CINECA
for HCP and e-infrastructures

June 5, 2012

CINECA
Casalecchio di Reno (BO)
Via Magnanelli 6/3, 40033 Casalecchio di Reno | 051 6171411 | www.cineca.it

Next future

- MIUR Ministry entrustment for the incorporation into CINECA of the CILEA Consortium in Milan and CASPUR Consortium in Rome

CASPUR
Consorzio interuniversitario per le Applicazioni di Supercalcolo Per Università e Ricerca

cilea
CONSORZIO INTERUNIVERSITARIO
CALCOLO ALTE PRESTAZIONI
TECNOLOGIE DELL'INFORMAZIONE E DELLA COMUNICAZIONE

- Create a unique entity that for critical mass will reinforce the Italian commitment for the development of e-infrastrcuture in Italy and in Europe for HCP and HCP technologies, scientific data repository and management, cloud computing for industries and Public administration, for the development of computing intensive and data intensive methods for science and engineering

- From beginning of 2013
  - Unique offer for open access of integrated tier0 and tier1 HCP national infrastructure
  - Unique offer of education and training activities under the umbrella of PRACE Training advanced center action
  - Integrated help desk and scale up process for HCP users support

Partnership agreement with

Consiglio Nazionale delle Ricerche

CINECA

DEMOCRITOS
DEmocritos MOdeling Center for Research In aTOmistic Simulation
IOM

DEMOCRITOS will expand the current activities towards computational molecular biophysics and new areas of materials science. This rapidly advancing field requires the deployment of dedicated hardware and human resources in a similarly way as large facilities are necessary to even conceive certain experimental enterprises. In order to optimize the use of available resources, as well as the synergies that DEMOCRITOS will create, our action will be deployed by implementing a 'distributed virtual laboratory'. Such a laboratory will be strongly supported by SISSA and by CINECA.

CNR NANO

Istituto Nanoscienze is a new institute of CNR devoted to frontier research in nanoscience and nanotechnology, established in 2010 from three former laboratories of INFM: NEST-Pisa, NNL-Lecce, and S3-Modena. CNR NANO interests range from fundamental science to emerging technologies as well as applied projects of direct industrial and societal interest.

Partner CASPUR. The goal of the EPIGENOMICS FLAGSHIP PROJECT (Progetto Bandiera Epigenomica 2012 - 2015) is to understand, using a wide range of experimental models and approaches, how epigenetic mechanisms regulate biological processes, determine phenotypic variation and contribute to the onset and progression of diseases.

Partnership agreement with

The INFN - the National Institute of Nuclear Physics - is an organization dedicated to the study of the fundamental constituents of matter, and conducts theoretical and experimental research in the fields of sub nuclear, nuclear, and astroparticle physics. Fundamental research in these areas requires the use of cutting-edge technologies and instrumentation, which the INFN develops both in its own laboratories and in collaboration with the world of industry. These activities are conducted in close collaboration with the academic world.

With CINECA is in place a framework agreement for the development of applications for theoretical simulations of physical phenomena in the domain of fundamental constituents of matter.

Particularly development joint activity will be conducted in the development of new prototype solution for advanced computing based on massively parallel accelerated architecture.

INFN will co-invest for the development of the HCP infrastructure in CINECA looking head in the prospective of the follower of the FERMI system.

Partnership agreement with

CINECA

European Laboratory for Non-Linear Spectroscopy is a European reference point for research with light waves, based on a fundamental multi-disciplinary approach. LENS since its birth in 1991 is a center of excellence of the University of Florence.

Objectives of the partnership:

data processing, data mining, knowledge extraction/acceleration, “critical mission system management, assistance guided decision

Data mining Connectome Project (Consortium Connettoma)

Data mining cardiology (predictive cardiopatics e-networks)

Data mining neuro-omics (Besta)

Data mining Human Brain Project

Data mining pharmaceutical data (Novartis)

Stream Computing

DeepQA (Watson)

Specific agreement with LENS for the Human Brain Project

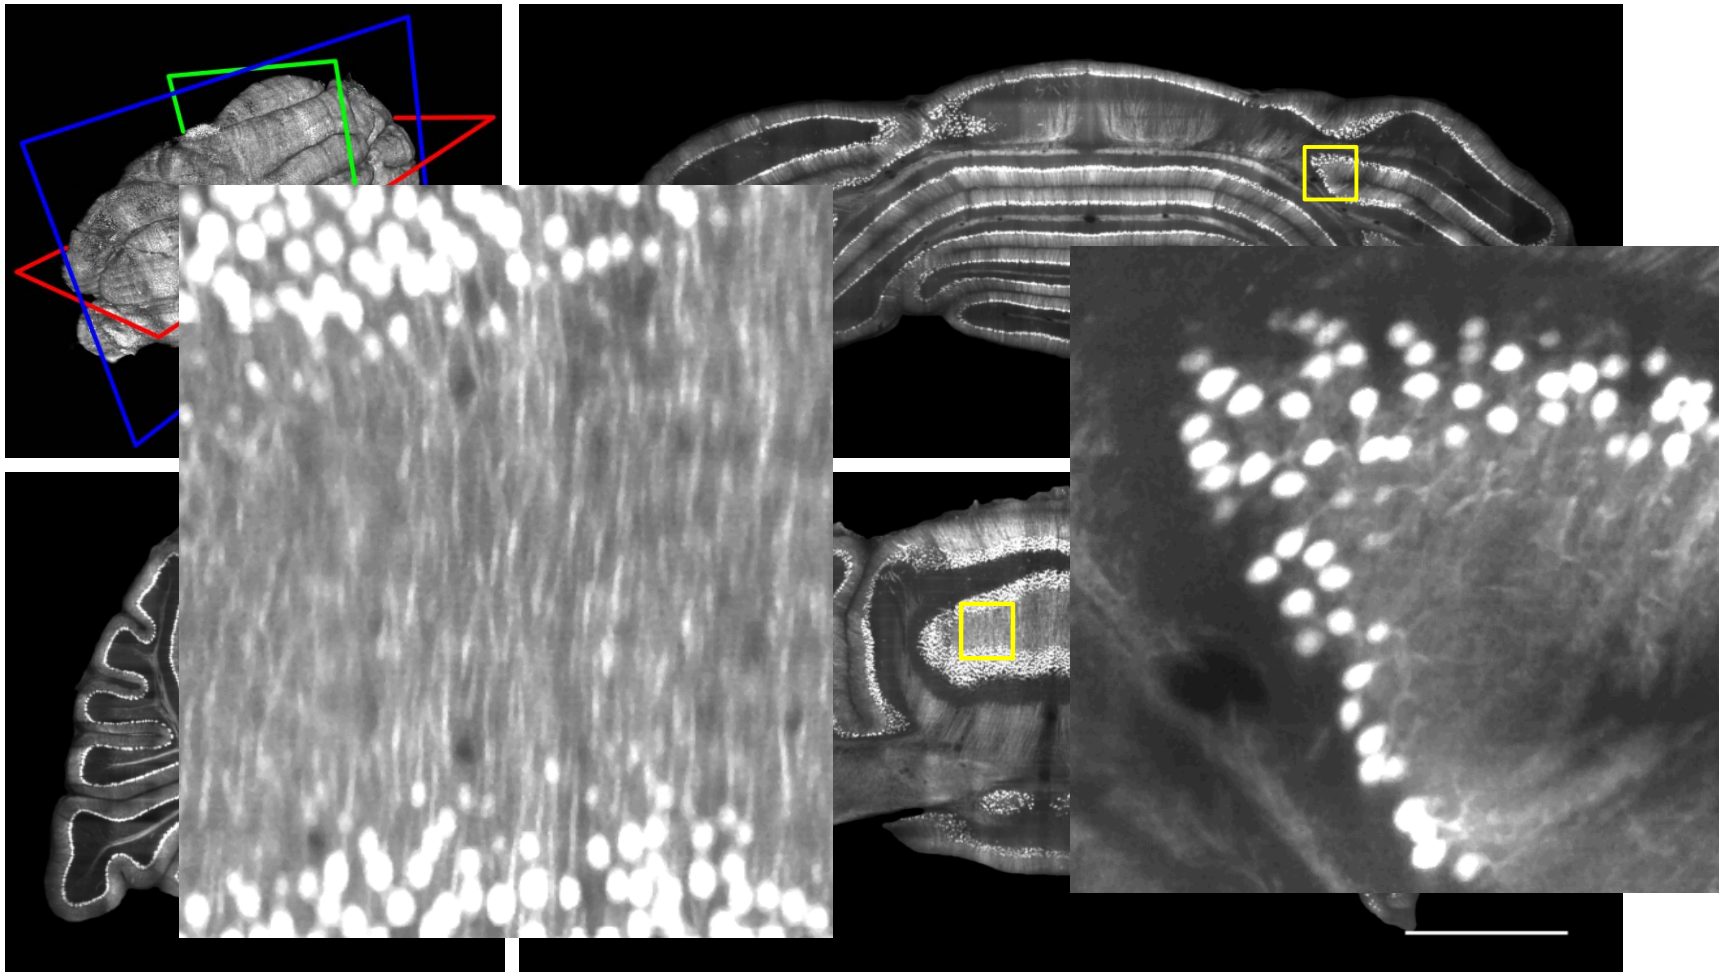

Scale bars: 1 mm

Partnership agreement with

The Abdus Salam
International Centre for Theoretical Physics

HPC Data Centric Italian Exascale Lab

Initial activities

- Validation of the new “EURORA” PRACE Prototype architecture.
- Evaluation of new programming models (ArBB, TBB, Cilk, OpenMP).
- Evaluation of MIC-to-MIC MPI communications.
- Evaluation and comparison between MIC to GPU
- Porting of full applications:
  - SISSA QuantumEspresso, Gromacs,
  - ICTP HP-SEE
  - OGS Solid Earth model and Ocean/biochemistry circulation model
  - INFN QCD, Lattice model, grid model
- Hot Water Cooling efficiency
- Specific measurement of power consumption by specific sensors
- Independent Power consumption measurements for each main component of the system
- Evaluation of TCO and total energy efficiency

Funded by the MIUR Minister with 1 Million Euro for the participation of Italy to the PRACE action on top of the fund assured to CINECA. All the partners cofound in similar amount either in cash of in kind.

Added value pay per use agreement with:

- **INAF, ESA, ASI for the post processing of the observed data of the Plank Mission and Gaia Mission;**
- **IIT for the Italian Drug Discovery Network initiative lead by IIT**
- **INGV for the EU ESFRI RI project EPOS lead by INGV**

The partecipation to the EU funded projects

HPC-Europa2
EC-funded collaborative research visits and access to Europe's biggest supercomputers for scientist of all levels, in all disciplines, form all EU states

Project coordinator

PRACE

Hosting Member

EUDAT

Principal Infrastructure provider

MMM@HPC

HPC for Molecular modeling

Works with EMI

European Middleware Infrastrcucture

V-MusT

Virtual archaeology

HPCWORLD

Handbook of HPC e-Science Infrastructure Allocation Reviewing, Selection and Management

HPC for Geophysic

MONT BLANC

DEEP Dynamical Exascale Entry Platform

ECESI

VERCE

Towards EXASCALE R&D (3 projects)

RISC

Latin American

- Production management
- Application development

BOMBARDIER

ORACLE RACING

FERRETTIGROUP
PASSION. INNOVATION. EXCELLENCE.

LAPCOS Progettazione collaborativa e simulazione

OpenFOAM
The Open Source CFD Toolbox

Computational fluid dynamic

Dompé

Drug design

PROTEZIONE CIVILE
Presidenza del Consiglio dei Ministri
Dipartimento della Protezione Civile

arpa SIM Servizio Idro Meteo

PIEMONTE Arpa Agenzia Regionale per la Protezione Ambientale

JRC CCU

Meteo climatology

FONDAZIONE CASSA DI RISPARMIO IN BOLOGNA

APA Stereo 3D Movie
Award Siggraph Asia 2011

HCP infrastructure for technical computing and added value service for industries

Logical Name	Production (April 2010)	Production (April 2011)	Production (April 2012)	Special Appl. (june 2008)
Model	IBM LS21 / HS21	HP Proliant	**IBM LS21 / HS21**	IBM HS21
Architecture	Linux Cluster	Linux Cluster	**Linux Cluster**	Linux Cluster
Processor	Intel Nehalem Qc 2.6 Ghz	Westmere Ec 2.6 Ghz	**Intel Sandybridge 8c 2.8 Ghz**	Intel Xeon Qc 2.8 Ghz
# of core	10240	15360	**2560**	3280
# of node	1280	1280	**320**	512
# of rack	25	25	**4**	12
Total RAM	45 Tera Byte	45 Tera Byte	**10 Tera Byte**	4 Tera Byte
Interconnection	CISCO Infiniband SDR 4x	Mellanox Infiniband QDR 4x	**Infiniband QDR 4x**	CISCO Infiniband SDR 4x
Operating System	RedHat	RedHat	**RedHat**	RedHat
Total Power	~ 800 Kwatts	~ 800 Kwatts	**~ 200 Kwatts**	~ 380 Kwatts
Peak Performance	> 100 Tera Flops	> 120 Tera Flops	**> 57 Tera Flops**	40 Tera Flops

Procurement 2008 terms of reference

- 100 Tflop/s system at initial 2009 with option to upgrade to Pflop/s size
- a prototype to demonstrate and facilitate technological transition
- 1 PF in 2012 at a not to exceed price

2009 IBM selected vendor

- 100Tflops SP6 + BlueGene/P proto
  - 12 SP6 p575 frames + I/O nodes + 6 DCS9900 FC Storage servers
  - 1 BlueGene/P rack + I/O nodes + 1 DCS9900 FC Storage server

CINECA

BlueGene/Q option (2.1 PFlop/s)

- Reuse of the water cooling infrastructure + free cooling
  - Better PUE: from 1.3 (worst case measured) to 1.15 (expected)
- Reuse of the power distribution infrastructure with same power consumption (~900Kw) – UPSs and generators included
- Rough transition (stop services during installation)
  - room issue: no separate room space to host the new system during transition (due to other businesses)

Main (technological) requirements

- Transition to water cooling (better PUE)
  - the 100 Tflop/s plant should be reused for the next (current) technological step (for a system of 2 PF peak)
- Power consumption: 100Tflop/s hits the Megawatt (not over)
  - reuse of UPSs, power panels etc. and same or lower power consumption for the next (current) technological step
- Smooth transition
  - business continuity
  - data migration

Hardware: the big blocks picture

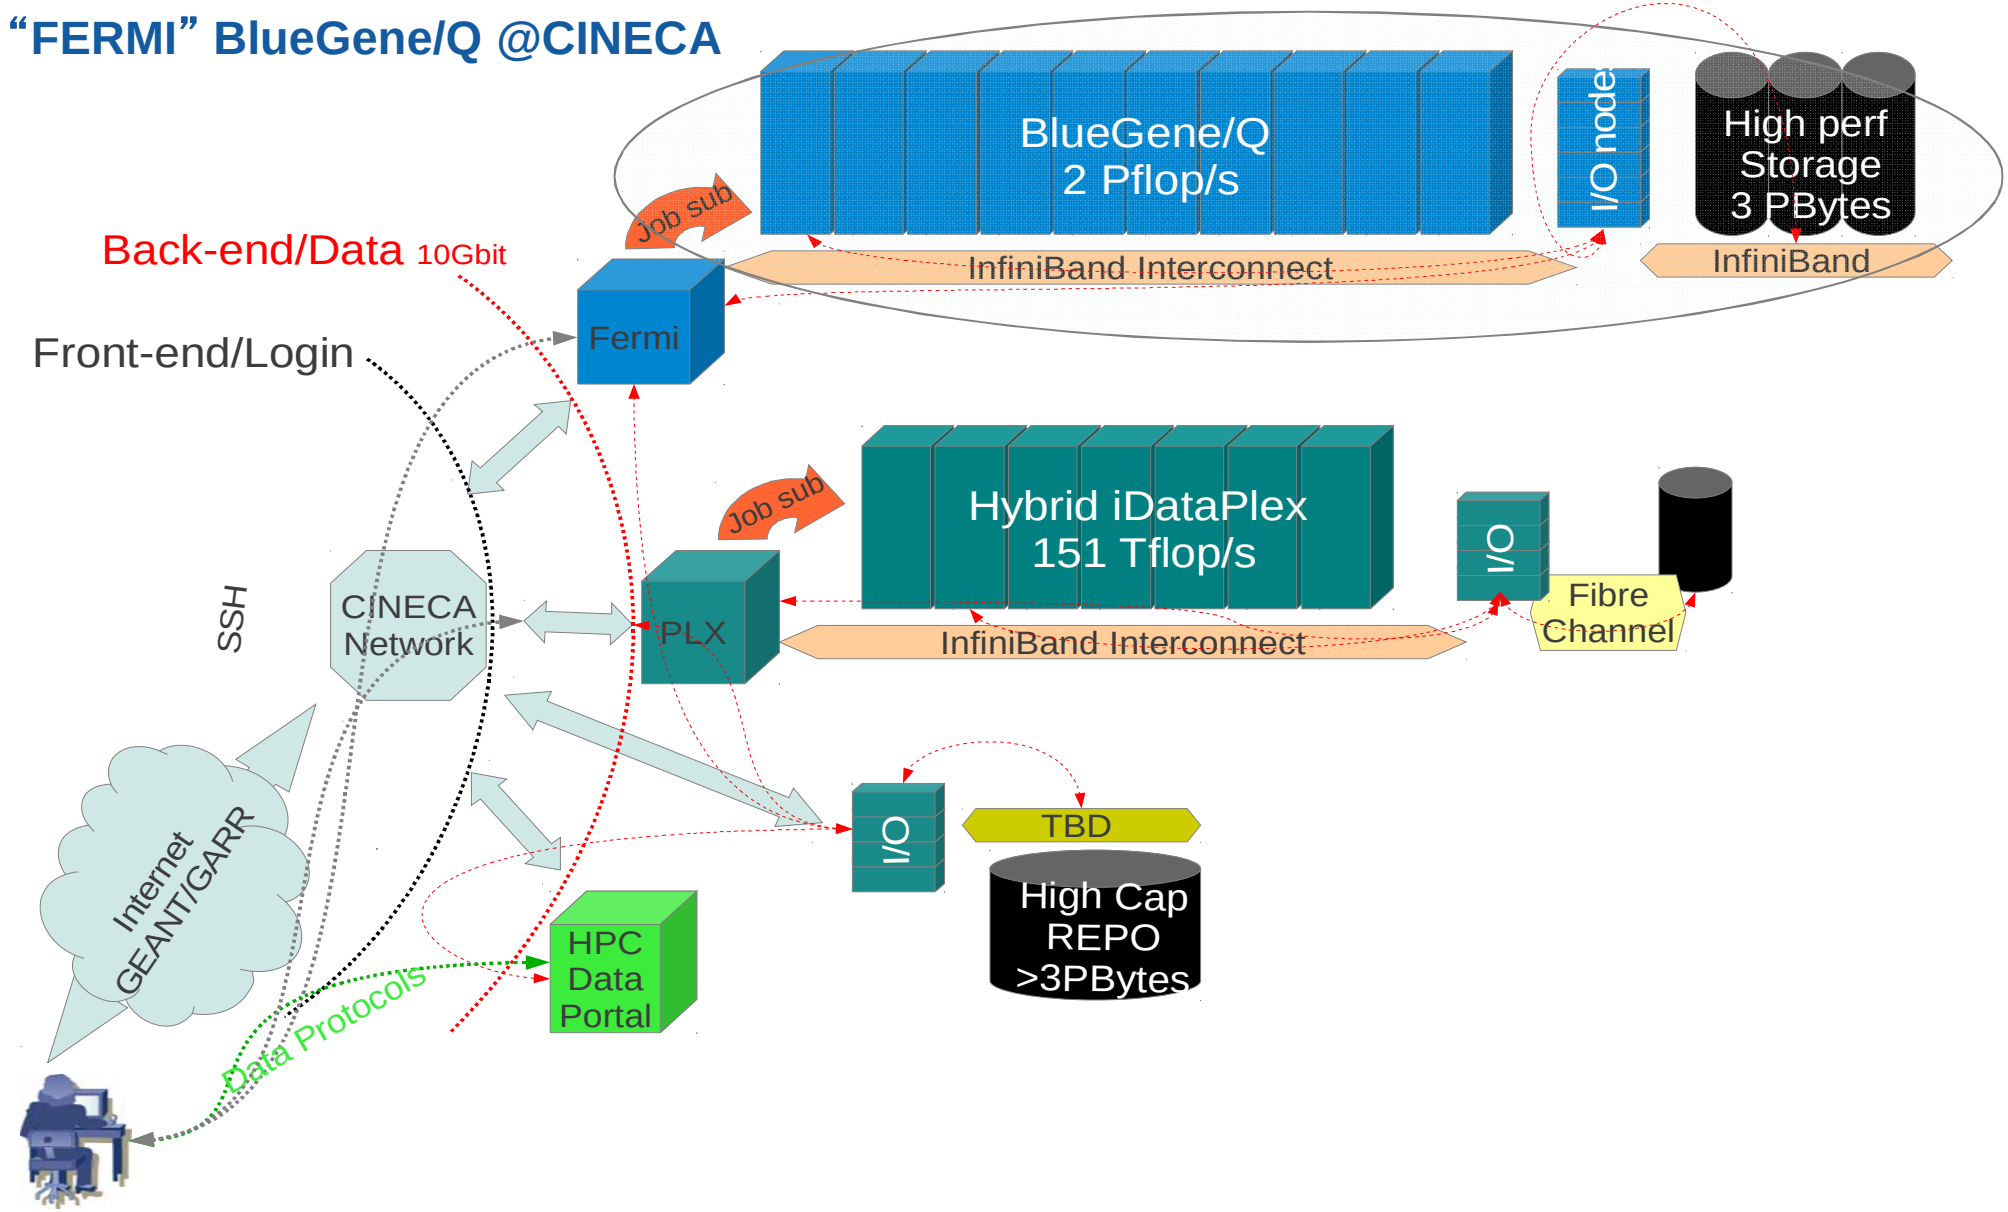

HPC infrastructure for scientific computing

Logical Name	FERMI (June 2012)	SCS PLX (june 2011)
Model	**IBM BG / Q**	IBM IDATAPLEX
Architecture	**MPP**	Linux Cluster
Processor	**IBM Power2A 1,6 GHz**	Intel Westmere Ec 2.4 Ghz
# of core	**163840**	3288
# of node	**10240**	274 + 548 Nvidia Fermi GPGPU
# of rack	**10**	14
Total RAM	**163 Tera Byte**	~ 16 Tera Byte
Interconnection	**IBM 5D Torus**	Qlogiq QDR 4x
Operating System	**RedHat**	RedHat
Total Power	**~ 1000 Kwatts**	~ 350 Kwatts
Peak Performance	**~ 2100 Tera Flops**	~350 Tera Flops

Storage infrastructure

System	Frontend bandwidth (GB/s)	Capacity (TB)	LAN / SAN tech.	Disks Tech.	
2 x S2A9500	3,2	140	FCP 4Gb/s	FC	
4 x S2A9500	3,2	140	FCP 4Gb/s	FC	
6 x DCS9900	5,0	540	FCP 8Gb/s	SATA	
4 x DCS9900	5,0	720	FCP 4Gb/s	SATA	
3 x DCS9900	5,0	1500	FCP 4Gb/s	SATA	
Hitachi Ds	3,2	360	FCP 4Gb/s	SATA	
4 x SFA10000	10,0	2700	QDR	SATA	June 2012
1 x 5100	3,2	66	FCP 8Gb/s	FC	
3 x SFA12000	40,0	3600	QDR	SATA	June 2012
		>9,7 PB			

BlueGene/Q “FERMI” Italian PRACE tier0

- Technology architecture: Blue Gene / Q, 2.1PF peak; 10 full rack, 163.840 core processor, 1GB of memory RAM per core.
- Storage system: 3PByte of scratch capacity with a front end bandwidth throughput in excess of 120 GByte /sec.
- File system: GPFS.
- Dedicated Cluster front-end for storage system management
- Full production July 2012.