http://ifisc.uib-csic.es/nuredduna

NUREDDUNA GRID
IFISC's Contribution to e-Science

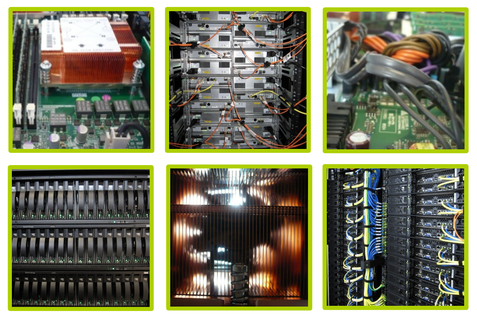

Nuredduna, the main tool for High Throughput Computing at IFISC, is composed of two parts: a cluster designed and build at IFISC which with 200 computational cores, and a grid cluster with the Grid-CSIC initiative to promote e-science. The grid cluster includes an IBM iDataplex with 512 computational cores and 96TB of raw storage.

THE GRID

- The term grid computing originated in the early 1990s as a metaphor for making computer power as easy to access as an electric power grid.
- It is the technology allowing the coordinated use of computational resources in a decentralized way.
- Security is achieved by using digital certificates.

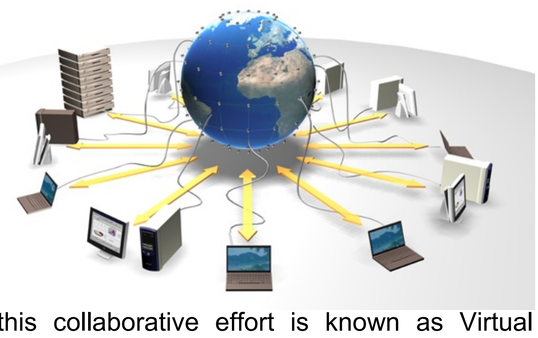

- Bringing together many different groups in this collaborative effort is known as Virtual Organizations (VOs). These VOs are flexible, can remain for a long period or can disappear just as quickly as they were created.
- Grids are used for solving scientific, technical or business problems that require a great number of computer processing cycles.

GRID-CSIC

- The Spanish National Research Council (CSIC) is deploying a distributed computing infrastructure to facilitate CSIC researchers the accomplishment of scientific projects requiring computing resources beyond the capabilities of a single user or research group. The Grid-CSIC infrastructure intents to foster multidisciplinarity and joint projects between CSIC centers.
- IFISC became a member of Grid-CSIC in early 2010. The Nuredduna grid cluster is a resource shared among the collaborating CSIC centers.
- Researchers from CSIC will have access to over 8000 cores and a total storage of more than 1PB (1000TB). It is built in a completely interoperable way, which allows the resources to be shared with similar initiatives in Europe and World-Wide
- Sites:
  - Cantabria - Instituto de Física de Cantabria
  - Valencia - Instituto de Física Corpuscular
  - Granada - Instituto de Astrofísica de Andalucía
  - Illes Balears - Institut de Física Interdisciplinar I Sistemes Complexos
  - Barcelona – CIN2 and ICMAB
  - Madrid - Centro de Física "Miguel Antonio Catalán"

- Grid-CSIC participates on the Spanish National Grid Initiative (NGI) that is part of the European Grid Initiative (EGI).

HOW GRID WORKS

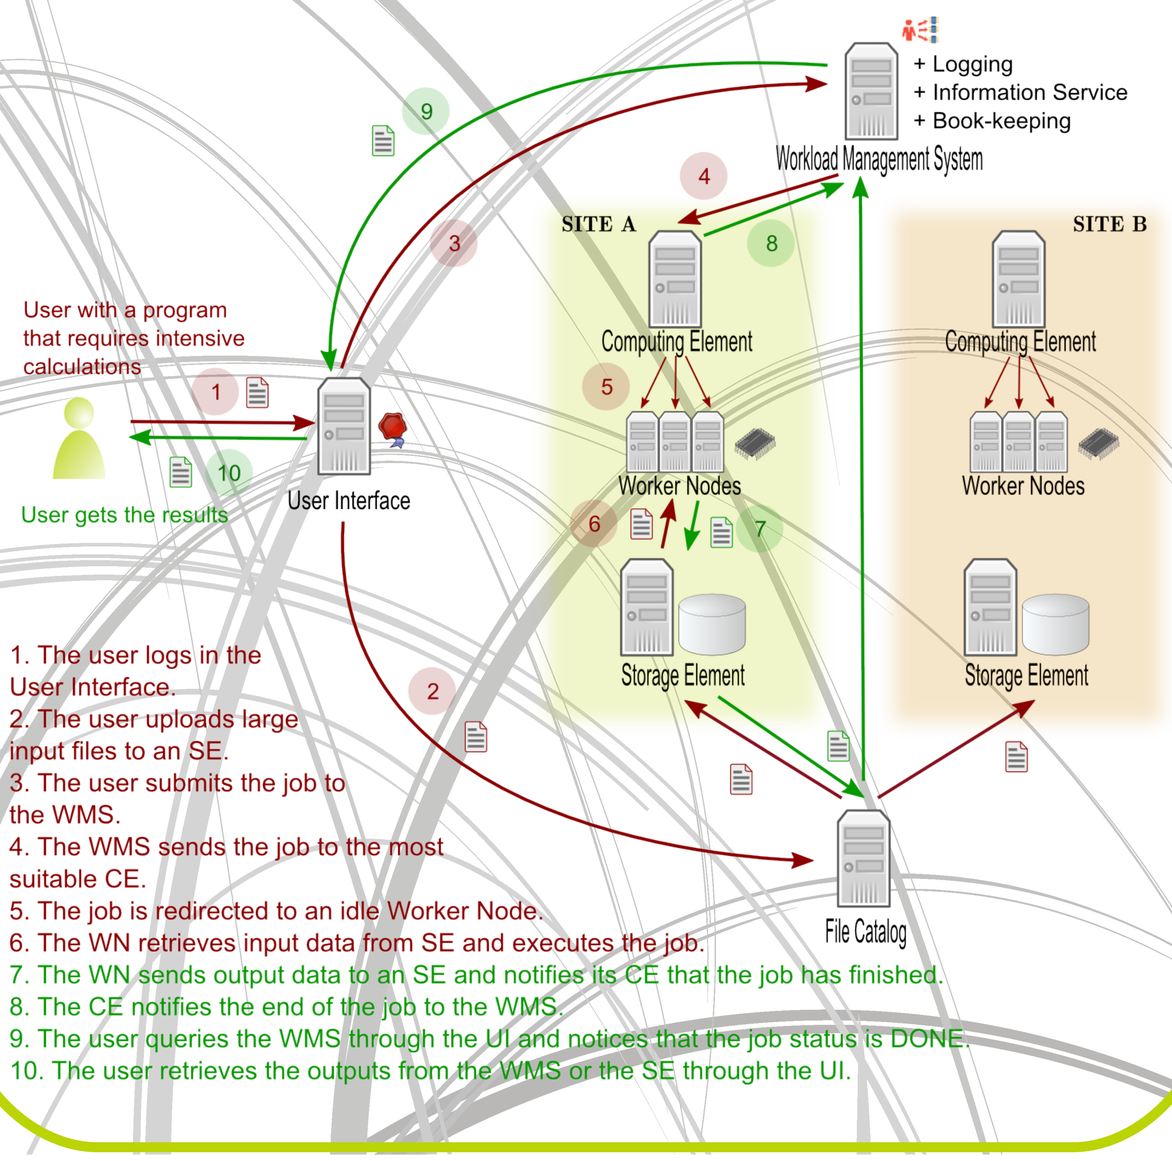

1. The user logs in the User Interface.
2. The user uploads large input files to an SE.
3. The user submits the job to the WMS.
4. The WMS sends the job to the most suitable CE.
5. The job is redirected to an idle Worker Node.
6. The WN retrieves input data from SE and executes the job.
7. The WN sends output data to an SE and notifies its CE that the job has finished.
8. The CE notifies the end of the job to the WMS.
9. The user queries the WMS through the UI and notices that the job status is DONE.
10. The user retrieves the outputs from the WMS or the SE through the UI.

HARDWARE CONFIGURATION

• Green computing

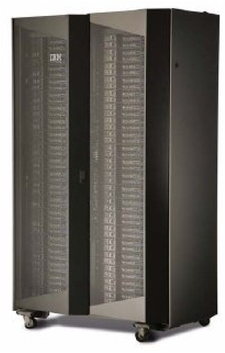

- Flexible and modular IBM iDataPlex system. Front cabling and front access keeps service and maintenance costs down.
- Reduced depth and rational wiring layout allows for efficient cooling.
- Uses less space. Each rack admits up to 84 1U dual processor nodes plus space for switches.
- Uses less power, up to 40% less energy than usual racks.
- The iDataplex has four power distribution units (PDU), each one of them can provide electricity up to 16 dual-node boxes. PDUs can be monitored externally and balance efficiently the load between the three electrical phases.

• Nodes (computational units)

- 64 top rated for both performance and energy efficiency dx360M2 nodes.
- Each pair of nodes share the power supply.

- Each node integrates two low consumption Intel Xeon L5520 quad-core processors running at 2.26MHz and it has 16GB of DDR-3 memory at 1066MHz.
- Each node incorporates a 15000 rpm SAS 160GB hard drive for the Operating System and the swap space.

• Storage

- The total amount of RAM for the 512 cores is 1024GB.

- IBM iDataPlex computational cluster is completed with four storage servers. Each one integrates two L5520 quad-core processors running at 2.26MHz, has 12GB of DDR-3 memory at 1066MHz and an integrated Qlogic Fibre Channel HBA Card.

- A total amount of 96TB is distributed between two high performance IBM DS4700 and four DS4000 EXP810 storage expansion units.
- RAID5 is used to get storage reliability with optimal input/output performance, consequently, IFISC has 54TB of redundant storage.

• Connectivity

- Computational units and storage nodes are connected through two dedicated Gigabit Ethernet lans, one for management and the other for communications between nodes.
- Discs are connected to storage servers through 4Gb multimode Fibre Channel cables. Redundant disk controllers and redundant fibre channel connections make IFISC grid data a high reliable system.

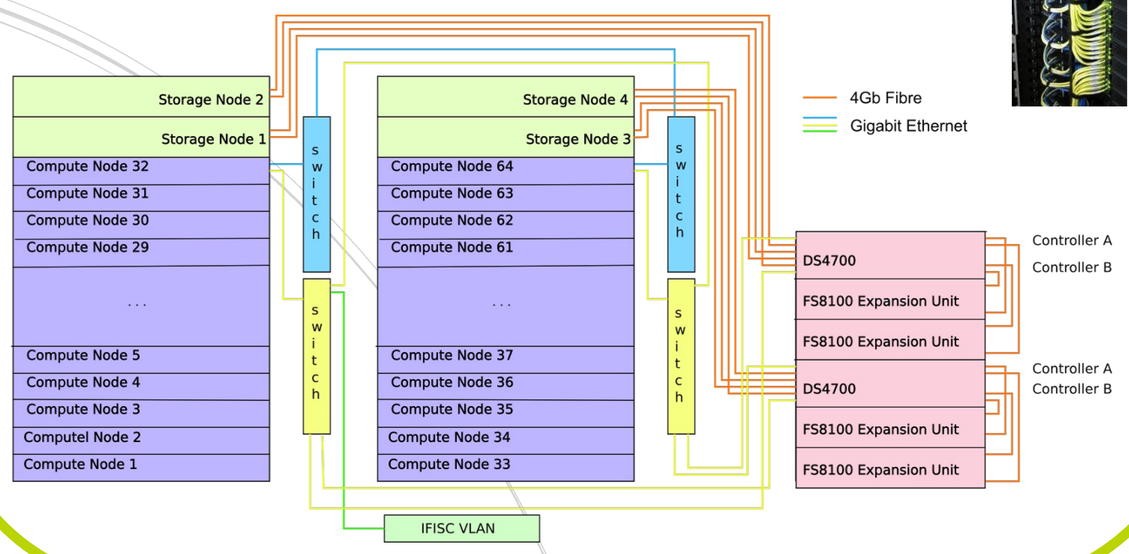

SOFTWARE CONFIGURATION

- The Operating System for nuredduna grid cluster is Scientific Linux 5.4 for 64 bit, which is commonly used in european grids.
- The Batch Queue System used is Sun Grid Engine 6.2, an open source system developed and supported by Oracle.
- A high-performance shared-disk clustered file system developed by IBM, GPFS, is the filesystems integrated in nuredduna.
- gLite 3.2 is the grid middleware running on nuredduna. The node types installed at IFISC are: User Interface, Site-BDII, Storage Element, Computing Element and Worker Nodes

In order to simplify users access to grid, IFISC has designed and developed a set of suitable scripts that are available at the User Interface.

CSIC Consejo Superior de Investigaciones Científicas

Universitat de les Illes Balears