



Access to HPC resources in Italy and Europe


Paolo RAMIERI

p.ramieri@ Cineca.it

SuperComputing Applications and Innovation Department



www.hpc.cineca.it




CINECA
is a non profit Consortium,
made up of 70 Italian
Universities, 4 research
Institutions and MIUR.



CINECA is now the largest Italian computing centre, one of the most important worldwide. .
The High Performance Systems department (SCAI: SuperComputing Applications and Innovation) offers support to scientific and technological research through supercomputing and its applications.

The Story

- 1969: CDC 6600 1st system for scientific computing
- 1975: CDC 7600 1st supercomputer
- 1985: Cray X-MP / 4 8 1st vector supercomputer
- 1989: Cray Y-MP / 4 64
- 1993: Cray C-90 / 2 128
- 1994: Cray T3D 64 1st parallel supercomputer
- 1995: Cray T3D 128
- 1998: Cray T3E 256 1st MPP supercomputer
- 2002: IBM SP4 512 1 Teraflops
- 2005: IBM SP5 512
- 2006: IBM BCX 10 Teraflops
- 2009: IBM SP6 100 Teraflops
- 2012: IBM BG/Q 2 Petaflops




TOP500

#7 june 2012



Rank	Site	Computer
1	DOE/NNLS/LLNL United States	Sequoia - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM
2	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu
3	DOE/SC/Argonne National Laboratory United States	Mira - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM
4	Leibniz Rechenzentrum Germany	SuperMUC - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR IBM
5	National Supercomputing Center in Tianjin China	Tianhe-1A - NUDT YH MPP, Xeon X5670 6C 2.93 GHz, NVIDIA 2050 NUDT
6	DOE/SC/Oak Ridge National Laboratory	Jaguar - Cray XK6, Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA 2090
7	CINECA Italy	Fermi - BlueGene/Q, Power BQC 16C 1.60C IBM
8	Forstungswissenschaftszentrum Berlin (FZB) Germany	SuperMUC - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR IBM
9	CEA/TGCC-GENCI France	Curie thin nodes - Bullx B510, Xeon E5-2680 8C 2.700GHz, Infiniband QDR Bull
10	National Supercomputing Centre in Shenzhen (NSCS) China	Nebulae - Dawning TC3600 Blade System, Xeon X5650 6C 2.66GHz, Infiniband QDR, NVIDIA 2050 Dawning



FERMI



*High-end system, only
for extremely scalable
applications*

Name: Fermi

Architecture: BlueGene/Q (10 racks)

Processor type: IBM PowerA2 @1.6 GHz

Computing Nodes: 10.240

Each node: 16 cores and 16 GB of RAM

- **Computing Cores:** 163.840

RAM: 1 GByte / core (163 TByte total)

Internal Network: 5D Torus

Disk Space: 2 PByte of scratch space

Peak Performance: 2 PFlop/s

Power Consumption: 820 kWatts

N. 12 in Top 500 rank (June 2013)

National and PRACE Tier-0 calls



GALILEO



Name: Galileo

Model: IBM NeXtScale

- **Architecture:** IBM NeXtScale

Processor type: Intel Xeon Haswell@ 2.4 GHz

Computing Nodes: 516

Each node: 16 cores, 128 GB of RAM

- **Computing Cores:** 8.256

RAM: 66 TByte

Peak Performance: 1.2 PFlops

Internal Network: Infiniband 4xQDR switches (40 Gb/s)

Accelerators: 768 Intel Phi 7120p (2 per node on 384 nodes)

+ 80 Nvidia K80 (2 per node on 40 nodes)

X86 based system for production of medium scalability applications



• ~~National and PRACE Tier-1 calls~~

MARCONI (2016)



- Complessivamente:
 - 16,5 PFs peak performance,
 - 10 PB storage
 - 3 MW di potenza assorbita

 - Caratteristiche tecniche:
 - Rete interna: Intel OmniPath
 - Architettura: Lenovo NeXtScale
 - A1: Broadwell 2x18 core, 2.3 GHz; 1500 nodi, 2 PFlops
 - A2: KnightsLanding 68 core, 1.4 GHz; 3600 nodi, 11 PFlops
 - A3: SkyLake 2x20 core, 2.3 GHz; 1500 nodi, 4,5 PFlops
-

BigData - PICO

Storage and processing of large volumes of data

Name: Pico

Model: IBM NeXtScale

Processor type: Intel Xeon Ivy Bridge@2,5Ghz

Computing Nodes: 66+

Each node: 20 cores, 128 GB of RAM

Computing Cores: 1.320+

RAM: 6,4 GB/core

plus

2 Visualization nodes (with 2 GPU NVIDIA K40)

2 Big Mem nodes (512 GB RAM)

4 BigInsight nodes (32 TB of local disk)


2 Very Large Mem nodes (1TB RAM)





The European HPC-Ecosystem

PRACE Research Infrastructure (www.prace-ri.eu): the top level of the European HPC ecosystem



Creation of a European HPC ecosystem involving all stakeholders

- ✓ HPC service providers on all tiers
- ✓ Scientific and industrial user communities
- ✓ The European HPC hw and sw industry

- **CINECA:**
- - represents Italy in PRACE
- - hosting member in PRACE
 - - Tier-0 system
 - **BG/Q 2 PFlop/s**
 - - Tier-1 system
 - **> 5 % GALILEO**
- involved in PRACE 1IP, 2IP, 3IP
- PRACE 2IP prototype Eol

Access to HPC resources: CINECA aims and basic principles



Our objectives:

- ✓ Providing Italian and European researchers with an advanced computational environment
- ✓ Supporting Italian researcher for increasing their competitiveness
- ✓ Following Italian researchers in their path towards Tier 0
- ✓ Soliciting large-scale and computationally intensive projects

Basic principles:

- ✓ Transparency
- ✓ Fairness
- ✓ Conflict of Interest management
- ✓ Confidentiality

L'offerta HPC



HPC in Italy



HPC offer in Italy: ISCRA Italian SuperComputing Resource Allocation



The aim of ISCRA is to ensure adequate support to Italian scientists and engineers.

The allocation is of 600M core hours per year on

- FERMI

The access is by

- Online submission of proposals
- Peer-review process by an international panel

The proposal are scientifically evaluated by international reviewers and technically evaluated by Cineca experts.

Applications and codes are evaluated on the basis of their computational readiness.

HPC offer in Italy: ISCRA



Class B: Standard Projects; two calls / year
1-10M core hours
duration: 12 months
FERMI only



Class C:

Small Projects

continuous submission, 12 selections per year

>50K core hours

up to 1M core hours on FERMI

up to 200K core hours on GALILEO

only two C projects approved per year per user

duration: 9 months

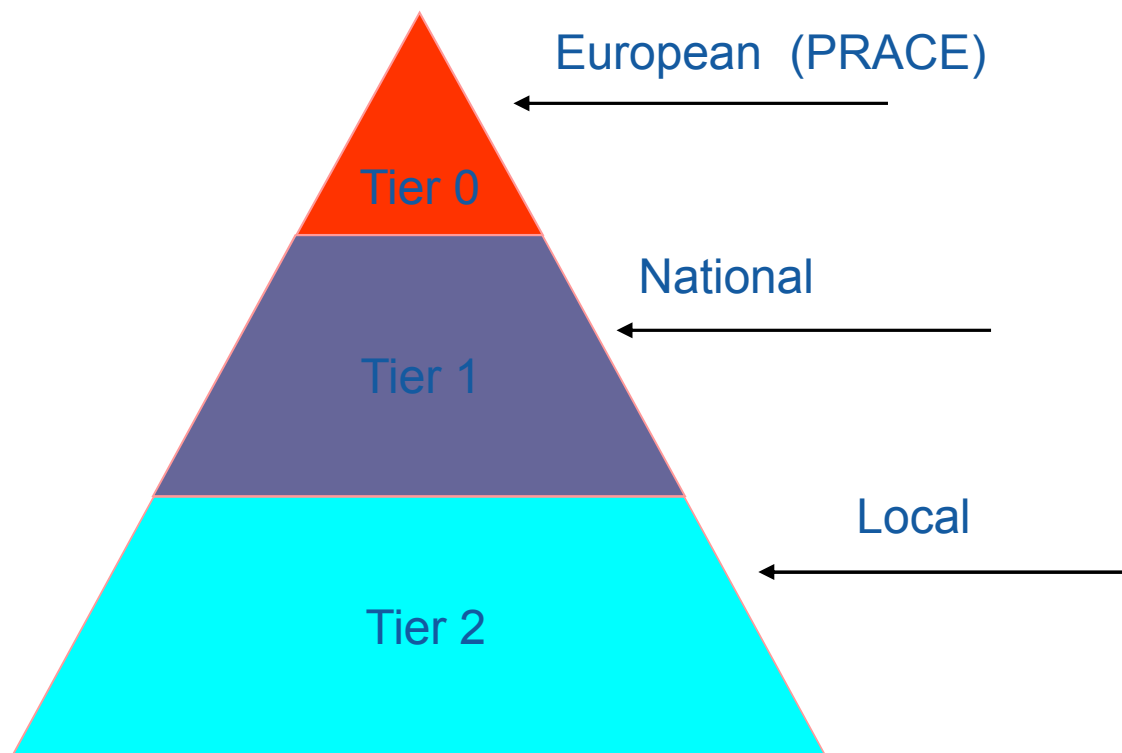
Trial:

on demand

L'offerta HPC in Europa: PRACE



Computing provisioning pyramid



PRACE offers access to Tier 0 (and Tier 1) systems

- **Preparatory Access**

- Intended for preliminary resource use required to prepare proposals for Project Access
- Technical review
- Continuously open calls

- **Project Access**


- Intended for individual researchers and research groups including multi-national research groups
- Technical and Scientific review

PRACE Tier 0 Access: CURRENT CALL



- Applications accepted on call
- For projects which use codes that have been previously tested and must have demonstrated high scalability and optimisation

Projects dimension



typ 20K

typ 1M

typ 3M

typ 10-70M



Useful links



- ISCRA: <http://www.hpc.cineca.it/services/iscra>
- PRACE: www.prace-ri.eu/hpc-access?lang=en