

CESGA

**Advanced computing services for the research
community**

(September 2009)

Ignacio López Cabido

(Technical Deputy Subdirector, Galicia Supercomputing Centre)





CERTIFICATED FIRM

Outline

- Who we are
- Services
- Some plans
- Applications
- Challenges

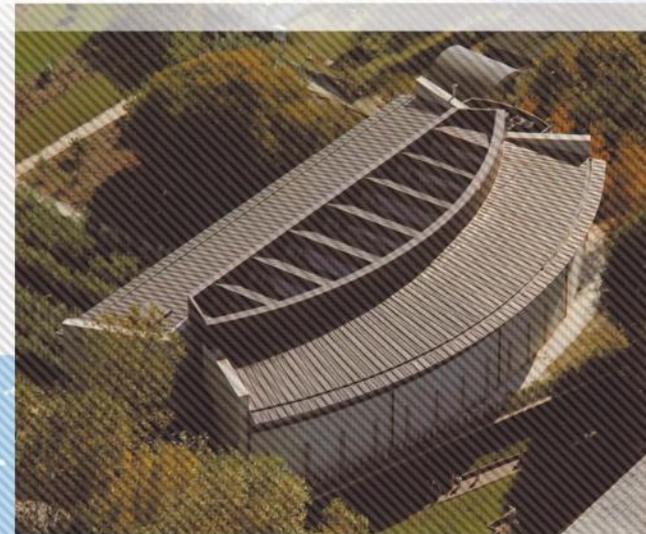
ESTABLISHED IN 1993 IN SANTIAGO DE COMPOSTELA (SPAIN)



CESGA



SANTIAGO DE COMPOSTELA



MISSION STATEMENT

To provide high performance computing and advanced communications **resources and services** to the scientific community of Galicia and to the Spanish National Research Council (CSIC), as well as, to institutions and enterprises with R&D activity.

To promote high quality research in **Computational Science** in close collaboration with the research community from Galicia as well as from other regions or countries all over the world; contributing in this way to the advancement of science, to **transfer technology** to industry and administrations , and as consequence, to the welfare of society as a whole.

LEGAL ENTITIES

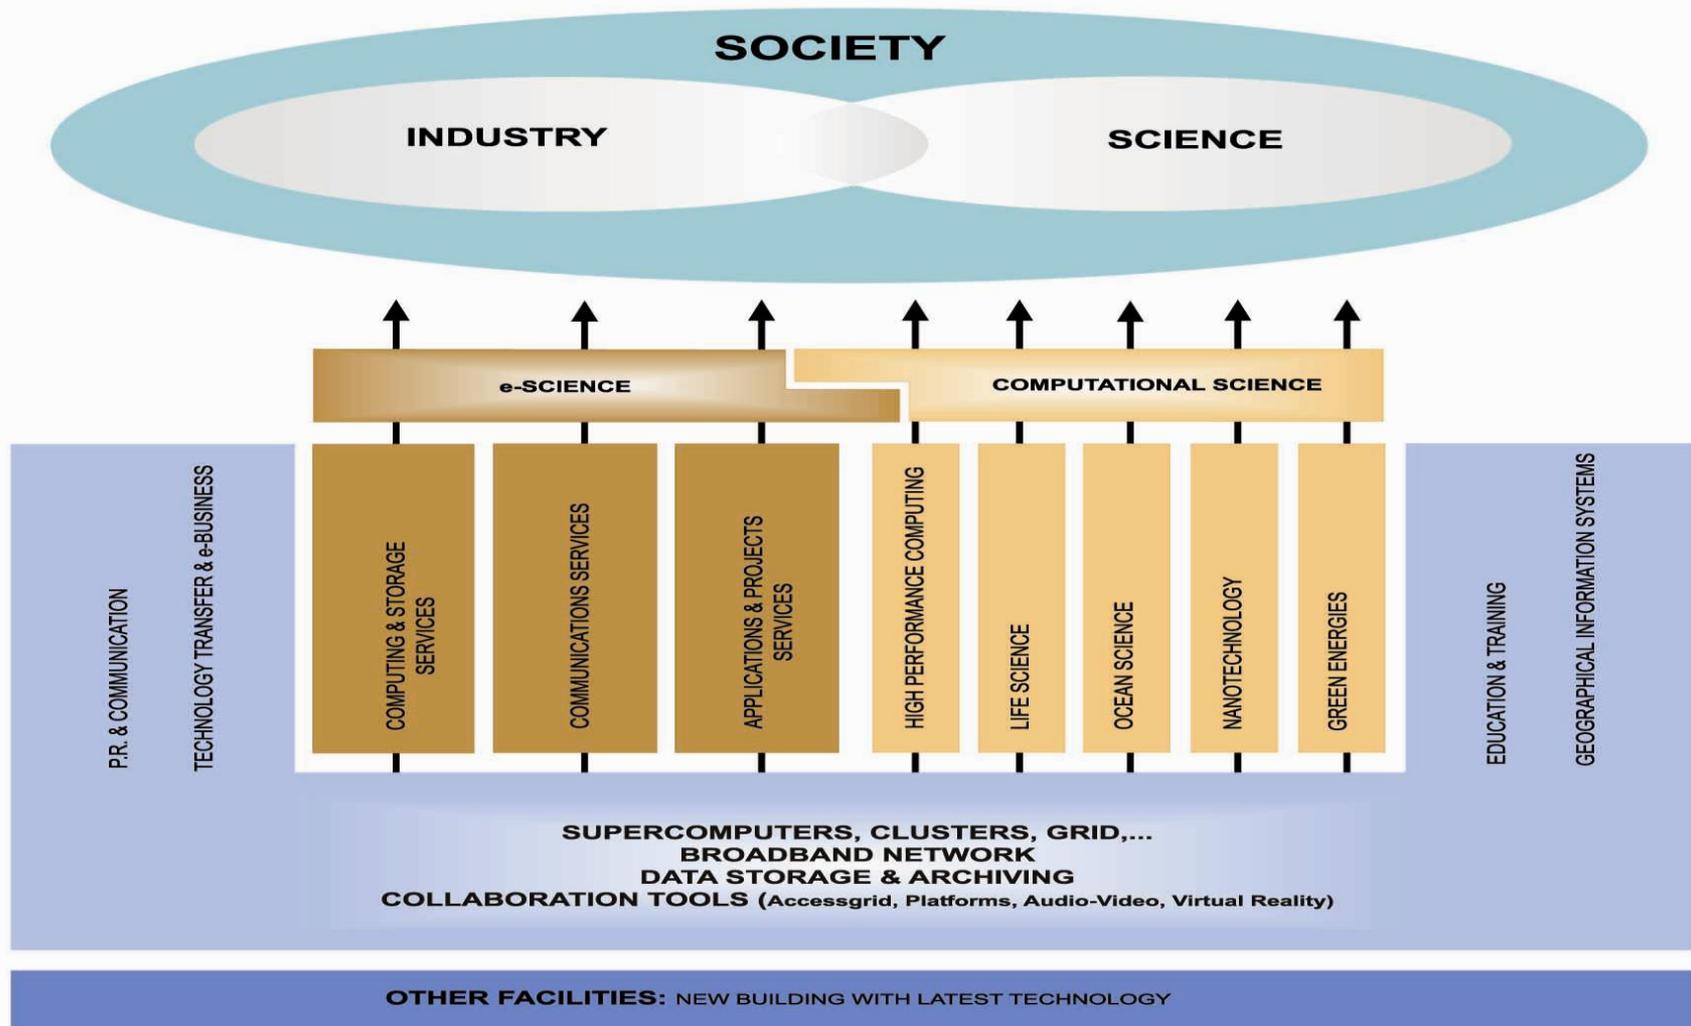
- **Public Company**
- **Public Foundation**

PARTNERS

- **Regional Government of Galicia**
- **National Research Council of Spain**



FUNCTIONAL SCHEME



AREAS OF ACTIVITY

RESEARCH

CESGA - C²SRC

(CESGA – Computational Science Research Centre)

- Starting date: 2009
- Staff: 115 Researchers & Technicians(recruiting)

ADVANCED SERVICES & INNOVATION

- Starting date: 1.993
- Staff: 60, currently

CESGA - C²SRC

CESGA COMPUTATIONAL SCIENCE RESEARCH CENTRE



- **Application Areas:**
 - Strategic for Galicia.
 - High demand of HPC resources.
 - Research experience in Galicia and CSIC.
 - Singular in Galicia and CSIC.



CENTER STRUCTURE

R&D Galician Plan 2006-2010

- Life Sciences
- Nanotechnology
- Ocean Sciences
- New Energies

Strategic Actions R&D Spanish National Plan 2008-2011

- Biotechnology.
- Nanotechnology.
- ICT.
- New Energies.
- Health.
- Biotechnology.
- New Energies and Climate Change.
- ICT.
- Nanoscience and Nanotechnology.

R&D Excellence Centers

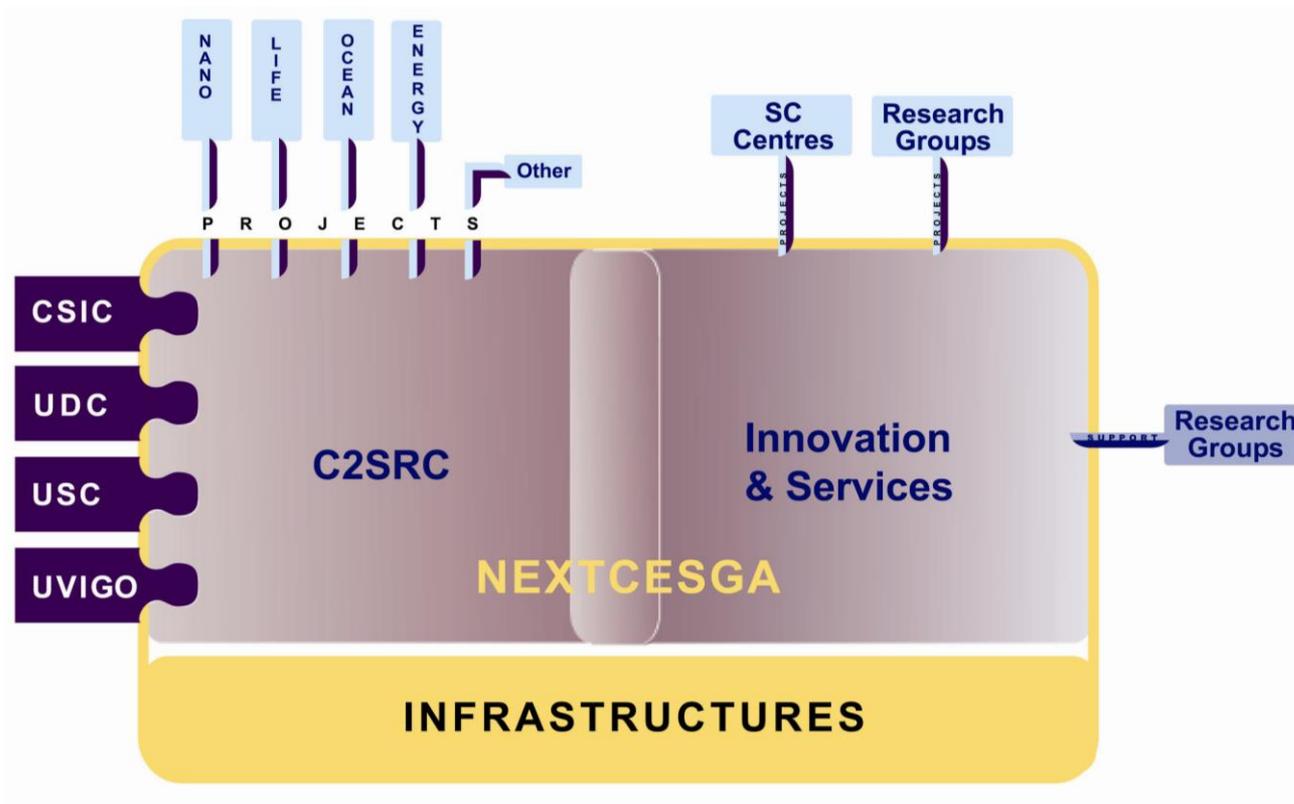
- Centro de Investigación en Ciencias del Mar.
- Centro de Investigación en Ciencias y Tecnologías de la Vida.
- Centro de Electrónica para Vehículos Inteligentes.
- Centro Hispano-Portugués de Investigación en Nanotecnología.

HPC + Simulation

Application Areas
CESGA - C²SRC



CENTER STRUCTURE FOR THE FUTURE



CENTRO DE SUPERCOMPUTACIÓN DE GALICIA

CESGA-ADVANCED SERVICES And INNOVATION



CURRENT CESGA's COMMUNITY OF USERS

- **Galician Universities**
- **Galician Regional Government Research Centres**
- **Spanish National Research Council (CSIC) Centres**
- **Other public or private organizations worldwide**
 - Hospital R&D Departments
 - Industries R&D Departments
 - Technological & Research Centres
 - Other Universities worldwide
 - Non-profit R&D organizations



- **HPC, HTC & GRID Computing**
- **User Data Storage**
- **Advanced Communications Network**
- **e-Learning & Collaboration Infrastructures**
- **GIS (Geographical Information Systems)**
- **Technology Transfer to Industry and e-Business**
- **Public Relations Department**

TECHNOLOGY

CESGA'S TECHNOLOGICAL EVOLUTION: INSTALLED SERVERS

1993
VP 2400



2.5 GFLOPS

1998
VPP 300 **AP 3000**



14.1 GFLOPS **12 GFLOPS**

1999
HPC 4500



9.6 GFLOPS

2001
SVG



9.9 GFLOPS

2002
HPC 320 **BEOWULF**



64 GFLOPS **16 GFLOPS**

2003
SUPERDOME



768 GFLOPS

2004, 2005, 2006
SVG



3,142 GFLOPS

2007
FINISTERRAE



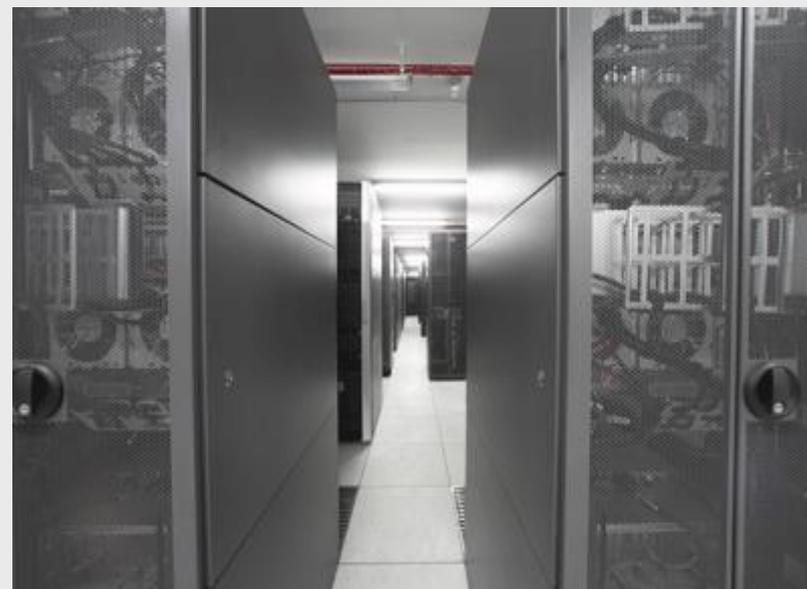
16,000 GFLOPS

Installation Year	1993	1998	1999	2001	2002	2003	2004	2005	2006	2007
Capacity				SVG			SVG	SVG	SVG	
Capability	VP2400	VPP300E AP3000	HPC4500		HPC320	SUPERDOME				FINIS TERRAE

FINIS TERRAE

EXPANDING
THE
FRONTIERS OF KNOWLEDGE

FINIS TERRAE



Spanish National Singular Scientific & Technological Infrastructure

More than: **16,000 GFLOPS**

2,580 CPUs

19,640 GB Memory

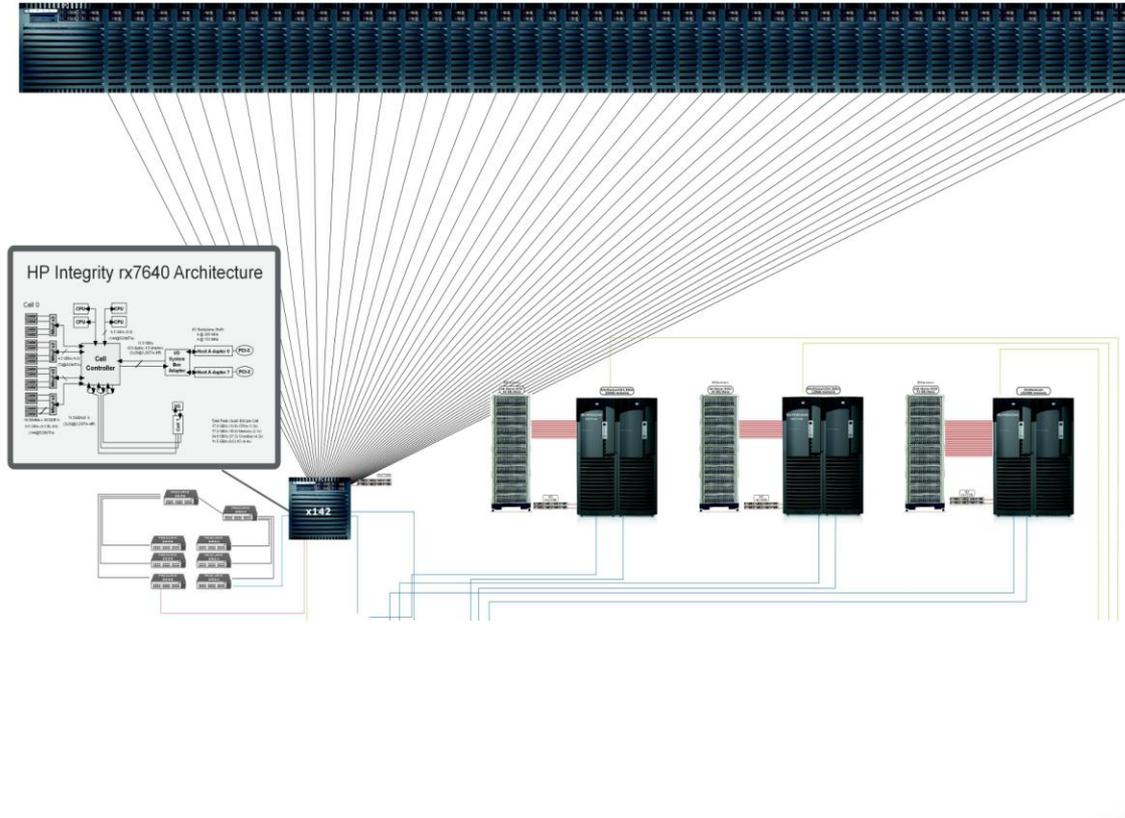
LINUX, UNIX, WINDOWS

ICTS: Unique S & T Infrastructure Required to conduct research that strives for excellence



FINIS TERRAE – COMPUTING NODES

- Fiber Chanel 2Gbits/s
- Fiber Chanel 4Gbits/s
- Infiniband 10 Gbits/s
- 1 GbitEthernet
- FastEthernet
- USCSI 320 MB/S

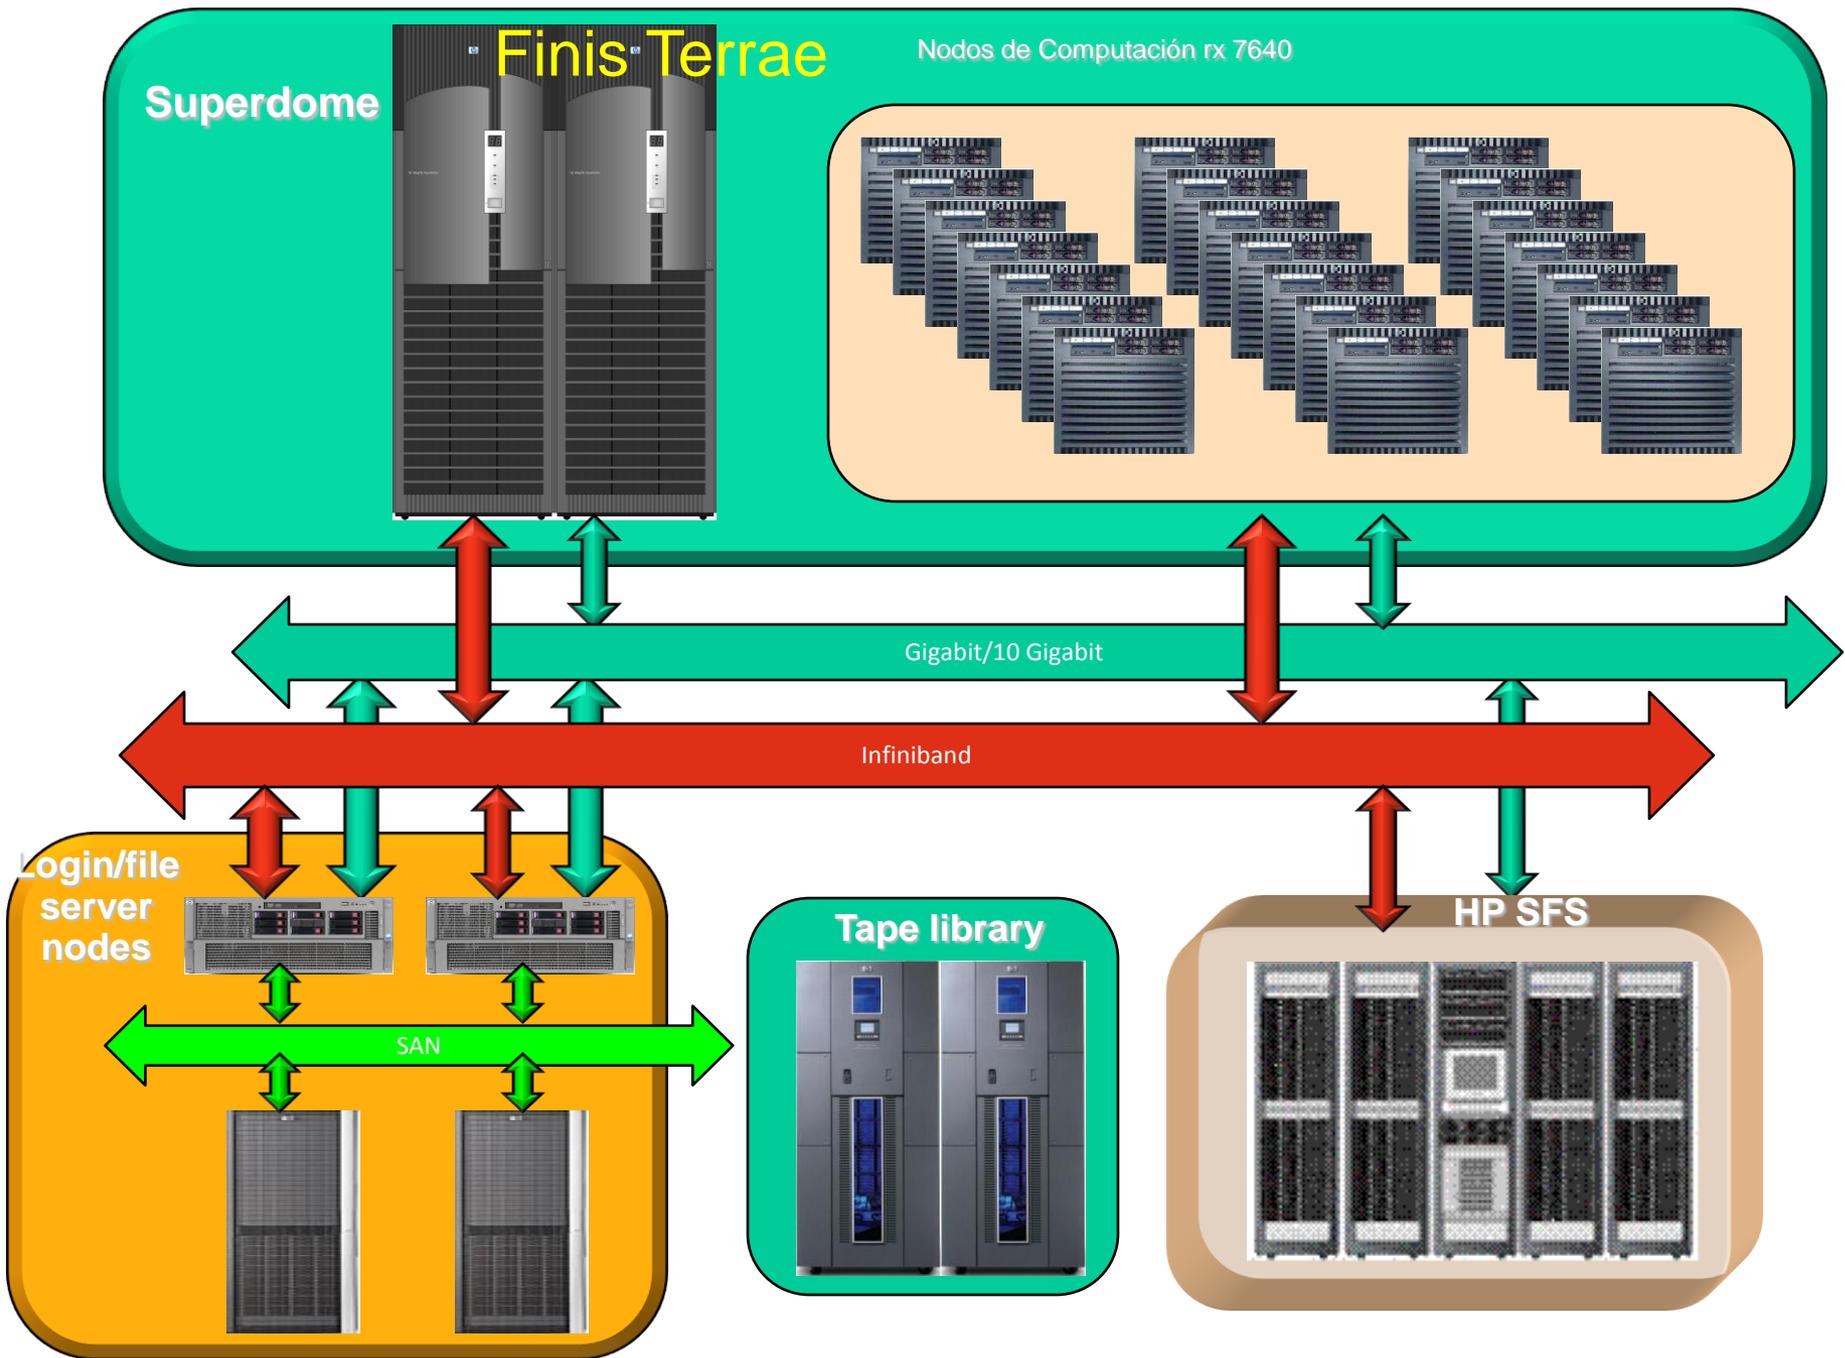


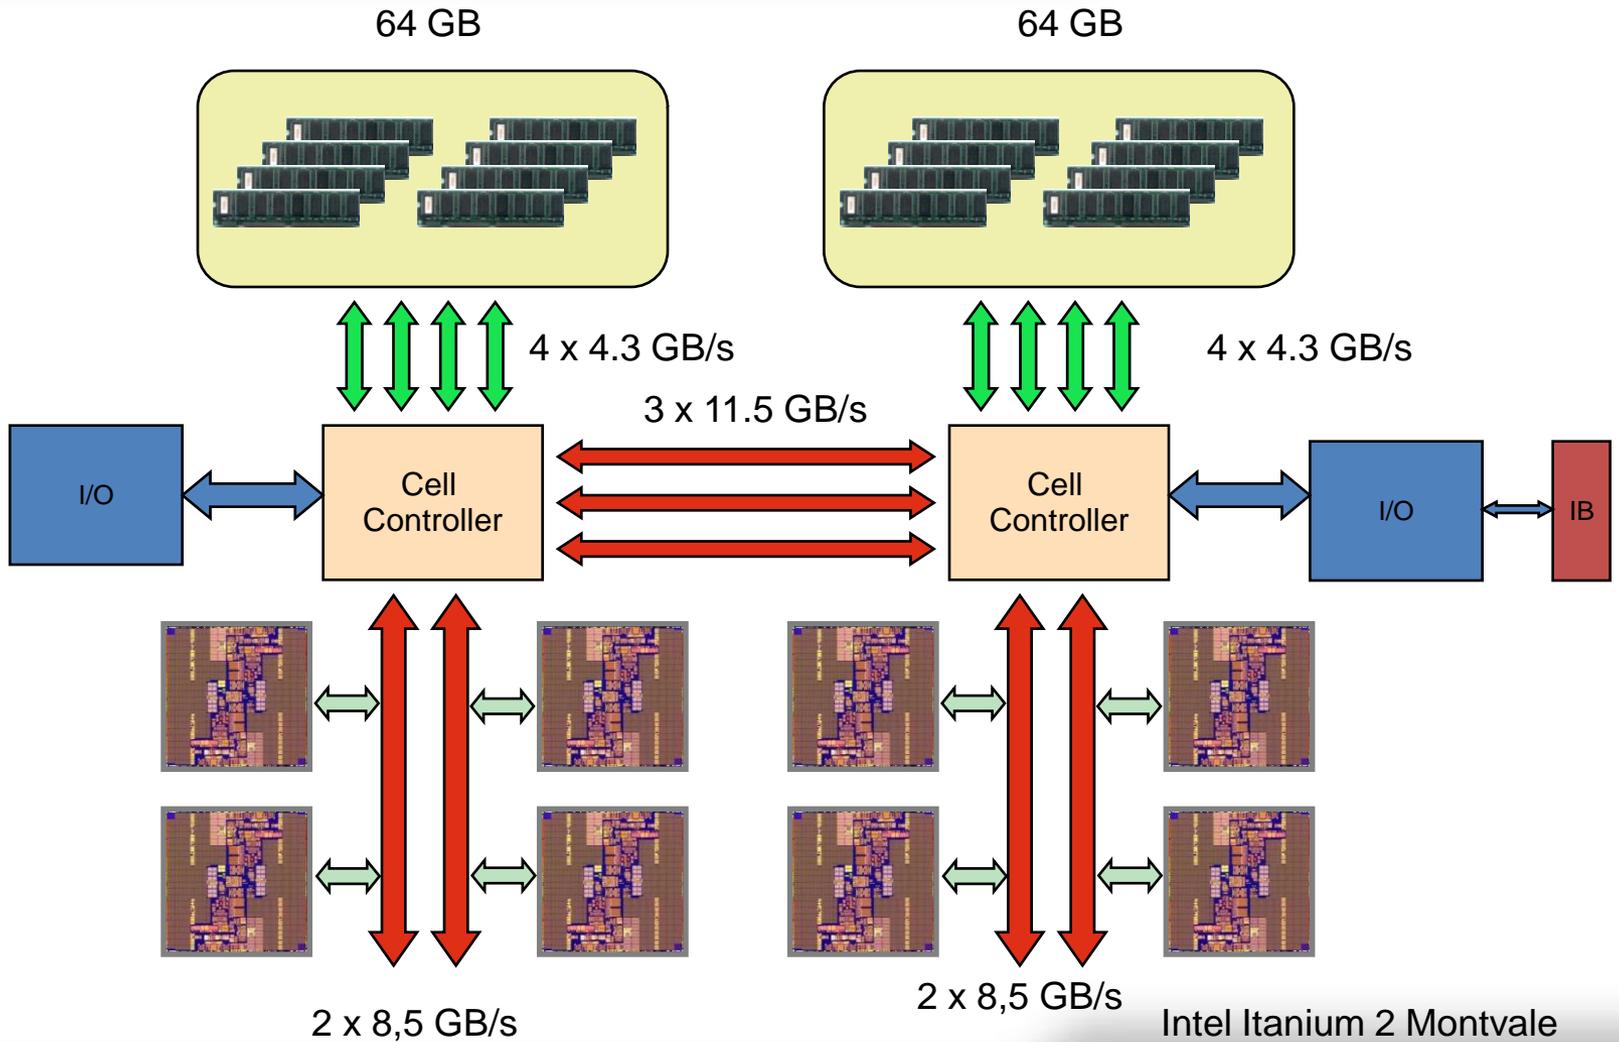
SUPERCOMPUTING NODES:

147 cc-NUMA Nodes with Itanium CPUs connected through a high efficiency INFINIBAND network

- 1 node: 128 cores, 1,024 GB memory
- 2 nodes: 64 CPUs, 128+256 GB memory
- 142 nodes: 16 cores, 128 GB memory
- 2 nodes: 4 cores, 4 GB memory for testing

SANTIAGO DE COMPOSTELA, SPAIN,

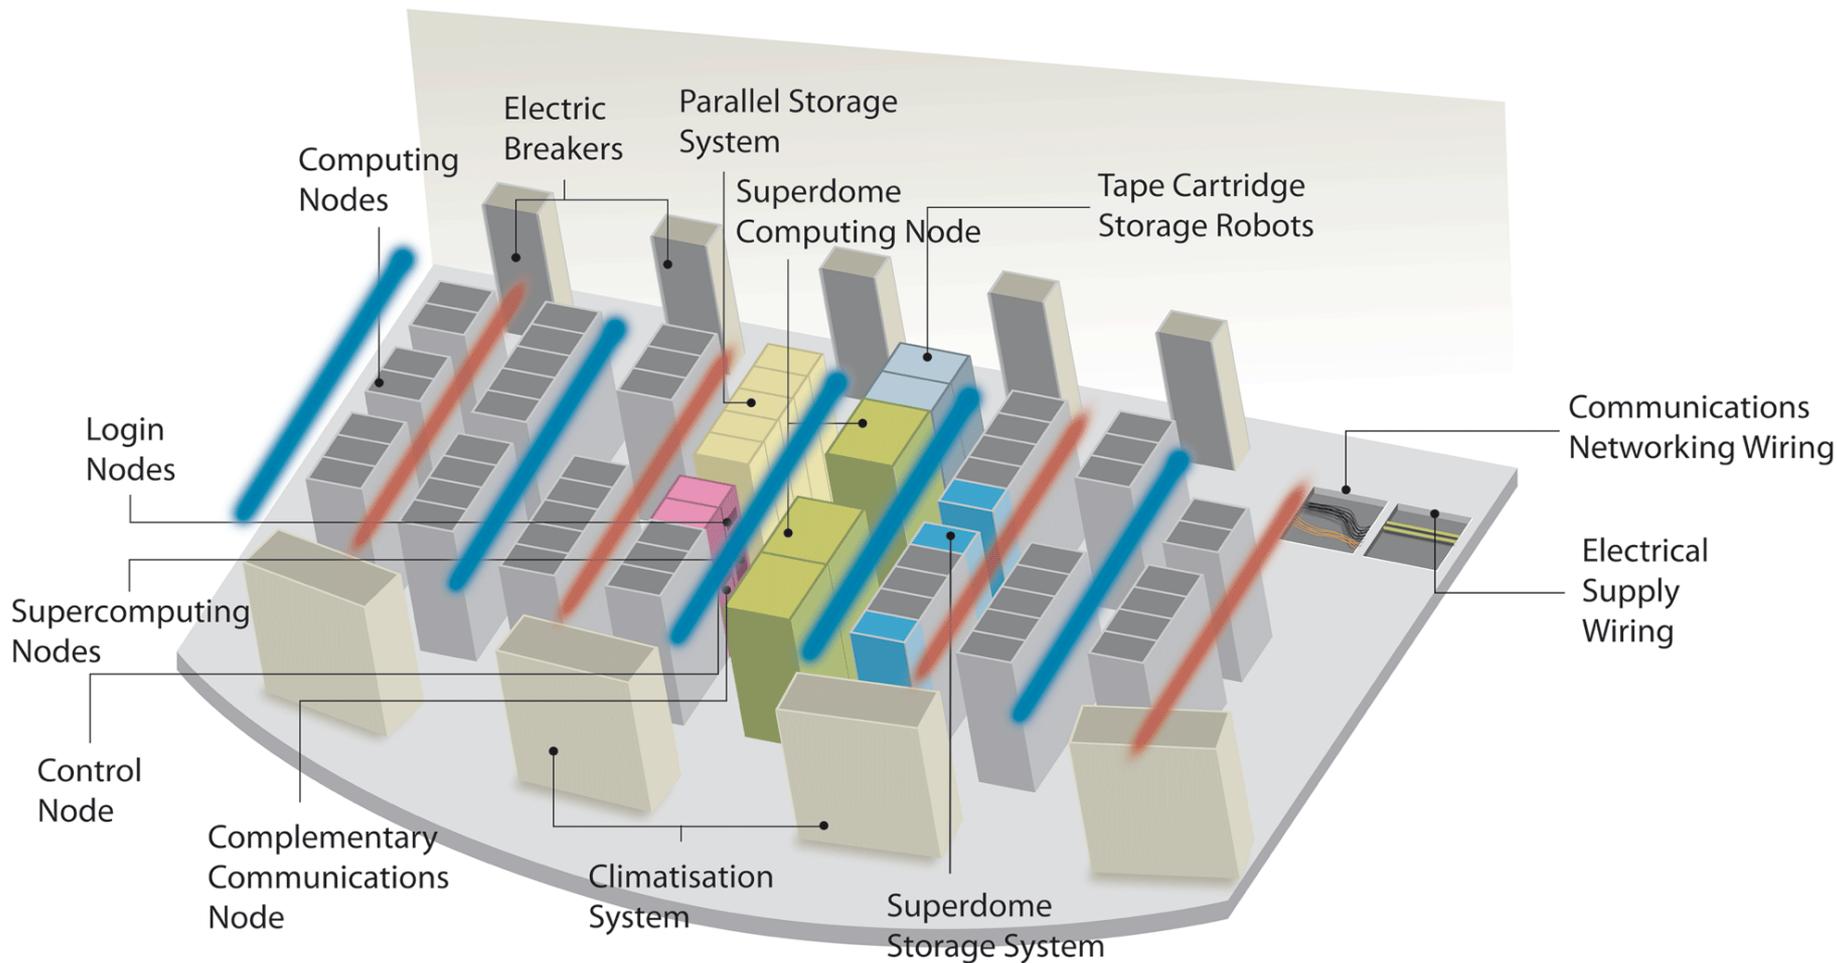




Intel Itanium 2 Montvale



The Finis Terrae Cluster



Technical Specs.

Surface Area: 140 m²

Weight:



35.000 Kg

Storage:



2.200.000 GB on tape

390.000 GB on disk

Memory:



19.670 GB

2.528 Processing Cores

142 nodes, each with 16 cores & 128 GB memory

1 node with 128 cores & 1.024 GB memory

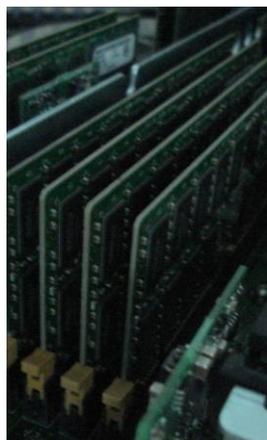
1 node with 128 cores & 384 GB memory

Node Interconnect INFINIBAND
4x DDR at 20 Gbps

85 Km of interconnect cable

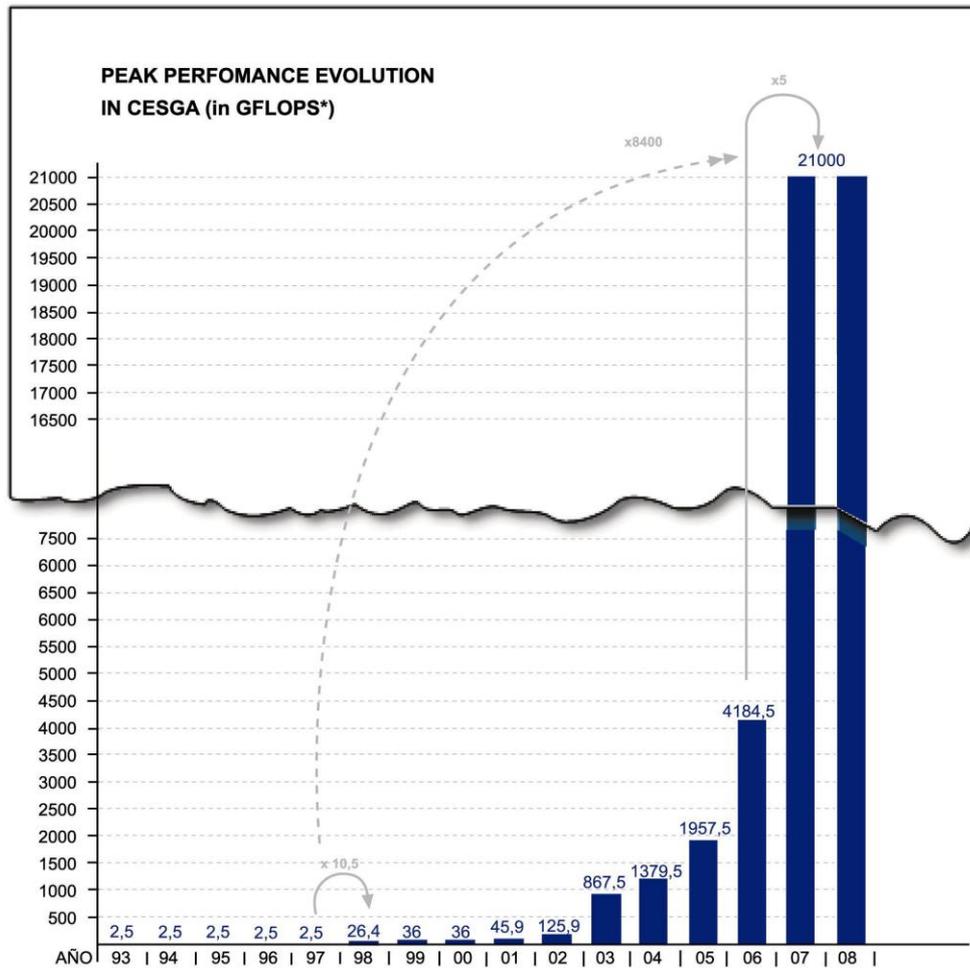
Open Software: Linux, Lustre, Globus...

FINISTERRAE





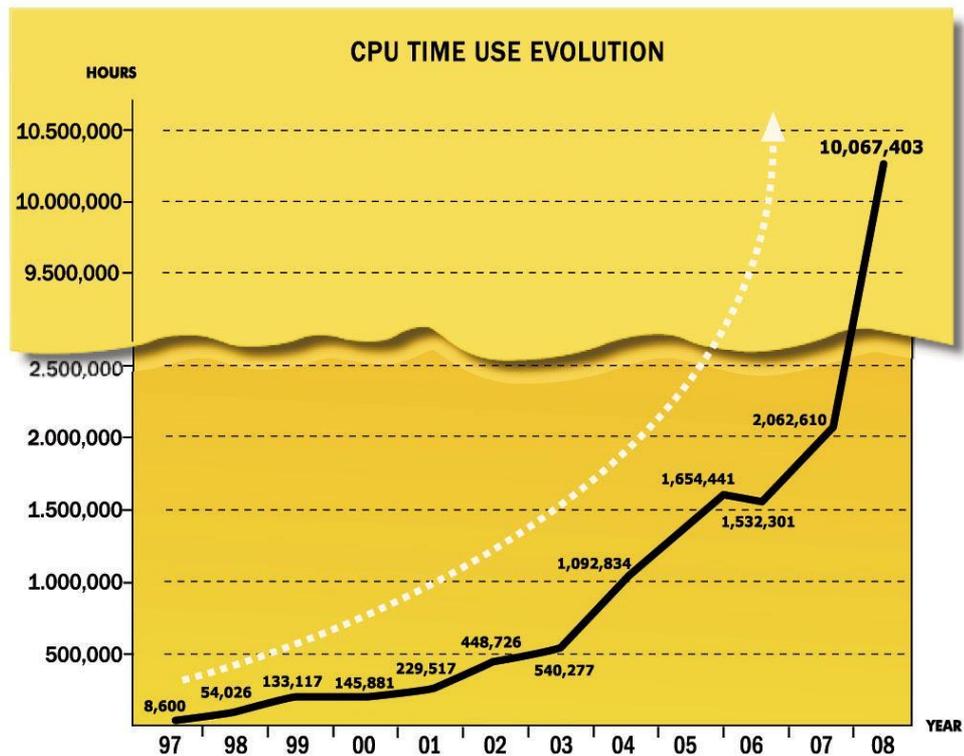
CESGA'S PEAK PERFORMANCE EVOLUTION



* 1GFLOPS= 1.073.741.824 (floating point operations per second)



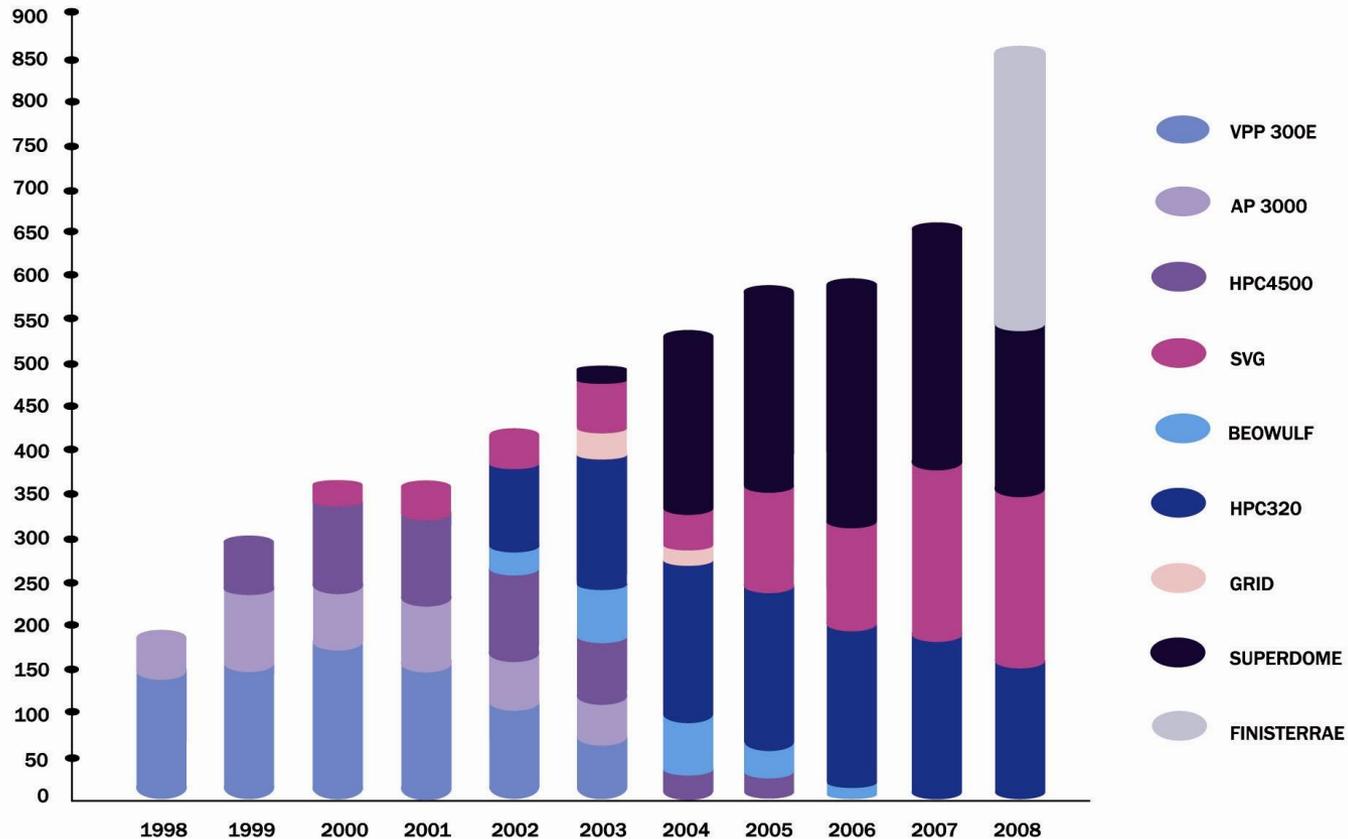
CPU TIME USE EVOLUTION AT CESGA



DEMAND OF COMPUTING RESOURCES AT CESGA

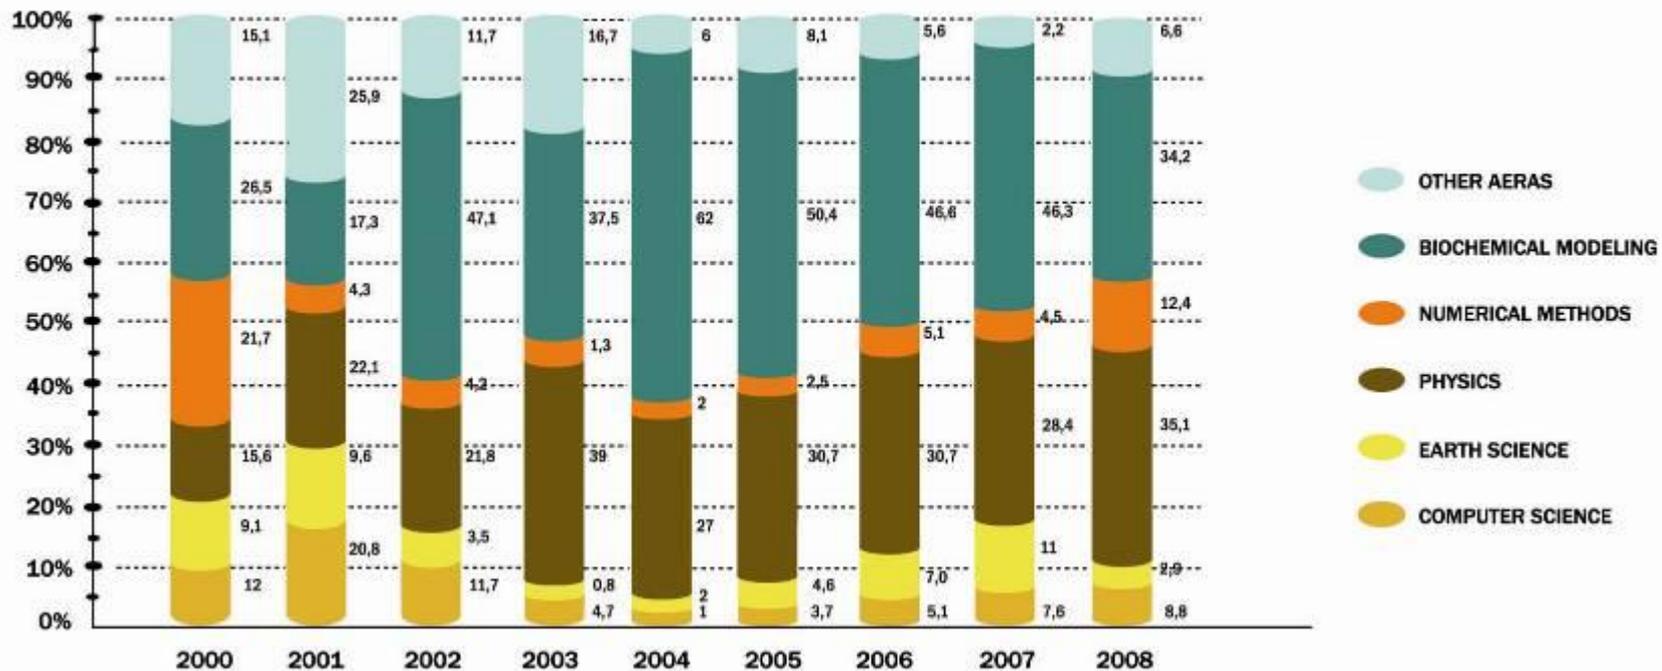


USER ACCOUNT EVOLUTION PER SYSTEM PER YEAR

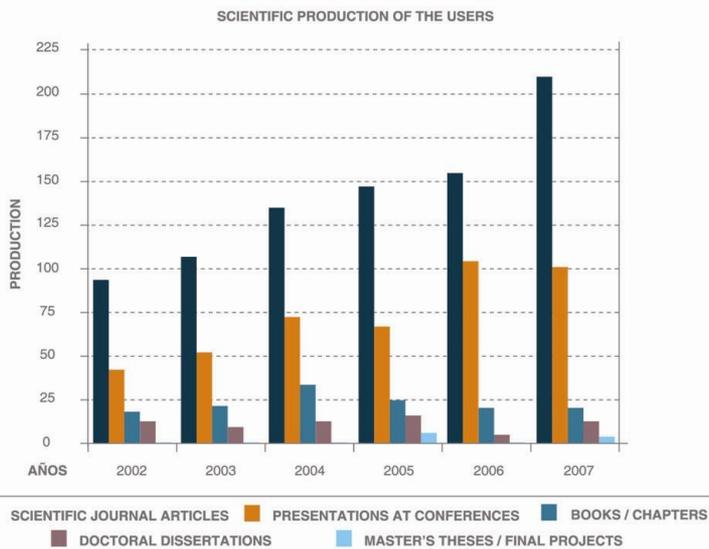


CPU USE DISTRIBUTION BY RESEARCH AREAS

CPU USE DISTRIBUTION BY RESEARCH AREA



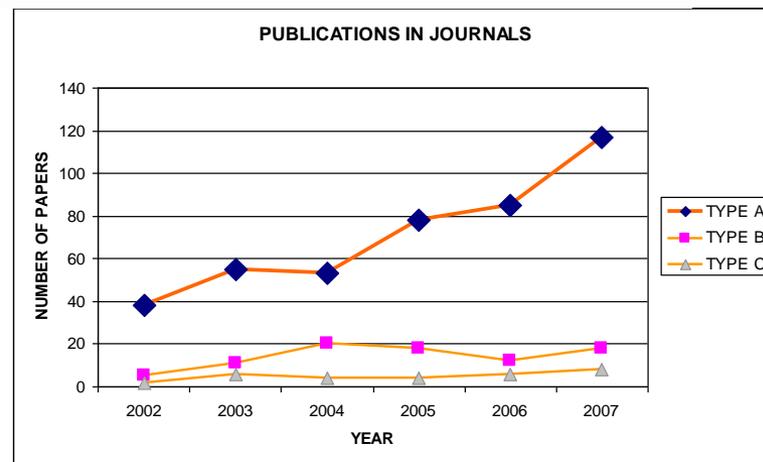
CESGA USERS'S SCIENTIFIC PRODUCTION 2002-2007



DISTRIBUTION OF SCIENTIFIC PRODUCTION PER YEAR

PRODUCT	2002	2003	2004	2005	2006	2007
SCIENTIFIC ARTICLES	93	107	134*	147	154	208
ACCEPTED	4	8	30	18	28	19
SUBMITTED	16	29	28	27	23	34
PUBLISHED	73	70	91	104	108	185
PRESENTATIONS IN CONFERENCES	42	52	72*	65	105	101
PRESENTATIONS	33	40	53	37	53	46
POSTERS	9	12	19	28	52	47
DOCTORAL THESES	17	21	32	24	18	18
DEFENDED	9	4	5	9	10	11
PRESENTED	0	0	0	0	3	5
IN PROCESS	8	17	27	15	5	3
MASTERS THESES GRADUATE PROJECTS	12	9	12	14	5	10
DEFENDED	6	4	5	13	3	8
IN PROCESS	6	5	7	1	2	2
BOOKS / CHAPTERS	NA	NA	NA	5	4	3
ACCEPTED / IN PRESS	NA	NA	NA	4	1	1
SUBMITTED	NA	NA	NA	1	0	2
PUBLISHED	NA	NA	NA	0	3	0
TOTAL	169	194	256	260	301	351

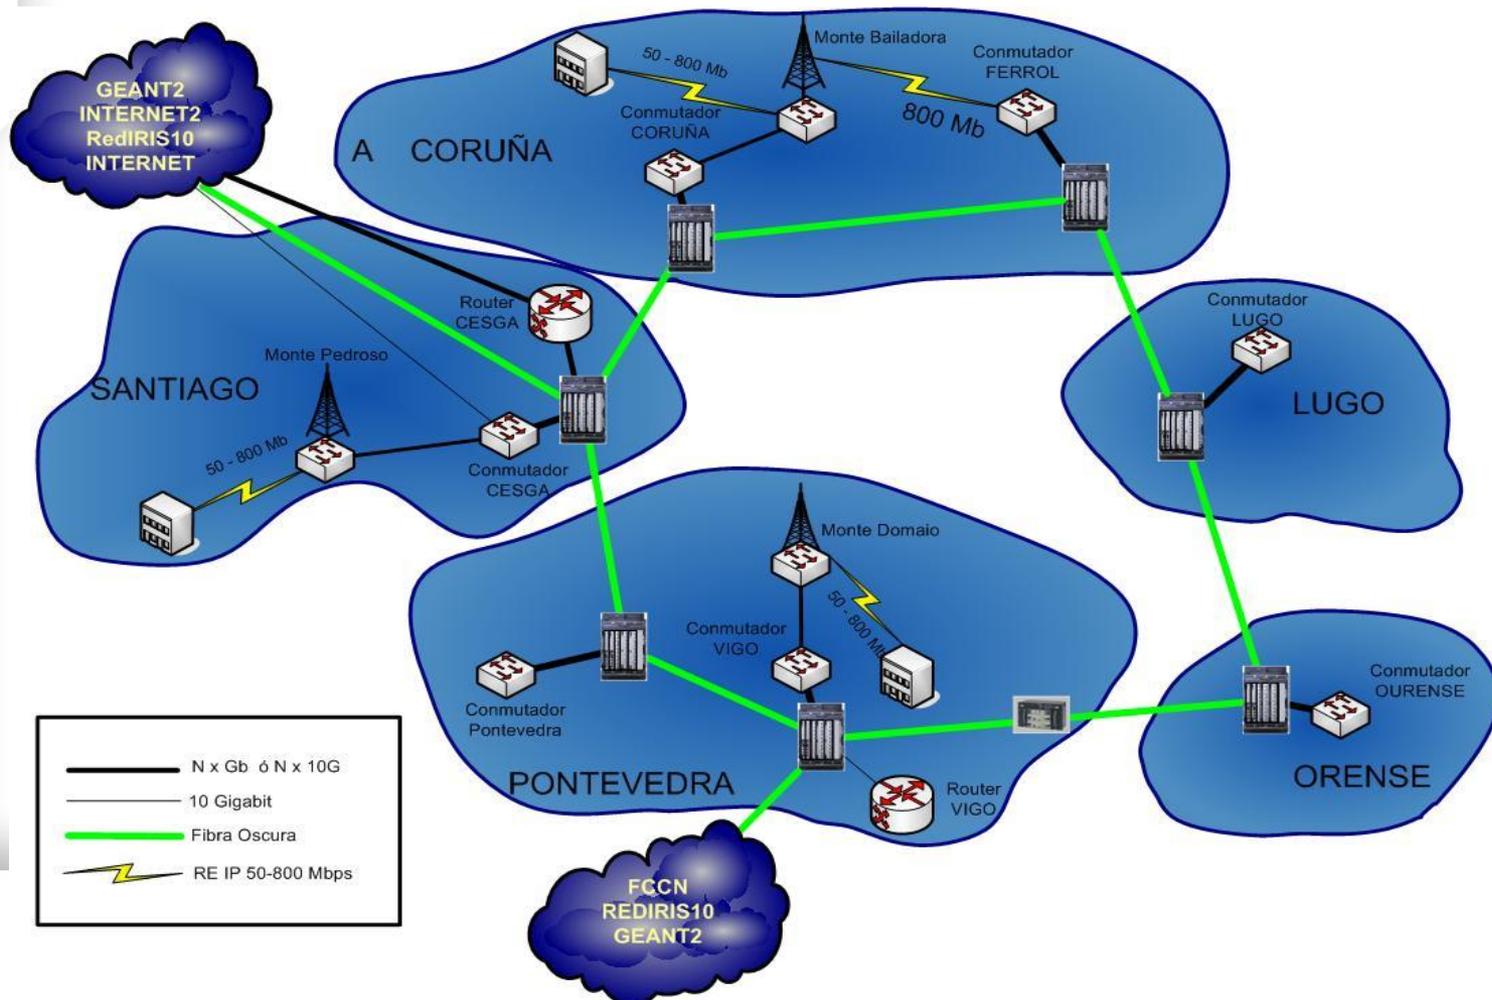
* 6 PUBLICATIONS & 16 CONFERENCES WITH AUTHORS FROM MORE THAN ONE INSTITUTION



RECETGA NOVA

58 points of presence in Galicia

RECETGA-Rediris Link: Dark fiber (Increase ~ 80 times)



CESGA - C²SRC

CESGA STRATEGIC PLAN 2010-13

Achieving and maintaining **international recognition** as a high level and high quality supercomputing facility and as a Research Centre of Excellence in the area of **Computational Science**

KEY ACTIONS

☐ Upgrade the computing and associated facilities

- Design, building and installing Finis Terrae II;
- Upgrade of SVG (the CESGA HTC computer);
- Design and deployment of a new on-line storage;
- Actualization of RedIris Nova lines;
- New visualization and collaboration facilities;
- Security infrastructure to guarantee a good usage of the systems.

☐ Start and foster new high added-value services for scientists:

- Specialized storage and archiving services,
- Improve scientific development application support,
- Specialized visualization services

☐ Improve IT departments' research and support activities:

- E-learning and other collaboration tools,
- Innovation and technological support.

CHANGES IN ORGANIZATION

❑ **Improvement of and new services**

- Scientific application support
- Storage and archiving support .

❑ **Focusing current services toward computational science:**

- E-Learning
- E-Business and technological transfer
- GIS

TECHNOLOGICAL PLAN: INVESTMENTS

SUPERCOMPUTING. FINIS TERRAE II:

□ Year 2H2010-11

- Constitution of a scientific and technical committees to determine the best architecture of Finis Terrae II:
 - Gather information about users current and future users requirements.
 - Gather information about technology: industry and other Centres.
 - Proposal of the new supercomputer: scientific, technical, logistic, and economical conditions.

□ Year 2012:

- Call for tenders.
- Deployment of Finis Terrae II

□ Year 2013:

- Finis Terrae II: open production.

HTC/GRID:

□ Year 2010:

- Upgrade of HTC/GRID servers: 1.000 cores

CARE AND DATA STORAGE:

□ Year 2010:

- Update of storage subsystem: Very High throughput, permanent storage of large datasets, multiprotocol

COLLABORATION AND VISUALIZATION TOOLS:

□ Year 2012:

- Deployment of two AccessGrid Rooms
- Deployment of one Auditorium in the new building
- Virtual reality: one AccessGrid room with VR.

LOCAL AREA NETWORK:

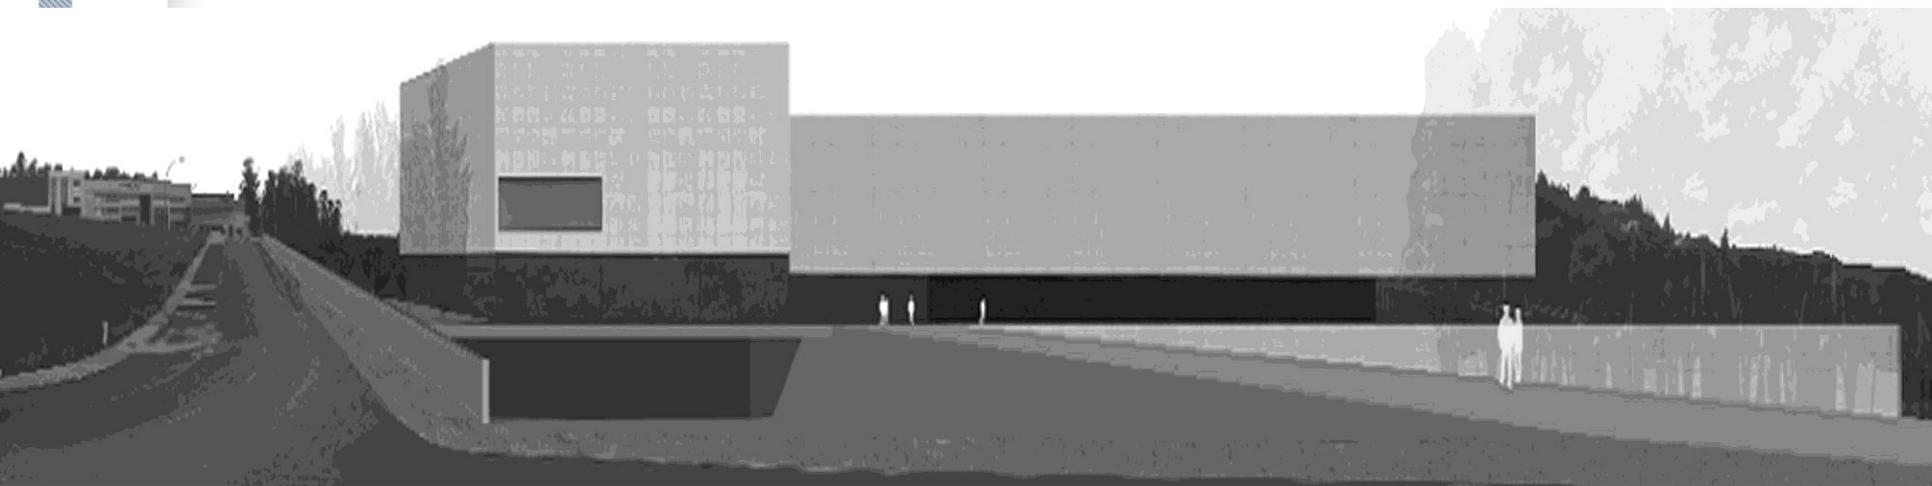
□ Year 2012:

- Deployment of new internal local area network. Improving: bandwidth, availability, serviceability, reliability, security, integrity and quality



CERTIFICATED FIRM

NEW FACILITIES



XUNTA DE GALICIA
CONSELLERÍA DE INNOVACIÓN,
E INDUSTRIA



CONSEJO SUPERIOR
DE INVESTIGACIONES
CIENTÍFICAS



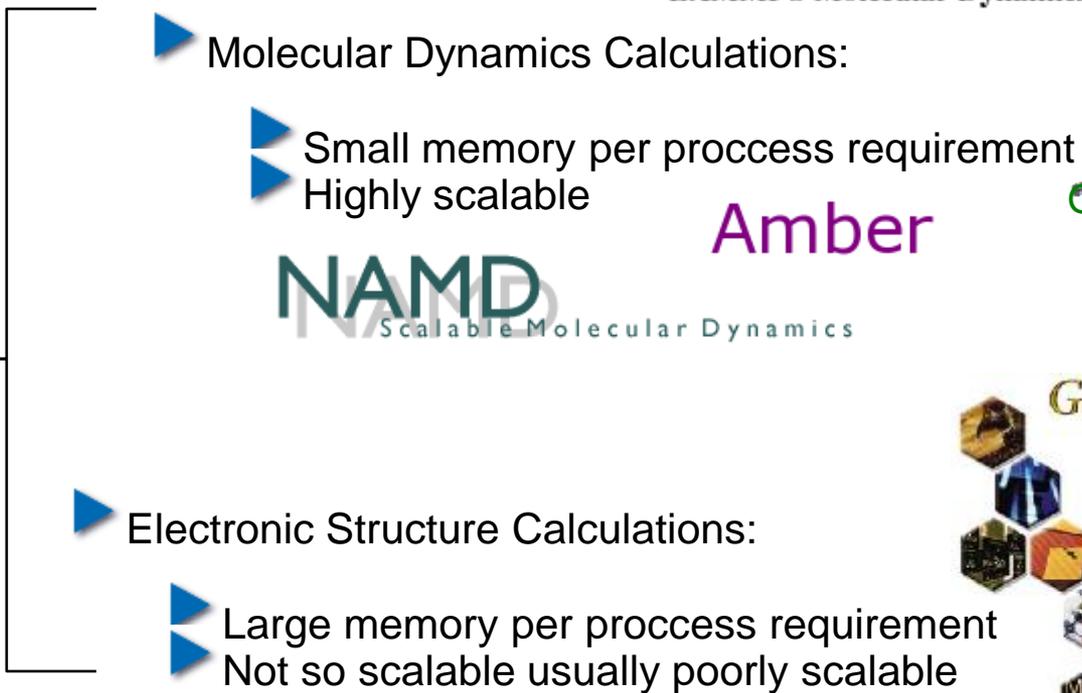
MINISTERIO
DE CIENCIA
E INNOVACIÓN



FEDER
FONDO EUROPEO DE
DESENVOLVEMENTO REGIONAL

Molecular Simulation Applications

abinit
Amber
CPMD
Dalton
Gamess
Gaussian 03
Gaussian 98
Gromacs
LAMMPS
Molden
NAMD
NWChem
Octopus
Schrodinger Suite
SIESTA



LAMMPS Molecular Dynamics Simulator

NAMD
 Scalable Molecular Dynamics

Amber



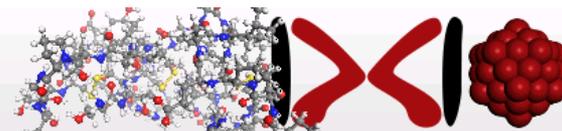
Gaussian 03



See up-to-date list in WWW.CESGA.ES

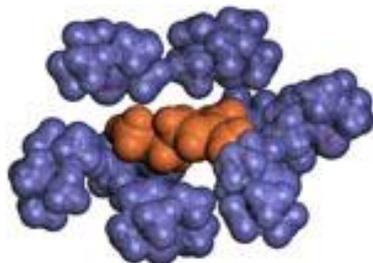
Users Licensed Applications installed at CESGA

TURBOMOLE GmbH



Molpro quantum chemistry package

Materials Studio



It is possible in many cases run private applications on CESGA servers (the license must admit it)

Application Use

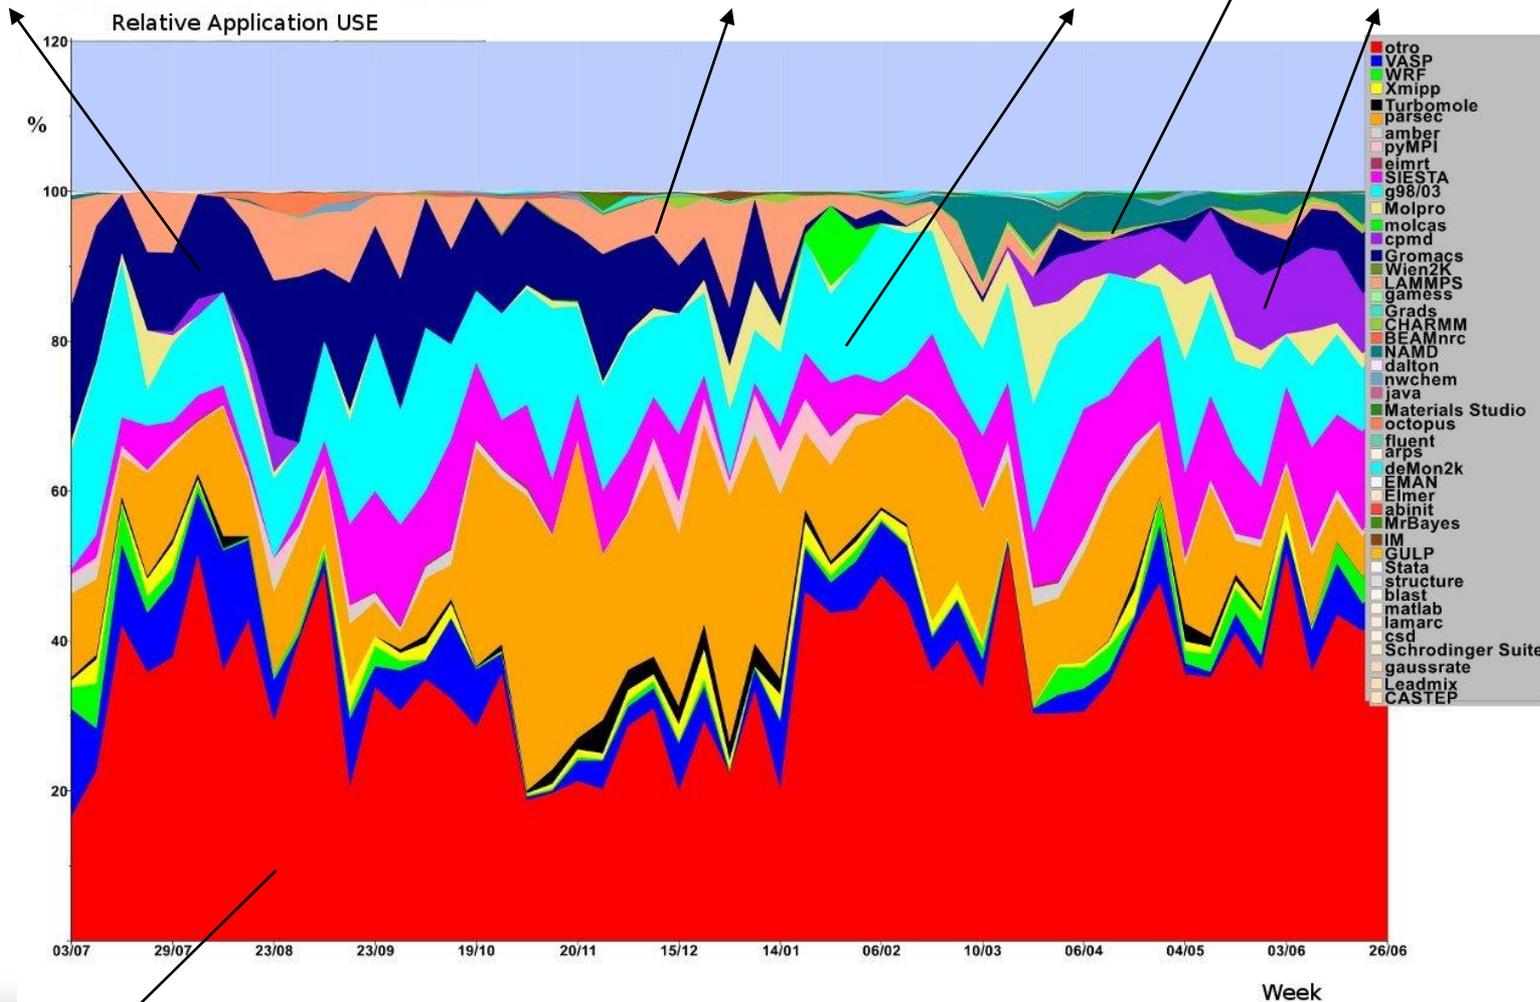
Gromacs

LAMMPS

G98/03

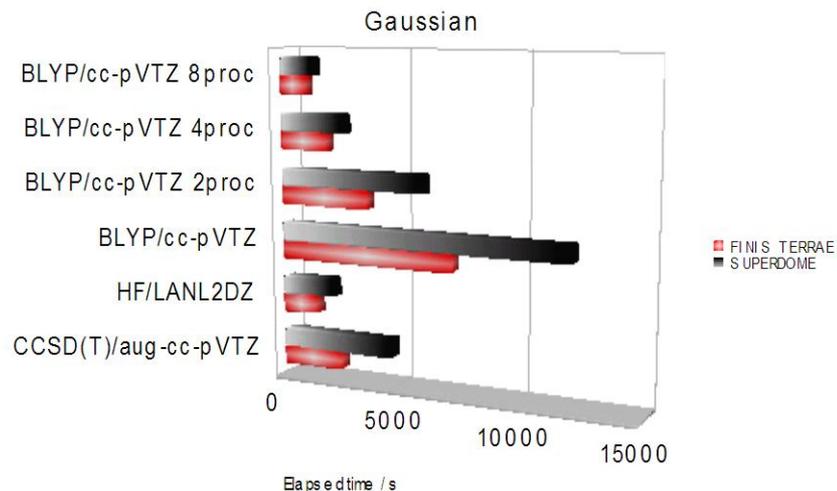
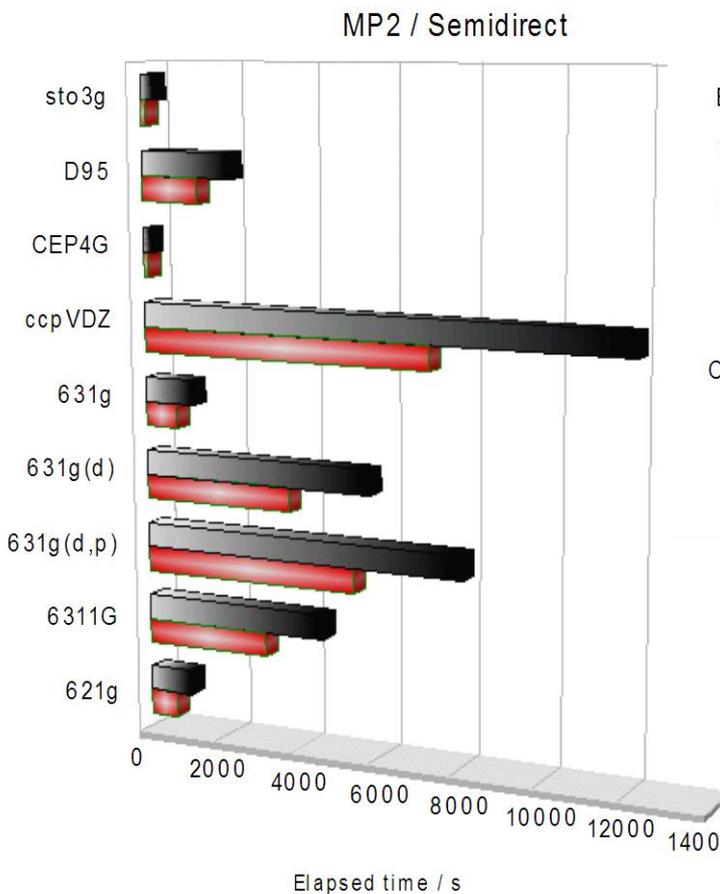
NAMD

CPMD



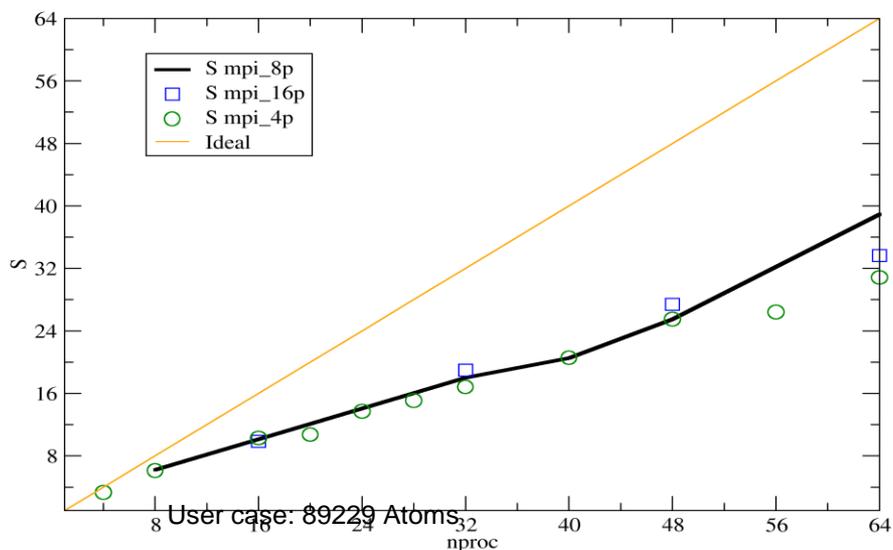
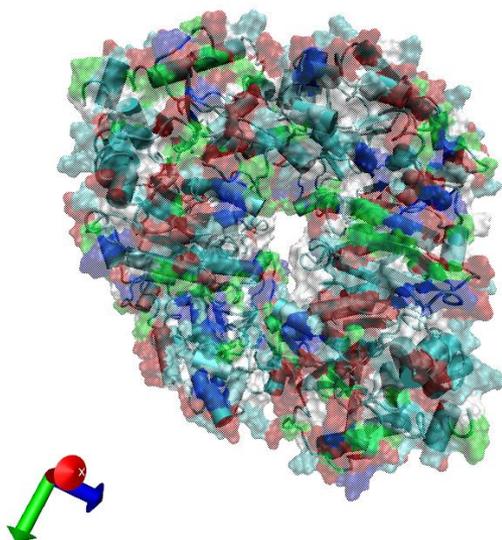
User codes

Electronic Structure – Gaussian Performance



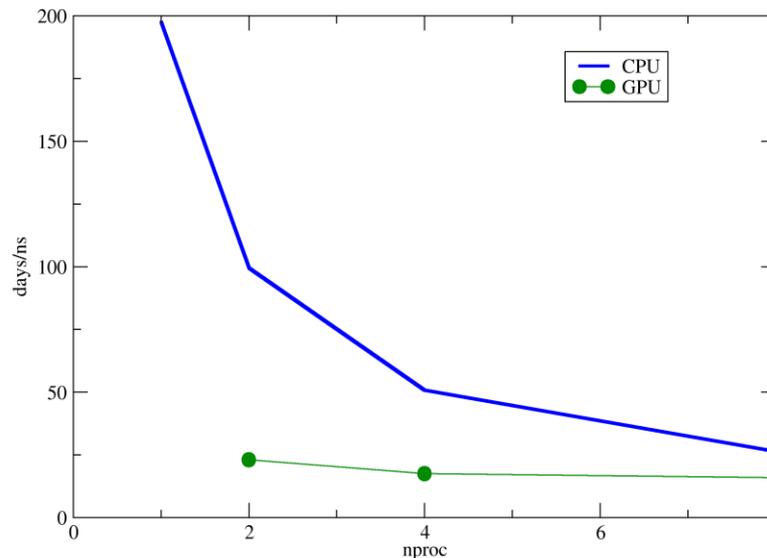
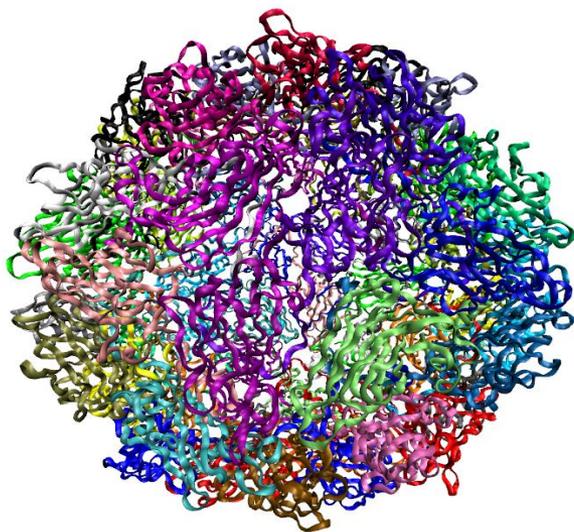
Improved performance over 1 core ia64
 Good scaling inside a node (OpenMP) using direct calculations

- ▶ Gromacs: Highly Efficient Molecular Dynamics
 - ▶ Usually poor scaling. Improved in version 4.0
 - ▶ Optimum number of PME nodes (all to all communications)
 - ▶ Thick nodes problem (how many processes per node?)



New Architectures (GPUs): NAMD Performance

NANoscale Molecular Dynamics (NAMD) is a highly scalable molecular dynamics (MD) code.
 NAMD is especially geared towards the simulation of large biomolecular systems.
 Benchmark: Satellite Tobacco Mosaic Virus (STMV) (1,066,628 atoms)
 Platform: A 8 cores Nehalem E5520 with a NVIDIA Tesla C1060



Single GPU acceleration: 1GPU = 4-6 cores

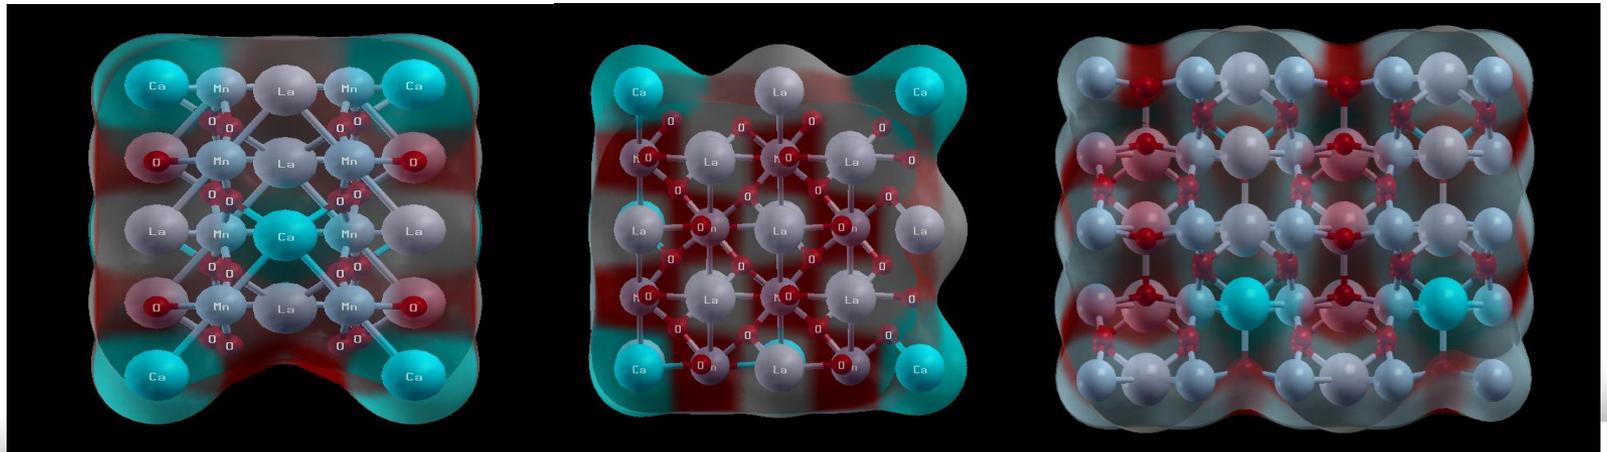
Finis Terrae : Grand challenges

- Different scientific fields / know-how
- Coming from different groups / Institutions / research projects
- Computationally highly demanding

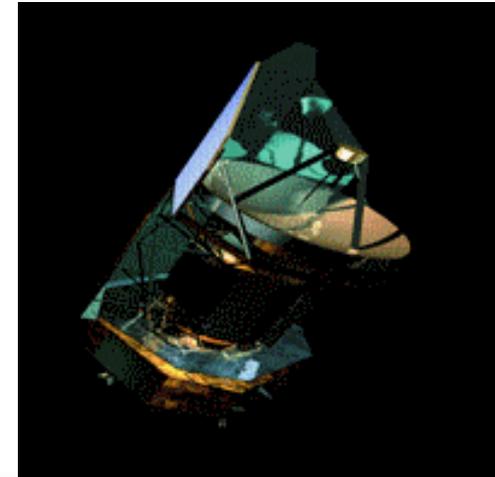


"PHASE SEPARATION"

- **Condensed Matter Physics: Materials Design**
- **3rd/4th physics most important problem of decade (AIP)**
- **Wien2k: electronic structure calculations of solids using density functional theory (DFT).**
- **Two levels parallelism:**
 - **Coarse grain: k points**
 - **Fine grain: ScaLapack**
- **24 nodes: 384 cores, 100G memory, 300G disk -> 68.000h**

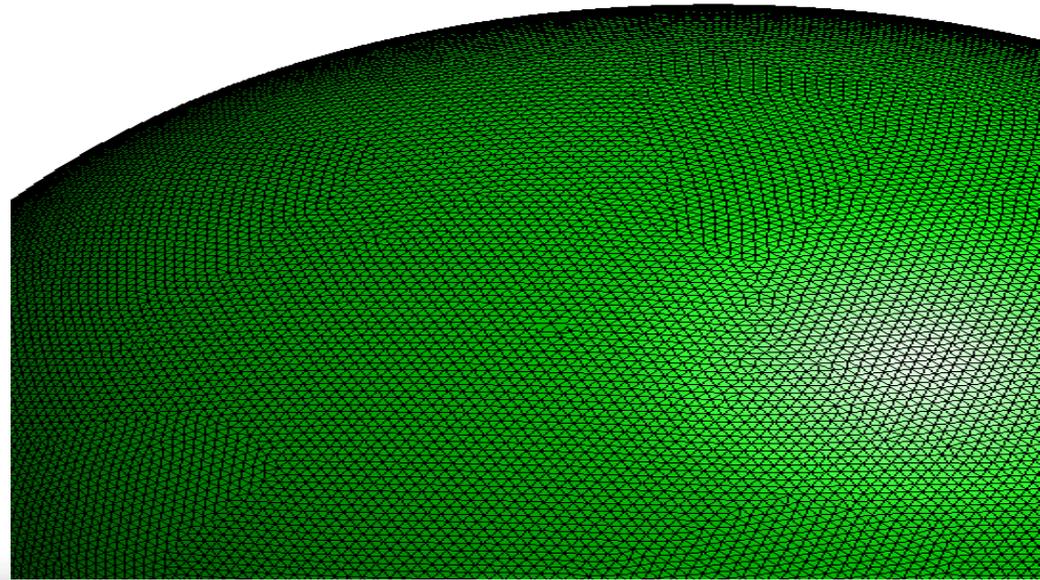
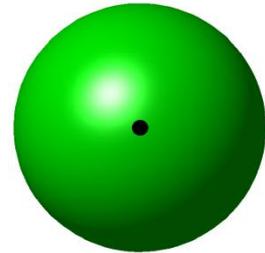


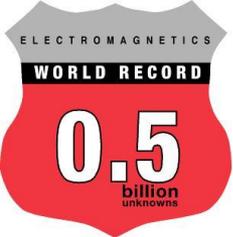
- **Molecular and Infrared Astrophysics**
- **Genetic Algorithms: PIKAIA multimodal optimization problems / FASTWIND**
- **MPI master slave schema:**
 - **master task took care of the GA-related operations**
 - **slave tasks to perform the model calculations**
- **20.000 – 40.000 models (15 min each)**
- **80 cores**



“FEKETE POINTS”

- **I-MATH: Applied Maths (Potential theory/Numerical methods)**
- **Forces Method -> Estimation of Fekete points**
- **7th Smale: Mathematical problems for 21st century**
- **Several parallelizations paradigms:**
 - **MPI**
 - **OpenMP**
 - **MPI/OpenMP**
 - **High Throughput**
- **1024 cores**
- **350.000 hours**
- **50 million executions**





HEMCUVE"

- Electromagnetic problems in large structures
- HEmCUVE++: Electromagnetic calculations based on FAST MULTIPOLE methods
- MPI/OpenMP
- Highly demanding on memory per process: Balance between CPU time and memory
- 1024 cores, 6T of memory

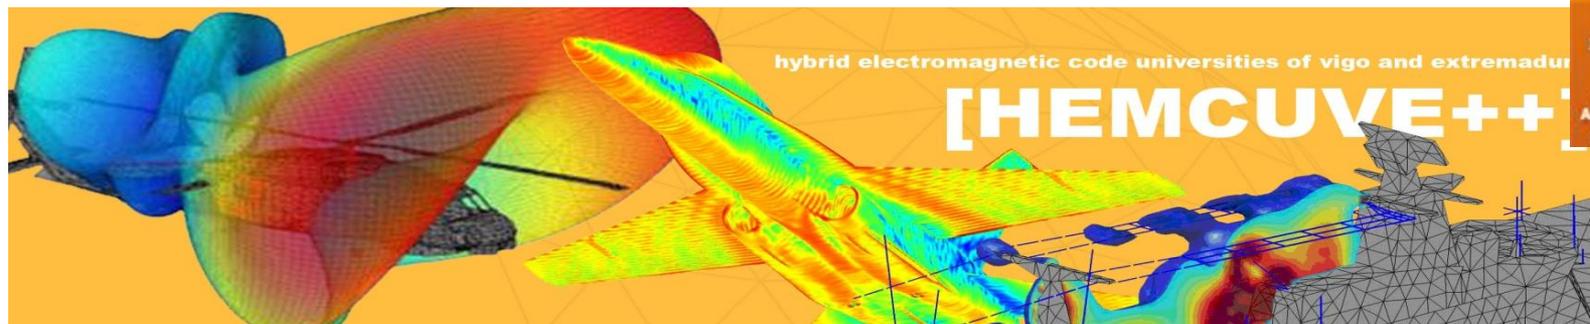


Award 2009

presented at

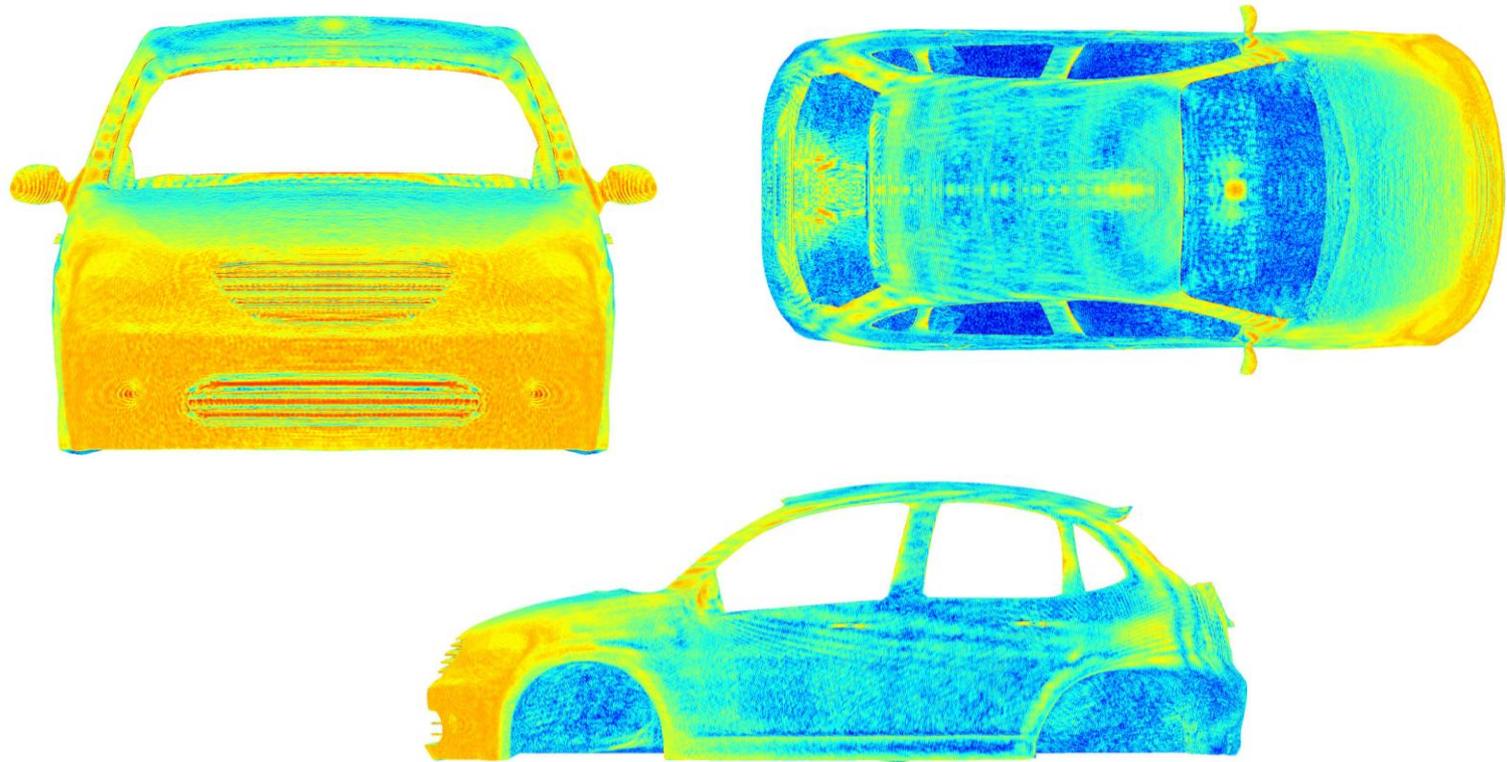


INTERNATIONAL
SUPERCOMPUTING
CONFERENCE



“HEMCUVE”

Citröen C3 at 24.125 GHz (radar frequency) 40M unknowns



In progress 79 GHz: 300M – 400M unknowns