INFN PISA

SCIENTIFIC COMPUTATION ENVIRONMENT (GRID HPC AND INTERACTIVE ANALYSIS)

E.MAZZONI A.CARBONI S.AREZZINI G.CARUSO A.CIAMPA S.COSCETTI S.PIRAS

INFN-PISA 2013 resources summary:

total number of cores: ~6700
total disk space: ~1,3 PB
WAN bandwith: 2 x 10 GE

All the computing nodes are grouped in 15 farms and 3 clusters, each one housing identical or similar hardware. Access to the facility supports both GRID and local submission to resources. Clusters are equipped with high speed interconnections for parallel computing.

~3750 cores Grid Computing:
all INFN VO are supported, with CMS and THEOPYS as stockholders. The centre is a WLCG LHC-T2 for CMS Experiment.

~240 cores THEONUC farm:
aims at the computing requirements of the INFN Theoretical Physics Community. High memory is a key feature, total available RAM is more than 800 GB.

~130 cores FAI (Facility for Interactive Analysis): scheduler managed, scientifical software provider.

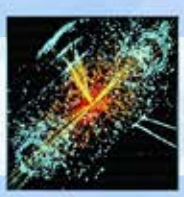

Clusters

~230 core Diablo cluster with Myrinet interfaces. This cluster is born in collaboration with the Engineering Department of Pisa University and is mainly used for Fluent, Starcd and Foam computations.

~1000 core Tramontana cluster with Infiniband DDR 20 Gb/s interfaces. The cluster was intended as centralized service for serial and parallel computation for the national theoretical community.

~1600 core Theocluster cluster with Infiniband QDR 40 Gb/s interfaces. Each node has 4x16 cpu cores and 512 GB of memory in order to fulfill both cpu/memory bound computations.

Advanced coprocessors GPU & Intel MIC

Seven servers hosts a total number of 36 Nvidia GPU:

- 24 Tesla T10
- 4 Tesla M2090
- 8 k20 GPU

The k20 GPUs are currently tested in cluster configuration using Infiniband QDR 40 Gb/s interconnections.

Two Intel Xeon Phi corprocessors

Data Storage and Filesystem

1,3 PB of available disk space are housed in 5 DDN systems: 2 S2A9900 , 1 SFA12K and 2 Metadata server.
The filesystem is GPFS, DDN systems are accessed by 10 GPFS servers (2 of them are used exclusively for metadata access): all the link are 8 Gb/s fibre channel and grant a total bandwidth of 160 Gb/s toward the worker nodes. Many access protocol supported: GRID-srm, Posix, xRootD.

Access to resources and job submission

Access to all the computing resources of the centre is granted by a set of user interfaces to all INFN Theoretical Physics Community and CMS staff and is based upon INFN-AAI access credentials. Both GRID and local submission are possible. Job submission is based over batch queues, the scheduler is LSF, both version 7 and 9 are currently used.

INFN-AAI: project that aims to create a unique Authentication and Authorization Infrastructure for INFN, with these unique credentials users can access the facility avoiding complex registration procedures.

Net Architecture

The centre is geographically connected with a 2x10 Gb/s net. Internal flat architecture is based on a Force 10 E1200 exascale switch housing 8x90 1 Gb/s ports and 2x10 10 Gb/s ports. 1 Gb/s ports are used to connect worker nodes while 10 Gb/s are for GPFS nsd servers and StoRM gridftp servers

SUMA Project and Collaborations

SUMA (SUperComputer Massive) is a major project funded by the Italian Ministry of education.
Theocluster is part of SUMA, in synergy with a new cluster that will be shortly installed at the SISSA (International School for Advanced Studies) computing center, in Trieste. The latter will have 5000 computing cores with 2 GB RAM each, and a peak performance of at least 100 TFlops. We are planning to make the queues of the two installations mutually visible via federated accesses. In this way INFN users can submit jobs to the SISSA cluster while SISSA users can submit jobs to Theocluster.

The work has been partially funded under contract 20108T4XTM of 'Programmi di Ricerca Scientifica di Rilevante Interesse Nazionale (Italy)'