- A tool for Genome-Wide Association Studies
- Examples of scientific application:
  - Genotype imputation and association analysis of type 2 diabetes cases and controls with 70K subjects
  - Genotype imputation of 0.5 Million patients and controls suffering 44 genetic diseases using the 1000 whole genome sequences as reference panel

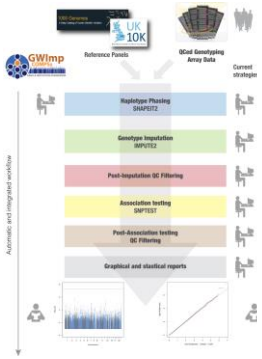

Can machines perform complex cognitive tasks?

Can they make music?

Difficulty: **Subjectivity**

We need human support

ALGORITMO

Minotauro (because GPUs)

Recurrent neural network

Generated songs

Feedback and raw material

Music festival (because people)

Voting and creation

Contact us!

Visit our website: www.res.es

Subscribe to our newsletter

Follow us in Twitter:
@RES_HPC

applications@res.es
dissemination@res.es


BSC
Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación


EXCELENCIA
SEVERO
OCHOA

THANK YOU!


R E S
RED ESPAÑOLA DE
SUPERCOMPUTACIÓN


Infraestructuras
Científicas y Técnicas
Singulares