


RED ESPAÑOLA DE  
SUPERCOMPUTACIÓN



# RES and PRACE



RED ESPAÑOLA DE  
SUPERCOMPUTACIÓN

- RES is a distributed infrastructure of Spanish HPC systems
- The objective is to provide a unique service to HPC users in Spain

- Services included:

- HPC resources
- Storage
- Support to submit your application
- Benchmarking
- Technical support
- User Training
- Scientific seminars
- Dissemination activities



RED ESPAÑOLA DE  
SUPERCOMPUTACIÓN



- Created by BSC in 2006, and supported by Spanish Government and so far, around 45 Mhours per period distributed.

# RES members and resources



## BSC-CNS (MareNostrum) +1PFlop/s +48.000 cores

Processor: Intel Sandy-Bridge processor  
Memory: 2GB/core  
Disk: 2PB + 5PB  
Network: Infiniband FDR10



## UPM (Magerit II) +3900 cores, 103Tflops

Processor: 3.920 (245x16) Power7 3.3 GHz  
Memory: 8700 GB  
Disk: 190 TB  
Network: IB QDR



## Universidad de Malaga - UMA (picasso), +2312 cores, 31Tflops

Processor: Intel Xeon  
Memory: +1 TB /node



## IAC, UV (LaPalma, Tirant) - 9.4 and 18 Tflops

Processor: 1024 and 2048 PowerPC 970 2.2 GHz  
Memory: 1 TB  
Disk: 14 + 10 TB  
Network: Myrinet



## CSUC (Pirineus) 14,3 Tflops

Processor: 1344 cores (224 Intel Xeon X7542 with 6 cores)  
Memory: 6.14 TB  
Disk: 112 TB



## CenitS (Lusitania) 1,6 Tflops

Processor: 2800 cores (608 Xeon E5450 with 4 cores)  
Memory: 2 TB



## BSC-CNS (MinoTauro) +180TFlops/s, with + 1500 cores and 256 GPUS M2090 Nvidia \*

Processor: 128 nodes: 2 Nehalem + 2 M2090  
Memory: 3 TB  
Network: IB QDR

## UC (Altamira) + 3800 cores, 79.84 Tflops

Processor: Intel Sandy-Bridge processor  
Memory: 4GB/Core  
Network: Infiniband FDR

## Universidad de Zaragoza - UZ (caesaraugusta), + 3072 cores, 25.8 Tflops

Processor: AMD Opteron 6272, 2.1 GHz (Interlagos)  
Memory: 256 GB RAM memory  
Network: Infiniband QDR

## FCSC (Calendula) 328 Tflops

Processor: 2800 cores (608 Xeon E5450 with 4 cores)  
Memory: 8.5 TB  
Network: Infiniband

## CESGA (Finis Terrae2) 330 Tflops

Processor: 7712 cores (306 Haswell 2680v3- 12 cores)  
Memory: 44 TB  
Disk: 1 PB  
Network: Infiniband

# Institutions and machines in RES (April 2016)



# Technical Support offered by RES

## “(Code optimization and porting for efficient use of supercomputers:

- Enhancement of parallelism and scalability
- I/O optimization
- Code porting (New architecture, accelerators, programming models, ...)
- Code debugging and profiling (Totalview, Paraver, DDT...)

## “(But also:

- User trainings
- Management of technical and administrative issues:
  - Account creation and user accounting
  - Help in the submission of proposals
  - Weekly reports regarding the activity's progress

## “(And support to submit test activities

- Allowing users to execute and optimize their code in a RES site, with a limited amount of hours.
- Via web form in : [www.bsc.es/RES](http://www.bsc.es/RES)



# How to access RES?

Via RES Área ([www.bsc.es/RES](http://www.bsc.es/RES))

☞ Call announcements for RES access

☞ Terms and usage conditions


☞ User login

☞ Accounting, reporting, etc.

☞ Proposal submission

- Title of the activity
- Description of the project and the activity
- Software and libraries
- Description of PI and members
- Requested resources
- Summary of the activity for dissemination

☞ Maximum of 2 periods per activity ( it means only 1 continuation per application).



The screenshot shows the login interface for RES-Red Española de Supercomputación. At the top, it says "RES-Red Española de Supercomputación: Login". Below this is a dark blue button labeled "New User". The main text asks the user to "Please, enter your username (e\_mail) and your password:". There are two input fields: "Username:" and "Password:". Below the password field is a checkbox labeled "Check here if you forgot your password". At the bottom right is a blue button labeled "Submit".

# How to access RES? Online form

## 5.- Resources

### a) To which machine(s) are you requesting access?

☐ MareNostrum 3 (IBM System X iDataplex with Infiniband / >40000 cores) [More info](#)

☐ Magerit-2 (Power 7 with Infiniband / 3920 cores) [More info](#)

☐ Altamira (IBM System X iDataplex with Infiniband / 2200 cores)

☐ Picasso (Intel Xeon/sandybridge with Infiniband / 1328 cores ) [More info](#)

☐ Tirant (IBM powerPC cluster with myrinet / 2048 cores)

☐ Memento (Cluster of AMD opteron 6272 with Infiniband / 3072 cores)

☐ La Palma (IBM powerPC cluster with myrinet / 1024 cores)

☐ MinoTauro (BULL cluster with Intel processors + NVIDIA CUDA accelerators) [More info](#)

☐ Caléndula (Intel Xeon with Infiniband / 1800 cores) [More info](#)

☐ Pirineus (SGI Altix UV1000 with Numalink v5, Xeon X7542 / 1344 cores) [More info](#)

☐ Lusitania (HP SuperDome with 256 cores Itanium2 Montvale and 2TB of main memory in the computational node) [More info](#)

☐ Cibeles (Intel Xeon with Infiniband / 560 cores) [More info](#)

☐ Finis Terrae 2 (cluster BULL con 7712 cores, Infiniband Fat-Tree FDR, 44TB memoria, 1 PB disco) [More info](#)

**INFORMATION: The estimated cost of the requested hours, considering only the electricity cost, is 0 euros.**

### b) Please select:

☐ The required resources have to be executed in the selected machines, the other architectures do not fit the requirements to execute the proposal.

\*\* this option implies that if no hours in this machine/these machines are available, the acces committee will reject the full application.

☐ The architectures selected for the requested resources are only a suggestion. If no hours in this machine/these machines are available, please grant resources in any other similar architecture where the codes used for the application may run efficiently.

**c) If the Activity is a 'Long Term Activity' (which will extend over several Application Periods), give an estimate of the total resources that the Activity will require until it is completed (including all Periods)**

Number of application periods expected to complete this Activity

Total Requested Time: (Thousands of hours) expected to complete this Activity

### a) To which machine(s) are you requesting access?

☒ MareNostrum 3 (IBM System X iDataplex with Infiniband / >40000 cores) [More info](#)

Please, try to estimate in terms of resource consumption, the typical run as well as the largest case (the most resource demanding experiment) for your simulation.

Interprocess communication

Null

Typical job run:

Number of processors needed for each job

Estimated number of jobs to submit

Average job durations (hours) per job

Total memory used by the job (GBytes)

Largest job run:

Number of processors needed for each job

Estimated number of jobs to submit

Average job durations (hours) per job

Total memory used by the job (GBytes)

Total Requested time: (Thousands of hours)

If this activity is asking for more than 2Million CPU hours, you need to justify the amount of resources requested for the activity (max 1000 characters)  
/ 1000 characters

Total disk space (GBytes)

Minimum

desirable

Total scratch space (GBytes)

Minimum

desirable

Total tape space (GBytes)

Minimum

desirable



# Access Committee

## Core Team

## 4 Expert Panels (10 people per panel):

- Astronomy, Space, and Earth Sciences
- Biomedicine and Life Sciences
- Mathematics, Physics and Engineering
- Chemistry and Material Science.

## Designated by MICINN, independent from BSC-CNS

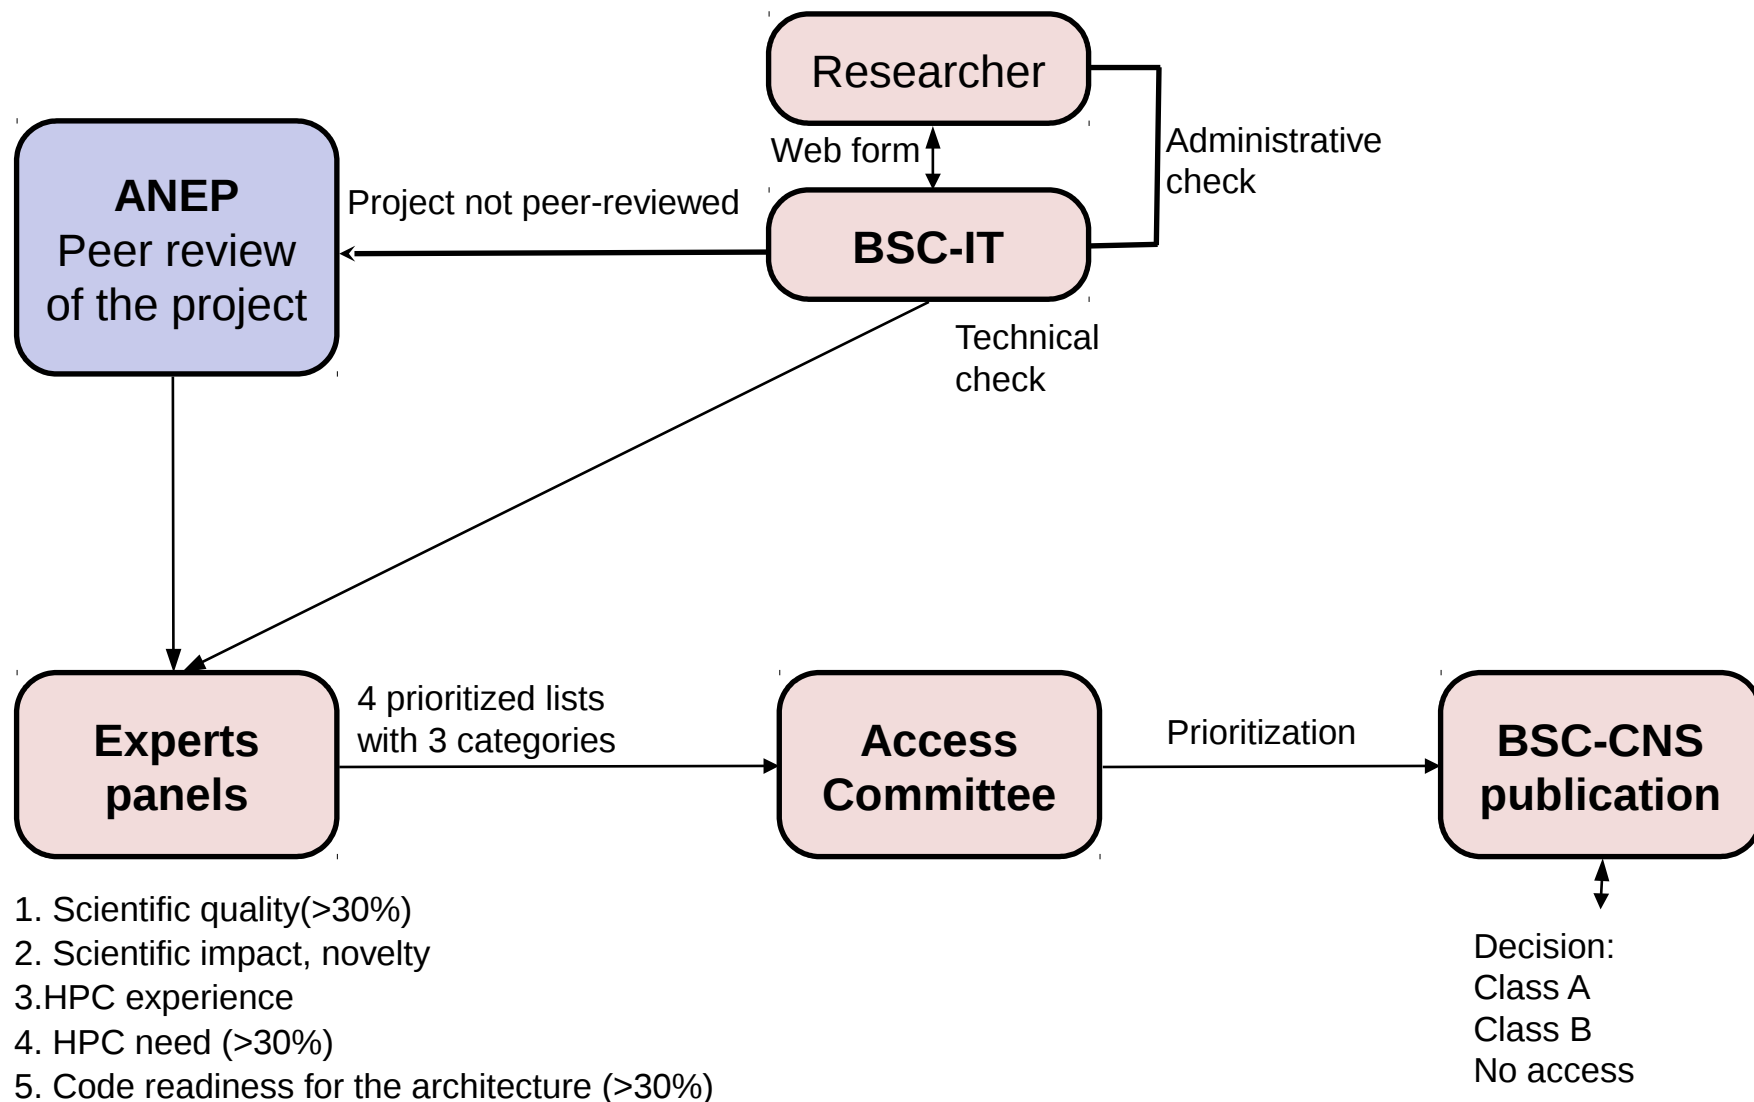
- Renewed every three years

## Access allocation every 4 months

## Resources distribution

- Depending on the machine, minimum of 20% of the machine provided for the RES users. In total, currently more than 40Mhours per period (hours with priority) + 10% of overbooking (hours without priority)


# Peer Review Process





# ESFRI: European Infrastructure Roadmap

- **PRACE** is an international not-for-profit association under Belgian law, with its seat in Brussels.
- **PRACE** counts 25 members and 2 observers.
- The **PRACE** Hosting Members are France, Germany, Italy and Spain





# PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

## 4 Countries offering computing resources on 6 world-class machines





**MareNostrum:** IBM  
BSC, Barcelona, Spain

**JUQUEEN:** IBM  
BlueGene/Q  
GAUSS/FZJ  
Jülich, Germany



**SuperMUC:** IBM  
GAUSS/LRZ  
Garching, Germany

**Hazel Hen:** Cray  
GAUSS/HLRS,  
Stuttgart, Germany




**Marconi:** Lenovo  
CINECA, Bologna, Italy



**CURIE:** Bull Bullx  
GENCI/CEA  
Bruyères-le-Châtel, France




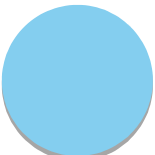
## PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE


 **467** scientific projects **enabled**

 **12.5 thousand million** core hours awarded since 2010 with peer review, main criterion is **scientific excellence**. Open R&D access for **industrial users** with **>50 companies** supported

 **>7 500** people trained by **6 PRACE Advanced Training Centers** and others events

 **24 Pflop/s** of peak performance on **6 world-class systems**


 **>500 M€** of funding for **2010-2015**, access **free at the point of usage**

 **25 members**, including **4 Hosting Members** (France, Germany, Italy, Spain)





# Access through PRACE Peer-Review



**Free-of-charge** required to publish results at the end of the award period



**Preparatory Access** (2 to 6 months)



**SHAPE Programme** (2 to 6 months)



**Project Access** (12, 24 or 36 months)



**Centers of Excellence** : 0,5 % of the total resources available for CoE

**Criterion:  
Scientific  
Excellence**

# Intel iDataPlex – MareNostrum, BSC, Barcelona, Spain

☞ MareNostrum – hosted by BSC-CNS, BarcelonaSpain

- MareNostrum is based on the Intel Xeon-Architecture and it provides a peak performance of 1.1 Petaflops
- 3,056 compute nodes + 42 nodes with 2 Xeon Phi
- 2 PB of disk storage



<http://www.bsc.es/marenostrum-support-services/mn3>

# Lenovo cluster – Marconi, CINECA, Casalecchio di Reno, Italy

- MARCONI: Lenovo System supercomputer 2PFlops and 54,432 Intel Broadwell cores
  - Model: Lenovo NeXtScale
  - Architecture: Intel OmniPath Cluster
  - Nodes: 1.512
  - Processors: 2 x 18-cores Intel Xeon E5-2697 v4 (Broadwell) at 2.30 GHz
  - Cores: 36 cores/node, 54.432 cores in total
  - RAM: 128 GB/node, 3.5 GB/core
  - Internal Network: Intel OmniPath
  - Disk Space: 17PB (raw) of local storage
  - Peak Performance: 2 PFlop/s

Near future -> KNL (3600 nodes) around 11PFlops



For preparatory access only

# IBM Blue Gene/Q – JUQUEEN, GCS, Jülich, Germany

Composed of 458,752 processing cores using 16 cores forming a node with 16 GB of memory for a total of 448 TB.

Performance

-IBM PowerPC A2, 1.6 GHz

-Peak: 5.9 PFlop/s

<http://www.fz-juelich.de/jsc/juqueen>





# Cray XC40 – Hazel Hen, GCS, HLRS, Stuttgart, Germany

⌘ The Cray XC40 Hazel Hen is based on Intel Haswell Processor and the Cray Aries network. The system is an upgrade / follow on project of Hornet and has a peak performance of over 7 Petaflops.

It is composed of 7712 dual socket nodes equipped with INTEL Sandy Bridge 12 core Processors leading to overall 185,088 processing cores. Nodes are equipped with 128 GB main memory.

Hazel Hen features 965 Terabyte of Main Memory and a total of 11 Petabyte of storage capacity

⌘ Peak Performance 7.420 PFlop/s

⌘ <https://www.hlr.de/systems/platforms/cray-xc40-hazel-hen/>





# Intel iDataPlex – SuperMUC, GCS, LRZ, Garching, Germany

SuperMUC – hosted by GCS in LRZ, Garching, Germany. Currently 155,000 cores.

- SuperMUC is based on the Intel Xeon-Architecture and it provides a peak performance of 3.2 Petaflops
- Will be upgraded to > 6 Petaflops in 2015



# Cray XC30 – Piz Daint, CSCS, Lugano, Switzerland

## 5272 Compute Nodes with

- One Intel® Xeon® E5-2670 (8-core)
- One NVIDIA® Tesla® K20X

## Theoretical Peak Floating-point Performance per node

- 166.4 Gigaflops (Intel® Xeon® E5-2670)
- 1311.0 Gigaflops (NVIDIA® Tesla® K20X)

## Theoretical Peak Performance 7.787 Petaflops

## Memory Capacity per node:

- 32 GB (DDR3-1600)
- 6 GB non-ECC (GDDR5)

## Interconnect Configuration Aries routing and communications ASIC, and Dragonfly network topology

## System storage capacity 2.5 PB

# Support from BSC-CNS

- Technical help in applications
  - In order to increase the *success ratio* of applications
- Scalability tests at local machines
- Help in access and execution
  - Code porting
  - I/O transfers from/to BSC
  - Batch system access
  - Code performance
- Data transfer
  - Grace period after project allocation





**Barcelona  
Supercomputing  
Center**

*Centro Nacional de Supercomputación*

Thank you!

Questions ?