

MareNostrum



MareNostrum was a supercomputer of the BSC-CNS based on PowerPC processors, the architecture BladeCenter, a Linux system and a Myrinet interconnection. These four technologies configure the base of the architecture and design.

See below a summary of the system:

- Peak Performance of 42.3 TFLOPS
- 4812 IBM PowerPC 970FX processors at 2.2 GHz (2406 JS20 blades)
- 9.6 TB of main memory
- 236 TB of disk storage
- Interconnection networks:
 - Myrinet and Gigabit Ethernet
- Linux: SuSe Distribution
- 650 KW of power consumption

Myrinet racks

The 2406 blade nodes JS20 are interconnected through a high speed interconnection network called Myrinet. The different nodes are interconnected via fibre optic cables.

Four in MareNostrum are dedicated to network elements which allow to interconnect the different nodes connected to the Myrinet network. These four racks are located in the center of the room and each node has a fibre optic cable. The network elements connect the different cables allowing the interconnection from one point to another from the different nodes.

The total Myrinet network is connected through:

- 10 elements Clos256+256
- 2 elements Spine 1280

Storage servers

Further to the local disk of each node with a 36GB capacity, MareNostrum has 20 storage servers arranged in 7 racks. These have a total of 560 disks of 512GB and each one provide a total capacity of 280 TB external storage. These disks are working with GPFS (Global Parallel File System) which offers a global vision of the file system and also allows a parallel access.

The 2406 nodes access the disks through the Gigabit network.

Each one of the 20 storage nodes has two nodes p615 in charge of the disk requests, a controller type FAStT100 and one unit EXP100. Each node p615 Gbit connections and Ethernet 10/100 connection.

Operations rack

One of the racks is the operation rack where the system can be managed. This rack is located in the machine console.

The content of each rack is the following one:

- 1 Monitor 7316-TF3
- 2 nodes p615
- 2 HMC (consoles) 7315-CR2
- 3 remote nodes asynchronous
- 1 Chasis BCIO BladeCenter IO
- 4 Switches Cisco 3550

Gigabit switches

One of the racks of MareNostrum is dedicated to the interconnection of the Gigabit network and one part of the interconnection elements of the Ethernet 10/100 network.

The elements are the following ones:

- **1 switch Force10 E600 Gigabit Ethernet.** It contains 7 slots in total, 6 of them have cards 10/100/1000 Base-T of 48 ports offering a total of 288 ports GigabitEthernet/IEEE 802.3
- **4 Switches Cisco 3550 48-port Fast Ethernet**

Source URL (retrieved on 6 febr 2023 - 07:57): <https://www.bsc.es/ca/mareostrum/mareostrum/mn1>