Barcelona Supercomputer Center
National Supercomputing Facility

Objectives

- Research in Supercomputing and Computer Architecture
- Collaborate in R&D e-Science projects with prestigious scientific teams
- Manage BSC supercomputers to accelerate relevant contributions to research areas where intensive computing is an enabling technology

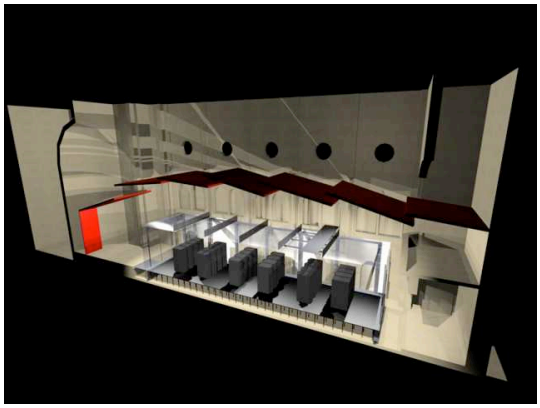

Barcelona Supercomputer Center
National Supercomputing Facility

Structure

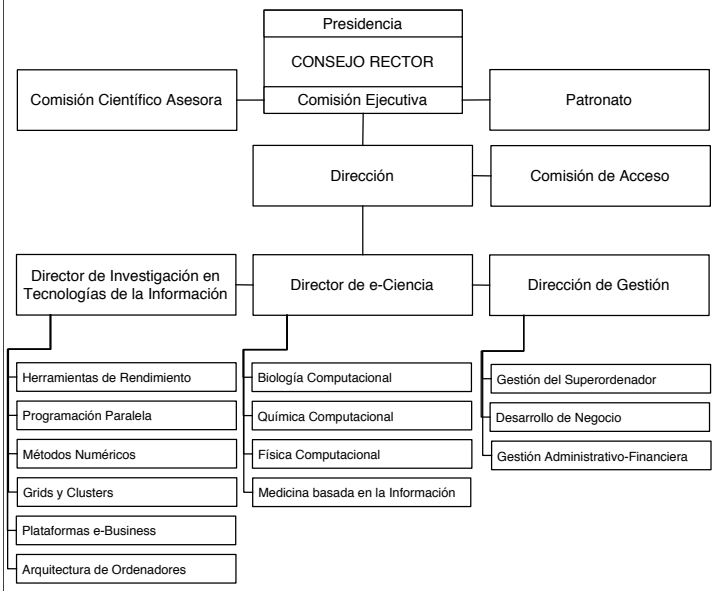

Jornadas Técnicas RedIRIS 2004 Toledo, 29 de Octubre de 2004

Barcelona Supercomputer Center
National Supercomputing Facility

IT research and development projects

- Continuation of CEPBA (European Center for Parallelism in Barcelona) research lines
- Deep Computing
  - Performance Tools
  - Parallel programing
  - Grid
  - Code Optimization
- Computer Architecture
  - Vector processors
  - Network processors

Barcelona Supercomputer Center
National Supercomputing Facility

e-Science projects

R&D collaborations

- Computational Biology
- Computational Chemistry
- Computational Physics
- Information based Medicine

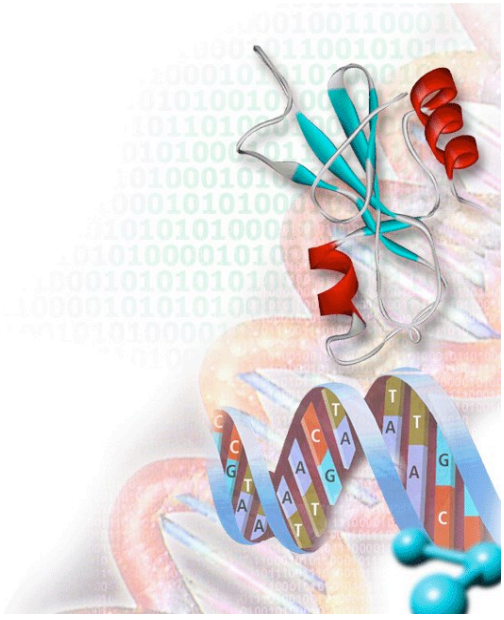

Management projects

- Supercomputer Management
  - System Administration
  - Users support
- Business Development
  - External Relations
  - Promotion
  - Technology Transfer
  - Education
- Administration
  - Accounting and Finances
  - Human Resources

Barcelona Supercomputer Center
National Supercomputing Facility

MareNostrum / Iberus

- PowerPC 970 FX processors (dual processors)
- 4GB ECC 333 DDR memory per node
- 3 networks
  - Myrinet
  - Gigabit
  - 10/00 Ethernet
- Diskless network support
- Linux cluster

Barcelona Supercomputer Center
National Supercomputing Facility

MareNostrum: Overall system description

27 Compute Racks (RC01-RC27)
- 162 BC chassis w/OPM and gigabit ether switch
- 2268 JS20+ nodes w/myrinet daughter card

4 Myrinet Racks (RM01-RM04)
- 10 clos256+256 myrinet switches
- 2 Myrinet spines 1280s

7 Storage Server Racks (RS01-RS07)
- 40 p615 storage servers 6/rack
- 20 FastT 100 3/rack
- 20 EXP100 3/rack

1 Operations Rack (RH01)
- 7316-TF3 display
- 2 p615 mgmt nodes
- 2 HMC model 7315-CR2
- 3 Remote Async Nodes
- 3 Cisco 3550
- 1 BC chassis (BCIO)

1 Gigabit Network Racks
- 1 Force10 E600 for Gb network
- 4 Cisco 3550 48-port for 10/100 network

Barcelona Supercomputer Center
National Supercomputing Facility

Processor: PowerPC 970FX

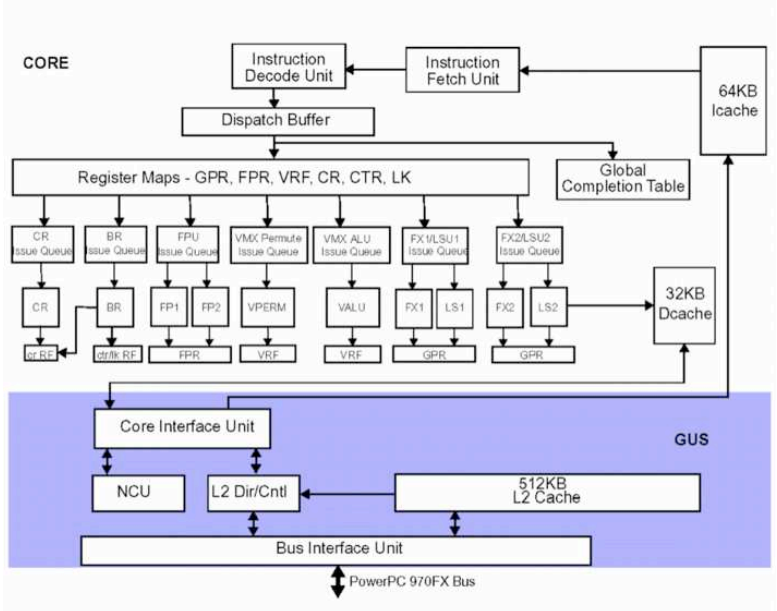

Jornadas Técnicas RedIRIS 2004 Toledo, 29 de Octubre de 2004

Barcelona Supercomputer Center
National Supercomputing Facility

Blades, blade center and blade center racks

Blade Center
- **14 blades per chassis (7U)**
  - **28 processors**
  - **56GB memory**
- **Gigabit ethernet switch**

6 chassis in a rack (42U)
- **168 processors**
- **336GB memory**

Barcelona Supercomputer Center
National Supercomputing Facility

MareNostrum: Floor plan with cabling

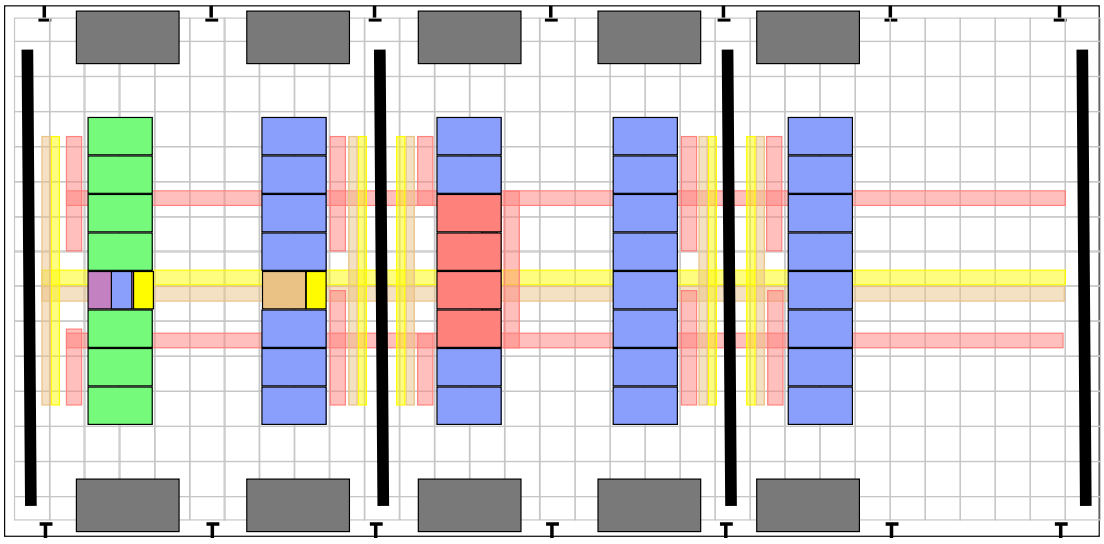

27 racks + 1 BC chassis
- **4564 processors**
- **9TB memory**

Blade centers
Storage servers
Gigabit switch
Myrinet racks
Operations rack
10/100 switches

Barcelona Supercomputer Center
National Supercomputing Facility

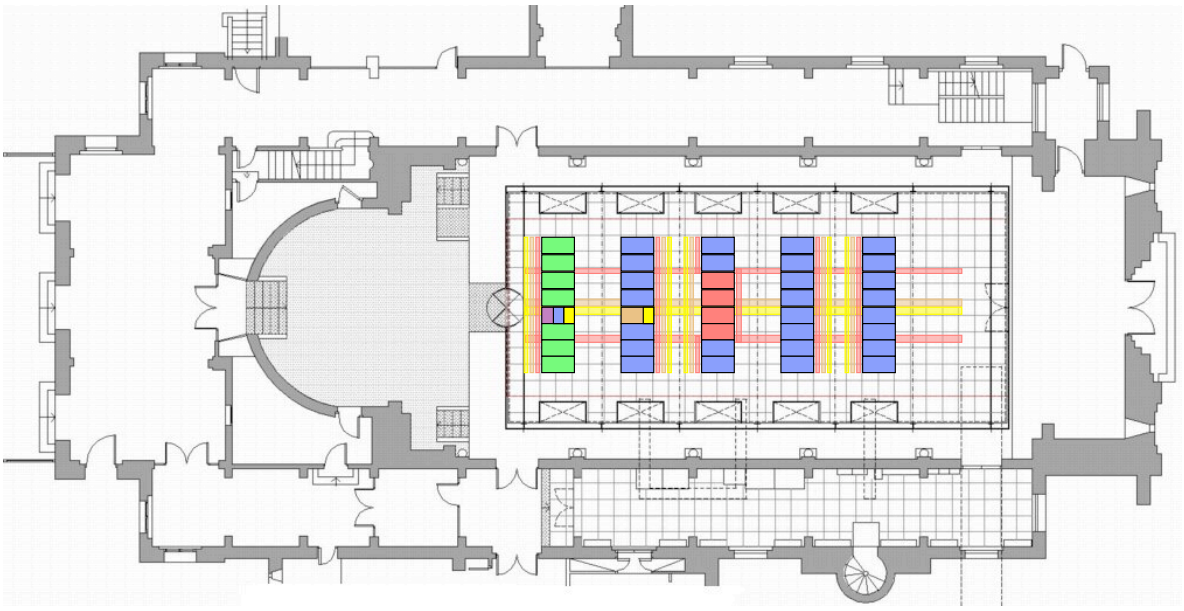

Barcelona Supercomputer Center
National Supercomputing Facility

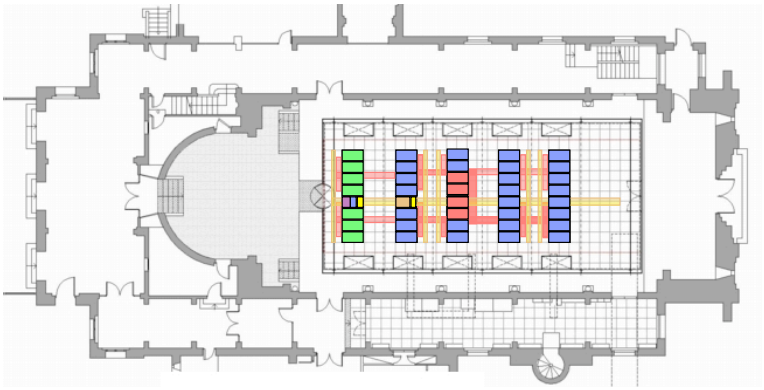

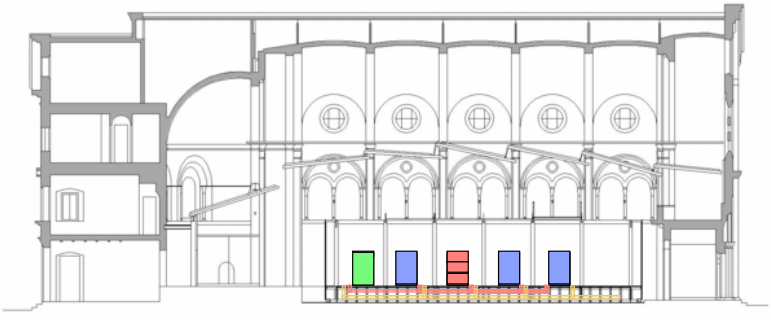

256 blades connected to
1 clos 256+256 Myrinet

1280 blades connected to
5 clos 256+256 Myrinet
and 1 spine 1280

2282 blades connected to
10 clos 256+256 Myrinet
and 2 spine 1280

20 x 7TB storage nodes

Management rack,
Force 10 Gigabit,
10/100 Cisco switches

Blade center racks

- 6 Blade Centers per rack
- 27 racks + 1 Blade Center
- Cabling per rack
  - 84 fiber cables to myrinet switch
  - 6 Gb to Force10 E600
  - 6 10/100 cat5 to Cisco

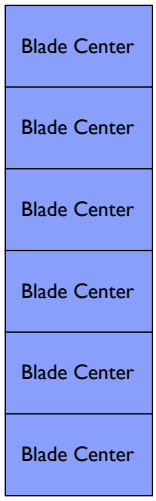

Barcelona Supercomputer Center
National Supercomputing Facility

Myrinet racks

- 10 Clos 256x256 switches
  - Interconnect up to 256 Blades
  - Connect to Spine (64 ports)
- 2 Spine 1280
  - Interconnect up to 10 Clos 256x256 switches
- Monitoring using 10/100 connection

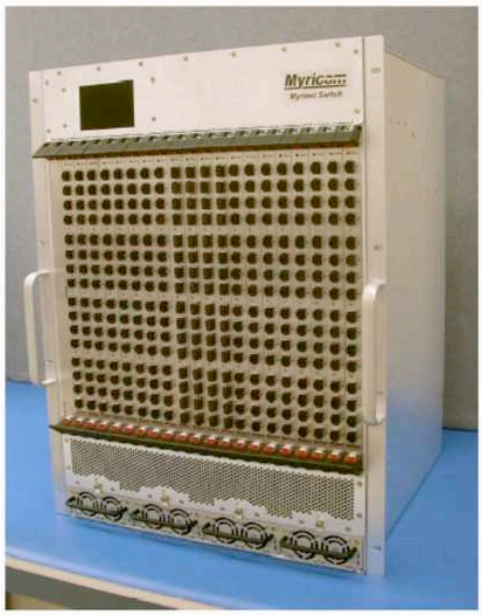

Barcelona Supercomputer Center
National Supercomputing Facility

Myrinet racks

Clos 256x256	Spine 1280	Spine 1280	Clos 256x256
Clos 256x256	Clos 256x256	Clos 256x256	Clos 256x256
Clos 256x256	Clos 256x256	Clos 256x256	Clos 256x256

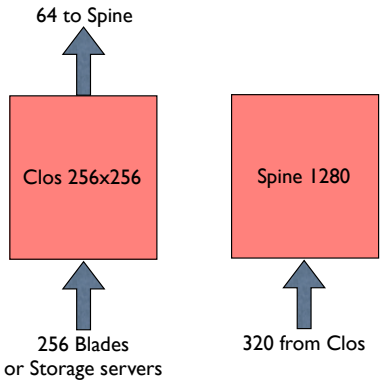

Barcelona Supercomputer Center
National Supercomputing Facility

Gb Subsystem: Force 10 E600

- Interconnection of Blade Centers
- Used for system boot of every blade center
- 212 internal network cables
  - 170 for blades
  - 42 for p615
- 76 connection available to external connect

Storage nodes

- Total of 20 storage nodes, 20 x 7 TBytes
- Each storage node
  - 2xP615
  - FastT100
  - EXP100
- Cabling per node
  - 2 Myrinet
  - 2 Gb to Force10 E600
  - 2 10/100 cat5 to Cisco
  - 1 Serial

Barcelona Supercomputer Center
National Supercomputing Facility

Management rack

- Contents
  - BCIO
  - Display
  - 2 HMC
  - 2 x p615
  - 3 x Cisco
  - 3 x 16-port Remote Async Nodes

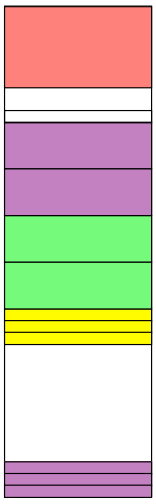

Scientific Applications

= Recent Computations by NSF Grand Challenge Research
= Next Step Projections by NSF Grand Challenge Research Teams
= Long Range Projections from Recent Applications Workshop

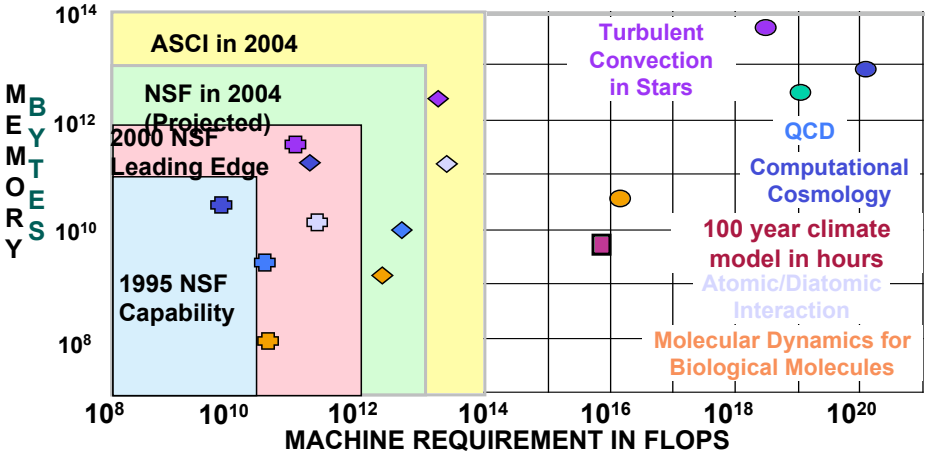

Source: Bob Voigt , NSF

Barcelona Supercomputer Center
National Supercomputing Facility

Top500 - Performance

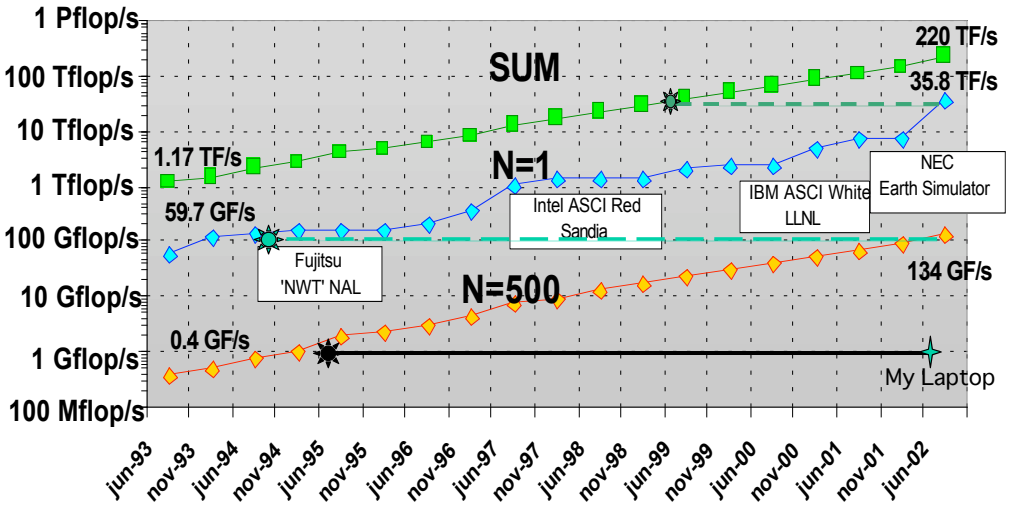

Jornadas Técnicas RedIRIS 2004
Toledo, 29 de Octubre de 2004

Top 500 list, November 2004 (upcoming)

Rank	Site Country/year	Computer/processors Manufacturer	Rmax Rpeak
1	NASA Ames United States/2004	Project Columbia 16x504proc SGI Altix 3000 1.5GHz Itanium2 w/ Infiniband / 8064 SGI	42707 48384
2	IBM Thomas V United States	ork / 16384	36010 45875
3	Earth Simulat Japan/2002		35860 40960
4	BSC Spain/2004	et/ 3564	20530 31363
5	Lawrence Live United States	cs / 4096	19940 22938

MareNostrum: number comparison

	Project Columbia	BlueGene /L	Earth Simulator	MareNostrum
Number of racks		>64	640	40
Number of processors	10240	128k	5120	4500
System peak performance	65 TF	360 TF	41 TF	40 TF
Total system Memory		16 TB	10 TB	9 TB
Power		< 1.6 MW	5.1 MW	600 KW
Floor space		400 m^2	3250 m^2	120 m^2
Bisectional bandwidth		0.7 TB/s	7.8 TB/s	0.6 TB/s
Number of cables		5000	83200	2600

Barcelona Supercomputer Center
National Supercomputing Facility

Thank you !

Prof. Mateo Valero
BSC Director

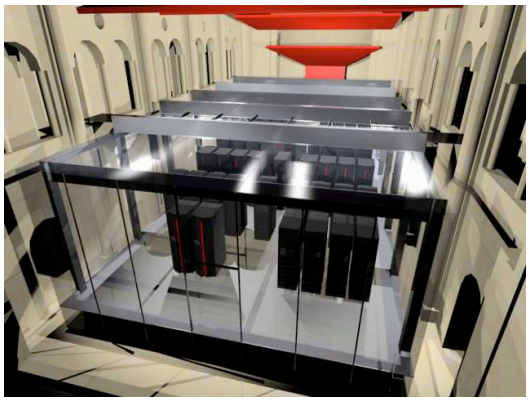

Barcelona Supercomputer Center
National Supercomputing Facility