



# CEA Computing Complex

## TERA and CCRT

### 100 Tflops of Itanium Montecito

*CEA/DAM/DIF/DSSI*

# CEA Computing Center

---

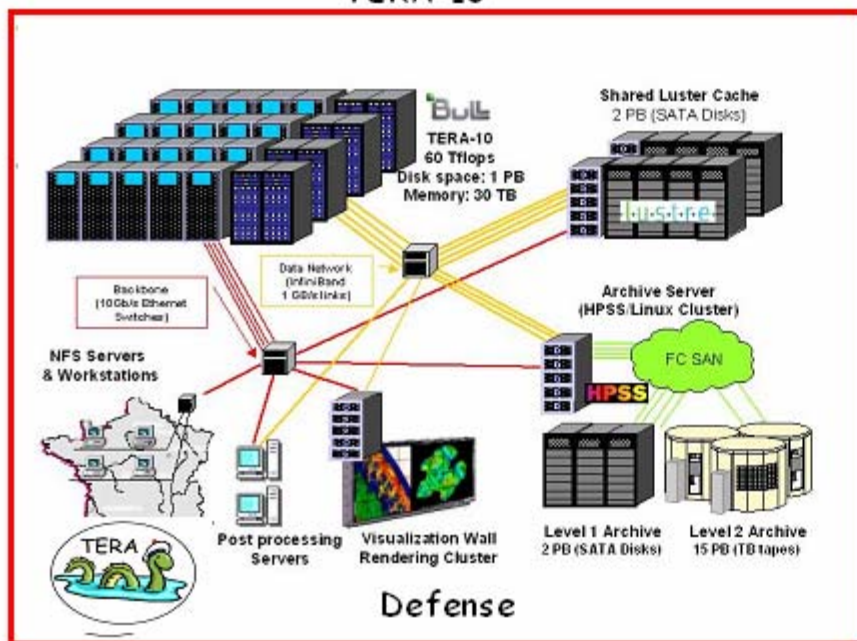


**CEA DIF**  
(South of Paris)

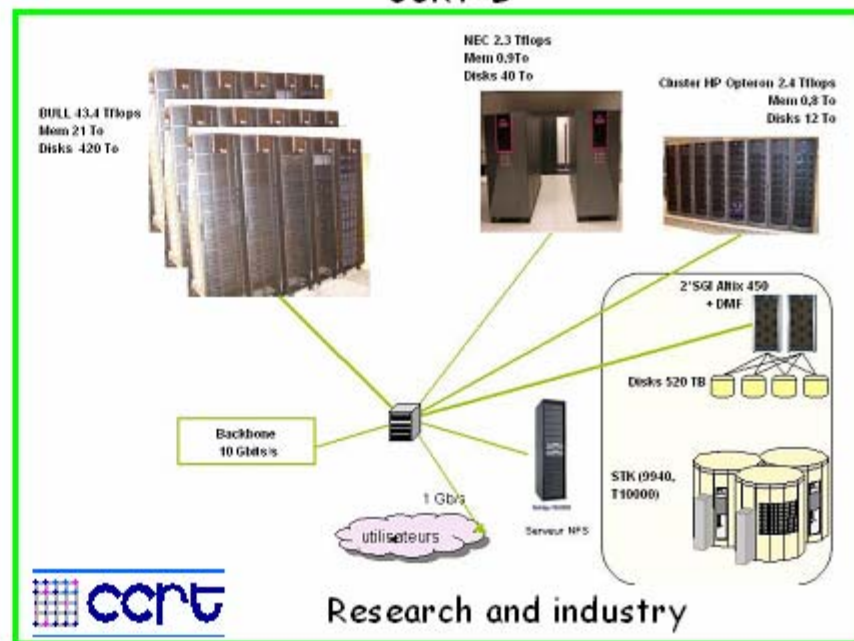


# Current supercomputers at CEA/DIF

TERA-10



CCRT-B





TERA-10 system

Main characteristics

Installation December 2005

Compute nodes 565\*16 proc.

Peak performance (sustained) > 60 Tflops (12,5 Tflops)

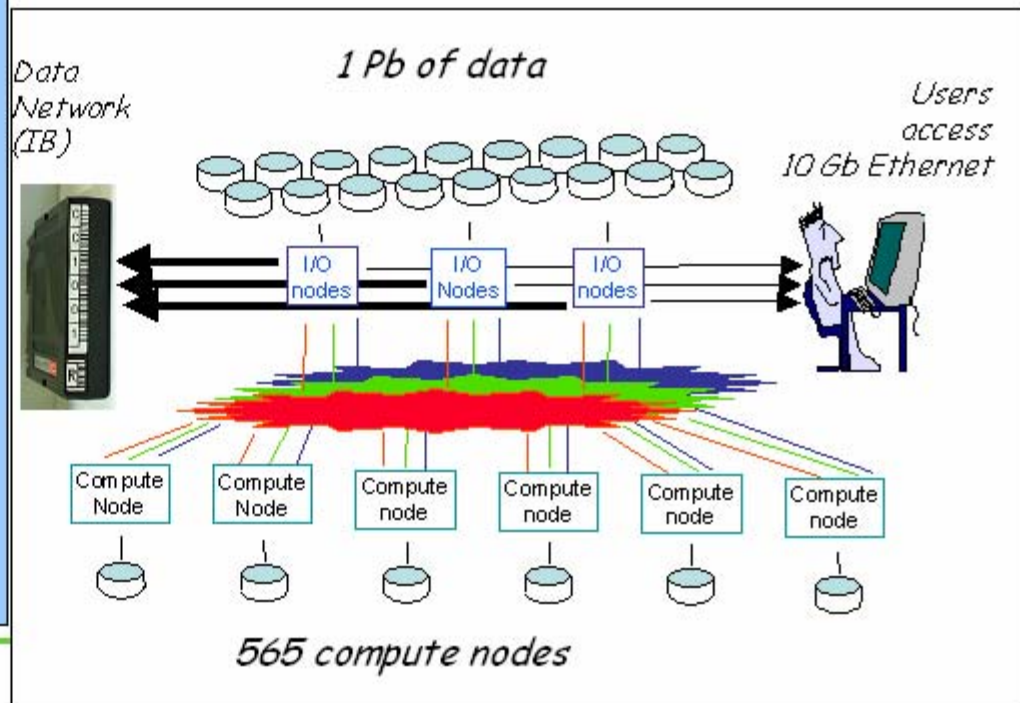
Total memory 30 Tb

Disk storage 1 Pb - 58 nodes  
100 Gb/s  
(54 OSS+2 MDS)

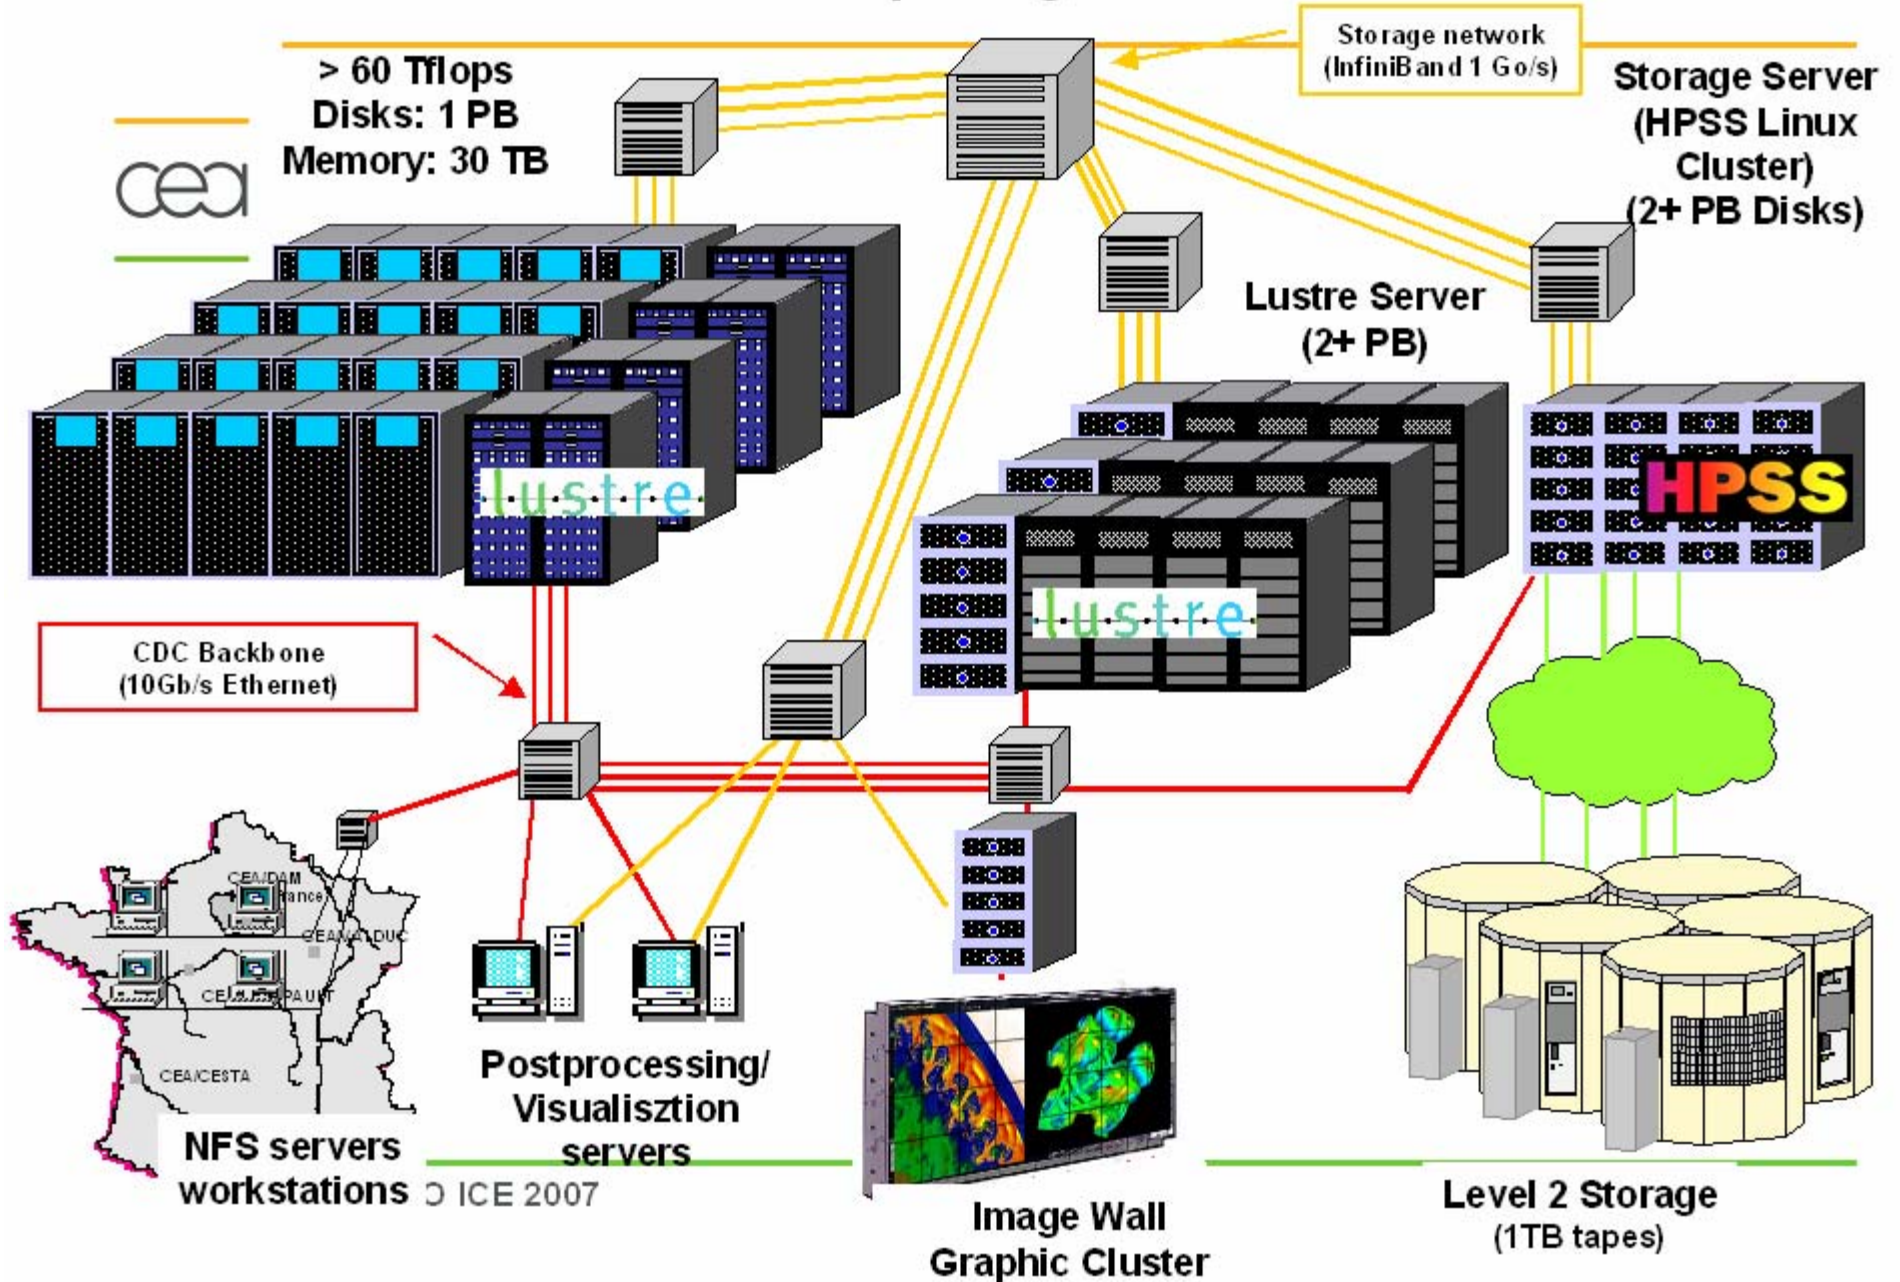
Power 1800 kW

272 racks  
60 miles of cables

GELATO ICE 2007



# TERA-10 Computing Center Architecture



## CCRT-B Compute System

---



- **Computing Power 43.4 Tflops:**
  - 840 compute nodes with 24 GB of memory
  - 8 compute nodes with 128 GB of memory
  - 26 service nodes (management, Lustre MDS and OSS)
  - MPI interconnect based on Infiniband
  - 422 TB of DDN disks (S<sup>2</sup>A9550)
  - Software stack:
    - ✉ Linux Bull Advanced Server (Operating system)
    - ✉ Intel development environment
    - ✉ LSF/Slurm
  
- **Machine ready to grow to up to 57 Tflops**

# CCRT-B Computing Center Architecture




**BULL 43.4 Tflops**  
Mem 21 To  
Disks 420 To



**NEC 2.3 Tflops**  
Mem 0.9To  
Disks 40 To



**Cluster HP Opteron 2.4 Tflops**  
Mem 0.8 To  
Disks 12 To



# Linux Itanium Ports

---



## ● System Tools

- Meta Scheduler (KMS)
  - ✉ Automatic Management of multi months simulations on multiple computer
- Parallel HPSS copy tool
- Token Server
  - ✉ Allows storage resources control (by user, file type)
- Linux Cluster Administration tools
- User Space NFSV2/3/4 server (Ganesha)
- Log Analysis tool (Casimir)

## ● Codes

- All codes are ported
- From Molecular Dynamics to Computational Fluid Dynamics



- **Experience sharing on system administration**
  - Day to day life with large montecito cluster
  - Data Sharing with Lustre
  - Cluster Administration tools
  
- **Experience sharing on porting codes**
  - Use of large SMP nodes
  - Use of large Clusters
  - I/O intensive codes
  - Code development tools
  
- **Feedback on CEA contributions**
  - Linux Kernel for large HPC cluster (tuning, debugging)
  - Lustre-HSM binding
  - Distributed Lustre