

## TERA-100 supercomputer

Designed by the French manufacturer Bull and put into production in 2011, Tera-100 is a cluster of 4370 Bull Server MESCA bullx S Series servers, with four Intel® octo-core 2.27 GHz Xeon® 7500 each. The NovaScale servers are interconnected by an InfiniBand QDR network (40 Gb/s). The system offers the processing capacity of its 17,480 processors with 300 TB of memory. The machine also encompasses 22 frames containing each nine B505 hybrid nodes: 2 processors (4 cores @ 2.4Ghz), 3GB/core and a high performance graphics card per node. This partition is suited for various new codes as well as optimized remote visualization.



← Tera-100, 1.25 petaflops (more than a million billion operations per second), first European computer to sustain a one petaflop/s performance on Linpack test.

Tera-100 exploitation requires huge storage capacity for the data produced by CEA/DAM simulation codes. Tera-100 has a local storage space of 8 PB (~ 9600 disks) with a 300 GB/s bandwidth, beside a communication bandwidth of 200 GB/s to 15 PB of global storage, provided by 42 router nodes dedicated to data transfer.

Tera-100 is operated through the Bull HPC software stack that encompasses Linux® operating system and Lustre parallel file system. This software stack is based on open source components integrated and optimized by Bull.

Some figures may help to understand such a cluster: 90 km of cables, an internal storage capacity equivalent to 240 times the volume of digitized supports from Bibliothèque Nationale de France. The machine consists of more than 200,000 elements integrated into 200 water-cooled cabinets.

Computers are increasingly powerful and their power consumption increases as well. Today electricity represents a significant part of the budget of a computing center. The electrical power required to power and cool Tera-100 is in the order of 7 MW (peak). Energy optimization at different levels (computer, cooling, usage of passive devices instead of UPS) has significantly reduced the power consumption of the facility. The computing power of Tera-100 is 20 times higher than Tera-10 for a 3-fold only increase of the power consumption. A simulation on Tera-100 thus consumes 7 times less electricity than the equivalent computation on Tera-10.

### An optimized production environment

Experience of data centres management acquired for many years has led CEA/DAM to select open source products to handle all its computing and storage systems. By improving software components such as Lustre, Slurm and developing others ones (like Lustre-HSM, clustershell, Robinhood, NFS-Ganesha, shine, ...), CEA/DAM has been able to optimize the operation of its computing centres.

TERA 100 was elected Top Supercomputing Achievement in 2011 by HPC Wire readers.