EUROFEL-Report-2005-DS2-006

EUROPEAN FEL Design Study

Deliverable N°:	D2.1
Deliverable Title:	A Parallel-Computing Resource for the EC
Task:	DS-2
Author:	Luca Gianessi, ENEA
Contract N°:	011935

Project funded by the European Community
under the "Structuring the European Research Area" Specific Programme
Research Infrastructures action

Description

The budget available (~40k€) suggested to choice an architecture based on **Intel/AMD** processors and **Linux OS**. An agreement was achieved with **ENEA's Computer Division**, which expressed availability to **host the computer** in Frascati's Research Center, **provide support** and - if possible - to increase the cluster size, provided some requirements (basically space occupancy, ease of maintenance and a reasonable cost of ownership) were met.

For both the above reasons the possibility of assembling a system with Computer Off The Shelf (COTS) hardware has been abandoned, and only "blade" of "rack-mounted" systems have been considered thereafter. From the architectural point of view, during 2005 the obvious choice based on Intel has been Xeon CPUs with *HyperThreading* technology (HT). On the AMD side the *dual-core* (essentially the embedding of two – or more - processors on the same chip) seemed very promising.

A benchmark resembling the weighted calculation expected from a TREDI simulation has been prepared with the intention to test a number of different computing platforms.

The benchmark has been extensively tested on several available systems (some of them available at ENEA, some made available temporarily by vendors) ***dual-core*** **proved by far more cost/performance effective than** ***HyperThreading*** **technology**, the former being able to execute to concurrent processes with virtually no performance penalties and the latter yielding a mere 25% speed up - for the system booted with HT on - with respect to the same system run with HT disabled. Moreover, the Intel/Xeon HT platform seemed generally more expensive in absolute terms.

A careful analysis of the *pro* and *con* suggested that rack-mounted systems were more appropriate - and cheaper than blade based ones, more suited to bigger parallel configurations for the higher density achievable (essentially due to superior heat dissipation capabilities).

The final choice has been a rack-mounted (APC rack/SuperMicro MB) system, composed of:

- 1X Dual Core/Dual Opteron265 (4 cores) with
  2GB Ram, 1x160GB+4X300GB HD (FrontEnd node);

- 12x Dual Core/Dual Opteron265 (48 cores) with
  2GB Ram, 1X160GB HD each.

The order has been issued in July 05, but the system has been delivered at the end of Oct 05 (one month later that expected). The short delay was justified by a mistake of the vendor in shipping the right cabinets.

The cluster was assembled/tested at the beginning of November and become operational by mid December.

A few minor details as the queuing system and IPMI functionality remain to be completed but the system is already actively being used for simulation work and is available for the qualified users.

The installed software includes:

Operating system: Scientific Linux X86_64 4.1

- CERNLIB 64 bit
- HDF 5 1.6.5
- TREDI
- LAM MPI 7.1.1
- MPICH 1.2.7
- MPICH2 1.0.3
- CERN ROOT 5.06.00

A picture of the cluster is shown in fig. 1

Fig. 1 The 53 nodes parallel cluster, deliverable 2.1