

Computing

computing

Different architectures of high-performance computing systems are available for CESGA's community of users. Depending upon the algorithms involved, users choose the computing architecture that is most appropriate for their needs.

The highlights of 2008 regarding computing servers were:


- The installation and the initiation of operation of Finis Terrae,
- The integration of the HP-Superdome servers into the Finis Terrae cluster,
- The HPC-320 server was removed from service, and
- The users employed more than 10 million CPU hours.

In 2008, CESGA centered its efforts on putting Finis Terrae into operation and supporting it. This system ranked number 100 on the TOP500 list of November 2007 and began operation in February 2008 in order to begin to work on relevant computational challenges.

Finis Terrae was made available to all users on April 1, 2008. In 2008, the following computing servers were available to users:



COMPUTING SERVERS IN 2008			
HIGH PERFORMANCE COMPUTING SERVERS			
SERVERS	YEAR INSTALLED	ARCHITECTURE	PROCESSORS, MEMORY, PEAK PERFORMANCE
COMPAQ HPC320	2002	SMP CLUSTER	32 CPU's, 80 GB MEMORY, 64 GFLOPS
FINIS TERRAE	2008	SMP (NUMA) CLUSTER	2,580 CORES, 20 TB, 16 TFLOPS
HIGH THROUGHPUT COMPUTING SERVERS INTEGRATED IN THE GALICIAN VIRTUAL SUPERCOMPUTER (SVG)			
SERVERS	YEAR INSTALLED	ARCHITECTURE	PROCESSORS, MEMORY, PEAK PERFORMANCE
SVG	2001-2006	DISTRIBUTED PC CLUSTER	50 CPU's, 0,5 -1 GB MEMORY CPU, 9,9 GFLOPS , 110 CPU, 300 GFLOPS (2004)
COMPAQ BEOWULF	2002	BEOWULF CLUSTER	16 CPU, 8 GB MEMORY, 16 GFLOPS
SVG DELL	2004	PC CLUSTER	80 CPU, 80 GB MEMORY, 512 GFLOPS
SVG BLADES	2006	BLADE CLUSTER	292 CORES, 148 GB MEMORY, 2,227 GFLOPS
SERVERS HOUSED AT CESGA			
SERVERS	YEAR INSTALLED	ARCHITECTURE	PROCESSORS, MEMORY, PEAK PERFORMANCE
METEOGALICIA	2004	2 SMP NODES	20 CPU, 40 GB MEMORY, 48 GFLOPS
LHCb-USC	2002-2008	PC CLUSTER	339 CORES, 312 GB MEMORY, 1600 GFLOPS




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



Evolution of CPU Consumption


During 2008, the new Finis Terrae system and the SVG cluster were the most utilised servers. In order to facilitate the transition to Finis Terrae for the users, one of the HP Superdome systems was maintained in operation for the duration of the year. The number of hours consumed increased significantly, multiplying the total by 5, going from 2 million hours (2,062,610 actual hours) for 2007 to a total of 10,067,403 hours for 2008.



Number of Jobs Executed

The number of jobs executed represents the quantity of simulations that the users performed on each one of the computing servers. This value does not only depend on the available computing capacity but also on the resources necessary for the execution of the simulations. The SVG cluster and Finis Terrae were the most utilised systems. The SVG was especially important given that it is a server that is oriented to productivity, that is, the performance of a great number of independent projects with relatively low computing needs. In summary, the 2008 total number of jobs of 1,214,079 was much higher than that of 2007 (215,883).


JOBS EXECUTED PER SYSTEM



Average in-queue time


The average in-queue time represent the average amount of time that users must wait from the moment that they request CESGA resources until their simulations begin running on the systems. This time varies as a function of the quantity of simulations that are being performed and desirably it should be the closest possible to zero in order to avoid that users have to wait much time until they receive the results of their simulations.

Logically, the higher the level of occupation of the computing systems, the more time it is necessary to wait until the necessary resources are available. This average is a good indicator of the level of saturation that exists in the computational resources. Normally, the waiting times decrease in summer and during the Christmas holidays, and they increase significantly as new users are incorporated. In summary, the waiting times have decreased in great measure thanks to the incorporation of the Finis Terrae server, reducing the waiting time from an average of more than 7 hours



in 2007 to an average of only 2 hours for all of the computing servers in 2008. Even more important is the decrease in waiting time in the server with the greatest capacity. In 2007, the average for the HP Superdome was 43 hours and it decreased to less than 2 hours for Finis Terrae. This allows researchers to run more simulations and to receive their results in less time, thereby increasing their competitiveness.

**AVERAGE IN QUEUE TIME FOR ALL PROCESSES IN ALL SYSTEMS
(JANUARY 2001-DECEMBER 2008)**



HPC: high performance computing servers

Different server architectures for High Performance Computing (HPC) are available for CESGA users. These involve machines designed to solve a reduced number of problems, but of large dimensions, in a limited time. These architectures incorporate scalar high performance processors with access to large memory size, utilizing internal networks with a low latency time and a high transfer capacity.

During 2008, the users had access to three different high performance systems.

Finis Terrae

The supercomputer FINIS TERRAE was installed at CESGA in December, 2007. It is an integrated system with shared memory nodes, with SMP NUMA architecture. A list of the components of FINIS TERRAE follows.

A total of 144 computing nodes:

- 142 HP Integrity rx7640 nodes with 16 Itanium Montvale cores and 128 GB of memory each,
- 1 HP Integrity Superdome node with 128 Itanium Montvale cores and 1,024 GB of memory, and
- 1 HP Integrity Superdome node with 128 Itanium 2 cores and 384 GB of memory.

A hierarchical storage system with:

- 22 nodes for storage management with a total of 96 processing cores,
- 390,000 GB on disk, and
- 2,200,000 GB on tape (cartridge robot).

An interconnect Infiniband 4x DDR at 20 Gbps.


An external network at 10 Gbps.

The system concurrently supports multiple operating systems as demanded by the research community, such as, Unix, Linux, and Windows.

FINIS TERRAE includes open software, such as, Linux, Lustre, Grid Engine, and Globus.

The system has the following compilers, libraries, and development tools: Intel C/C and Fortran, Intel MKL, Vtune, HP-MPI, and HP UPC.

FINISTERRAE



FINISTERRAE CLUSTER SMP NUMA TECHNICAL SPECIFICATIONS	
COMPUTER	Integrity rx7640/Superdome Itanium 2 Cluster
APPLICATION AREAS	Computational Science Applied to: Nanotechnology, Health & Life Sciences, Ocean Sciences, Energy, HPC
MULTIPLE CONCURRENT O.S.	Unix, Linux, Windows
OPEN SOFTWARE	Linux, Lustre, Globus etc...
COMPILERS, LIBRARIES & DEVELOPMENT TOOLS	Intel C/C++ and Fortran, Intel MKL, Vtune, HP-MPI and HP UPC
NODE INTERCONNECT	Infiniband 4x DDR at 20 Gbps
EXTERNAL NETWORK CONNECTION	10 Gbps
PROCESSOR	Intel IA-64 Itanium 2 Montvale Dual Core 1600 MHz (6.4 Gflops)
COMPUTING NODES	- 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
PROCESSING CORES	2,580
MEMORY	19,670 GB
STORAGE NODES	22 nodes with 96 cores
HIERARCHICAL STORAGE	390,000 GB on disk & 2,200,000 GB on tape



HP Integrity Superdome

This SMP NUMA Cluster is formed by two nodes with a total of 128 Intel Itanium2 processors, 1500 MHz, 6 Mbytes cache, 384 GB of memory, and 4.6 Terabytes for temporary storage on disk. An Infiniband network between the two nodes is available. For permanent data, the user had an HP EVA-3000 disk rack for massive storage with FibreChannel technology. During 2008, this server was integrated into Finis Terrae which means that its operating system is Linux SLES10SP1 and can be accessed using the same batch system.

COMPAQ HPC 320

This cluster was made up of 8 SMP machines with 4 Alpha EV68 processors of 1 GHz each and 80 GB of total memory. The nodes of this cluster were interconnected by means of a Memory Channel Dual Rail network. After 5 years of extraordinary service to CESGA's community of researchers, COMPAQ HPC 320 reached the end of its technological life and was finally retired from service in June, 2008.

HTC: high throughput computing & grid servers

CESGA offers server architectures designed to solve a large number of problems of reduced computational complexity in a short period of time. The High Throughput Computing servers incorporate a large quantity of scalar processors with fast access to a local memory of limited size. These servers have an internal interconnect network with medium-high latency.

This type of architecture is ideal for iterative processes with low dependency between themselves, such as, genetic algorithms, massive processing of experimental data, the rendering of images, parametric computing, etc.



Supercomputer (SVG). In 2007, 32 processors and 32 GB of memory were added.

SVG reflects the wager that the Centre made back in 1999 for low cost cluster systems as an ideal solution to the increase in throughput jobs.

HTC systems, as well as other clusters, are increasingly more abundant in the laboratories and departments of research groups. These groups use the services of CESGA as a complement to their own, without the necessity to port their applications or to learn new operating systems.

In 2006, CESGA acquired a DELL Blade Cluster with 292 processing cores, 148 GB of memory, and a peak performance of 2,227 GFLOPS. This cluster was integrated into the Galician Virtual

Due to their characteristics, the HTC servers installed at CESGA are used in GRID experiences and projects.

SIGNIFICANT COMPUTING INICIATIVES IN WHICH CESGA PARTICIPATES

housing of computing equipment




SUN NODES
METEOGALICIA

Meteogalicia

In CESGA's Data Centre, there is a SUN Microsystems computing server composed of two SMP nodes with a total of 20 CPUs, 40 GB of memory, and 48 GFLOPS of peak performance that was acquired by Meteogalicia in 2004. Since that date, the Centre has taken care of the housing and administration of this equipment which is used in the development of meteorological research projects.

Since its creation, Meteogalicia has used CESGA's servers in order to produce daily meteorological forecasts. In addition to this SUN machine, Meteogalicia regularly used the following servers in 2008: Finis Terrae, HP Integrity Superdome, HPC 320, and the Galician Virtual Supercomputer (SVG).

LHCb-USC CLUSTER



LHCb-USC

Since 2002, the Experimental Group of High Energy Physics of the University of Santiago de Compostela (GAES) has had a PC Cluster housed in CESGA's Data Centre. It participates as a Tier-2 Centre in the LHCb experiment coordinated by CERN. CESGA systems technicians have been responding to the necessities of equipment administration. This equipment is available for all users when it is not employed by the project. Since it was first housed at CESGA, this cluster has steadily grown with the addition of new computing cores.



XUNTA DE GALICIA
CONSELLERÍA DE MEDIO AMBIENTE
E DESENVOLVEMENTO SOSTIBLE
Dirección Xeral de
Desenvolvemento Sostible

METEOGALICIA



Worldwide LHC Computing Grid
Distributed Production Environment for Physics data Processing